



HAL
open science

Large Deviations Theory in Statistical Physics: Some Theoretical and Numerical Aspects

Grégoire Ferré

► **To cite this version:**

Grégoire Ferré. Large Deviations Theory in Statistical Physics: Some Theoretical and Numerical Aspects. Mathematics [math]. Université Marne La Vallée, 2019. English. NNT: . tel-02441535

HAL Id: tel-02441535

<https://hal.science/tel-02441535>

Submitted on 15 Jan 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Ecole Doctorale Mathématiques et Sciences et Technologies de l'Information et de la
Communication (MSTIC)

THÈSE DE DOCTORAT

Discipline : Mathématiques Appliquées

présentée par

Grégoire FERRÉ

Large Deviations Theory in Statistical Physics: Some Theoretical and Numerical Aspects

Thèse dirigée par Gabriel STOLTZ
et préparée au CERMICS, École des Ponts ParisTech

Soutenue le 27 novembre 2019 devant le Jury composé de :

<i>Président du jury</i>	M. Thierry BODINEAU	École Polytechnique
<i>Rapporteurs</i>	M. Arnaud GUILLIN	Université Blaise Pascal
	M. Denis TALAY	INRIA Sophia-Antipolis
<i>Examineurs</i>	Mme Laure DUMAZ	Université Paris-Dauphine
	M. Adrien HARDY	Université de Lille
<i>Directeur de Thèse</i>	Gabriel STOLTZ	École des Ponts ParisTech

À mes parents

À D.

On dit que le pouce opposable est ce qui différencie le plus nettement l'homme du singe. Il faut joindre à cette propriété cette autre que nous avons, de nous diviser contre nous-mêmes, notre faculté de produire de l'antagonisme intérieur. Nous avons *l'âme opposable*. Peut-être le JE et le ME de nos expressions réfléchies sont-ils comme le pouce et l'index de je ne sais quelle main de... Psyché? Alors les mots *comprendre* ou *saisir* s'expliqueraient assez bien.

Paul Valéry

REMERCIEMENTS

Mes premiers remerciements vont bien sûr à Gabriel Stoltz, pour lequel j'aurai du mal à exprimer toute ma gratitude. S'il a encadré ma thèse pendant ces trois dernières années, je connais Gabriel depuis bien plus longtemps, car notre rencontre date peut-être de six ans ! L'expérience forte de la thèse a cependant été l'occasion de profiter pleinement de ses exceptionnelles qualités scientifiques et humaines, et je lui suis infiniment redevable de son encadrement, des opportunités dont il a su me faire bénéficier, de ses conseils désintéressés et honnêtes, de son aide dans les moments difficiles et surtout de tout ce qu'il m'a appris.

Je tiens aussi à remercier les membres du jury, Thierry Bodineau, Laure Dumaz et Adrien Hardy d'avoir accepté de participer à la soutenance. Je suis particulièrement reconnaissant envers les rapporteurs de cette thèse, Arnaud Guillin et Denis Talay, pour avoir pris le temps d'effectuer une lecture si consciencieuse du manuscrit.

Au cours de ces trois années, j'ai eu la chance de croiser plusieurs personnes dont les compétences scientifiques et la gentillesse ont sans conteste permis la réalisation de ce travail. Je pense notamment à Hugo Touchette, qui m'a initié à la théorie des grandes déviations et ses applications en physique, et dont la rencontre (heureusement au début de la thèse !) a été un déclencheur. Je pense également à Djalil Chafaï, avec lequel la collaboration sur les gaz de Coulomb a été un sujet d'enrichissement important, et dont la curiosité mathématique a été source de nombreuses discussions et de beaucoup d'enthousiasme. Enfin, le contact de Mathias Rousset a été l'occasion d'apprendre énormément en théorie des probabilités, et j'ai de très bons souvenirs de nos échanges (souvent techniques !), tandis qu'un travail avec Marielle Simon a permis de nombreuses discussions stimulantes. Je remercie grandement ces personnes qui m'ont tant appris durant ces années.

Je tiens également à remercier Eric Vanden-Eijnden de m'avoir accueilli pendant deux mois au Courant Institute à New York. Sa curiosité intellectuelle et sa prévenance ont permis de faire de ce séjour une forte expérience scientifique et humaine.

Parce qu'une thèse ne va pas sans organisation, je dois remercier Isabelle Simunic pour son accompagnement sur les plans logistique et personnel dans ce parcours. Son engagement auprès des étudiants et sa gentillesse constante m'auront aidé à bien des moments, et en bien des points du globe !

Plus généralement, je tiens à remercier tous les membres du CERMICS, dont l'enthousiasme donne à ce laboratoire une coloration particulière. Je remercie en particulier Antoine Levitt et Julien Reygner pour plusieurs discussions intéressantes, ainsi que Tony Lelièvre, Eric Cancès et Benjamin Jourdain pour des interactions toujours sympathiques et stimulantes (assez souvent dans un aéroport ou à l'étranger !). Je remercie également Stefano Olla pour avoir encadré le stage précédent ma thèse et pour plusieurs échanges à la suite de celui-ci.

Les doctorants du laboratoire sont également l'élément clé qui permet de s'épanouir et d'être soutenu pendant le travail, tant par des discussions informelles qu'à travers une aide technique précieuse. J'exprime donc ma reconnaissance envers les anciens (Boris, Alexandre, Marc, Julien, Athmane...) les contemporains (Lingling, Pierre-Loïk, Frédéric...) et les suivants (Adrien, Robert, Sami, Adel, Inass, Rafaël, Clément, Victor, Oumaima...) pour avoir aussi fait de cette vie quotidienne un plaisir.

Je remercie également mes amis Louis et Vincent, qui m'apportent toujours autant malgré les années, et que le hasard des voyages m'a amené à retrouver d'un bout à l'autre des États-Unis.

Ce travail a aussi été un engagement personnel, et je remercie mes parents, qui m'ont toujours soutenu dans les tournants qu'a pris ma vie, ainsi que mon frère et ma soeur qui supportent mon caractère depuis bien plus longtemps que tous ceux qui liront ces remerciements, et en méritent bien pour cela.

Enfin, je remercie D. pour tout ce que je n'ai pas dit et qui est sûrement l'essentiel.

Résumé

Cette thèse s'intéresse à différents problèmes de grandes déviations en rapport avec la physique statistique, qu'elle aborde sous l'angle théorique aussi bien que numérique. La première partie concerne l'étude de grandes déviations en temps long pour les processus de diffusion. Tout d'abord, de nouveaux résultats d'ergodicité sont montrés pour les dynamiques de Feynman–Kac, en temps discret et en temps continu. Ceci conduit à de nouveaux résultats fins (au sens de la topologie considérée) sur les grandes déviations de mesures empiriques de processus de diffusion. Divers aspects numériques sont ensuite abordés. Tout d'abord, des estimées d'erreur précises sont fournies pour les discrétisations de processus de Feynman–Kac, la non-linéarité de la dynamique demandant le développement de nouveaux outils. Afin de réduire la variance des estimateurs classiques de grandes déviations, un algorithme adaptatif est ensuite présenté, qui utilise les techniques dites d'approximation stochastique. Enfin, nous abordons un problème numérique concernant les systèmes à basse température, et présentons une méthode pour construire une approximation du contrôle optimal à partir de la théorie du chemin de réaction. La dernière partie de cette thèse porte sur un sujet légèrement différent, celui des gaz de Coulomb, qui apparaissent en physique mais aussi dans la théorie des matrices aléatoires. Nous présentons d'abord une méthode efficace pour la simulation de tels gaz, avant de nous tourner vers l'étude des gaz sous contrainte. Pour ceux-ci, nous prouvons de nouveaux résultats de concentration dans la limite d'un grand nombre de particules, sous certaines conditions sur la contrainte. Nous présentons également un algorithme de simulation qui confirme les attentes théoriques.

Abstract

This thesis is concerned with various aspects of large deviations theory in relation with statistical physics. Both theoretical and numerical considerations are dealt with. The first part of the work studies long time large deviations properties of diffusion processes. First, we prove new ergodicity results for Feynman–Kac dynamics, both in continuous and discrete time. This leads to new fine results (in the sense of topology) for large deviations of empirical measures of diffusion processes. Various numerical problems are then covered. We first provide precise error estimates on discretizations of Feynman–Kac dynamics, for which the nonlinear features of the dynamics demand new tools. In order to reduce the variance of naive estimators, we provide an adaptive algorithm relying on the technique of stochastic approximation. We finally consider a problem concerning low temperature systems. We present a new method for constructing an approximation of the optimal control from the instanton (or reaction path) theory. The last part of the thesis is concerned with the different topic of Coulomb gases, which appear both in physics and random matrix theory. We first present an efficient method for simulating such gases, before turning to gases under constraint. For such gases, we prove new concentration results in the limit of a large number of particles, under some conditions on the constraint. We also present a simulation algorithm, which confirms the theoretical expectations.

Préambule

Théorie des grandes déviations en physique statistique : Quelques aspects théoriques et numériques.

Le langage courant associe généralement à la notion d'aléatoire celle d'imprévisible, et renvoie à un évènement dont on ne pourrait prévoir l'issue. Si la modélisation aléatoire sert en effet à décrire une succession de phénomènes dont on ne peut prévoir *exactement* le résultat, la théorie des probabilités s'emploie à extraire de ce hasard une forme de *régularité*. Ainsi, on ne pourra dire si une pièce jetée retombera sur pile ou face, mais que l'issue aura telle ou telle probabilité (un demi pour chaque face si la pièce n'est pas biaisée). Dans la seconde moitié du XXème siècle, les mathématiciens ont poussé très loin ces raisonnements en introduisant de nombreux objets aléatoires, faisant au passage des connexions fructueuses avec d'autres branches des mathématiques comme l'analyse ou la géométrie.

Parallèlement à ces travaux, les physiciens ont incorporé depuis la fin du XIXème siècle l'aléatoire comme outil fondamental de la *modélisation*. Ainsi le mouvement brownien aura d'abord servi à décrire de petites particules dans un bain de pollen, avant d'être utilisé pour représenter l'agitation thermique en physique, puis l'incertitude en finance. D'une certaine façon, la physique statistique illustre la recherche de régularité dans ces modèles aléatoires, au sein d'une représentation *microscopique*, aléatoire, de la matière. Considérant par exemple un ensemble d'atomes, on ne saura décrire précisément le mouvement de chacun, mais bien plutôt leurs *propriétés moyennes*. Les échelles macroscopiques de la matière comportant un nombre extrêmement grand d'atomes (de l'ordre de 10^{23} pour une goutte d'eau), ce procédé de moyennisation est souvent une approximation raisonnable.

Cependant, les multiples échelles mises en jeu rendent ces limites (*e.g.* grand nombre de particules, temps long) difficiles à comprendre, tant du point de vue théorique que de la simulation numérique. Les systèmes étudiés n'auront ainsi jamais vraiment atteint leur état d'équilibre, qui représente leur comportement moyen, mais seront sujets à des *fluctuations*. Ces comportements anormaux, au sens "loin de la moyenne", seront généralement rares à l'échelle d'une simulation numérique, mais peuvent être beaucoup plus fréquents à l'échelle macroscopique. Il est donc important de les caractériser et, si possible, de les quantifier. Notons dès à présent que la question du calcul numérique lié aux évènements rares est délicate, pour la simple raison qu'un évènement de probabilité très faible se produira aussi rarement lors d'une simulation.

Une façon de s'attaquer à ces problèmes est de recourir à la théorie des grandes déviations qui, comme son nom l'indique, tente de décrire les états improbables d'un système. Ce cadre d'étude peut cependant être utilisé à d'autres fins, notamment afin de décrire l'état le plus probable d'un système, ou de mettre au point des méthodes numériques efficaces pour l'estimation d'évènements rares. Comme pour le développement de la physique statistique en générale, de nombreux physiciens et mathématiciens ont contribué au développement de cette théorie qui est encore un sujet de recherche actif, notamment pour les systèmes dits irréversibles.

De façon plus modeste, nous nous intéressons dans cette thèse au comportement de trois types de systèmes entrant dans le cadre des grandes déviations :

1. Les propriétés ergodiques et les fluctuations en temps long ;
2. Les systèmes métastables à basse température ;
3. Les gaz en interaction coulombienne dans la limite d'un grand nombre de particules.

Dans tous les cas, il s'agira de comprendre le comportement moyen d'un système aléatoire, et la façon dont il peut en dévier. Nous abordons ces questions sous trois angles complémentaires :

1. Fournir des preuves rigoureuses de théorèmes abstraits (ergodicité, concentration, conditionnement, etc.) ;
2. Mettre au point des algorithmes efficaces afin de quantifier numériquement les résultats théoriques ;

3. Donner un cadre à l'analyse d'erreur des algorithmes employés, afin de prouver leur efficacité.

Il est intéressant de noter que les résultats théoriques obtenus ont souvent été motivés par des calculs numériques surprenants, et nous espérons retranscrire (de façon malheureusement bien incomplète) cette interaction stimulante entre simulation et théorie. Nous pouvons maintenant décrire plus précisément le contenu de la thèse au vu de la discussion ci-dessus.

Partie I : Introduction

Nous commençons par une présentation générale en Partie I. Au-delà de la simple introduction des modèles en Section 1.1, reposant souvent sur des équations différentielles stochastiques, ce chapitre cherche à donner une vision intuitive des trois problèmes mentionnés plus haut. En particulier, nous verrons que les phénomènes limites s'expriment naturellement à travers le principe de Laplace auquel la théorie des grandes déviations, décrite en Section 1.2, permet de donner un sens rigoureux. La Section 1.3 montre comment utiliser ces outils dans le cadre de la physique statistique, en essayant d'insister sur la caractère physique des objets abstraits (entropie, énergie électrostatique, etc.). Nous aborderons finalement quelques problèmes numériques liés au calcul des fonctions de grandes déviations en Section 1.4.

Partie II : résultats théoriques d'analyse en temps long

Les contributions apportées par ce travail débutent en Partie II, et concernent l'étude théorique du comportement en temps long de systèmes markoviens (chaînes de Markov et processus de diffusion).

En Chapitre 2, nous prouvons des résultats d'ergodicité pour les dynamiques de Feynman–Kac, qui apparaissent aussi bien dans la théorie des grandes déviations qu'en mécanique quantique et dans la théorie du filtrage non linéaire. De façon plus spécifique, nous étudions la compétition entre la dynamique sous-jacente et son poids statistique à travers une technique de fonction de Lyapunov, ce qui permet de caractériser de façon précise la stabilité.

Le Chapitre 3 concerne le problème dual : les fluctuations de moyennes empiriques de processus de diffusion. À partir d'outils abstraits développés dans le Chapitre 2, nous dégageons un nouveau principe de grandes déviations (LDP) pour la moyenne empirique d'une diffusion (nouveau en cela qu'il considère une topologie plus fine que celle utilisée dans les résultats connus). La préoccupation de cette partie était de comprendre sous quelles conditions, pour une dynamique donnée, la moyenne d'une fonction satisfait un LDP. Autrement dit, nous obtenons l'équivalent d'une condition de Cramer pour les diffusions, qui n'a rien d'évident a priori (les conditions obtenues sont bien différentes d'une condition de moment exponentielle par exemple). Dans un second temps, nous nous intéressons à la fonctionnelle d'entropie dynamique associée au principe de grandes déviations. En particulier, nous utilisons et généralisons des travaux récents afin de décomposer l'entropie en deux parties :

- Une entropie *réversible* correspondant à une information de Fisher ;
- Une entropie *irréversible* exprimée par la norme duale d'un espace de Sobolev adéquat.

Si le premier résultat est bien connu pour les systèmes réversibles, nous croyons le second assez original et relié à des travaux récents sur les déviations de courants empiriques, notion découverte en physique théorique il y a une dizaine d'années. De manière plus générale, il s'agit d'une manifestation du second principe de la thermodynamique pour les systèmes irréversibles.

Partie III : analyse numérique et algorithmes

Si la Partie II apporte certaines réponses à des questions abstraites, elle reste relativement éloignée de l'estimation *quantitative* des quantités considérées. Nous abordons donc en Partie III différents aspects numériques liés à ces problèmes de grandes déviations¹.

Le Chapitre 2 ayant montré l'utilité des dynamiques de Feynman–Kac pour l'estimation de fonctions de grande déviations, nous entreprenons en Chapitre 4 une analyse d'erreur de leur discrétisation. Plus précisément, étant donné un semigroupe continu ayant certaines propriétés d'ergodicité en temps long, nous considérons sa discrétisation par un semigroupe discret avec un pas de temps Δt , et étudions l'erreur commise sur les propriétés ergodiques. Nous utilisons pour cela une analyse d'erreur à

¹Je tiens à cependant signaler que les Chapitres 4 et 5 ont *motivé* les questions de la Partie II, et non l'inverse.

la Talay–Tubaro, qui se révèle cependant difficile à mettre en oeuvre du fait de la non linéarité de la dynamique. Ces difficultés demandent d’introduire de nouvelles techniques de preuve, en particulier la construction d’un problème spectral approché. Un résultat intéressant de ce travail est l’introduction d’une nouvelle notion d’erreur, celle concernant la fonction génératrice des cumulants associée à l’intégrale en temps long du poids de Feynman–Kac. Cette erreur ne peut pas être considérée pour les processus markoviens, car dans ce cas la fonction génératrice est toujours nulle pour le processus et sa discrétisation (en l’absence de poids de Feynman–Kac). Un espoir est donc que les techniques développées dans ce travail puissent être utiles à l’analyse d’erreur d’autres dynamiques non linéaires.

Les chapitres suivants concernent la mise au point d’algorithmes efficaces pour le calcul de fonctions de grandes déviations et la réduction de variance. En effet, comme il sera détaillé en Partie I, les estimateurs naïfs de fonctions de grandes déviations ont une grande variance, ce qui conduit à des coûts de calcul prohibitifs pour une erreur donnée. Il est donc utile de développer des méthodes permettant de réduire ces coûts. Le Chapitre 5 concerne l’étude des grandes déviations en temps long. L’objectif de ce travail est de mettre en place un algorithme *adaptatif* de résolution d’un problème de contrôle, afin de réduire la variance de façon optimale. Nous adaptons pour cela des algorithmes utilisés en finance pour le contrôle du risque, que nous enrichissons de nouveaux estimateurs issus de travaux récents en physique théorique.

Le Chapitre 6 s’intéresse aux dynamiques à basse température. Ces systèmes souffrent généralement de ce que l’on appelle *métastabilité*, qui décrit ce comportement erratique où un système reste longtemps piégé dans un état avant de sauter brusquement dans un autre. Ces transitions rares et brutales conduisent à une variance importante des estimateurs de quantités moyennes du système. Nous développons en Chapitre 6 une méthode pour construire de façon systématique des contrôles réduisant la variance du système en petite température (ce que nous montrons formellement). Nous utilisons pour cela la théorie des *chemins de transitions*, que nous affinons afin d’améliorer le contrôle. Nous montrons formellement que cette approximation n’est rien d’autre qu’un développement de Taylor du contrôle optimal autour du chemin de réaction.

Partie IV : gaz de Coulomb et matrices aléatoires

La Partie IV s’est développée autour d’une idée numérique purement fortuite, reposant sur le lien entre matrices aléatoires et gaz de Coulomb. En effet, la densité de probabilité des valeurs propres de certains modèles de matrices (Wigner, Ginibre, etc.) peut s’écrire sous la forme d’une mesure de Gibbs, qui décrit l’état d’équilibre d’un système où deux actions s’opposent :

- Une force de confinement qui ramène les valeurs propres vers l’origine ;
- Une répulsion singulière (coulombienne ou logarithmique) incitant les particules à se repousser les unes les autres.

Sous des changements d’échelles appropriés, la mesure empirique d’un tel gaz converge vers un état déterministe appelé *mesure d’équilibre*. Étant donnés les travaux sur le comportement en temps long des systèmes aléatoires mentionnés ci-dessus, il semblait naturel de considérer cette mesure d’équilibre comme la mesure invariante d’une dynamique (de Langevin par exemple), et de l’échantillonner par une discrétisation de ladite dynamique. Le Chapitre 7 montre que c’est en effet possible, à condition toutefois d’employer un algorithme adapté à la singularité des interactions. Cet outil peut ensuite être utilisé pour tester numériquement des conjectures, comme l’universalité des fluctuations au bord du support de la mesure empirique pour les gaz de Coulomb.

Au cours de ce projet est apparu un prolongement intéressant, consistant à ajouter des *contraintes* au modèle. Il se trouve que des échantillonneurs de mesures de probabilité restreintes à une sous-variété ont été développés en dynamique moléculaire et en statistique computationnelle pour différentes raisons (notamment le calcul d’énergie libre). Nous avons donc étudié numériquement le comportement de gaz conditionnés avec un de ces algorithmes. Le comportement en a été si surprenant que nous avons souhaité mieux le comprendre et fournir des preuves des résultats observés. Le Chapitre 8 décrit cette démarche, qui repose finalement sur le *principe de conditionnement de Gibbs*. Nous prouvons que ce principe s’applique dans le contexte des gaz de Coulomb, et étudions finement les problèmes d’optimisation sous contrainte auxquels il conduit afin de décrire l’état limite.

Preamble

Large Deviations Theory in Statistical Physics: Some Theoretical and Numerical Aspects.

Common knowledge generically associates randomness to unpredictability, or an event whose issue cannot be forecast. Random modelling is indeed used for describing phenomena whose result cannot be predicted *exactly*, but probability theory strives to extract from this randomness a form of *regularity*. In a coin toss game for instance, one cannot say what the result will be, but one can describe the probability for each side to come out (one half for each side in an unbiased situation). In the second half of the XXth century, mathematicians have pushed forward this theory very far by introducing a series of random objects, obtaining fine results together with connections to other branches of mathematics such as geometry or analysis.

Since the end of the XIXth century, in parallel, physicists associated embedded randomness as a fundamental tool for *modelling*. A famous example is the Brownian motion, first used to describe small particles in a pollen bath, which then became a standard tool for modelling thermal agitation in classical physics. In some way, statistical physics illustrates the search for regularity in random models in a *microscopic*, random, description of matter. Consider for instance a set of atoms interacting together. We are in general not able to describe precisely the motion of each of those atoms, but rather some *average property*. This averaging effect can be justified by the very large amount of particles in a real-life system (say around 10^{23} for a droplet of water), and often is a reasonable approximation.

However, the multiple scales at stake make the limiting procedure difficult to understand, both from a theoretical and a numerical standpoint. The systems under consideration do not reach exactly their equilibrium state, which indeed represents their average behaviour, but are subject to *fluctuations* around it. These abnormal behaviours, in the sense “far from the average”, are typically rarely observed in numerical simulations, but they may well not be rare on a macroscopic scale. It is therefore important to characterize them and quantify them if possible. Let us mention right now that the numerical estimation of rare events is a difficult area, for the simple reason that those events are rarely observed during a simulation.

One way to address these problems is to resort to the theory of large deviations that, as the name suggests, tries to describe the unlikely states of a system. This framework can also be used for other purposes such as finding the most probable state of a system, or designing efficient algorithms for rare event estimates. Like the development of statistical physics in general, many physicists and mathematicians contributed to the theory, which is still an active area of research, in particular for irreversible systems.

At a more humble level this thesis considers the behaviour of three types of systems entering the framework of large deviations:

1. Ergodic properties and fluctuations in the long time limit;
2. Metastable systems at low temperature;
3. Gases interacting through Coulomb potentials in the mean field limit.

In any case, we are interested in the average behaviour of a random system and the way it may deviate from it. We address these problems in three ways:

1. Give rigorous proofs of abstract theorems (large deviations, conditioning, etc.);
2. Design efficient algorithms for quantifying the abstract results;
3. Provide an error analysis framework for the algorithms employed.

Interestingly, the theoretical results we obtain were often motivated by surprising numerical simulations. We hope to reproduce (in an uncomplete way) this stimulating interaction between simulation and theory, and we describe below in more details the material of the thesis.

Part I: Introduction

We start with a general introduction in Part I. Section 1.1 presents most of the models (often based on stochastic differential equations) and problems in an informal way. This highlights the importance of the Laplace principle for studying limiting procedures for probability measures, which is given a precise meaning by the large deviations theory described in Section 1.2. Section 1.3 shows how to use these tools in a statistical physics context, by insisting on the physical interpretation of the abstract objects at hand (entropy, electrostatic energy, etc.). We finally present some numerical problems related to the computation of large deviations functions in Section 1.4.

Part II: Theoretical analysis for long time problems

The contributions of the present thesis start in Part II, which is concerned with the long time behaviour of Markovian dynamics (Markov chains and diffusion processes).

In Chapter 2, we prove new ergodicity results for Feynman–Kac dynamics, which appear in large deviations theory but also in quantum physics and non-linear filtering. To be more specific, we study the trade-off between the underlying dynamics and its statistical weight through Lyapunov function techniques, which allows to characterize precisely the stability of the dynamics.

Chapter 3 turns to the dual problem of the fluctuations of empirical measures of diffusions. Using the abstract tools developed in Chapter 2, we prove a new large deviations principle (LDP) for such empirical measures (new in that it concerns a finer topology than usual results). The preoccupation of this part was to understand, for a given dynamics, which functions have an empirical mean satisfying a LDP. In other words, we obtain an equivalent to Cramer’s condition for diffusions, which is not obvious a priori (in particular this is not an exponential moment condition). We next consider the dynamical entropy associated with the LDP, and show that it can be decomposed into two parts:

- A *reversible* part which is a Fisher information;
- An *irreversible* part expressed as the dual norm of a relevant Sobolev space.

If the first result is well-known for reversible systems, we believe the second to be quite original and related to recent works on large deviations for empirical currents, an active field in nonequilibrium statistical physics. From a more general perspective, this is an illustration of the second law of thermodynamics for irreversible systems.

Part III: Numerical analysis and algorithms

While Part II is mostly concerned with theoretical problems, we turn in Part III to numerical and algorithmic issues related to large deviations ².

Chapter 2 proved the importance of Feynman–Kac dynamics in the context of large deviations. We thus consider in Chapter 4 the problem of their numerical discretization. More precisely, given a time continuous semigroup with certain ergodic properties in the long time limit, we consider its discretization with a discrete semigroup of time step $\Delta t > 0$ and study the error made on the ergodic properties. For this we use techniques à la Talay–Tubaro, which is made difficult by the non-linearity of the dynamics. This requires to introduce new techniques of proof, in particular the construction of an approximate spectral problem. An interesting output of this work is a new notion of error concerning the cumulant generating function. This error does not make sense for Markovian dynamics, for which the cumulant is always zero for the process and its discretization. A hope is that our technique could be used for studying other non-linear dynamics.

The two following chapters treat some aspects of the computation of large deviations functions and variance reduction. Indeed, as made precise in Part I, naive estimators of large deviations functions have a dramatically large variance, leading to prohibitive computational costs for a fixed error. It is thus useful to develop methods to reduce this variance and compute more efficiently the quantities of interest. Chapter 5 is concerned with long time large deviations. The goal of this work is to design an *adaptive* algorithm, solving on the fly an optimal control problem to reduce the variance in an optimal way. For this we use an algorithm from finance (risk-sensitive control), which we adapt to our context and enhance with new estimators based on recent theoretical results.

²Let us mention that the study of Part II was *motivated* by the works of Chapters 4 and 5, and not conversely.

Chapter 6 turns to low temperature problems. Systems at low temperature often suffer from *metastability*, a notion which describes the behaviour of a system remaining trapped for a long time in a state before hopping abruptly to another one. Those rare but important transitions lead to a high variance of the quantities of interest. We develop in Chapter 6 an approach to build variance reduction controls for computing expectations at low temperature (which we prove formally). The strategy relies on the *transition path theory*, from which we build higher order terms to improve the control. We show formally that the so-built expansion is actually the Taylor expansion of the optimal control around the transition path.

Part IV: Coulomb gases and random matrices

Part IV has grown around a numerical idea relying on the relation between random matrices and Coulomb gases. For this we note that the probability density of the eigenvalues of some matrix models (Wigner, Ginibre, etc.) can be written as a Gibbs measure, describing the equilibrium state of a system under the two competing actions:

- A confinement force pushing the eigenvalues towards the origin ;
- A singular repulsion (Coulombic or logarithmic) repelling the particles from each other.

Under an appropriate rescaling, the empirical measure of such a gas converges to a deterministic state called *equilibrium measure*. According to the above mentioned works on long time analysis, it seemed natural to consider the Gibbs measure associated to the gas as the invariant measure of a dynamics (for instance the Langevin dynamics) and to sample it numerically. The Chapter 7 shows that this is indeed possible, provided one uses an algorithm adapted to the singularity of the interactions. This numerical tool can then be used to explore conjectures, such as the universality of the fluctuations at the edge of the support of the empirical measure for a Coulomb gas.

An interesting sequel to this project is the addition of *constraints* to the dynamics. Simulating diffusions under constraints is actually an active field in molecular dynamics and computational statistics for various reasons (*e.g.* for free energy computations). We therefore studied numerically conditioned Coulomb and log-gases with one of these algorithms³. The results were so surprising that we wished to provide proofs of what we observed. The Chapter 8 describes our results, which rely in the end on a *Gibbs conditioning principle*. After proving this principle, we provide a fine analysis of the constrained optimization problems to which it leads in the context of Coulomb gases, and compare the results with numerical simulations.

³We mention that sampling the spectrum of random matrices with a dynamics is not efficient from a computational perspective. However, conditioning the spectrum of a random matrix is difficult and costly: our constrained algorithm is then interesting in a random matrix context.

List of publications

Here is a list of articles that were written during this thesis:

- [185] Ferré, G., Vanden-Eijnden, E and Gräffe, T. (2019+). *Approximate controls for low temperature systems*.
In preparation.
- [183] Ferré, G. and Stoltz, G. (2019). *Large deviations of empirical measures of diffusions in fine topologies*.
Preprint arXiv:1906.0941.
- [83] Chafaï, D., Ferré, G. and Stoltz, G. (2019). *Coulomb gases under constraint: some theoretical and numerical results*.
Preprint arXiv:1907.05803 .
- [181] Ferré, G. , Rousset, M. and Stoltz, G. (2018). *More on the long time stability of Feynman-Kac semigroups*.
Preprint arXiv:1807.00390.
- [82] Chafaï, D. and Ferré, G. (2018). *Simulating Coulomb and Log-Gases with Hybrid Monte Carlo Algorithms*.
The Journal of Statistical Physics, 174(3), 692-714.
- [184] Ferré, G. and Touchette, H. (2018). *Adaptive sampling of large deviations*.
The Journal of Statistical Physics, 172(6), 1525-1544.
- [182] Ferré, G. and Stoltz, G. (2019). *Error estimates on ergodic properties of Feynman-Kac semigroups*.
Numerische Mathematik, 143(2), 261-313.

Below is a list of articles published before the thesis:

- [179] Ferré, G., Haut, T. and Barros, K. (2017). *Learning molecular energies using localized graph kernels*.
The Journal of Chemical Physics, 146(11), 114107.
- [431] Zentner, I., Ferré, G., Poirion, F. and Benoit, M. (2016). *A biorthogonal decomposition for the identification and simulation of non-stationary and non-Gaussian random fields*.
Journal of Computational Physics, 314, 1-13.
- [180] Ferré, G., Maillet, J.-B. and Stoltz, G. (2015). *Permutation-invariant distance between atomic configurations*.
The Journal of Chemical Physics, 143(10), 104114.

Oral presentations in conferences

Below is a list of conferences where I presented the above mentioned works:

- International Conference on Scientific Computation and Differential Equations (SciCADE), Innsbruck, *Error estimates on ergodic properties of Feynman–Kac semigroups*, 2019.
- International Congress on Industrial and Applied Mathematics (ICIAM), Valencia, *Conditioned random matrices and Coulomb gases: some numerical and theoretical aspects*, 2019.
- Journées de Probabilités 2019, Dourdan, *Large deviations for the empirical measure of diffusion: revisiting Cramer’s condition with Lyapunov functions*, 2019.
- Franco-German Meeting Workshop on Mathematical Aspects in Computational Chemistry, Aachen, *Feynman-Kac models: stability and further issues*, 2018.
- SIAM meeting on Mathematical Aspects of Material Science, Portland (Oregon), *An Adaptive Algorithm for Sampling Large Deviation Functions*, 2018.
- SIAM meeting on Mathematical Aspects of Material Science, Portland (Oregon), *Error Estimates and Stability for Diffusion Monte Carlo Algorithms*, 2018.
- International Conference in Monte Carlo and Quasi-Monte Carlo Methods in Scientific computing, Rennes, *Long time stability of Feynman-Kac dynamics*, 2018.
- Congrès National d’Analyse Numérique, Cap d’Agde, *Error estimates and ergodic properties of Markov chains and Feynman-Kac models*, 2018.
- IHP, Paris, Young Probabilists’ Day (Les probabilités de demain), *Long time stability of Feynman-Kac models*, 2018.
- ICTS, Bangalore, Large deviations theory in statistical physics, recent advances and future challenges, *Error estimates for Feynman-Kac semi-groups*, 2017.
- IHP, Paris, Stochastic Dynamics Out of Equilibrium, Young Researchers’ Talks, *Error estimates for Feynman-Kac semi-groups*, 2017.
- Institute for Pure and Applied Mathematics, UCLA, Los Angeles, Understanding Many-Particle Systems with Machine Learning, *Learning potential energy landscapes with localized graph kernels*, 2016.

Talks in seminars

Below is a list of seminars talks which I gave:

- Université de Marseille, Séminaire de probabilité, *Large deviations of diffusion processes for unbounded observables*, 2019.
- Université de Lille, Séminaire, *Long time behaviour of linear and nonlinear semigroups in probability*, 2019.
- Université Paris-Dauphine, Young researchers’ seminar, *A gentle introduction to ergodic theory for Markov chains and Feynman-Kac dynamics*, 2019.
- CERMICS, ENPC, Paris, MsMath working group, presentation of the paper *Scaling limit of the Stein variational gradient descent: the mean field regime*, by J. Lu, Y Lu and J. Nolen, 2018.
- Courant Institute of Mathematical Science, New York, Student Probability seminar, *Ergodicity for Markov chains, coupling probabilities and analysis*, 2018.
- CERMICS, ENPC, Paris, PhD seminar, *Crash course on ergodicity for Markov chains: doing probabilities like an analyst*, 2018.

Poster presentations in conferences

Finally, below is a list of conferences where I presented posters:

- CIB-CECAM, Lausanne, Computational Mathematics for model reduction and predictive modelling in molecular and complex systems, 2019.
- CIRM, Marseille, Advances in Computational Statistical Physics, 2018.
- INRIA Rennes, Simulation aléatoire : problèmes actuels, 2018.
- Alan Turing Institute, London, Data-Driven Modelling of Complex Systems, 2018.
- Duke University, Durham, Quasi Monte-Carlo and High-Dimensional Sampling Methods for Applied Mathematics, 2017.
- IPAM, Los Angeles, Complex High-Dimensional Energy Landscapes, 2017.
- IHP, Paris, Stochastic Dynamics Out of Equilibrium, 2017.
- IPAM, Los Angeles, Understanding Many-Particle Systems with Machine Learning, 2016.

I	Introduction	23
1	Introduction	25
1.1	Models and motivations	26
1.1.1	Stochastic differential equations and computational statistical physics	26
1.1.2	Ergodicity and fluctuations	32
1.1.3	Systems at small temperature	37
1.1.4	Random matrices and Coulomb gases	40
1.2	Abstract large deviations theory	44
1.2.1	A natural framework for the concentration of probability measures	44
1.2.2	Important results	46
1.3	Large deviations in statistical physics	51
1.3.1	Long time fluctuations	51
1.3.2	Small temperature regime	54
1.3.3	Singular gases	55
1.4	Some numerical aspects of rare events analysis	56
1.4.1	Discretization of SDEs	57
1.4.2	Estimating rare events	59
1.5	Contributions	62
1.5.1	Feynman–Kac dynamics and large deviations	63
1.5.2	Numerical analysis and algorithms	65
1.5.3	Coulomb gases and random matrices	68
	Bibliographical comments	71
	Notational comments	72
II	Feynman–Kac dynamics and large deviations: Theoretical analysis	73
2	Long time analysis of Feynman–Kac dynamics	75
2.1	Motivation	76
2.2	Results	77
2.2.1	Framework	77
2.2.2	Results in discrete time	78
2.2.3	Results in continuous time	83
2.3	Applications	87
2.3.1	Examples in discrete time	87
2.3.2	Applications to diffusion processes	90
2.3.3	Convergence results uniform with respect to the time step	91
2.4	Discussion	94
2.5	Complementary elements	95
2.5.1	Stability of Markov chains	95
2.5.2	Useful theorems	95
2.5.3	Proof of Lemma 2.4	97

2.5.4	Proof of Lemma 2.5	98
2.5.5	Proof of Lemma 2.6	100
2.5.6	Proof of Lemma 2.20	101
3	Large deviations of empirical measures of diffusions in fine topologies	103
3.1	Presentation	104
3.2	Large deviations principle	106
3.2.1	Setting	106
3.2.2	Statement of the main results	109
3.3	Decomposition of the rate function	113
3.3.1	Donsker–Varadhan variational formula	113
3.3.2	Entropy decomposition: symmetry and antisymmetry	114
3.4	Applications	116
3.4.1	Overdamped Langevin dynamics	116
3.4.2	Underdamped Langevin dynamics	118
3.5	Conclusion and perspectives	122
3.6	Proofs	123
3.6.1	Proof of the large deviations principle	123
3.6.2	Proofs of Section 3.3	135
3.6.3	Proof of Proposition 3.8	139
3.6.4	Proof of Lemma 3.23	139
III	Numerical analysis and algorithms	141
4	Error estimates for ergodic properties of discretized Feynman–Kac semigroups	143
4.1	Presentation	144
4.2	Convergence properties of Feynman–Kac semigroups	146
4.2.1	Continuous dynamics	146
4.2.2	Discretization	148
4.3	Numerical analysis of the discretization	151
4.3.1	Error estimates on the invariant probability measure	151
4.3.2	Alternative error estimate for the principal eigenvalue	156
4.3.3	TU Lemma	157
4.3.4	Second order schemes	158
4.4	Numerical application	161
4.4.1	Monte Carlo discretization	161
4.4.2	Galerkin discretization	162
4.4.3	Numerical results	163
4.5	Discussion	166
4.6	Proofs	167
4.6.1	Proof of the results of Section 4.2.1	167
4.6.2	Proof of Theorem 4.9	169
4.6.3	Proofs related to Theorem 4.16	171
4.6.4	Proof of Proposition 4.23	176
4.7	Markov contractions and Dobrushin coefficients	177
5	Adaptive sampling of large deviations functions	179
5.1	Motivation	179
5.2	Framework	180
5.2.1	Model and notation	180
5.2.2	Large deviations functions and driven process	182
5.3	Adaptive algorithm	184
5.3.1	Time discretization	184
5.3.2	Stochastic approximation and annealing	185
5.3.3	Spatial projection	186
5.3.4	Algorithm	186
5.3.5	Remarks	187
5.4	Applications	188

5.4.1	Ornstein–Uhlenbeck process	188
5.4.2	Periodic diffusion	191
5.5	Conclusion	192
6	Constructing approximate controls for small temperature systems	195
6.1	Motivation	195
6.2	Optimal control and low temperature limit	196
6.2.1	A direct approach	196
6.2.2	Optimal tilting on path space	197
6.3	Instanton and optimal bias	198
6.3.1	First order control	198
6.3.2	Second order control	200
6.3.3	Consistency with optimal control	201
6.4	Numerical applications	203
6.5	Conclusion	205
IV	Coulomb gases and random matrices	207
7	Sampling Coulomb gases with Hamiltonian Monte Carlo	209
7.1	Presentation	209
7.1.1	Boltzmann–Gibbs measures	210
7.1.2	Static energy and equilibrium measures	211
7.1.3	Two remarkable gases from random matrix theory	211
7.2	Simulating log-gases and Coulomb gases	213
7.2.1	Standard sampling methods	214
7.2.2	Hybrid Monte Carlo algorithm	215
7.3	Numerical experiments on remarkable models	218
7.3.1	Case study: 1D	218
7.3.2	Case study: 2D	218
7.3.3	Case study: 3D	218
8	Coulomb gases under constraint: some theoretical and numerical results	227
8.1	Motivation	228
8.2	From conditioning to shifting: quadratic confinement with linear constraint	231
8.3	A general conditioning framework	233
8.3.1	Gibbs conditioning	233
8.3.2	Linear statistics	235
8.3.3	Quadratic statistics	237
8.4	Numerical illustration	239
8.4.1	Description of the algorithm	239
8.4.2	Numerical results	243
8.5	Proofs	248
8.5.1	Proof of Theorem 8.1	248
8.5.2	Proofs of Section 8.3.1	250
8.5.3	Proof of Theorem 8.9	252
8.5.4	Proof of Theorem 8.16	256

Part I

Introduction

Contents

1.1 Models and motivations	26
1.1.1 Stochastic differential equations and computational statistical physics	26
1.1.2 Ergodicity and fluctuations	32
1.1.3 Systems at small temperature	37
1.1.4 Random matrices and Coulomb gases	40
1.2 Abstract large deviations theory	44
1.2.1 A natural framework for the concentration of probability measures	44
1.2.2 Important results	46
1.3 Large deviations in statistical physics	51
1.3.1 Long time fluctuations	51
1.3.2 Small temperature regime	54
1.3.3 Singular gases	55
1.4 Some numerical aspects of rare events analysis	56
1.4.1 Discretization of SDEs	57
1.4.2 Estimating rare events	59
1.5 Contributions	62
1.5.1 Feynman–Kac dynamics and large deviations	63
1.5.2 Numerical analysis and algorithms	65
1.5.3 Coulomb gases and random matrices	68
Bibliographical comments	71
Notational comments	72

This chapter introduces the main models studied along the manuscript, the mathematical framework of large deviations theory and the numerical issues raised by practical applications. Section 1.1 is a heuristic presentation of the systems at hand, with an emphasis on physical motivations. It covers in particular three classes of problems:

1. Long time behaviour and ergodicity for stochastic differential equations (Section 1.1.2);
2. Systems at small temperature (Section 1.1.3);
3. Coulomb gases in the mean field limit (Section 1.1.4).

A common feature of these models is their relation to the Laplace principle, which can be used to understand their limiting behaviour in different regimes: long time, small temperature and large number of particles.

We next turn to the mathematical formalism necessary to rigorously handle this reasoning by presenting the framework of large deviations theory in Section 1.2. Once the basic elements are set (Section 1.2.1), we present some important results in Section 1.2.2. In particular, we introduce the mathematical notions of *entropy* and *free energy* that are used throughout this thesis. Section 1.3 then

comes back to the three above mentioned problems by recasting them in the words of large deviations. In addition to introducing the subjects treated in the next parts, we hope that this presentation illustrates the flexibility of the methods used here, and provides nice examples of application of this rather abstract theory.

An interesting feature of the problems at hand is that they are often motivated by practical situations, and the abstract solutions provided by the large deviations theory are not always satisfactory for the practitioner. This motivates several numerical problems, some of which are presented in Section 1.4. More precisely, our focus is on discretization issues for stochastic differential equations (Section 1.4.1), and variance reduction for computing large deviations functions (Section 1.4.2).

Finally, we close this introduction by presenting the main contributions of this work in Section 1.5. This is followed by brief comments on the bibliography and notation.

1.1 Models and motivations

In this section, we first describe the main models studied in the present work together with their physical motivations. We next introduce the problems treated throughout in an informal way, which we hope to be a gentle introduction to the settings of Parts II, III and IV. In particular, we wish to emphasize the importance of the Laplace principle in models from probability theory and statistical physics.

Section 1.1.1 introduces models based on stochastic differential equations, in particular overdamped and underdamped Langevin dynamics, insisting on reversibility and irreversibility. We then discuss the long time behaviour and fluctuations of these stochastic models in Section 1.1.2, which is the opportunity to present the concept of free energy and dynamical entropy in this context. Next, Section 1.1.3 is an introduction to the study of systems at low temperature, a common situation in physics. Finally, Section 1.1.4 is concerned with a quite different model of particles under singular interactions, which is motivated here by the analysis of the spectrum of random matrices.

The presentation below also allows to introduce the main notation used in the different parts. The mathematical style of the presentation is rather informal to highlight the main ideas, and the reader is assumed to be familiar with basic notions of analysis and probability. The presentation of the mathematical tools necessary for the precise analysis of the phenomena described below is postponed to Section 1.2.

1.1.1 Stochastic differential equations and computational statistical physics

1.1.1.1 Modelling with random equations

We generally consider stochastic differential equations (SDE) of the form

$$dX_t = b(X_t) dt + \sigma(X_t) dB_t, \quad (1.1)$$

where the process $(X_t)_{t \geq 0}$ has values in a state space $\mathcal{X} \subset \mathbb{R}^d$ and $d \geq 1$ is a physical dimension (one may think of $\mathcal{X} = \mathbb{R}^d$ or $\mathcal{X} = \mathbb{T}^d$, the d -dimensional torus). Here, $b : \mathcal{X} \rightarrow \mathbb{R}^d$ is a drift, $\sigma : \mathcal{X} \rightarrow \mathbb{R}^{d \times m}$ is a diffusion field (both assumed to be smooth here) and $(B_t)_{t \geq 0}$ is an m -dimensional Brownian motion for some $m \geq 1$. This setting encompasses many situations in statistical physics, where the process $(X_t)_{t \geq 0}$ typically represents a set of particles evolving randomly in time. The drift b is a force, which may derive from a potential energy function or can be induced by an external forcing such as an electric field. The Brownian motion is a thermal noise whose intensity is prescribed by the possibly inhomogeneous temperature field proportional to σ .

The generator of the dynamics [262] reads

$$\mathcal{L} = b \cdot \nabla + \frac{\sigma \sigma^T}{2} : \nabla^2, \quad (1.2)$$

where \cdot denotes the scalar product in \mathbb{R}^d and, for two matrices $A, B \in \mathbb{R}^{d \times d}$ we write $A : B = \text{Tr}(A^T B)$, where A^T is the transpose of A . Moreover, ∇ and ∇^2 are the gradient and Hessian operators respectively. Although details on the definition of \mathcal{L} will be given later, the operator can be defined to act at least on smooth and compactly supported functions as soon as b and σ are smooth. We insist on the fact that, under this convention, \mathcal{L} prescribes the evolution of *observables* of the system. In other words, if φ is a smooth function over \mathcal{X} and if we set, for all time $t \geq 0$,

$$(P_t \varphi)(x) = \mathbb{E}_x[\varphi(X_t)], \quad (1.3)$$

then

$$\partial_t P_t \varphi = \mathcal{L} P_t \varphi, \quad P_0 \varphi = \varphi. \quad (1.4)$$

In (1.3), \mathbb{E}_x denotes the expectation with respect to all the realizations of the trajectories $(X_t)_{t \geq 0}$ satisfying (1.1) and starting at $X_0 = x$. More generally, if ν is a probability measure, we denote by \mathbb{E}_ν the expectation when the initial condition X_0 is distributed according to ν . In this setting, (1.4) is a manifestation of the semigroup structure of the Markovian dynamics (1.1) [361, 286].

A dual viewpoint of (1.4) is to consider the evolution of the *probability density* of $(X_t)_{t \geq 0}$. More precisely, it is possible to associate to the dynamics the probability measure over \mathcal{X} defined by $\mu_t(\cdot) = \mathbb{P}_x(X_t \in \cdot)$, which satisfies when $X_0 = x$ and under mild assumptions detailed later on the equation

$$\partial_t \mu_t = \mathcal{L}^\dagger \mu_t, \quad \mu_0 = \delta_x, \quad (1.5)$$

where \mathcal{L}^\dagger is the adjoint of \mathcal{L} in $L^2(\mathcal{X})$, the space of Lebesgue square integrable functions over \mathcal{X} . More precisely

$$\mathcal{L}^\dagger \mu = \nabla \cdot \left[-b\mu + \nabla \cdot \left(\frac{\sigma \sigma^T}{2} \mu \right) \right],$$

where $\nabla \cdot$ denotes the divergence operator. The evolution of μ_t describes the probability of presence of X_t in some region of space at time $t \geq 0$. Alternatively, (1.4) is a weak formulation of (1.5) for a class of test functions φ . Note that the convention in physics differs from the one used here, since in this field the generator \mathcal{L} is often used to describe the evolution of the probability measure μ_t . We thus emphasize that, in our notation, \mathcal{L} prescribes the evolution of observables while \mathcal{L}^\dagger is associated with the evolution of the probability density of $(X_t)_{t \geq 0}$.

1.1.1.2 Invariant measure

A crucial point of statistical physics is that one is often not interested in the particular motion of the particles, but rather in their *average behaviour*. For instance, if a test function φ as above represents some observable in statistical physics such as a pressure, one would like to know the average value of this quantity rather than its value for a particular configuration. This point of view is justified by the following scaling arguments:

- in general, the number of particles under consideration is large and the observable of interest is somewhat local, so that there is a mean field averaging;
- the time scale of interest is generally much larger than the relaxation time of the system, hence a long time averaging.

The first aspect is related for instance to the fields of hydrodynamic, thermodynamic and mean field limits depending on the system under consideration (see the very good books [392, 268, 386] for more details in different contexts). An instance of such a scaling limit is illustrated in Part IV, and we will come back to this aspect in Section 1.1.4. From a physical perspective, this limit is motivated by the large number of particles in realistic systems: for instance, a droplet of water contains of the order of 10^{23} water molecules, which is pretty much $+\infty$ for a mathematician! On the other hand, the time averaging effect is generally referred to as *ergodicity*. This is the main preoccupation of this thesis, which we motivate by the following considerations. Typically, an atomic system of particles evolves on time scales of the order of picoseconds (10^{-12} s), which corresponds to the typical period of atomic vibrations. For a molecule, dihedral rotations evolve on a scale of roughly 10^{-11} seconds, while a human eye observes phenomena of the order of a second. As a result, for applications, the relaxation time of the system can be considered to be reached¹ and we can consider a long time limit of (1.1), which we formalize with the notion of *invariant measure*.

If we assume that after a long time, the probability to find the system in some state does not change anymore, *i.e.* $\partial_t \mu_t = 0$, (1.5) becomes

$$\mathcal{L}^\dagger \mu = 0, \quad (1.6)$$

¹Our motivations stem from atomistic physics, but the same reasoning can be applied to other systems involving different time scales. For instance, the large deviations theory described next has recently been used for studying rare events in climate models, see [355] and references therein.

for some probability measure μ called the *invariant measure* (or stationary measure) of the process $(X_t)_{t \geq 0}$. The terminology is clear since, by (1.5), it holds

$$\int_{\mathcal{X}} P_t \varphi(x) \mu(dx) = \int_{\mathcal{X}} \varphi(x) \mu(dx). \quad (1.7)$$

In other words, $\mathbb{E}_\mu[\varphi(X_t)] = \mathbb{E}_\mu[\varphi(X_0)]$, which indeed means that μ is invariant under the flow of $(X_t)_{t \geq 0}$. The above relation can also be written in differential form as

$$\int_{\mathcal{X}} (\mathcal{L}\varphi) d\mu = 0, \quad (1.8)$$

for any smooth test function φ . This is actually the dual formulation of (1.6). Finally, another consequence of (1.6) is that, for any measurable set $A \subset \mathcal{X}$ it holds

$$\mathbb{P}_\mu(X_t \in A) = \mathbb{P}_\mu(X_0 \in A), \quad (1.9)$$

where \mathbb{P}_μ denotes the probability measure under all the realizations of the dynamics (1.1) when X_0 is distributed according to μ . We insist on the fact that, in (1.9), the position X_t of the process at time $t > 0$ is in general different from X_0 for a particular realization of (1.1), but the overall probability measure is unchanged. This is the crux of such a statistical description of matter.

For a given dynamics such as (1.1), it is generally a difficult task to prove the existence of an invariant measure μ and the relaxation of the law of the process towards this stationary state. Interestingly, the proofs generally rely on energetic criteria, for which the existence of a “decreasing in average” energy ensures stability. We will come back to this remark in Section 1.1.2. However, the existence of an invariant measure or, say, stationary behaviour, is not the only interest of physical systems. In particular, dynamical properties play a crucial role in the understanding of various physical systems; we mention for instance autocorrelations, entropy production, transport coefficients, kinetic transition rates and macroscopic fluctuations theory [99, 422, 36, 123, 171, 392, 34, 34]. In particular, dynamics with the same invariant measure may have very different dynamical properties, as illustrated in the sequel.

1.1.1.3 Reversibility and irreversibility

In order to understand more precisely the dynamical features of a physical system, it is interesting to introduce the notion of *reversibility*. There is a long history and literature on the subject, and we will content ourselves with a brief mathematical description in our context. For this, we assume that there exists an invariant measure μ for the SDE at hand and introduce the Hilbert space of square integrable functions with respect to μ ,

$$L^2(\mu) = \left\{ \varphi : \mathcal{X} \rightarrow \mathbb{R} \text{ measurable} \mid \int_{\mathcal{X}} |\varphi|^2 d\mu < +\infty \right\}.$$

We next introduce the adjoint \mathcal{L}^* of \mathcal{L} with respect to μ , which is defined by the following equality: for any smooth functions φ, ψ ,

$$\int_{\mathcal{X}} (\mathcal{L}\varphi)\psi d\mu = \int_{\mathcal{X}} \varphi(\mathcal{L}^*\psi) d\mu. \quad (1.10)$$

Similarly, for the evolution operator $(P_t)_{t \geq 0}$, it holds

$$\forall t \geq 0, \quad \int_{\mathcal{X}} (P_t \varphi)\psi d\mu = \int_{\mathcal{X}} \varphi(P_t^* \psi) d\mu. \quad (1.11)$$

With this notation and using (1.10), the invariance relation (1.6) defining μ reads

$$\mathcal{L}^* \mathbb{1} = 0,$$

where $\mathbb{1}$ is the constant function equal to 1. The reversibility of a dynamics entails an invariance of the evolution of the law of the process with respect to time reversal, when the process is initialized according to the invariant measure. More precisely, we say that $(X_t)_{t \geq 0}$ is *reversible* with respect to μ if

$$\mathcal{L}^* = \mathcal{L}.$$

This terminology can be understood by reformulating (1.11) as: for any $x, y \in \mathcal{X}$ and infinitesimal elements $dx \subset \mathcal{X}$, $dy \subset \mathcal{X}$:

$$P_t(x, dy)\mu(dx) = P_t(y, dx)\mu(dy). \quad (1.12)$$

This means that, if $X_0 \sim \mu$ (and hence $X_t \sim \mu$ by definition of the invariant measure), the law of $(X_s)_{0 \leq s \leq t}$ coincides with that of $(X_{t-s})_{0 \leq s \leq t}$, which shows some invariance with respect to the arrow of time. If $\mathcal{L}^* \neq \mathcal{L}$, we say that $(X_t)_{t \geq 0}$ is *irreversible*. This situation generally characterizes the presence of a *current* (typically particles or energy), which may be due to an electric field or an inhomogeneous heat bath [194, 35, 49, 96]. We further distinguish two types of irreversibility in Remark 1.1 below.

In general, it is quite clear that studying irreversible dynamics is more difficult than studying reversible ones. This can be seen for instance at the level of the generator by noting that the theory of self-adjoint operators can be used for \mathcal{L} considered over $L^2(\mu)$ in the reversible case [304]. Note that we will sometimes use the terminology of *equilibrium* and *nonequilibrium* dynamics for the above notions when no confusion arises.

1.1.1.4 Two important examples: overdamped and underdamped Langevin dynamics

In order to further motivate modelling with SDEs and give more substance to the above definitions, we present two particular models that will be frequently come across through the various chapters.

Overdamped Langevin dynamics. Our first case study is the overdamped Langevin dynamics [304, 350]. An instance of such an equation is obtained from (1.1) by choosing $\mathcal{X} = \mathbb{R}^d$, $b = -\nabla V$ for some smooth potential V , $m = d$ and $\sigma = \sqrt{2\beta^{-1}}I_d$ where I_d is the identity matrix of size d , and $\beta > 0$. In this case, (1.1) becomes

$$dX_t = -\nabla V(X_t) dt + \sqrt{2\beta^{-1}} dB_t. \quad (1.13)$$

The interpretation of the above equation is as follows: the system strives to minimize the potential energy V through the drift term $-\nabla V$, while experiencing the influence of a thermal noise represented by a Brownian motion whose intensity is dictated by the inverse temperature $\beta > 0$ (in general it holds $\beta = 1/(k_B T)$ where T is the temperature and $k_B \approx 10^{-23}$ is the Boltzmann constant). This leads to a random behaviour where $(X_t)_{t \geq 0}$ spends more time in regions of low energy, while less time is spent in the regions of high energy. In particular, one can easily check that the invariant measure of $(X_t)_{t \geq 0}$ is given, under mild assumptions on V , by

$$\mu(dx) = Z^{-1} e^{-\beta V(x)} dx, \quad Z = \int_{\mathcal{X}} e^{-\beta V(x)} dx, \quad (1.14)$$

where Z is a normalization constant often called *partition function* or *free energy*, which is assumed to be finite for any $\beta > 0$. The measure (1.14), called *Gibbs measure*, is a pivotal element of equilibrium statistical mechanics [22]. Moreover, one can check that the generator $\mathcal{L} = -\nabla V \cdot \nabla + \beta^{-1} \Delta$ (where Δ denotes the Laplacian operator) is such that $\mathcal{L}^* = \mathcal{L}$, so the dynamics is reversible with respect to μ . As an illustration, we plot in Figure 1.1 the potential

$$V(x) = 3|x|^4 - 3|x|^2 + x \quad (1.15)$$

in dimension one, with its corresponding Gibbs measure at inverse temperature $\beta = 1$ as given by (1.14).

Underdamped Langevin dynamics. At small time scales, the dynamics (1.13) may not seem reasonable from a physical perspective because the system has no inertia. This is why practitioners consider instead the more realistic (underdamped) Langevin dynamics, in which particles are described by a position and a momentum. In this case, the state space is $\mathcal{X} = \mathbb{R}^d \times \mathbb{R}^d$ and we write $X_t = (q_t, p_t)$ where $q \in \mathbb{R}^d$ is the position and $p \in \mathbb{R}^d$ the momentum variable. The drift and diffusion coefficients are respectively set to

$$b(q, p) = \begin{pmatrix} p \\ -\nabla V(q) - \gamma M^{-1} p \end{pmatrix} \in \mathbb{R}^{2d}, \quad \sigma = \sqrt{\frac{2\gamma}{\beta}} \begin{pmatrix} 0 & 0 \\ 0 & I_d \end{pmatrix} \in \mathbb{R}^{2d \times 2d}, \quad (1.16)$$




Figure 1.1 – Representation of the potential (1.15) (left) with its associated Gibbs measure (1.14) at inverse temperature $\beta = 1$ (right).

where $M \in \mathbb{R}^{d \times d}$ is a mass matrix assumed to be invertible and $\gamma > 0$ is a friction. This leads to the equation

$$\begin{cases} dq_t = M^{-1} p_t dt, \\ dp_t = -\nabla V(q_t) dt - \gamma M^{-1} p_t dt + \sqrt{\frac{2\gamma}{\beta}} dB_t. \end{cases} \quad (1.17)$$

Typically, we consider a set of N particles in \mathbb{R}^3 , so that $d = 3N$ and the full dimension of (1.17) is $6N$.

Relation with Hamiltonian dynamics. The Langevin dynamics can be introduced as a perturbation of Newton's equation of motion [301, 350], which reads

$$\begin{cases} dq_t = M^{-1} p dt, \\ dp_t = -\nabla V(q_t) dt. \end{cases} \quad (1.18)$$

A non-equilibrium forcing could be added to the gradient force $-\nabla V$, but we assume for now that the force derives from the potential energy V . It is easily checked that the dynamics (1.18) conserves the Hamiltonian

$$H(q, p) = V(q) + \frac{p^T M^{-1} p}{2}, \quad (1.19)$$

where p^T denotes the transpose of $p \in \mathbb{R}^d$. By this we mean that, for all $t \geq 0$, it holds $H(q_t, p_t) = H(q_0, p_0)$. In the definition (1.19) of the Hamiltonian, the first term is a *potential energy* while the second is a *kinetic energy*; a Hamiltonian which is the sum of two such energies is often called *separable*. The dynamics (1.18) is actually a particular case of Hamiltonian system. A feature of such dynamics is that they preserve an energy, here the Hamiltonian defined in (1.19) (this is only one property out of many, and we refer to [215, 296] and references therein for more insight).

However, most physical systems conserve temperature rather than energy (think of a gas in a box in contact with a thermostat). This motivates adding to the force $-\nabla V$ a noise corresponding to an exchange of energy with a heat bath at fixed temperature. To compensate this input of energy, the particles are assumed to be damped by a friction effect. This way of fixing the temperature is related to the notion of *canonical ensemble* [22, 405]. In (1.17), the thermal noise is then represented by a Brownian motion at inverse temperature $\beta > 0$, while the damping is given by $-\gamma M^{-1} p_t$. This is why, in (1.17), the part $-\gamma M^{-1} p_t dt + \sqrt{2\gamma/\beta} dB_t$ is generally referred to as *fluctuation-dissipation* part, while the part corresponding to (1.18) is called *Hamiltonian*. In particular, we may decompose the generator of the dynamics according to

$$\mathcal{L} = \mathcal{L}_{\text{ham}} + \gamma \mathcal{L}_{\text{FD}},$$

with

$$\mathcal{L}_{\text{ham}} = M^{-1} p \cdot \nabla_q - \nabla V \cdot \nabla_p \quad \text{and} \quad \mathcal{L}_{\text{FD}} = -M^{-1} p \cdot \nabla_p + \beta^{-1} \Delta_p.$$

In the above decomposition, the operator \mathcal{L}_{ham} (sometimes called Liouville operator) corresponds to the Hamiltonian part of the dynamics, while \mathcal{L}_{FD} stands for the fluctuation dissipation part². The above choice of parameters ensures that the Gibbs measure

$$\mu(dqdp) = Z^{-1} e^{-\beta H(q,p)} dqdp, \quad Z = \int_{\mathcal{X}} e^{-\beta H(q,p)} dqdp \quad (1.20)$$

is an invariant probability measure of the dynamics, assuming again sufficient regularity and integrability for any $\beta > 0$. One can prove, under mild assumptions, that (1.20) is actually the only invariant measure of the dynamics, see [420, Proposition 15] for a proof.

We observe now that the marginal of μ in the position variable is the invariant measure (1.14) of the overdamped dynamics (1.13). However, (1.17) has very different dynamical properties. In particular it holds

$$\mathcal{L}_{\text{ham}}^* = -\mathcal{L}_{\text{ham}}, \quad \mathcal{L}_{\text{FD}}^* = \mathcal{L}_{\text{FD}},$$

where adjoints are taken in $L^2(\mu)$ for the Gibbs measure μ defined in (1.20). This particular structure shows that, in terms of operators, the dynamics is irreversible but *reversible up to momentum reversal*. In other words, if \mathcal{R} denotes the momentum reversal operator $\mathcal{R}\varphi(q,p) = \varphi(q,-p)$, it holds

$$\mathcal{L}^* = \mathcal{R}\mathcal{L}\mathcal{R}.$$

This property is used in Part IV to design efficient sampling algorithms from (1.17), and also appears in Chapter 3 when dealing with the dynamical entropy.

Another particular feature of the Langevin dynamics (1.17) is that it does not directly dissipate energy in position, but only in momenta through the friction term $-\gamma M^{-1}p_t$. The Hamiltonian part \mathcal{L}_{ham} serves (among other things) as a transmission of the energy dissipation from momenta to positions, a phenomenon detailed in Part II.

Relation between overdamped and underdamped Langevin dynamics. Before closing this section, we mention that the overdamped dynamics (1.13) can be obtained from (1.17) when $\gamma \rightarrow +\infty$ under an appropriate rescaling. To see this, we integrate the second line in (1.17) to obtain

$$p_t - p_0 = - \int_0^t \nabla V(q_s) ds - \gamma(q_t - q_0) + \sqrt{2\gamma} B_t.$$

By introducing now $Q_t^\gamma = q_{\gamma t}$ and $P_t^\gamma = p_{\gamma t}$, this becomes

$$Q_t^\gamma - Q_0^\gamma = \frac{P_0^\gamma - P_t^\gamma}{\gamma} - \int_0^t \nabla V(Q_s^\gamma) ds + \sqrt{2} B_t.$$

When $\gamma \rightarrow +\infty$, we observe that Q_t^∞ converges formally towards the solution of (1.13), see [350, Section 6.5]. This short computation explains the denomination *overdamped* for (1.13), since the dynamics (1.13) can be derived as a large friction limit of the finite friction Langevin dynamics (1.17), provided one considers the relevant time scale.

Remark 1.1 (Reversibility and irreversibility). *We have seen that the overdamped and underdamped Langevin dynamics presented above both admit a Gibbs measure as an invariant measure. While the overdamped dynamics (1.13) is reversible, the Langevin dynamics (1.17) is irreversible but reversible up to momentum reversal. In the general situation when $(X_t)_{t \geq 0}$ satisfies (1.1), we may prove under some conditions that the process has an invariant measure μ satisfying (1.6) without further information on the reversibility of the dynamics. Although there is no clear classification of irreversibility, we would like to say that the Langevin dynamics is “not so much irreversible”, because the symmetry breaking of the operator is only due to the direction of the momenta. This motivates the following personal distinction between degrees of irreversibility:*

1. *The invariant measure is a Gibbs measure and the dynamics is reversible with $\mathcal{L}^* = \mathcal{L}$ (for instance (1.13));*

²The reader familiar with stochastic processes will notice that the fluctuation-dissipation part is an Ornstein-Uhlenbeck process in momenta, with variance prescribed by β^{-1} and where γ is a time scale parameter.

2. The invariant measure is a Gibbs measure but the dynamics is irreversible, i.e. $\mathcal{L}^* \neq \mathcal{L}$. Local modifications in the drift lead to local deformations of the invariant measure. The dynamics (1.17) enters this definition, but also modified versions of the overdamped dynamics where a divergence-free drift is added, see [249, 363, 146];
3. The invariant measure is unknown and a priori $\mathcal{L}^* \neq \mathcal{L}$, which is the general situation (1.1). Local modifications of the drift lead to nonlocal deformations of the invariant measure, which is typically due to the presence of currents.

Although the overdamped and underdamped Langevin dynamics are particular cases of the more general SDE (1.1), because of their physical significance we will repeatedly come back to them in order to illustrate our theoretical and numerical results. Moreover, many generalizations of these equations, and even of (1.1) are available. Concerning the Langevin equation, we can be interested in generalized kinetic energies, which can be used *e.g.* for designing more efficient sampling schemes [315, 393]. Another possible way to make the Langevin dynamics more general is to add a *memory kernel* to the evolution of the momentum. In this case the dynamics is not Markovian any more but appropriate changes of variables can, in some important cases, turn it into an SDE like (1.1). We refer for instance to [435, 4, 78, 297] for more background on this situation. As a last example, we mention Langevin-like dynamics where the fluctuation-dissipation mechanism is replaced by a resampling of momenta dictated by a Poisson process. These processes have been introduced in biology to model the motion of bacteria with random velocity flips, but found interesting applications in computational statistics, see [168, 64, 41, 152] for more information. Although these models do not strictly enter the framework described by (1.1), they can also be studied by the techniques used in this work. Yet other models can be discussed (Hawkes processes or fractional Brownian motion for instance) but we stop the discussion here for conciseness.

In all cases, the long time behaviour of the system at hand is an important feature on which we would like to provide more information. The next section presents some of these aspects, which are related to the crucial notion of *ergodicity*.

1.1.2 Ergodicity and fluctuations

1.1.2.1 Ergodicity

We consider in this section the long time behaviour of (1.1) starting some definitions. From the expression (1.4) of the semigroup associated with the dynamics, we expect that for any φ in a class of sufficiently smooth and integrable functions it holds

$$\forall x \in \mathcal{X}, \quad \mathbb{E}_x[\varphi(X_t)] \xrightarrow{t \rightarrow +\infty} \int_{\mathcal{X}} \varphi d\mu,$$

where μ solves (1.6). This is a result of *convergence in law*. Although it provides an important information, this convergence may be inconvenient because the rate of convergence may depend on the observable φ . This is why convergence results generally take the form [362, 219]

$$\sup_{\|\varphi\| \leq 1} \|P_t \varphi - \mu(\varphi)\| \leq C e^{-ct}, \quad (1.21)$$

for some constants $C, c > 0$, where $\|\cdot\|$ is an appropriate functional norm. As discussed below, $\|\cdot\|$ is typically a weighted supremum norm³, and the convergence (1.21) is called in this case convergence in weighted total variation norm.

A problem with (1.21) is that, in practice, we cannot compute $\mathbb{E}_x[\varphi(X_t)]$. In a simulation, we rather use the long time average

$$L_t(f) = \frac{1}{t} \int_0^t f(X_s) ds,$$

for a function f belonging to some functional space, and for which we expect

$$L_t(f) \xrightarrow{t \rightarrow +\infty} \int_{\mathcal{X}} f d\mu, \quad (1.22)$$

³Another standard approach is to consider functions belonging to an appropriate Lebesgue or Sobolev space.

in some probabilistic sense, typically almost surely with respect to $(X_t)_{t \geq 0}$ (we use different notation for the test functions in (1.21) and (1.22) for reasons that will appear clearer later on). Such convergence results can be obtained under recurrence assumptions when $f \in L^1(\mu)$, see [323, 269] for details. We emphasize that, in (1.22) the temporal average becomes a spatial average, a key feature of *ergodicity*: this entails that the proportion of time that the process spends in a region of space is dictated by the invariant measure.

The convergence result (1.22) is the most important one for practical applications since it genuinely corresponds to what is computed during a numerical simulation. However, one may be interested in the convergence of the empirical mean (1.22) independently of the observable f at hand. To do so, we consider the empirical average

$$L_t = \frac{1}{t} \int_0^t \delta_{X_s} ds, \quad (1.23)$$

where δ_x is the Dirac mass at $x \in \mathcal{X}$. Therefore, for any $t \geq 0$, L_t is a random probability measure. In what follows, we denote by $\mathcal{P}(\mathcal{X})$ the space of probability measures, for which a topology should be specified (this topology issue will be treated at length in Chapter 3). If $\mu \in \mathcal{P}(\mathcal{X})$ is the invariant measure of $(X_t)_{t \geq 0}$, we typically expect

$$L_t \xrightarrow[t \rightarrow +\infty]{} \mu, \quad (1.24)$$

almost surely in the topology on $\mathcal{P}(\mathcal{X})$, or in another appropriate probabilistic sense. The convergence (1.24) then characterizes the evolution of the empirical average independently of the choice of observable, and thus the relaxation of the process towards its stationary state.

There are in general three types of conditions to fulfill in order to obtain the convergence (1.21) or (1.22) (see [362] for an overview):

1. Regularity of the transition density, as provided by Hörmander's condition [244, 420, 161];
2. Controllability or irreducibility, *i.e.* the system should be able to visit all the state space [396, 394, 362]. This is often stated as $\mathbb{P}_x(X_t \in A) > 0$ for all $x \in \mathcal{X}$ and all non-empty open set A ;
3. Confinement of the dynamics, to avoid probability mass to go to infinity [330, 20, 219].

In the situations considered here, the first two conditions are generally satisfied, and the most important problem is to prove some confinement of the dynamics. When considering the convergence in law, since the density of $(X_t)_{t \geq 0}$ satisfies the partial differential equation (PDE) (1.5), there is a vast literature on the long time behaviour of PDEs for convergence in Sobolev spaces [238, 220, 420, 304, 129, 295], which also has an entropic counterpart [52, 420, 308, 172]. These methods are useful not only because they allow to derive convergence results such as (1.22), but because they often come with rather explicit convergence rates. This is helpful in a number of situations (see [308] for an example of application to hydrodynamic limits of nonequilibrium evolutions).

However, methods based on PDEs sometimes lack flexibility and do not apply to all the situations of interest, which motivates considering Lyapunov function techniques. In the following, we say that a dynamics with generator \mathcal{L} satisfies a Lyapunov condition if there exist a function $W : \mathcal{X} \rightarrow [1; +\infty)$ (going to infinity at infinity) and constants $a > 0$, $C \in \mathbb{R}$ such that

$$\mathcal{L}W \leq -aW + C. \quad (1.25)$$

This energetic condition has a number of consequences (return time to compact sets, existence of an invariant measure and ergodicity, compactness of the evolution operator, large deviations, tightness of the evolution operator, etc.), and we refer to [330, 427, 324, 217, 362, 219] for more insight. The analysis carried out in Part II relies on this technique to exploit the confinement of the dynamics.

In order to make these ideas more precise, we introduce the space⁴

$$B^\infty(\mathcal{X}) = \left\{ \varphi : \mathcal{X} \rightarrow \mathbb{R} \text{ measurable} \mid \sup_{x \in \mathcal{X}} |\varphi(x)| < +\infty \right\}, \quad (1.26)$$

⁴The standard notation in probability textbooks would be \mathcal{B}_b , the space of bounded measurable functions [124]. However, since we use a weight W , we do not want b to be interpreted as a Lyapunov function, and we prefer the notation B^∞ , which is reminiscent of the space L^∞ of essentially bounded functions. On the other hand, we keep the notation \mathcal{B} for the space of bounded operators over a Banach space.

with norm $\|\varphi\|_{B^\infty} = \sup_{\mathcal{X}} |\varphi|$. When a Lyapunov function W is available, it is natural to introduce the space of functions growing at most like W at infinity, namely

$$B_W^\infty(\mathcal{X}) = \left\{ \varphi : \mathcal{X} \rightarrow \mathbb{R} \text{ measurable} \mid \sup_{x \in \mathcal{X}} \frac{|\varphi(x)|}{W(x)} < +\infty \right\},$$

with norm $\|\varphi\|_{B_W^\infty} = \|\varphi/W\|_{B^\infty}$. In order to understand why the space $B_W^\infty(\mathcal{X})$ is appropriate, we recall that the semigroup of the dynamics is defined by $(P_t\varphi)(x) = \mathbb{E}_x[\varphi(X_t)]$. Therefore, integrating (1.25) with a Grönwall lemma shows that, for all $t \geq 0$,

$$\forall x \in \mathcal{X}, \quad P_t W(x) \leq e^{-at} W(x) + \frac{C}{a},$$

where C and a are as in (1.25). The meaning of this condition is that the energy W remains bounded in expectation along the dynamics. If W has compact level sets, this implies that the process $(X_t)_{t \geq 0}$ spends most of its time in a compact set. This condition shows for instance that $(P_t)_{t \geq 0}$ is a semigroup of bounded operators on $B_W^\infty(\mathcal{X})$ [351, 259]. The method can also be applied to prove the existence of a strong solution to (1.1) for arbitrary long times, see [51] for an interesting application to systems interacting via a singular potential.

In order to provide an example of such Lyapunov function, we come back to the overdamped dynamics over $\mathcal{X} = \mathbb{R}^d$ presented in the previous section, namely

$$dX_t = -\nabla V(X_t) dt + \sqrt{2} dB_t,$$

with generator $\mathcal{L} = -\nabla V \cdot \nabla + \Delta$, for a smooth potential V . Bearing in mind that a Lyapunov function should correspond to some energy in the system, it is natural to use the potential energy V to construct it. As a matter of example (which will be used in Part II) we may set, for all $x \in \mathbb{R}^d$,

$$W(x) = e^{\theta V(x)},$$

for some $\theta \in (0, 1)$. In this case, a simple computation shows that (see Chapter 3)

$$-\frac{\mathcal{L}W}{W} = \theta(1-\theta)|\nabla V|^2 - \theta\Delta V. \quad (1.27)$$

Moreover, if we set $\theta = 1/2$, we obtain

$$-\frac{\mathcal{L}W}{W} = \frac{1}{4}|\nabla V|^2 - \frac{1}{2}\Delta V,$$

which is the potential of the Witten Laplacian associated to V , see [425, 236, 20, 304] and references therein. From (1.27), we obtain (1.25) under appropriate *growth condition* at infinity on V , typically when V grows like $|x|^q$ at infinity with $q > 1$. This explains the terminology of *confinement* (see also Remark 3.11 in Chapter 3).

Building a Lyapunov function for the Langevin dynamics

$$\begin{cases} dq_t = M^{-1}p_t dt, \\ dp_t = -\nabla V(q_t) dt - \gamma M^{-1}p_t dt + \sqrt{\frac{2\gamma}{\beta}} dB_t, \end{cases}$$

is a more difficult task, because the damping only acts in the momentum variable. As a result, the natural choice $W(q, p) = H(q, p)$ fails to provide a Lyapunov function when the position space is unbounded. In order to transfer the damping from momenta to positions, a common strategy is to introduce a cross term by setting, for all $(q, p) \in \mathbb{R}^d \times \mathbb{R}^d$,

$$W(q, p) = H(q, p) + \varepsilon p \cdot q,$$

or

$$W(q, p) = \exp(\theta H(q, p) + \varepsilon p \cdot q),$$

for an appropriate choice of parameters $\theta, \varepsilon > 0$. To our knowledge, this strategy was first introduced in the early 2000s independently by [427, 399, 324], and allows to deduce (1.25) under appropriate growth condition on V (see Lemma 3.23 in Chapter 3 for a precise statement).

1.1.2.2 Small fluctuations and Central Limit Theorem

In practice (1.22) is computed with a finite simulation time $t > 0$, and it is therefore natural to quantify the statistical error arising from this estimator. By the central limit theorem, we expect errors of order $1/\sqrt{t}$, and we thus magnify the difference between the empirical average and the ergodic average at scale \sqrt{t} . We can then prove, for $f \in L^2(\mu)$ and provided the Poisson equation $-\mathcal{L}\psi = f - \mu(f)$ admits a solution in $L^2(\mu)$, that the following central limit theorem holds [38]:

$$\sqrt{t} \left(\frac{1}{t} \int_0^t f(X_s) ds - \int_{\mathcal{X}} f d\mu \right) \xrightarrow[t \rightarrow +\infty]{} \mathcal{N}(0, \sigma_f^2), \quad (1.28)$$

in law, where $\mathcal{N}(0, \sigma_f^2)$ is a centered normal random variable with variance $\sigma_f^2 > 0$. This *asymptotic variance* is defined by

$$\sigma_f^2 = \lim_{t \rightarrow +\infty} t \mathbb{E}_x \left[\left(\frac{1}{t} \int_0^t f(X_s) ds - \int_{\mathcal{X}} f d\mu \right)^2 \right]. \quad (1.29)$$

An application of the Itô formula actually leads to the alternative expression

$$\sigma_f^2 = -2 \int_{\mathcal{X}} \psi \mathcal{L}\psi d\mu, \quad \text{where} \quad -\mathcal{L}\psi = f - \mu(f), \quad (1.30)$$

where we used the shorthand notation

$$\mu(f) = \int_{\mathcal{X}} f d\mu.$$

Let us explain briefly why (1.30) holds. First, under appropriate conditions on \mathcal{L} , the Poisson equation in (1.30) admits a solution in $L^2(\mu)$ for $f \in L^2(\mu)$, see [349, 277, 304] and references therein. In this case, we can formally apply Itô's formula to ψ for the dynamics $(X_t)_{t \geq 0}$, which gives

$$d\psi(X_s) = \mathcal{L}\psi(X_s) ds + \nabla\psi(X_s) \cdot \sigma(X_s) dB_s.$$

By integrating in time we obtain

$$\psi(X_t) - \psi(X_0) = \int_0^t \mathcal{L}\psi(X_s) ds + \int_0^t \nabla\psi(X_s) \cdot \sigma(X_s) dB_s.$$

Using that ψ is solution to $-\mathcal{L}\psi = f - \mu(f)$ leads to

$$\sqrt{t} \left(\frac{1}{t} \int_0^t f(X_s) ds - \int_{\mathcal{X}} f d\mu \right) = \frac{\psi(X_t) - \psi(X_0)}{\sqrt{t}} + \frac{1}{\sqrt{t}} \int_0^t \nabla\psi(X_s) \cdot \sigma(X_s) dB_s.$$

The first term on the right hand side of the above equation can be shown to converge to zero in the sense of square integrability in probability as $t \rightarrow +\infty$. The second term can be treated with theorems of convergence of martingales [178], leading to the convergence to a normal distribution with variance given by

$$\int_{\mathcal{X}} |\sigma^T \nabla\psi|^2 d\mu.$$

One can check that this expression indeed matches (1.30) (with computations similar to that performed at the beginning of Chapter 3), and this strategy allows to obtain the central limit theorem (1.28). This framework can be generalized when f does not belong to $L^2(\mu)$ but alternatively to the dual of an appropriate Hilbert space, and we refer to [273] for details (note that the framework developed in this book will be used in Chapter 3).

In order to provide more intuition on the asymptotic variance, we believe that it is worthwhile to express it in terms of an autocorrelation. For this we first note that, when the semigroup $P_t = e^{t\mathcal{L}}$ has appropriate decay estimates [304, 129], the inverse of the generator \mathcal{L} can be written as

$$\mathcal{L}^{-1} = - \int_0^{+\infty} e^{t\mathcal{L}} dt.$$

Plugging this expression in (1.30), we obtain

$$\sigma_f^2 = 2 \int_0^{+\infty} \mathbb{E}_\mu [\Pi f(X_t) \Pi f(X_0)] dt,$$

where we denoted by $\Pi f = f - \mu(f)$. The above formula can also be obtained by directly expanding (1.29), see [304] for more details. This shows that an observable with a rapidly decaying autocorrelation has a small variance, while high positive correlations lead to a large variance. This heuristic can be used to design more efficient sampling algorithms, for example by adding an irreversible drift to the dynamics [248, 249, 146, 363]. We also mention that it is possible to control the convergence towards the central limit theorem by higher order techniques such as Berry–Esseen bounds, see *e.g.* [343] and references therein.

1.1.2.3 Large fluctuations

A particular issue with physical systems is, as we mentioned in Section 1.1, the very different time scales at stake in the dynamics. As a result, it is often interesting to consider very large fluctuations from the expected mean, since these fluctuations may not be so rare on another scale. From a probabilistic point of view, this suggests considering fluctuations beyond the \sqrt{t} -scaling of (1.28). One way to do so is to consider the whole probability distribution of $L_t(f)$, namely

$$\mathbb{P}(L_t(f) \in [a; a + da]),$$

for $a \in \mathbb{R}$. The general framework of large deviations presented in Section 1.2 below suggests that these probabilities should decay exponentially fast, *i.e.*

$$\mathbb{P}(L_t(f) \in [a; a + da]) \approx e^{-tI_f(a)+o(t)} da, \quad (1.31)$$

as $t \rightarrow +\infty$. The function $I_f : \mathbb{R} \rightarrow [0, +\infty]$ is generally called *rate function* (or dynamical entropy in this context), and governs the rate of decay of the probability of excursions from the average behaviour. The scaling (1.31) proves to be important in physics, for instance for studying irreversibility or phase transitions, see [48, 197, 198, 406, 170, 411] for recent accounts. From a more pragmatic perspective, computing I_f provides asymptotic estimates for the probabilities of fluctuations.

From (1.31), we obtain the following heuristic expression for I_f :

$$\forall a \in \mathbb{R}, \quad I_f(a) = - \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}(L_t(f) \in [a; a + da]),$$

to which a precise sense can be given through the large deviations theory presented in Section 1.2. Interestingly, a scaling such as (1.31) generalizes the central limit theorem (1.28). To see this, we observe that in the «standard» situation, I_f is convex and such that

$$I_f(\mu(f)) = 0, \quad I_f'(\mu(f)) = 0, \quad I_f''(\mu(f)) > 0.$$

This is because we expect I_f to vanish only at the ergodic average for (1.31) to make sense, if μ is the unique invariant measure. Therefore, a second order Taylor expansion in (1.31) shows that, for small $\varepsilon > 0$, it holds

$$\mathbb{P}(L_t(f) \in [\mu(f) - \varepsilon; \mu(f)]) \approx e^{-t \frac{I_f''(\mu(f))}{2} \varepsilon^2 + o(\varepsilon^2 t)} d\varepsilon,$$

meaning that small fluctuations are Gaussian with variance $1/I_f''(\mu(f))$, see [363] for details. A result such as (1.31) is called a *large deviations principle*, and illustrates the fact that large fluctuations are *a priori* non-Gaussian. Further explanations and examples are provided in Sections 1.2 and 1.3.

From a physical perspective and interpreting I_f as an entropy, it seems natural to introduce a free energy function λ_f which is the Legendre–Fenchel transform of I_f , namely

$$\forall k \in \mathbb{R}, \quad \lambda_f(k) = \sup_{a \in \mathbb{R}} \{ak - I_f(a)\}. \quad (1.32)$$

We can actually prove in many cases that

$$\lambda_f(k) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right]. \quad (1.33)$$

This expression is motivated by the Laplace principle⁵, for which we assume that the scaling (1.31) holds. We then rewrite the expectation in (1.33) with

$$\mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right] = \mathbb{E} \left[e^{tkL_t(f)} \right] = \int_{\mathbb{R}} e^{tka} \mathbb{P}(L_t(f) \in da) = \int_{\mathbb{R}} \exp \left(t(ka - I_f(a)) + o(t) \right) da.$$

By applying the Laplace principle, the above formula becomes, as $t \rightarrow +\infty$,

$$\mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right] \approx \exp \left(t \sup_a \{ak - I_f(a)\} + o(t) \right).$$

Plugging this estimate into (1.33) leads to (1.32). As a result, the classical Legendre duality between free energy and entropy appears here as an application of the Laplace principle at an exponential scale [406, 408]. Once again, the large deviations theory presented in Section 1.2 provides a mathematical framework to justify this heuristic while allowing for important generalizations.

Remark 1.2. *We motivate studying the long time behaviour of SDEs from physical consideration, but it is also an important problem in computational statistics. For Bayesian inference, one is interested in sampling the posterior distribution π of a set of parameters given data, which is obtained from the likelihood and the prior distribution. In order to sample from the target π one can run the dynamics (1.13) or (1.17) with the potential $V = -\log \pi$, i.e. minus the log-likelihood (possibly with an additional term given by the prior). In this perspective, a Gaussian prior amounts to adding a quadratic confinement to the dynamics. In this context, kinetic dynamics such as (1.17) are generally called Hamiltonian Monte Carlo methods, and we refer for instance to [367, 70, 368] for more details on the applications in statistics of the dynamics discussed here.*

We also mention that we present continuous systems in the introduction in order to avoid repetitions, but most of the objects have a discrete time counterpart (i.e. we can consider Markov chains instead of diffusion processes). We will extensively study discrete time dynamics in Chapter 2 for Feynman–Kac models as well as in Chapters 4 and 5 when considering discretizations of SDEs.

1.1.3 Systems at small temperature

The previous section was motivated by the long time scale at which physical phenomena occur. In practice, another important feature of real life systems is *metastability*. Roughly speaking, this phenomenon is produced by the many basins of attractions around local minima of the interaction potential of the dynamics. The dynamics then tends to spend a long time in a region of low energy, before making a rare transition to another region of low energy, where it remains trapped for a long time.

In order to illustrate this discussion we consider the dynamics (1.1) for a smooth function b and $\sigma = \sigma I_d$ a constant temperature field. If $b = \nabla V$ with a potential V which has many local minima, the dynamics $(X_t)_{t \geq 0}$ tends to stay close to one of those, before the noise induced by the Brownian motion transfers enough energy in the system so as to drive it outside of this region; the same reasoning applies when b possibly has a non-gradient part.

1.1.3.1 Low temperature and concentration

An additional difficulty arises in the low temperature regime, where the transitions become rarer and rarer. This motivates considering the following version of (1.1) over $\mathcal{X} = \mathbb{R}^d$:

$$dX_t^\varepsilon = b(X_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma dB_t, \quad (1.34)$$

where $b : \mathcal{X} \rightarrow \mathbb{R}^d$ is smooth and $\sigma \in \mathbb{R}^{d \times d}$ is constant, while $\varepsilon > 0$ denotes a small temperature parameter. Moreover we write $D = \sigma \sigma^T$ for the diffusion matrix, assumed to be positive definite for

⁵The method of Laplace states that, when integrating the exponential of a function f , the main contribution comes from the maximum of f . Assume for instance that $a < b$ and $f \in C^2(\mathbb{R})$ admits a maximum at $x_0 \in (a, b)$. Then one can prove with a Taylor expansion that

$$\int_a^b e^{tf(x)} dx \underset{t \rightarrow +\infty}{\sim} e^{tf(x_0)}.$$

In other words

$$\lim_{t \rightarrow +\infty} \frac{1}{t} \log \left(\int_a^b e^{tf(x)} dx \right) = f(x_0).$$

The large deviations framework presented in Section 1.2 is a generalization of this idea, which is useful in various situations.

simplicity. There is a very large literature on small temperature systems studying different aspects of the problem, for instance: typical rare trajectories [193, 413], diffusions on graphs for Hamiltonian systems [192, 193, 220], exit time from a metastable domain [69, 298], spectral analysis for small eigenvalues [236, 331], quasi-stationary distributions [101, 299, 125]. Since this is not the principal issue of the present work, we focus on the particular situation considered in Chapter 6.

For this, we fix a time $T > 0$ and consider the dynamics (1.34) over the time window $[0, T]$ and starting at an arbitrary position $x \in \mathcal{X}$. Our goal is to numerically compute free energy functions of the form⁶

$$Z_\varepsilon = \varepsilon \log \mathbb{E}_x \left[e^{\frac{f(X_T^\varepsilon)}{\varepsilon}} \right], \quad (1.35)$$

in particular in the small temperature regime $\varepsilon \rightarrow 0$. This is motivated by applications in physics where we expect this quantity to have a finite limit as $\varepsilon \rightarrow 0$ (see [157, 413, 210, 211] and references therein). In order to understand the small temperature behaviour of Z_ε , it is natural to study the typical trajectories of $(X_t^\varepsilon)_{t \in [0, T]}$ as $\varepsilon \rightarrow 0$. By setting $\varepsilon = 0$ in (1.34), we expect the dynamics to concentrate on a path $(\bar{\phi}_t)_{t \in [0, T]}$ solution to the ODE

$$\bar{\phi}_0 = x, \quad \frac{d}{dt} \bar{\phi}_t = b(\bar{\phi}_t), \quad (1.36)$$

and we expect that

$$\mathbb{P} \left(\lim_{\varepsilon \rightarrow 0} \|X^\varepsilon - \bar{\phi}\|_{C^0([0, T])} = 0 \right) = 1, \quad (1.37)$$

where

$$\|X^\varepsilon - \bar{\phi}\|_{C^0([0, T])} = \sup_{t \in [0, T]} |X_t^\varepsilon - \bar{\phi}_t|.$$

Actually, (1.37) is a consequence of [193, Theorem 1.1] in our setting, and there are quite precise results on the concentration rate of probabilities like (1.37).

As explained in [193, Chapter 3], one can actually prove, using Girsanov's theorem and Tchebychev's inequality, that the rate of decay of the left hand side of (1.37) is exponential and controlled by a so-called *action functional*. This functional $I : C^0(\mathcal{X}) \rightarrow [0; +\infty]$ is defined, for any absolutely continuous path $(\varphi_t)_{t \in [0, T]} \in \mathcal{X}$, by

$$I(\varphi) = \frac{1}{2} \int_0^T |b(\varphi_s) - \dot{\varphi}_s|_D^2 ds, \quad (1.38)$$

and set to $+\infty$ if φ is not absolutely continuous or if the above integral diverges. In the above formula, we denote by $|v|_D^2 = v \cdot D^{-1}v$ for $v \in \mathbb{R}^d$ the modified norm induced by the diffusion matrix $D = \sigma\sigma^T$, see [211]. It is then possible to derive the following kind of scaling, for any path $(\varphi_t)_{t \in [0, T]}$ and sufficiently small $\varepsilon, \delta > 0$:

$$\mathbb{P} \left(\|X^\varepsilon - \varphi\|_{C^0([0, T])} < \delta \right) \approx e^{-\varepsilon^{-1}I(\varphi)}. \quad (1.39)$$

The functional I can be interpreted as an energetic cost for not following the most probable path (1.36). Since the functional I vanishes on the path $(\bar{\phi}_t)_{t \in [0, T]}$ defined in (1.36), we see that (1.39) generalizes (1.37). In particular, $(X_t^\varepsilon)_{t \in [0, T]}$ concentrates exponentially fast on $(\bar{\phi}_t)_{t \in [0, T]}$.

With the aim of computing the small temperature limit of (1.35), we now turn to the probability for $(X_t^\varepsilon)_{t \in [0, T]}$ to reach a set $A \subset \mathcal{X}$. Because of the definition (1.35), we typically consider A to be a level set of the function f , namely $A = \{x \in \mathcal{X} \mid f(x) = z\}$ for some $z \in \mathbb{R}$. Intuitively, a path ending in A should minimize the functional I with the constraint that the endpoint lies in A , *i.e.* the minimization should be performed over continuous paths $(\varphi_t)_{t \in [0, T]}$ such that $\varphi_T \in A$. In particular, if $z \in \mathbb{R}$ and A is a z -level set of f , we introduce

$$C_z^0 = \left\{ \varphi \in C^0([0, T]; \mathcal{X}), \varphi_0 = x, f(\varphi_T) = z \right\},$$

to denote those paths whose endpoint takes the value z for the observable $f : \mathcal{X} \rightarrow \mathbb{R}$. For any $z \in \mathbb{R}$, we can therefore define the function

$$J(z) = \inf_{\varphi \in C_z^0} I(\varphi), \quad (1.40)$$

⁶We could also consider $Z_\varepsilon = \varepsilon \log \mathbb{E}_x \left[\exp \left(\varepsilon^{-1} \int_0^T f(X_s^\varepsilon) ds \right) \right]$ or suppose that T is a random time upon appropriate modifications.

which characterizes the minimum energy path for which f takes the value $z \in \mathbb{R}$ at final time. In other words we formally have, in the small ε limit,

$$\mathbb{P}(f(X_t^\varepsilon) \in dz) \approx e^{-\varepsilon^{-1}J(z)}. \quad (1.41)$$

The minimum energy path (*i.e.* the minimizer in (1.40), assuming uniqueness) is generally called *instanton* or *reaction path* [155, 156, 211]. It corresponds to the most likely path leading to a (possibly unlikely) value of f at final time.

We can now use the action functional to determine the small ε limit of (1.35) by characterizing the path contributing most to the expectation. Indeed, the expectation

$$\mathbb{E}_x \left[e^{\frac{f(X_T^\varepsilon)}{\varepsilon}} \right]$$

is dominated by the trajectories leading to large values of f , which are characterized by the minimization problem (1.40). Assuming (1.41) holds, we actually obtain in the small ε limit

$$\mathbb{E}_x \left[e^{\frac{f(X_T^\varepsilon)}{\varepsilon}} \right] = \int_{\mathbb{R}} e^{\frac{z}{\varepsilon}} \mathbb{P}(f(X_t^\varepsilon) \in dz) \approx e^{\varepsilon^{-1} \sup_z \{z - J(z)\}},$$

by the Laplace principle. This finally shows that

$$\lim_{\varepsilon \rightarrow 0} Z_\varepsilon = \sup_{z \in \mathbb{R}} \{z - J(z)\}. \quad (1.42)$$

This limit motivates the scalings chosen to define the process $(X_t^\varepsilon)_{t \in [0, T]}$ and the free energy Z_ε .

1.1.3.2 Instanton theory

Although (1.42) provides a formula for the zero temperature limit of the free energy, it may not be computationally tractable written this way. Our goal is now to present the so called *transition path theory*, which provides numerical methods for estimating the optimal path and free energy and found a number of interesting applications in physics [155, 156, 159, 158, 240, 211]. The idea is to rewrite the inf-sup problem (1.40)-(1.42) with a Lagrangian representation (in the sense of classical mechanics [211]). For this, we introduce the Lagrangian

$$\forall x, y \in \mathbb{R}^d, \quad L(x, y) = \frac{1}{2} |b(x) - y|_D^2, \quad (1.43)$$

so that I rewrites, for a path $(\varphi_t)_{t \in [0, T]}$,

$$I(\varphi) = \int_0^T L(\varphi_t, \dot{\varphi}_t) dt.$$

A minimizer of J in (1.40) is a path $(\phi_t)_{t \in [0, T]}$ solving the Euler–Lagrange equation

$$\frac{dL}{dx}(\phi_t, \dot{\phi}_t) - \frac{d}{dt} \frac{dL}{dy}(\phi_t, \dot{\phi}_t) = 0, \quad (1.44)$$

with appropriate boundary conditions (see below). The above equation and classical mechanics considerations motivate introducing the momentum variable $(\theta_t)_{t \in [0, T]}$:

$$\theta_t = \frac{dL}{dy}(\phi_t, \dot{\phi}_t), \quad (1.45)$$

which is associated to a path $(\phi_t)_{t \in [0, T]}$ solving (1.44). We also define the following Hamiltonian, for any $(x, z) \in \mathbb{R}^d \times \mathbb{R}^d$:

$$H(x, z) = \sup_{y \in \mathcal{X}} \{z \cdot y - L(x, y)\}. \quad (1.46)$$

This Hamiltonian is explicit when L is given by (1.43), and reads

$$\forall (x, z) \in \mathbb{R}^d \times \mathbb{R}^d, \quad H(x, z) = b(x) \cdot z + \frac{1}{2} z \cdot Dz. \quad (1.47)$$

Under appropriate convexity properties of L , it is also possible to invert (1.46) to obtain

$$\forall (x, y) \in \mathbb{R}^d \times \mathbb{R}^d, \quad L(x, y) = \sup_{z \in \mathcal{X}} \{y \cdot z - H(x, z)\}.$$

By optimizing the above equation in z , we find that the optimal z^* is provided by

$$y - \nabla_z H(x, z^*) = 0. \quad (1.48)$$

By (1.47), this equation is actually equivalent to

$$z^* = D^{-1}(y - b(x)) = \frac{d}{dy} L(x, y). \quad (1.49)$$

By applying the previous relations to the optimal path $(\phi_t)_{t \in [0, T]}$ (i.e. setting $x = \phi_t$, $y = \dot{\phi}_t$ and noting that $z^* = \theta_t$ by (1.45)-(1.49)), the optimality condition (1.48) provides a first equation for the optimal pathway $(\phi_t, \theta_t)_{t \in [0, T]}$:

$$\dot{\phi}_t = \nabla_z H(\phi_t, \theta_t) = b(\phi_t) + D\theta_t.$$

On the other hand, the Euler-Lagrange equation (1.44) provides the second relation (using (1.49) for the last equality)

$$\dot{\theta}_t = \nabla_x L(\phi_t, \dot{\phi}_t) = \nabla b(\phi_t) \cdot D^{-1}(b(\phi_t) - \dot{\phi}_t) = -\nabla b(\phi_t) \cdot \theta_t,$$

where $\nabla b \in \mathbb{R}^{d \times d}$. As a result, we obtain a coupled system for the optimal path and momentum $(\phi_t, \theta_t)_{t \in [0, T]}$ which reads, taking into account the boundary conditions:

$$\begin{cases} \dot{\phi}_t = b(\phi_t) + D\theta_t, & \phi_0 = x, \\ \dot{\theta}_t = -\nabla b(\phi_t) \cdot \theta_t, & \theta_T = \nabla f(\phi_T). \end{cases} \quad (1.50)$$

This shows that the maximization problem in (1.42) is attained for the value $z = f(\phi_T)$, while the minimization problem (1.40) is solved by the instanton. Therefore, plugging (1.50) into (1.40), we obtain

$$\lim_{\varepsilon \rightarrow 0} Z_\varepsilon = f(\phi_T) - \frac{1}{2} \int_0^T \theta_s \cdot D\theta_s ds. \quad (1.51)$$

In other words, the limiting free energy is the terminal value of f along the reaction path (final cost) minus the kinetic energy cost to reach it. Note that, similarly to Section 1.1.2, the above computations rely on an application of the Laplace principle at exponential scale.

Quite importantly, (1.50) not only provides the limit value (1.51) for the free energy, but also characterizes the path contributing most to the expectation in the small temperature regime. As mentioned above, this method found many interesting applications, because (1.50) can often be approximated efficiently in practical situations, see [211] for a review. This is because (1.50) is simply a system of ordinary differential equations with initial and terminal conditions (which may be high dimensional, though). A natural strategy to estimate the solution is to integrate forwards $(\phi_t)_{t \in [0, T]}$, then to integrate $(\theta_t)_{t \in [0, T]}$ backwards starting from the terminal condition $\nabla f(\phi_T)$, and to iterate the procedure. Other variants are also available, such as the so-called string method [155, 158].

It is important to note that (1.50) is a zero temperature limit, so that no randomness is involved. This greatly simplifies the computation of the optimal path, but makes the application of the method to finite temperature systems (1.34) difficult. We will come back to this point in Chapter 6 by relating the instanton theory presented here to a standard optimal control problem associated to (1.34)-(1.35).

Finally, we mention that the small temperature problem can be related to the long time setting of Section 1.1.2, see for instance [193, Chapter 6] and [67]. A very interesting application of this relation can be found in [342], where the small temperature instanton is used to characterize heavy tailed fluctuations in time.

1.1.4 Random matrices and Coulomb gases

In this section, we motivate the random gas models studied in Part IV, in particular the Coulomb gas setting. For this, we start with an experiment involving random matrices by considering a matrix $M \in$

$\mathbb{C}^{N \times N}$ for some $N \in \mathbb{N}^*$, such that $M_{i,j}$ are independent standard Gaussian random variables for all $i, j \in \{1, \dots, N\}$. These models arise in quantum physics, where M can be thought of as a discrete random Laplacian operator – we refer to [203, 9, 327, 188, 401] for more insight. We expect the spectrum of the matrix to provide information about the energy levels in the system when the number of elements becomes large. The behaviour of the eigenvectors of large random matrices is also important in various situations [282, 144] and is related to problems of spectral clustering in computational statistics [163, 102], but we do not consider these issues here.

1.1.4.1 Model: Ginibre Ensemble

As a numerical experiment, we draw one matrix M and plot in Figure 1.2 (left) the complex spectrum of M/\sqrt{N} for $N = 100$ (the scaling is motivated by the fact that $\mathbb{E}(\text{Tr}(M^*M)) = N$, where $M^* = \overline{M}^T$). We repeat the experiment with $N = 1000$ and plot the result in Figure 1.2 (right). We observe what seems to be a uniform distribution on the unit disk. Our goal is now to briefly explain with singular gas models why this surprising behaviour arises, and motivate the study of more general systems.




Figure 1.2 – Eigenvalues of M/\sqrt{N} where M has i.i.d. Gaussian entries, for $N = 100$ (left) and $N = 1000$ (right). The axes correspond to the real and imaginary parts of the eigenvalues. The right plot is shown together with the unit circle in red.

The law on the matrix M/\sqrt{N} (assuming the entries are independent complex Gaussian variables) is actually expressed in the following concise form:

$$e^{-N\text{Tr}(MM^*)} \prod_{i,j=1}^N dM_{i,j},$$

(we refer to [56, 188, 9] for more details). This model is generally called the (complex) *Ginibre Ensemble*, and has the particular feature of being invariant with respect to unitary transforms [327]. Therefore, denoting by $X_N = (X_{N,1}, \dots, X_{N,N}) \in \mathbb{C}^N$ the spectrum of M , we can show that X_N is distributed according to the following probability measure on \mathbb{C}^N (which we identify with $(\mathbb{R}^2)^N$):

$$P_N(dx) = Z_N^{-1} e^{-N \sum_{i=1}^N |x_i|^2} \prod_{1 \leq i < j \leq N} |x_i - x_j|^2 dx,$$

where Z_N is a normalization constant. In the above formula, the double product comes from a Vandermonde determinant in the change of variable, see [327, Chapter 15] for details, in particular Eq. (15.1.10). Writing this product in an exponential form, we obtain the distribution

$$P_N(dx) = Z_N^{-1} \exp \left[-N^2 \left(\frac{1}{N} \sum_{i=1}^N |x_i|^2 - \frac{1}{N^2} \sum_{i \neq j} \log |x_i - x_j| \right) \right] dx. \quad (1.52)$$

The reason for writing P_N as (1.52) is twofold. First, we see that the quantities inside parentheses are now intensive, so that we expect these quantities to have a finite limit when $N \rightarrow +\infty$. Moreover,

beyond this random matrix model, the measure P_N takes the form of a more general distribution over $(\mathbb{R}^d)^N$:

$$P_N(dx) = Z_N^{-1} e^{-\beta_N (\frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} K(x_i - x_j))} dx, \quad (1.53)$$

where V is a confinement potential, $K : \mathbb{R}^d \rightarrow \mathbb{R} \cup \{+\infty\}$ is an interaction kernel, and

$$Z_N = \int_{(\mathbb{R}^d)^N} e^{-\beta_N (\frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} K(x_i - x_j))} dx$$

is a normalizing constant assumed to be finite. The sequence $(\beta_N)_{N \geq 1}$ is a *cooling schedule*. Indeed, since P_N is a Gibbs measure, β_N plays the role of an inverse temperature. Assuming that $\beta_N \rightarrow +\infty$ then means that we consider a low temperature limit in addition to a large number of particles limit. The random matrix model described above (the *Ginibre Ensemble*) corresponds to

$$d = 2, \quad \beta_N = N^2, \quad V(x) = |x|^2, \quad K(x) = -\log |x|. \quad (1.54)$$

We shall assume below that the parameters of the models are as in (1.54), bearing in mind that more general situations are allowed.

1.1.4.2 Limiting distribution and electrostatics

From the observations of Figure 1.2, we expect the empirical measure of X_N to concentrate on a uniform distribution on the unit disc. This is actually a particular instance of a more general concentration property for the empirical average

$$\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i},$$

under P_N . Here, δ_x denotes the Dirac mass at $x \in \mathbb{R}^2$. A natural heuristic for understanding the concentration of μ_N is to note that

$$\frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} K(x_i - x_j) = \mathcal{E}_{\neq}(\mu_N),$$

where

$$\mathcal{E}_{\neq}(\mu) = \int_{\mathbb{R}^d} V(x) \mu(dx) + \iint_{x \neq y} K(x - y) \mu(dx) \mu(dy).$$

In this procedure, we see that a difficult problem is to deal with the singularity of K at 0, which makes the treatment of the diagonal terms in the double integral difficult. However, it is possible to address this issue [386, 84, 148, 314] and, by applying the Laplace principle in (1.53), we may guess that μ_N converges in some probabilistic sense to

$$\mu_{\star} \in \underset{\mathcal{P}(\mathbb{R}^d)}{\operatorname{argmin}} \mathcal{E}, \quad (1.55)$$

where

$$\mathcal{E} : \mu \in \mathcal{P}(\mathcal{X}) \mapsto \int V(x) \mu(dx) + \iint K(x - y) \mu(dx) \mu(dy) \quad (1.56)$$

is the electrostatic energy of a distribution of charges. A measure satisfying (1.55) is generally called an *equilibrium measure* (not to be confused with the notion of equilibrium discussed in Section 1.1.1). Once again, minimizing the energy \mathcal{E} is made difficult when K has singularities that make the energy \mathcal{E} blow up over Dirac masses, which is the case when $K = -\log |\cdot|$.

In order to identify the limiting measure μ_{\star} (assuming uniqueness for ease of exposition), a first order Euler–Lagrange analysis shows that [84, 85], formally,

$$2K * \mu_{\star} + V = C_{\star} \quad (1.57)$$

over the support of μ_{\star} , for the constant $C_{\star} = \mathcal{E}(\mu_{\star})$ (here, $K * \mu_{\star}$ denotes the convolution of K with μ_{\star} , see Chapter 8 for definitions). Therefore, this provides a characterization of μ_{\star} through

an integral equation. Note however that the support of μ_\star is unknown, which makes the resolution of (1.57) difficult apart from systems with very particular symmetries.

In the random matrix model leading to the choice (1.54), since $d = 2$ we observe that $K = -\log|\cdot| = g$, which is solution to

$$-\Delta g = 2\pi\delta_0, \quad (1.58)$$

in the sense of distributions, where Δ denotes the Laplacian operator. This means that $K = g$ is the Coulomb kernel in dimension 2 [386, 284]. Therefore, by differentiating (1.58) (under appropriate regularity of V) we obtain that

$$\mu_\star = \frac{\Delta V}{4\pi}$$

over the support of μ_\star . Since $V(x) = |x|^2$, the measure is constant and equal to $1/\pi$ over its support, which is a disk by symmetry (recall that $\Delta|x|^2 = 4$ when $d = 2$). Since the total mass of the measure μ_\star is one, the radius of the disk is equal to one. This explains the plots of Figure 1.2 showing the unit disk appearing with increasing dimension N . Note here the importance of the dimension: when $d = 1$ and $K = -\log|\cdot|$ as it arises from models of symmetric matrices leading to a real spectrum [327, Chapters 2 and 3], the interaction kernel K is not the Coulomb kernel g solution to (1.58). In this case, we may use the terminology of *one dimensional log-gas*, and we refer to [29] for a nice treatment of this problem.

The purpose of this short example was to provide a connexion between random matrix theory and systems of points defined by a singular Gibbs energy, which is one motivation for the study of Part IV. Obviously, there are (many!) other ways to deal with random matrices, for instance methods based on moments, complex analysis or orthogonal polynomials, see [327, 9, 401]. An advantage of the energetic point of view is to provide a natural candidate (the electrostatic energy \mathcal{E}) for controlling the convergence towards the equilibrium measure μ_\star . However, this approach is rather difficult to apply to general matrix models, see for instance the recent works [54, 17].

1.1.4.3 More general models

We may consider a more general gas of particles interacting through a Hamiltonian of the form

$$\forall x \in (\mathbb{R}^d)^N, \quad H_N(x_1, \dots, x_N) = \frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} K(x_i - x_j), \quad (1.59)$$

where K is typically singular at 0, and with associated probability distribution over $(\mathbb{R}^d)^N$

$$P_N(dx) = Z_N^{-1} e^{-\beta_N H_N(x_1, \dots, x_N)} dx. \quad (1.60)$$

Such models are interesting on their own, independently from random matrix considerations. This is motivated in particular in statistical physics by the Coulomb gas model, *i.e.* when $K = g$ is a Coulomb interaction. By this we mean that g is solution to

$$-\Delta g = c_d \delta_0, \quad \text{with} \quad c_d = \begin{cases} 2\pi & \text{if } d = 2, \\ (d-2)|\mathbb{S}^{d-1}| = \frac{d(d-2)\pi^{d/2}}{\Gamma(1+d/2)} & \text{if } d \geq 3, \end{cases}$$

in the sense of distributions, where Γ is the gamma function and $|\mathbb{S}^{d-1}|$ is the volume of the d -sphere. We know that g is actually given by

$$g(x) = \begin{cases} \log \frac{1}{|x|} & \text{if } d = 2, \\ \frac{1}{|x|^{d-2}} & \text{if } d \geq 3, \end{cases}$$

and we recover the situation discussed above when $d = 2$. Coulomb gases (and the more general Riesz gases) have been the subject of important recent works [386, 84, 287, 387, 85], and they will be at the heart of Part IV.

Finally, although these models are *static* in essence, we can relate them to the dynamics discussed in Section 1.1.1. For this, it suffices to note that (1.60) is a Gibbs measure. As a result, it can be considered for instance as the invariant probability measure of the processes (1.13) or (1.17). This is interesting for sampling purposes: when no random matrix model is available, we may sample the

measure μ_N with one of those dynamics. Although the numerical integration is made non-trivial by the singularity of K , this is possible and we detail this strategy in Chapter 7. On the other hand, this poses theoretical issues since the long time stability discussed in Section 1.1.2 is rather hard to prove in this case. There have been recent works on the subject, see for instance [239, 316] for interesting results.

1.2 Abstract large deviations theory

We have seen in Section 1.1 that problems arising from physics often lead to limiting behaviours because of the different scales involved, both in space and time. Interestingly, the three problems discussed above were related to the Laplace principle as some parameter becomes large (time, inverse temperature, number of particles). Actually, this feature can be cast in the framework of large deviations theory. In a nutshell, this theory provides an appropriate setting for characterizing the asymptotic concentration of probability measures.

In section 1.2.1, we first present the framework of large deviations theory. Section 1.2.2 then provides some important results and examples, such as the Gärtner–Ellis theorem. Our presentation closely follows the excellent books [124, 166, 119] to which we refer for complements.

1.2.1 A natural framework for the concentration of probability measures

We now go into more precise definitions by considering a topological space \mathcal{Y} with its Borel sigma-field \mathcal{B} . The goal of large deviations theory is to define a framework for the concentration of a sequence of probability measures $(\pi_n)_{n \geq 0}$ over $(\mathcal{Y}, \mathcal{B})$ when n becomes large (note that n can be discrete or continuous). The behaviour of the family of probability measures $(\pi_n)_{n \geq 0}$ is generally controlled by a rate function defined below.

Definition 1.3 (Rate function). *A lower semicontinuous functional $I : \mathcal{Y} \rightarrow [0, +\infty]$ is called a rate function. The function I is said to be good if the level sets*

$$\{y \in \mathcal{Y} \mid I(y) \leq M\} \quad (1.61)$$

are compact for any $M \in (0, +\infty)$. The domain of I is defined by

$$D_I = \{y \in \mathcal{Y} \mid I(y) < +\infty\}.$$

Note that the lower semicontinuity of I means that the level sets (1.61) are closed. Therefore, a good rate function is immediately a rate function. An important consequence of the goodness of I is that the infimum over closed sets is attained, *i.e.* it is a minimum; this will be an important element in the study of Coulomb gases in Chapter 8. We recall the convention that the infimum of a lower bounded function over the empty set is set to $+\infty$.

In what follows, for a measurable subset $B \in \mathcal{B}$, we respectively denote by $\overset{\circ}{B}$ and \bar{B} the interior and closure of B with respect to the topology on \mathcal{Y} , while $B^c = \mathcal{Y} \setminus B$ denotes the complement of B in \mathcal{Y} . The notation $\underline{\lim}$ and $\overline{\lim}$ stand for the limit inferior and limit superior respectively. We finally set $(a_n)_{n \geq 0}$ to be an increasing sequence of positive real numbers. We are then in position to define the large deviations principle (LDP).

Definition 1.4 (Large deviations principle). *The family of probability measures $(\pi_n)_{n \geq 0}$ satisfies a large deviations principle at speed $(a_n)_{n \geq 0}$ and with rate function I if, for all $B \in \mathcal{B}$,*

$$-\inf_{\bar{B}} I \leq \underline{\lim}_{n \rightarrow +\infty} \frac{1}{a_n} \log \pi_n(B) \leq \overline{\lim}_{n \rightarrow +\infty} \frac{1}{a_n} \log \pi_n(B) \leq -\inf_{\overset{\circ}{B}} I. \quad (1.62)$$

The Definition 1.4 provides a precise meaning to the asymptotic scaling

$$\pi_n(B) \approx e^{-a_n \inf_B I}$$

when $n \rightarrow +\infty$, by providing appropriate lower and upper bounds. One can see why the infimums on the right and left hand sides of (1.62) cannot be taken over the whole set B in general by considering a singleton $B = \{y\}$ for some $y \in \mathcal{Y}$, and assuming that the measures π_n are non-atomic (which is a natural situation for the models considered in Section 1.1). In this case, $\pi_n(\{y\}) = 0$, so that the infimum in the lower bound of (1.62) is $+\infty$. This is in contradiction with the upper bound

since $\pi_n(\mathcal{Y}) = 1$. As a result, the formulation (1.62) is an appropriate way to deal with concentration of probability measures: not too rigid so as to be correct, but still containing useful information on the exponential concentration⁷. Actually, a set B satisfying

$$I_B = \inf_{\tilde{B}} I = \inf_{\overline{B}} I \quad (1.63)$$

is called an I -continuity set, for which it holds

$$\lim_{n \rightarrow +\infty} -\frac{1}{a_n} \log \pi_n(B) = I_B.$$

This was the intuitive definition of the rate function presented in Section 1.1.2. We also mention I -continuity sets because they are actually useful for conditioning (see Proposition 1.13 below), a crucial element in Chapter 8.

Although stated for an abstract family of probability measures $(\pi_n)_{n \geq 0}$, we generally consider large deviations principles associated to a family of random variables. For this we simply remark that a sequence of random variables $(Y_n)_{n \in \mathbb{N}}$ taking values in \mathcal{Y} induces a family of probability measures on \mathcal{Y} through

$$\pi_n(\cdot) = \mathbb{P}(Y_n \in \cdot).$$

In this case, we refer in general to a large deviations principle for the random variables $(Y_n)_{n \geq 0}$ instead of the induced measures $(\pi_n)_{n \geq 0}$ over \mathcal{Y} , which we believe does not lead to any confusion. Moreover, our presentation uses a large parameter n but we could as well consider indices going to 0 by a clear change of notation.

We explained in Section 1.1.2 that, in order to show the convergence of the law of a process to a limiting distribution, one generally tries to prove that the dynamics remains most of the time in some bounded region of the state space. This is related to the notion of *tightness* in probability, which provides a compactness criterion for probability measures through Prohorov's theorem [42], see also [217] for an applications to the ergodicity of Markov chains. We recall the definition below.

Definition 1.5 (Tightness). *A subset of probability measures $\Gamma \subset \mathcal{P}(\mathcal{Y})$ is tight if, for any $\varepsilon > 0$, there exists a compact set $K_\varepsilon \subset \mathcal{Y}$ such that, for any $\nu \in \Gamma$, it holds $\nu(K_\varepsilon) \geq 1 - \varepsilon$.*

Since large deviations are concerned with exponentially small events, it is quite natural to consider tightness at the exponential scale.

Definition 1.6 (Exponential tightness). *A sequence of probability measures $(\pi_n)_{n \geq 0}$ is exponentially tight at speed $(a_n)_{n \geq 0}$ if, for any $N > 0$, there exists a compact set $K_N \subset \mathcal{Y}$ such that*

$$\overline{\lim}_{n \rightarrow +\infty} \frac{1}{a_n} \log \pi_n(K_N^c) < -N. \quad (1.64)$$

The above definition means that, if $(\pi_n)_{n \geq 0}$ is associated to random variables $(Y_n)_{n \geq 0}$, the following scaling holds as $n \rightarrow +\infty$:

$$\mathbb{P}(Y_n \notin K_N) \lesssim e^{-Na_n}.$$

In words, the probability of not being in a compact set decays at rate e^{-Na_n} . Note that, in order to prove exponential tightness, it is enough to find a precompact set such that (1.64) holds. Exponential tightness is important since it allows to reduce the study of the upper bound in (1.62) to compact sets.

Quite often, the rate function I plays the role of an *entropy functional*, a terminology which will appear clearer with the examples below. Based on physical considerations, it is natural to associate to I a *free energy* function, and to expect it to be related to I through a Legendre–Fenchel transform [166] (see for instance Sections 1.1.2 and 1.1.3). In the abstract framework described here, we can define a free energy when $a_n = n$, for which we introduce, for $f \in \mathcal{Y}'$ (the topological dual of \mathcal{Y}),

$$\Lambda_n(f) = \log \int_{\mathcal{Y}} e^{\langle f, y \rangle_{\mathcal{Y}', \mathcal{Y}}} \pi_n(dy),$$

⁷Being an analyst by training, I have to make a parallel with Schwartz's theory of distributions. In analysis, one generally cannot hope for a regular solution to a given PDE. However, it is possible to define weak solutions. There is again a trade-off between choosing a weak enough notion of solution so as to find one, but not too weak so as to say something about it (for instance uniqueness in the best case scenario).

where we use the notation $\langle \cdot, \cdot \rangle_{\mathcal{Y}', \mathcal{Y}}$ to denote the duality bracket between \mathcal{Y}' and \mathcal{Y} . The free energy is then defined by

$$\bar{\Lambda}(f) = \overline{\lim}_{n \rightarrow +\infty} \frac{1}{n} \Lambda_n(nf). \quad (1.65)$$

The quantity (1.65) is generally called *scaled cumulant generating function* (SCGF, or more simply cumulant function), or *free energy*. When (1.65) is an actual limit (not a limit superior), we use the notation Λ instead of $\bar{\Lambda}$. We will prove below (see Theorems 1.8 and 1.10) that, under certain conditions ensuring the existence and regularity of Λ in (1.65), a LDP holds for $(\pi_n)_{n \geq 0}$ at speed n and with rate function

$$I(y) = \sup_{f \in \mathcal{Y}'} \{ \langle f, y \rangle_{\mathcal{Y}', \mathcal{Y}} - \Lambda(f) \}. \quad (1.66)$$

This is an abstract formulation of the well-known Legendre transform in physics [166]. From a more mathematical perspective, this Legendre transform is generally obtained via optimizing a Tchebychev inequality at exponential scale, which provides a mathematical ground to the considerations of Section 1.1. These questions, dating back to Gärtner and Ellis [199, 165], are well explained in [119], and will be used in the context of fluctuations in time developed in Part II. Now that we presented our abstract setting, we provide some examples and useful results that will be used throughout the work.

Remark 1.7. *The large deviations framework is a convenient way to deal with concentration of probability measures, and can be used for various purposes. The first one is to understand at which speed concentration of measure takes place. If there is in general a natural scale, heavy tailed phenomena lead to anomalous concentration speeds, which are much more complex to study than the standard cases described below, see for instance [54, 342, 17]. Maybe more interestingly, large deviations allow to characterize what is the most probable way in which fluctuations occur. For instance, in the framework of Section 1.1.2, the fluctuating trajectories can be characterized through an optimal control problem [68, 147, 149]. This is also useful in the small temperature regime, as explained in Section 1.1.3. Once again, heavy tailed phenomena show a surprising behaviour. In [54], it is shown that random matrices with sub-Gaussian entries fluctuate through subgraphs with a particular structure. In [342], the authors show that heavy tailed large deviations in time induce localized fluctuations. From a more pragmatic perspective, large deviations can provide an expression of the limiting object through a minimization problem. For instance, in the setting of Section 1.1.4, the equilibrium measure can be identified as the minimizer of the electrostatic energy (1.56). In Chapter 8, we will use this strategy to identify conditioned equilibrium distributions for the Coulomb gas. When following this procedure, we see that the expression of I is very important. It is however difficult in general to express I in a way which allows computations, and it is one of the contributions of Chapter 3 to provide a useful representation of the rate function in the context of long time large deviations for irreversible dynamics.*

Large deviations theory however has some limitations for which various refinements are available. First, a LDP is an asymptotic result, in which all subexponential terms (in particular prefactors) are not taken into account. However, these prefactors can be of paramount physical importance, as in the Eyring–Kramers law, see for example [67, 298] and references therein. In this situation, more refined tools have to be developed on a case by case basis. Moreover, one can be interested in non asymptotic bounds, with a precise dependence on the decay with respect to a distance parameter. For instance, if $(Y_n)_{n \in \mathbb{N}}$ are random variables taking values in a metric space (\mathcal{Y}, d) and converging to a limit Y^ in probability, we expect bounds of the type*

$$\mathbb{P}(d(Y_n, Y^*) \geq r) \leq e^{-a_n |r|^p},$$

for some exponent $p > 0$ and a sequence $(a_n)_{n \in \mathbb{N}}$. There are many works on the subject, and for this thesis we may cite [317, 85] in the context of Coulomb gases, [195, 264, 43] in the context of long time behaviour of SDEs, and [53] in the context of interacting systems of particles.

1.2.2 Important results

Before coming back to the issues raised in Section 1.1, we present some examples and helpful results. This is the opportunity to get more intuition on the rather abstract setting just described, and to present some tools used throughout. Once again, these issues are discussed in great detail in [124, 119, 166, 121, 424].

1.2.2.1 Cramer's and Sanov's theorems

In order to give more intuition on the large deviations framework presented in Section 1.2.1, we start with examples concerning sums of independent random variables. For this, we consider independent and identically distributed (i.i.d.) variables $(X_n)_{n \geq 0}$ taking values in $\mathcal{X} = \mathbb{R}$. The empirical mean is defined by

$$S_n = \frac{1}{n} \sum_{i=1}^n X_i. \quad (1.67)$$

We assume for simplicity that X_1 has finite first and second moments, so that S_n converges almost surely to $\mathbb{E}[X_1]$ and a central limit theorem holds. We are then interested in the large fluctuations of (1.67). Since S_n is real-valued, the framework of Section 1.2.1 is instantiated here with $\mathcal{Y} = \mathbb{R}$ equipped with the usual Borel sigma-field and, for any measurable set $B \subset \mathbb{R}$,

$$\pi_n(B) = \mathbb{P}(S_n \in B).$$

Since the variables X_n are independent, one can easily check that the cumulant function (1.65) tensorizes (see also Remark 1.11 below) so that, since $\mathcal{Y}^* = \mathbb{R}$, it holds

$$\forall k \in \mathbb{R}, \quad \Lambda(k) = \log \mathbb{E} [e^{kX_1}]. \quad (1.68)$$

The quantity (1.68) corresponds to a standard partition function in statistical physics [166], and takes value in $(-\infty; +\infty]$. In particular it may take the value $+\infty$ but its domain $D_\Lambda = \{k \in \mathbb{R} \mid \Lambda(k) < +\infty\}$ is convex, since Λ is convex by Hölder's inequality. In this simple situation, the Fenchel transform (1.66) reads, for $a \in \mathbb{R}$,

$$\Lambda^*(a) = \sup_{k \in \mathbb{R}} \{ka - \Lambda(k)\}, \quad (1.69)$$

and the following standard theorem holds [119, Chapter 2.2].

Theorem 1.8 (Cramer's theorem). *The family of probability measures $(\pi_n)_{n \geq 0}$ satisfies a large deviations principle over \mathbb{R} at speed n and with convex rate function Λ^* , i.e. (1.62) holds with $a_n = n$ and $I = \Lambda^*$ is given by (1.68)-(1.69).*

Moreover, if $0 \in \mathring{D}_\Lambda$, Λ^* is a good rate function.

Cramer's theorem has a very wide range of applications, and a priori does not require Λ to be finite to hold true. However, when X_1 does not have enough exponential moments (for instance $D_\Lambda = \{0\}$ because the tails of X_1 decay too slowly), the rate function may be trivial. This is why the condition $0 \in \mathring{D}_\Lambda$ allows to obtain interesting information on the concentration of measure.

Two examples are useful to give more substance to this theorem. Consider first the case where X_1 follows a Gaussian distribution with variance σ^2 . A simple computation shows that $D_\Lambda = \mathbb{R}$ and

$$\forall a \in \mathbb{R}, \quad I(a) = \Lambda^*(a) = \frac{a^2}{2\sigma^2}.$$

In other words, in the Gaussian case, large fluctuations are Gaussian just like small fluctuations (as provided by the central limit theorem). This situation is specific to Gaussian random variables, and we will come across this particular Gaussian structure in two other situations (see Chapters 3 and 8). Note also that in this case $D_I = \mathbb{R}$. Suppose next that X_1 follows an exponential law with parameter $\theta > 0$. A similar computation shows that $0 \in \mathring{D}_\Lambda$ and that the rate function reads

$$\forall a \in \mathbb{R}, \quad I(a) = \Lambda^*(a) = \begin{cases} \theta a - 1 - \log(\theta a), & \text{if } a > 0, \\ +\infty, & \text{otherwise.} \end{cases}$$

The fact that $I(a) = +\infty$ for $a \leq 0$ stems from the fact that the exponential law cannot take negative values. On the other hand, I is convex, has compact level sets, and vanishes at the mean $\mathbb{E}[X_1] = 1/\theta$. The local behaviour around the mean is quadratic, as we expect from the central limit theorem. However, large deviations are not quadratic any more, since I asymptotically grows linearly. This is characteristic of random variables with exponential tails.

On the other hand, S_n only represents the mean of the variables and, as noted in Section 1.1.2, we may be interested in more general observables. Considering now that X_1 takes value in $\mathcal{X} = \mathbb{R}^d$, this suggests to consider the empirical average

$$L_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}, \quad (1.70)$$

where δ_x is the Dirac mass at $x \in \mathbb{R}^d$. To obtain the average of a particular measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$, one can compute

$$L_n(f) = \frac{1}{n} \sum_{i=1}^n f(X_i),$$

from which we recover the previous situation with the choice $f(x) = x$. Now L_n is taking value in $\mathcal{Y} = \mathcal{P}(\mathcal{X}) = \mathcal{P}(\mathbb{R}^d)$ and we may define, for $B \subset \mathcal{P}(\mathcal{X})$,

$$\pi_n(B) = \mathbb{P}(L_n \in B). \quad (1.71)$$

In order to study the exponential concentration of $(\pi_n)_{n \geq 0}$, we need to equip \mathcal{Y} with a topology. In other words, we have to choose the class of functions for which the convergence of $L_n(f)$ makes sense. The standard procedure [124] is to use the weak topology on $\mathcal{P}(\mathcal{X})$ (the topology of convergence with respect to continuous bounded functions) or the strong topology (convergence with respect to bounded measurable functions). A LDP for $(\pi_n)_{n \geq 0}$ under the weak or strong topologies is generally called Sanov's theorem [124, 121].

Since the observables under consideration are generally unbounded, we may be interested in topologies corresponding to unbounded functions. A possibility is to consider the Wasserstein topology for $p \geq 1$ on $\mathcal{P}_p(\mathcal{X})$, the set of probability measures ν with moments of order p , *i.e.* the probability measures such that

$$\int_{\mathcal{X}} |x|^p \nu(dx) < +\infty.$$

One way to define the Wasserstein topology τ^p is to define convergence through weak convergence plus convergence of the p^{th} moment, see [419, Chapter 7.2] (this is alternatively the topology induced by the p -Wasserstein distance). In order to identify the rate function we define the following entropy functional: for any $\nu \in \mathcal{P}(\mathcal{X})$,

$$H(\nu | \mu) = \begin{cases} \int_{\mathcal{X}} \frac{d\nu}{d\mu} \log \left(\frac{d\nu}{d\mu} \right) d\mu, & \text{if } \nu \ll \mu, \\ +\infty, & \text{otherwise,} \end{cases} \quad (1.72)$$

where μ is the distribution of X_1 and $\nu \ll \mu$ means that ν is absolutely continuous with respect to μ [42]. A stronger version of Sanov's theorem is as follows [424].

Theorem 1.9 (Sanov's theorem). *The family of probability measures $(\pi_n)_{n \geq 0}$ defined in (1.71) satisfies a large deviations principle on (\mathcal{P}_p, τ^p) at speed n if and only if*

$$\forall k > 0, \quad \log \left(\int_{\mathcal{X}} e^{k|x|^p} \mu(dx) \right) < +\infty. \quad (1.73)$$

In this case, the rate function is given by $I(\nu) = H(\nu | \mu)$ for any $\nu \in \mathcal{P}_p(\mathcal{X})$ and $+\infty$ otherwise.

The above Sanov theorem is very powerful, since it proves the LDP in the p -Wasserstein topology, meaning that the empirical measure (1.70) can accommodate functions growing at most like $|x|^p$ for $p \geq 1$. We will come back to this point in Chapters 3 and 8. Interestingly, the LDP in the p -Wasserstein topology also implies the exponential moment condition (1.73), which is the most difficult part of the proof in [424]. Without assumption on the exponential moment, we simply obtain a LDP in the weak topology [121], which does not allow to recover Theorem 1.8. Note that the domain of the rate function is now more complicated: it is a subset of the probability measures that are absolutely continuous with respect to μ . In the standard situation where μ has a Lebesgue density in \mathbb{R}^d , it holds in particular $H(\delta_x | \mu) = +\infty$ for any $x \in \mathcal{X}$.

Note that Sanov's theorem is a good reason for calling the rate function an *entropy*. This is because in the case of independent variables, the rate function (1.72) is the standard definition of

entropy in information theory and statistical physics [279], also called Kullback–Leibler divergence. In the context of large deviations for the empirical average $(L_t)_{t \geq 0}$ defined in (1.23) in Section 1.1.2, the rate function is actually a *Fisher information* for reversible dynamics, hence the name *dynamical entropy* (see Chapter 3).

In practical situations, the random variables X_i in the empirical measure (1.70) can describe the positions of particles interacting through some potential energy, like in the example presented in Section 1.1.4; or may be the values of a time dependent process, like in the diffusion setting of Section 1.1.2. Since the independence is lost, the theorems of Cramer and Sanov do not apply any more, which motivates considering a more general setting than the exponential moment condition (1.73) for independent variables. We present below a general strategy pioneered by Gärtner and Ellis [199, 165] to address this issue.

1.2.2.2 Gärtner–Ellis theorem

The idea of the Gärtner–Ellis theorem is to deduce large deviations principles from appropriate properties of the free energy $\bar{\Lambda}$ defined in (1.65). Let us first motivate why this quantity and its Legendre–Fenchel transform naturally appear in this context by providing a heuristic proof of the upper bound in (1.62) (see [119, Theorem 4.5.3] for details).

We come back to random variables $(X_n)_{n \geq 0}$ taking values in \mathbb{R} , for which we assume without loss of generality that $\mathbb{E}[X_1] = 0$. The empirical mean S_n is defined as in (1.67), but we drop the assumption that the X_n 's are independent. A natural strategy for proving the upper bound of (1.62) is to resort to Tchebychev's inequality. For this we write, for any $a, k > 0$,

$$\pi_n([a; +\infty)) = \mathbb{P}(S_n \geq a) = \mathbb{P}\left(k \sum_{i=1}^n X_i \geq akn\right) \leq e^{-akn} \mathbb{E}\left[e^{k \sum_{i=1}^n X_i}\right].$$

By taking the logarithm, normalizing by $1/n$ and passing to the limit superior we obtain

$$\overline{\lim}_{n \rightarrow +\infty} \frac{1}{n} \log \pi_n([a; +\infty)) \leq -(ak - \bar{\Lambda}(k)),$$

where $\bar{\Lambda}$ is defined in (1.65). By optimizing over $k > 0$, we obtain the Chernoff's bound:

$$\overline{\lim}_{n \rightarrow +\infty} \frac{1}{n} \log \pi_n([a; +\infty)) \leq -\sup_{k > 0} \{ak - \bar{\Lambda}(k)\},$$

which is nearly the Legendre transform (1.66) (because the supremum is over $k > 0$ instead of $k \in \mathbb{R}$). Some manipulations [119, Lemma 2.2.5] allow to recover the whole conjugate function $\bar{\Lambda}^*$, which is then a natural candidate for the rate function. However, the above computation only deals with the upper bound in (1.62).

The fruitful idea of the Gärtner–Ellis theorem is to deduce the lower bound from the upper bound provided the cumulant function Λ exists and is sufficiently regular. The need for regularity can be understood because a lack of regularity of Λ results into linear parts for the conjugate $\bar{\Lambda}^*$. In these regions, the actual rate function may be nonconvex, and the Fenchel transform $\bar{\Lambda}^*$ may be the convex hull of the actual rate function. This default is overcome by assuming sufficient regularity of the free energy, which covers cases in which the rate function is indeed convex⁸. These convexity considerations are very well explained with examples in [166], and we refer to [126, 407, 346, 423] and references therein for examples of nonconvex rate functions and singular entropies. The next abstract statement is a powerful application of these convexity considerations, see [119, Corollary 4.6.14].

Theorem 1.10 (Gärtner–Ellis theorem). *Consider the abstract setting of Section 1.2.1 with an exponentially tight family of probability measures $(\pi_n)_{n \geq 0}$ over a locally convex Hausdorff topological vector space \mathcal{Y} . Suppose that*

$$\Lambda(\cdot) = \lim_{n \rightarrow +\infty} \frac{1}{n} \Lambda_n(n \cdot) \tag{1.74}$$

exists on \mathcal{Y}' , takes finite values and is Gateau-differentiable. Then $(\pi_n)_{n \geq 0}$ satisfies a large deviations principle at speed n in \mathcal{Y} with good convex rate function defined by

$$\forall y \in \mathcal{Y}, \quad \Lambda^*(y) = \sup_{f \in \mathcal{Y}'} \{ \langle f, y \rangle_{\mathcal{Y}', \mathcal{Y}} - \Lambda(f) \}.$$

⁸This is a reason why the Gärtner–Ellis theorem does not apply in all situations and, in many cases, one has to prove the LDP from scratch.

At this level of generality, the proof relies on the projective limit technique reducing the analysis to the finite dimensional case, for which the strategy is as follows. One first obtains the upper bound over compact sets by reduction to a finite number of balls and using Tchebychev's inequality as above. This upper bound is extended to closed set thanks to the exponential tightness. Next, the lower bound relies on an appropriate change of measure (Esscher transform), under which the fluctuation becomes typical. The most difficult part of the proof is to make the latter argument precise through fine convex analysis results – this is explained at length in Chapters 1 and 4.6 of [119].

Remark 1.11 (Independent variables). *When considering fluctuations of the empirical mean S_n defined in (1.67), we can recover Cramer's theorem under a stronger condition from Theorem 1.10. Supposing that the variables X_n are i.i.d., the cumulant function (1.74) simply reads*

$$\forall k \in \mathbb{R}, \quad \Lambda(k) = \lim_{n \rightarrow +\infty} \frac{1}{n} \log \mathbb{E} \left[e^{k \sum_{i=1}^n X_i} \right] = \log \mathbb{E} \left[e^{k X_1} \right],$$

which is simply (1.68), as already seen above. We observe the very particular behaviour of the exponential with respect to independance. This seemingly innocuous remark actually makes the situation of subexponential variables (i.e. with tails decaying like $e^{-|x|^\alpha}$ for $\alpha \in (0, 1)$) much more difficult. In this case, the lack of «tensorizability» of the subexponential makes the typical fluctuations heavy tailed, in the sense that only a part of the variables fluctuates to produce an abnormal value of S_n . We refer to [333, 334, 335, 336, 59] and the more recent work [17] for more details. Similar considerations are important for stochastic differential equations, see [342] for an interesting contribution in the case of an Ornstein–Uhlenbeck process.

1.2.2.3 Contraction principle and Gibbs conditioning: most probable unlikely events

Although proving a large deviations principle is a complicated problem, there is a nice tool for deducing a LDP for the image of random variables satisfying a LDP. This result, called contraction principle [119, Theorem 4.2.1], will be useful in the context presented in Section 1.1.2.

Theorem 1.12 (Contraction principle). *Consider two Hausdorff topological spaces \mathcal{Y} and \mathcal{Z} and a continuous application $\Pi : \mathcal{Y} \rightarrow \mathcal{Z}$, together with a good rate function $I : \mathcal{Y} \rightarrow [0, +\infty]$. Then:*

1. *The function $J : \mathcal{Z} \rightarrow [0, +\infty]$ defined as*

$$\forall z \in \mathcal{Z}, \quad J(z) = \inf \left\{ I(y), y \in \mathcal{Y} \mid \Pi(y) = z \right\} \quad (1.75)$$

is a good rate function.

2. *If $(\pi_n)_{n \geq 0}$ satisfies a LDP with rate function I , then $(\pi_n \circ \Pi^{-1})_{n \geq 0}$ satisfies a LDP with rate function J .*

The first important point of the contraction principle is that LDPs are transferred by continuous applications, which turns out to be a useful tool. Second, the resulting rate function satisfies some maximum entropy principle⁹: the cost of a fluctuation of the image is the minimum of the costs over the pre-images. In other words, rare events happen in the most likely of all unlikely ways.

In order to understand how to use Theorem 1.12, we present an application related to Sanov's result, Theorem 1.9. If a LDP is available for the empirical average

$$L_n = \frac{1}{n} \sum_{i=1}^n \delta_{X_i},$$

we would like to obtain a LDP for instance for the empirical mean

$$S_n = \frac{1}{n} \sum_{i=1}^n X_i.$$

For this, we note that S_n is the image of L_n by the following mapping:

$$\Pi : \nu \in \mathcal{P}(\mathcal{X}) \longmapsto \int_{\mathcal{X}} x \nu(dx) \in \mathbb{R}.$$

⁹The term “maximum entropy” may sound confusing since (1.75) comes with an infimum. This is because the entropy in physics is the opposite of the entropy in mathematics.

An important remark is that the application Π is *not* continuous for the weak topology. However, it is continuous for the p -Wasserstein topology for any $p \geq 1$. As a result, Theorem 1.9 together with Theorem 1.12 lead to a LDP for S_n . The corresponding rate function J is given by (1.75), where I is the entropy (1.72). This strategy is yet another path to Cramer's theorem.

Before closing this section, we present a result often called *Gibbs conditioning principle* or *maximum entropy principle*. The purpose of this result is to understand the behaviour of random variables when they are *conditioned* on some rare event of interest. The principle states that the most likely way for a measure to concentrate on a set conditioned on a rare event is to minimize the rate function I under the constraint defining the conditioning. This idea is generally used with Sanov's theorem [357, 306, 305], and the so obtained conditioned limits are sometimes called *entropic projections*. We will prove in Chapter 8 the following more general result (stated for random variables in a metric space for the sake of simplicity), which relies on the notion of I -continuity sets discussed in Section 1.2.1.

Proposition 1.13 (A Gibbs conditioning principle). *Suppose that Y_1, \dots, Y_n are random variables taking values in a metric space (\mathcal{Y}, d) satisfying a large deviations principle at speed $(a_n)_{n \geq 1}$, and with good rate function I . Consider a closed set B such that*

$$\inf_B I = \inf_B I < +\infty.$$

Then, the set of minimizers

$$\mathcal{I}_B = \left\{ y \in \mathcal{Y}, I(y) = \inf_B I \right\}$$

is a non-empty closed subset of B . Moreover, for any $\varepsilon > 0$, setting

$$A_\varepsilon = \left\{ y \in \mathcal{Y}, d(y, \mathcal{I}_B) > \varepsilon \right\}, \quad (1.76)$$

there exists $c_\varepsilon > 0$ such that

$$\overline{\lim}_{n \rightarrow +\infty} \frac{1}{a_n} \log \mathbb{P} \left(Y_n \in A_\varepsilon \mid Y_n \in B \right) \leq -c_\varepsilon. \quad (1.77)$$

The idea of Proposition 1.13 is that, in the large n limit, it holds

$$\mathbb{P} \left(Y_n \in A_\varepsilon \mid Y_n \in B \right) \approx e^{-c_\varepsilon n}.$$

Because of the definition (1.76) of A_ε , we see that the conditioned variables concentrate on \mathcal{I}_B , *i.e.* on a minimizer of I over B . From a philosophical perspective, the conditioned rare event occurs in the less costly way, or equivalently in the more likely way, which is a feature similar to the contraction principle above. If the rate function is explicit and the set B is expressed through a natural constraint, there is a hope that a minimizer under constraint solves an equation involving a Lagrange multiplier. We will see an application of this idea in the context of conditioned Coulomb gases in Chapter 8.

Remark 1.14. *There are other techniques than those presented above to cope with large deviations problems. Let us mention in particular the so-called weak convergence method, a framework explained in [147]. An interesting feature of this approach is to characterize the large deviations functions through optimal control problems, which is useful for practical applications, see Chapter 5. Let us also quote [149] as a successful example of application of this approach.*

1.3 Large deviations in statistical physics

Section 1.2.2 has presented the general setting of large deviations together with a number of important results. We now come back to the various problems raised in Section 1.1 to show that they enter the large deviations framework. This will demonstrate the flexibility of large deviations techniques and raise interesting numerical problems, some of which will be highlighted in Section 1.4.

1.3.1 Long time fluctuations

1.3.1.1 Large deviations principle

In this section, we consider the setting of Section 1.1.2. Since we study long time problems, we have $n = t$ in the notation of Section 1.2.1, and we study the large deviations of the empirical average

$$L_t = \frac{1}{t} \int_0^t \delta_{X_s} ds \quad (1.78)$$

in $\mathcal{Y} = \mathcal{P}(\mathcal{X})$, where $(X_t)_{t \geq 0}$ is solution to (1.1). As mentioned in Section 1.2.1, we should typically speak of the LDP for the family of probability measures

$$\pi_t(\cdot) = \mathbb{P}(L_t \in \cdot)$$

over $\mathcal{Y} = \mathcal{P}(\mathcal{X})$, but we generally refer to the LDP for the empirical average $(L_t)_{t \geq 0}$ for conciseness. Three approaches have been considered in the past for proving a LDP for the empirical measure (1.78): the approach of Donsker and Varadhan [131, 132, 134, 135, 416], that of Gärtner and Ellis [199, 165], and the weak convergence approach [68, 149]. An example of result can be stated as follows in the case where $\mathcal{X} = \mathbb{T}^d$ (or a compact set without border). We recall that the strong topology on $\mathcal{P}(\mathcal{X})$ is the topology associated with the convergence against bounded measurable functions.

Theorem 1.15. *Suppose that $\mathcal{X} = \mathbb{T}^d$, $b \in C^\infty(\mathcal{X})$ and $\sigma \in \mathbb{R}^{d \times d}$ is constant with $\sigma\sigma^T$ positive definite. Then the functional*

$$f \in B^\infty(\mathcal{X}) \mapsto \lambda(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E} \left[e^{\int_0^t f(X_s) ds} \right]$$

is well-defined and finite, and the empirical mean $(L_t)_{t \geq 0}$ satisfies a LDP over $\mathcal{P}(\mathcal{X})$ equipped with the strong topology, at speed t and with rate function given by: for $\nu \in \mathcal{P}(\mathcal{X})$,

$$I(\nu) = \begin{cases} \sup_{f \in B^\infty} \{\nu(f) - \lambda(f)\}, & \text{if } \nu \ll \mu, \\ +\infty, & \text{otherwise.} \end{cases} \quad (1.79)$$

Recall that $\mu \in \mathcal{P}(\mathcal{X})$ is the invariant measure of $(X_t)_{t \geq 0}$. Note that we do not aim at the greatest level of generality in the above result, and we refer *e.g.* to [199] for less restrictive conditions on the regularity of the coefficients. We also mention that Theorem 1.15 is actually a particular case of Theorem 3.10 stated in Chapter 3.

However, in a statistical physics perspective, the Langevin dynamics (1.17) does not satisfy the assumptions of Theorem 1.15 for two important reasons:

- the dynamics is degenerate, *i.e.* $\sigma\sigma^T$ has not full rank, as can be seen from (1.16);
- the state space \mathcal{X} is genuinely unbounded because the momenta belong to \mathbb{R}^d .

These issues are treated in Chapter 3, and we postpone the discussion on this point to the contributions presented in Section 1.5. We also mention that the important work of L. Wu [427] brings elements of answer to these problems.

For practical purposes, when the average of a particular function f is of interest, we can also prove a LDP for

$$L_t(f) = \frac{1}{t} \int_0^t f(X_s) ds,$$

in the space $\mathcal{Y} = \mathbb{R}$. As mentioned in Section 1.2.2, a LDP can be obtained a priori by the contraction principle. On the other hand, when $\mathcal{X} = \mathbb{T}^d$, we can prove a LDP for $L_t(f)$ at speed t and with rate function

$$\forall a \in \mathbb{R}, \quad I_f(a) = \sup_{k \in \mathbb{R}} \{ka - \lambda_f(k)\}, \quad (1.80)$$

where the free energy is now a function over \mathbb{R} defined through

$$\forall k \in \mathbb{R}, \quad \lambda_f(k) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right]. \quad (1.81)$$

This is again a consequence of the Gärtner–Ellis theorem. In particular, we can compute numerically (1.80) from the cumulant function λ_f , which raises interesting numerical questions discussed in Section 1.4 below.

1.3.1.2 Relation to spectral theory and Feynman–Kac dynamics

As should be clear by now, the functional $f \mapsto \lambda(f)$ plays a crucial role in the large deviations analysis, so we would like to obtain more information about it. The cumulant $\lambda(f)$ is actually the principal eigenvalue (largest in modulus) of the operator $\mathcal{L} + f$ defined over $B^\infty(\mathcal{X})$ (recall that \mathcal{L} is the generator of the dynamics (1.1) defined in (1.2)). This can be intuited by considering the Feynman–Kac semigroup

$$(P_t^f \varphi)(x) = \mathbb{E}_x \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right],$$

which has generator $\mathcal{L} + f$ by the Feynman–Kac formula [361, Chapter VIII]. From spectral analysis, we expect

$$P_t^f \varphi \underset{t \rightarrow +\infty}{\sim} e^{tr(f)},$$

where $r(f)$ is the largest eigenvalue in modulus of $\mathcal{L} + f$. This formula is completely natural for matrices and we shall give a precise meaning to it for operators in Chapter 2. On the other hand, by taking the logarithm and dividing by t , we obtain the cumulant $\lambda(f)$: this motivates that $\lambda(f) = r(f)$ is actually the largest eigenvalue of $\mathcal{L} + f$.

We can now invert the Legendre transform (1.79) to obtain the representation

$$\lambda(f) = \sup_{\nu \in \mathcal{P}(\mathcal{X})} \{ \nu(f) - I(\nu) \}, \quad (1.82)$$

as was intuited at the end of Section 1.1.2. In order for the above expression to be of some use, we note that I may be rewritten as

$$I(\nu) = \sup \left\{ - \int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu, u \in \mathcal{D}_{\mathcal{L}}, u > 0 \right\}, \quad (1.83)$$

where $\mathcal{D}_{\mathcal{L}}$ denotes the domain of \mathcal{L} over $B^\infty(\mathcal{X})$. The equations (1.82)–(1.83) form the celebrated Donsker–Varadhan variational formula, whose strength is to provide a variational representation for the principal eigenvalue of an a priori *non-symmetric* second order operator associated with a diffusion [133] (note however that [133] demands less regularity on the diffusion by assuming that the semigroup $(P_t)_{t \geq 0}$ is Feller instead of strong Feller, in which case the generator has to be considered on the space $C^0(\mathcal{X})$ of continuous functions instead of $B^\infty(\mathcal{X})$). To get some intuition on (1.83), consider an eigenvector h_f associated to $\lambda(f)$, namely

$$(\mathcal{L} + f)h_f = \lambda(f)h_f.$$

The above equation rewrites

$$- \frac{\mathcal{L}h_f}{h_f} = f - \lambda(f). \quad (1.84)$$

By integrating both sides of (1.84) with respect to a measure ν we obtain the correspondance between (1.79) and (1.83), an idea made precise in [124, Lemma 4.1.36] and revisited in Chapter 3. Finally, it is an interesting problem to derive alternative representations for the rate function I . With the interpretation of I as a dynamical entropy, we expect different behaviours for reversible and irreversible dynamics, as discussed in Section 1.1.1. A significant contribution to this question was brought by [49] in the case of an atom chain thermostated by an inhomogeneous heat bath, and we will also come back to this point in Chapter 3.

Remark 1.16. *The Donsker–Varadhan formula is a significant representation result for the principal eigenvalue of non-symmetric operators. As will appear clear in the analysis of Part II, this formula is closely related to the Perron–Frobenius and Krein–Rutman theorems in the theory of positive operators, see for instance [383, 112]. In the case of a matrix $A \in \mathbb{R}^{d \times d}$ such that $A_{ij} > 0$ for any $i, j \in \{1, \dots, d\}$, we indeed know that the largest eigenvalue Λ is real and simple by the Perron–Frobenius theorem [383]. Moreover, Λ can be represented by the so-called Collatz–Wielandt formula:*

$$\Lambda = \sup \left\{ \inf_{x > 0} \sum_{i=1}^d \frac{(Ax)_i}{x_i} p_i \mid p \in \mathbb{R}^d, \sum_{i=1}^d p_i = 1 \right\}.$$

This is the exact analog of (1.83) for matrices¹⁰, see [329, Chapter 8].

¹⁰I could not find a reference to this formula in the papers of Donsker and Varadhan, and I warmly thank A. Levitt for pointing out this result.

Before concluding this section, we relate the cumulant (1.81) to the more general framework of Feynman–Kac semigroups. We call Feynman–Kac dynamics the nonlinear semigroup $(\Phi_t^f)_{t \geq 0}$ defined on the space of probability measures $\mathcal{P}(\mathcal{X})$ as: for $t \geq 0$,

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in B^\infty(\mathcal{X}), \quad \Phi_t^f(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]}{\mathbb{E}_\nu \left[e^{\int_0^t f(X_s) ds} \right]}.$$

From a practical perspective, this is a form of importance sampling: the trajectories with larger values of f are given a larger weight. This idea is crucial in nonlinear filtering methods [115, 139, 140]. In our context, we relate the long time convergence of Feynman–Kac dynamics [114] to large deviations theory, a central topic of Chapter 2.

1.3.2 Small temperature regime

In Section 1.1.3 we have presented some aspects of SDEs at small temperature from a heuristic perspective. The large deviations framework for this situation has been worked out by Freidlin and Wentzell in the 1970s and 1980s, and we refer to [193] for a very good introduction to these problems. Recall the notation of Section 1.1.3: $(X_t^\varepsilon)_{t \in [0, T]}$ is solution to

$$dX_t^\varepsilon = b(X_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma dB_t, \quad (1.85)$$

over $\mathcal{X} = \mathbb{R}^d$, where b is smooth and $D = \sigma \sigma^T$ is positive definite. We are interested in the fluctuations of the process around the continuous path $(\bar{\phi}_t)_{t \in [0, T]}$ solution to

$$\bar{\phi}_0 = x, \quad \frac{d}{dt} \bar{\phi}_t = b(\bar{\phi}_t),$$

for some $x \in \mathcal{X}$. Therefore, the natural space for studying large deviations is $\mathcal{Y} = C^0([0, T], \mathcal{X})$ equipped with the uniform norm topology.

In order to provide more rigorous estimates like (1.39), we consider the following family of probability measures over $C^0([0, T], \mathcal{X})$:

$$\pi_\varepsilon(\cdot) = \mathbb{P}(X_t^\varepsilon \in \cdot). \quad (1.86)$$

An action functional is also associated to the dynamics (1.85):

$$\forall \varphi \in C^0([0, T], \mathcal{X}), \quad I(\varphi) = \frac{1}{2} \int_0^T |b(\varphi) - \dot{\varphi}_s|_D^2 ds, \quad (1.87)$$

where $|v|_D^2 = v \cdot D^{-1}v$ for any $v \in \mathbb{R}^d$, and the above integral is given the value $+\infty$ if φ is not absolutely continuous or if the integral diverges. As the reader should expect, the functional I is a rate function controlling a LDP for the family of probability measures $(\pi_\varepsilon)_{\varepsilon > 0}$ defined in (1.86). Assuming that b is globally Lipschitz, we obtain the following result (in [193] we use Theorem 1.1 in Chapter 4 for proving that I is a rate function in the sense of Definition 1.3, paragraph 3 in Chapter 3 for the definition of the action functional, and Theorem 3.3 in Chapter 3 for the equivalence with the Definition 1.4 of the LDP in Section 1.2.1).

Theorem 1.17 (Freidlin–Wentzell). *Consider the family of probability measures $(\pi_\varepsilon)_{\varepsilon > 0}$ defined in (1.86) where $(X_t^\varepsilon)_{t \in [0, T]}$ is solution to (1.85), with b globally Lipschitz and $D = \sigma \sigma^T$ positive definite. Then the functional $I : C^0([0, T], \mathcal{X}) \rightarrow [0; +\infty]$ defined in (1.87) is a good rate function.*

Moreover, $(\pi_\varepsilon)_{\varepsilon > 0}$ satisfies a large deviations principle at speed ε with rate function I : for any $B \subset C^0([0, T], \mathcal{X})$, it holds

$$-\inf_{\bar{B}} I \leq \liminf_{\varepsilon \rightarrow 0} \varepsilon \log \mathbb{P}(X_t^\varepsilon \in B) \leq \overline{\lim}_{\varepsilon \rightarrow 0} \varepsilon \log \mathbb{P}(X_t^\varepsilon \in B) \leq -\inf_B I,$$

where the adherence and interior are taken with respect to the uniform norm topology.

Theorem 1.17 has many applications. For example, (1.42) is an instance of the Varadhan principle [119, Chapter 4], which makes precise the heuristic of Section 1.1.3. Since we will not prove new theoretical results on this kind of asymptotics and essentially use the instanton described in Section 1.1.3, we close the discussion here. Note however that the ideas of large deviations can be used for studying various aspects of systems at small temperature. The behaviour of the spectrum of the Witten Laplacian at low temperature and related quasi-stationary distributions are two important families of problems (related to metastability) in which the ideas of concentration at exponential scale in the low temperature regime play a crucial role, see e.g. [389, 299, 331, 298] and references therein.

1.3.3 Singular gases

We finally come back to the situation of Section 1.1.4 by considering a set of points $X_N = (X_{N,1}, \dots, X_{N,N})$ distributed according to the following Gibbs distribution over $(\mathbb{R}^d)^N$:

$$P_N(dx) = Z_N^{-1} e^{-\beta_N \left(\frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} K(x_i - x_j) \right)} dx, \quad (1.88)$$

where V is a confinement potential, K is a (typically singular) interaction kernel and $(\beta_N)_N$ is a cooling sequence. Each particle being distributed in \mathbb{R}^d , we consider the empirical distribution

$$\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i} \quad (1.89)$$

under P_N , which is a random element in $\mathcal{P}(\mathbb{R}^d)$. Large deviations theory once again provides natural tools for studying the concentration of μ_N as $N \rightarrow +\infty$, where now

$$\pi_N(\cdot) = \mathbb{P}(\mu_N \in \cdot).$$

The genuine space for the LDP is therefore $\mathcal{Y} = \mathcal{P}(\mathbb{R}^d)$. As intuited in Section 1.1.4, the electrostatic energy defined by

$$\mathcal{E}(\mu) = \int V(x) \mu(dx) + \iint K(x-y) \mu(dx) \mu(dy) \quad (1.90)$$

plays a crucial role in the analysis. Fine properties of the above electrostatic energy and its consequences in terms of concentration of (1.89) under (1.88) have known many developments since the seminal work of Ben Arous and Guionnet [29], see [84, 148, 386, 287, 85, 387, 314]. In order to provide an example of result in this direction, we present [84, Theorem 1.1], which applies in particular to the Coulomb gas setting of Section 1.1.4. We recall that the weak topology on $\mathcal{P}(\mathbb{R}^d)$ is metrized by¹¹ the Fortet–Mourier distance defined by

$$\forall \mu, \nu \in \mathcal{P}(\mathbb{R}^d), \quad d_{\text{FM}}(\nu, \mu) = \sup_{\max(\|f\|_{B^\infty}, \|f\|_{\text{Lip}}) \leq 1} \left(\int_{\mathbb{R}^d} f d\nu - \int_{\mathbb{R}^d} f d\mu \right),$$

where the B^∞ -norm is defined in (1.26) and

$$\|f\|_{\text{Lip}} = \sup_{x \neq y} \frac{|f(x) - f(y)|}{|x - y|}.$$

Theorem 1.18 (LDP for singular gas). *Suppose that the following assumptions hold true:*

1. *The cooling sequence satisfies $\beta_N \gg N \log(N)$.*
2. *The function V is continuous and satisfies $\lim_{|x| \rightarrow +\infty} V(x) = +\infty$ and*

$$\int_{\mathcal{X}} e^{-V(x)} dx < +\infty.$$

3. *The kernel $K : \mathcal{X} \rightarrow (-\infty; +\infty]$ is continuous on \mathbb{R}^d and takes finite values on $\mathbb{R}^d \setminus \{0\}$. Moreover, for any compact set $C \subset \mathbb{R}^d$ the function*

$$z \mapsto \sup \{ K(y), |y| \geq |z|, y \in C \}$$

is locally integrable with respect to the Lebesgue measure.

4. *There exists $c > 0$ and $\varepsilon_0 \in (0, 1)$ such that*

$$\forall x, y \in \mathbb{R}^d, \quad K(x-y) \geq c - \varepsilon_0 (V(x) + V(y)).$$

5. *If $\nu \in \mathcal{P}(\mathcal{X})$ is such that $\mathcal{E}(\nu) < +\infty$, there exists a sequence $(\nu_n)_{n \in \mathbb{N}}$ of probability measures absolutely continuous with respect to the Lebesgue measure, such that $\nu_n \rightarrow \nu$ weakly and $\mathcal{E}(\nu_n) \rightarrow \mathcal{E}(\nu)$ when $n \rightarrow +\infty$.*

¹¹There are many ways to metrize weak convergence of measures. In large deviations theory, the Lévy–Prokhorov metric is often used [124], but the Fortet–Mourier distance is easier to manipulate. Other distances can be used [354, 434], which metrize the weak convergence with different geometric properties.

Then the empirical average μ_N satisfies a large deviations principle under P_N in $\mathcal{P}(\mathbb{R}^d)$ equipped with the weak topology, at speed $(\beta_N)_N$ and with good rate function defined by

$$\forall \nu \in \mathcal{P}(\mathbb{R}^d), \quad I(\nu) = \mathcal{E}(\nu) - \inf_{\mathcal{P}(\mathbb{R}^d)} \mathcal{E},$$

where $\inf_{\mathcal{P}(\mathbb{R}^d)} \mathcal{E} > -\infty$. More precisely, for any measurable set $A \subset \mathcal{P}(\mathbb{R}^d)$, it holds

$$-\inf_A I \leq \liminf_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P}(\mu_N \in A) \leq \overline{\lim}_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P}(\mu_N \in A) \leq -\inf_A I,$$

where the interior and closure are taken with respect to the weak topology. Moreover, denoting by

$$\mathcal{I} = \left\{ y \in \mathcal{Y}, I(y) = \inf_{\mathcal{P}(\mathbb{R}^d)} I \right\}$$

the closed set of minimizers, it holds

$$d_{\text{FM}}(Y_n, \mathcal{I}) \xrightarrow{n \rightarrow +\infty} 0,$$

almost surely under P_N .

Theorem 1.18 calls for some comments. First, it applies to the case of Coulomb interactions under appropriate growth conditions on V , see [84]. It then provides a rigorous justification of the Laplace principle invoked in Section 1.1.4, where the electrostatic energy \mathcal{E} plays the role of a rate function. The speed of this LDP is dictated by the cooling sequence $(\beta_N)_N$ defining P_N , provided the temperature decreases fast enough (*i.e.* β_N diverges fast enough). This condition actually ensures that entropic effects can be neglected, as nicely discussed in [84, 148], but can be relaxed to $\beta_N \gg N$, see [148]. An important feature of the LDP in this situation is to identify the limiting empirical measure as a minimizer of \mathcal{E} . This justifies the importance of this minimization problem, which recently attracted a lot of interest, in particular for the situation of Coulomb gases (described in Section 1.1.4) and more generally Riesz and log-gases [84, 287]. Note that the study of this minimization problem for the one dimensional log-gas leads to an alternative proof of the celebrated Wigner theorem for random matrices [29]. We mention that these results can be strengthened, for instance by reinforcing the topology of the LDP if V satisfies sufficient growth conditions, see [148] for an interesting account. These refinements are used in Chapter 8 to apply the Gibbs conditioning principle (Proposition 1.13) in this context.

Remark 1.19. *We have illustrated three possible applications of the large deviations toolbox: long time behaviour of empirical averages, small temperature limit of diffusions and large size limit of singular gases. However, there exist many other fields in which large deviations can be applied. Let us mention for instance quantum mechanical models [6, 44, 267] and multifractals [229, 417, 307] for interesting applications.*

1.4 Some numerical aspects of rare events analysis

Although the large deviations aspects presented above form a wide mathematical playground for theoretical developments, they also raise difficult numerical questions. For instance, computing probabilities of rare events is important in various fields, and have led to numerous mathematical treatments, see for instance [75, 167, 79, 80] and references therein. On the other hand, identifying the limiting distribution of random matrices and their fine properties (such as the distribution of the largest eigenvalue [19, 108, 421, 16]) is important for instance in finance [65, 66]. It would be too long to present all the applications of these ideas, so we restrict ourselves to some questions treated in this work, mostly concerning stochastic differential equations.

We first review in Section 1.4.1 elements on the numerical analysis of SDEs, in particular the different notions of error. We next present in Section 1.4.2 some variance issues raised by the simulation of rare events, together with two types of strategies to address them.

1.4.1 Discretization of SDEs

In this section, we consider again the SDE

$$dX_t = b(X_t) dt + \sigma(X_t) dB_t, \quad (1.91)$$

where the process is defined over $\mathcal{X} = \mathbb{R}^d$. The functions $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times m}$ are assumed to be smooth for simplicity, and $(B_t)_{t \geq 0}$ denotes as usual an m -dimensional Brownian motion. In order to compute expectations and empirical means as described in Section 1.1.2, a natural strategy is to discretize (1.91) with a Markov chain of time step Δt . Such a Markov chain $(x_n)_{n \in \mathbb{N}}$ is defined by its (homogeneous) evolution operator¹²

$$\forall \varphi \in B^\infty(\mathcal{X}), \quad \forall x \in \mathcal{X}, \quad (Q_{\Delta t} \varphi)(x) = \mathbb{E}[\varphi(x_{n+1}) | x_n = x]. \quad (1.92)$$

For concreteness, one may think of the discretization given by the Euler–Maruyama scheme, which reads

$$x_{n+1} = x_n + b(x_n)\Delta t + \sigma(x_n)\sqrt{\Delta t}G_n, \quad (1.93)$$

where $(G_n)_{n \in \mathbb{N}}$ is a family of independent standard Gaussian random variables. There are much more sophisticated discretization procedures, and we refer to the textbooks [270, 332] for an overview. Instead of stating very precise results of numerical analysis, we prefer to insist on the various ways to measure the error induced by the discretization procedure. We believe this discussion is important in order to understand the error analysis performed in Chapter 4, in which we introduce a new notion of discretization error for Feynman–Kac dynamics.

Strong error estimates. Strong error estimates quantify the difference between the trajectories of the continuous process and its discrete counterpart over finite time intervals. They consist in practice in measuring the maximum L^q -norm for some $q \geq 1$ (in expectation) between the discretization and the continuous process over a finite time window. Such a result typically takes the following form: there exists $p_s > 0$ such that for any time $T > 0$, there is $C > 0$ and $\Delta t^* > 0$ for which, for any $\Delta t \in (0, \Delta t^*]$,

$$\sup_{0 \leq n \leq T/\Delta t} \left(\mathbb{E}[|X_{n\Delta t} - x_n|^q] \right)^{\frac{1}{q}} \leq C \Delta t^{p_s}.$$

We will not consider this type of convergence here, since we are interested in average properties rather than in the trajectories themselves. The following alternative definition of error is therefore more useful in our context.

Weak error estimates. Errors on the expectations are called weak errors (weak being understood in the usual sense “against test functions”). We denote by $C_c^\infty(\mathcal{X})$ the space of smooth functions with compact support. Weak error estimates typically take the following form: there exists $p_w > 0$ such that for any $\varphi \in C_c^\infty(\mathcal{X})$ and $T > 0$, there is $C > 0$ and $\Delta t^* > 0$ for which, for any $\Delta t \in (0, \Delta t^*]$,

$$\sup_{0 \leq n \leq T/\Delta t} \left| \mathbb{E}[\varphi(X_{n\Delta t})] - \mathbb{E}[\varphi(x_n)] \right| \leq C \Delta t^{p_w}. \quad (1.94)$$

Such an estimate can generally be obtained for test functions $\varphi \in C^k(\mathcal{X})$ with the constant C controlled by $\|\varphi\|_{C^k}$, see [332]. We also mention that in general it holds $p_s \leq p_w$.

Quite interestingly, estimates such as (1.94) can be deduced from a one time step expansion. The idea is to compare the exact flow $P_{\Delta t} = e^{\Delta t \mathcal{L}}$ with its approximation $Q_{\Delta t}$ (recall that \mathcal{L} is the generator of the diffusion (1.91)), for which one can generally prove an expansion like

$$Q_{\Delta t} = e^{\Delta t \mathcal{L}} + O(\Delta t^{p_w+1}). \quad (1.95)$$

It is then possible to derive estimates like (1.94), while providing a precise functional setting to the above equality [348, 400, 276], see also [332, Theorem 2.1]. We are however not completely satisfied by this definition of error, since the constant C in (1.94) typically grows exponentially with T , while we aim at estimating ergodic properties as $T \rightarrow +\infty$. This motivates the notion introduced next.

¹²We generally use the letter P for the evolution operator of time continuous dynamics, while Q is used for discrete time dynamics.

Error on the invariant measure. As explained in Section 1.1.2, we generally want to compute the following type of ergodic average

$$\mathbb{E}[\varphi(X_t)] \xrightarrow{t \rightarrow +\infty} \int_{\mathcal{X}} \varphi d\mu, \quad (1.96)$$

where μ is the invariant measure of the process. Let us assume that the discretization $(x_n)_{n \in \mathbb{N}}$ is also ergodic with respect to an invariant measure $\mu_{\Delta t}$, namely

$$\mathbb{E}[\varphi(x_n)] \xrightarrow{n \rightarrow +\infty} \int_{\mathcal{X}} \varphi d\mu_{\Delta t}. \quad (1.97)$$

Except in particular situations (for instance when correcting a reversible diffusion with a Metropolis–Hastings step [368]), it holds $\mu_{\Delta t} \neq \mu$ so the averages on the right hand side of (1.96) and (1.97) are different. It is then natural to quantify the error on the invariant measure by measuring the difference between averages of observables, a strategy pioneered by Talay [398, 399]. Such error estimates are written as follows: there exists $p_i > 0$ and $\Delta t^* > 0$ such that, for any $\Delta t \in (0, \Delta t^*]$ and $\varphi \in C_c^\infty(\mathcal{X})$,

$$\int_{\mathcal{X}} \varphi d\mu_{\Delta t} = \int_{\mathcal{X}} \varphi d\mu + O(\Delta t^{p_i}). \quad (1.98)$$

Moreover, the leading order term can be made precise as

$$\Delta t^{p_i} \int_{\mathcal{X}} \varphi \psi d\mu,$$

where ψ is solution to a Poisson equation. Perhaps surprisingly, the estimate (1.98) on the long time average can also be deduced from an expansion of the one step evolution operator $Q_{\Delta t}$, provided some ergodicity conditions on the evolutions $(X_t)_{t \geq 0}$ and $(x_n)_{n \in \mathbb{N}}$ are met. In practice, one can generally obtain by Taylor expansion of $Q_{\Delta t}$ in power of Δt the following type of expansion (see the example of the Euler scheme below): for some $p_i \in \mathbb{N}^*$,

$$\forall \varphi \in C_c^\infty(\mathcal{X}), \quad Q_{\Delta t} \varphi = \varphi + \Delta t \mathcal{A}_1 \varphi + \Delta t^2 \mathcal{A}_2 \varphi + \dots + O(\Delta t^{p_i+1}), \quad (1.99)$$

for differential operators \mathcal{A}_k . The crucial point is to observe until which order the propagator $Q_{\Delta t}$ conserves the invariant measure. This corresponds to the integer $p_i \in \mathbb{N}^*$ for which

$$\forall k \in \{1, \dots, p_i\}, \quad \forall \varphi \in C_c^\infty(\mathcal{X}), \quad \int_{\mathcal{X}} \mathcal{A}_k \varphi d\mu = 0, \quad (1.100)$$

which is the expansion for small times of the invariance relation (1.7). In general, in accordance with (1.95), schemes of weak order p_i are such that

$$\forall k \in \{1, \dots, p_i\}, \quad \mathcal{A}_k = \frac{\mathcal{L}^k}{k!}.$$

This condition is however not necessary for (1.100) to hold (for instance it may happen that $\mathcal{A}_k = c_k^{-1} \mathcal{L}^k$ with $c_k \neq k!$).

In order to deduce (1.98) from (1.100), the strategy is to compare the invariance relations corresponding to $P_{\Delta t}$ and $Q_{\Delta t}$ as introduced in Section 1.1 (recall that $P_{\Delta t} = e^{\Delta t \mathcal{L}}$), namely: for any $\varphi \in C_c^\infty(\mathcal{X})$,

$$\int_{\mathcal{X}} P_{\Delta t} \varphi d\mu = \int_{\mathcal{X}} \varphi d\mu, \quad \text{and} \quad \int_{\mathcal{X}} Q_{\Delta t} \varphi d\mu_{\Delta t} = \int_{\mathcal{X}} \varphi d\mu_{\Delta t}.$$

We next expand in powers of Δt through (1.99) inside the second expression. In fact, we construct an approximation $\tilde{\mu}_{\Delta t}$ of $\mu_{\Delta t}$ for which the second equality holds up to a small error term. This strategy is adapted in Chapter 4 to obtain error estimates on the cumulant function and the Feynman–Kac semigroup presented in Section 1.3.1. The difficulty in this situation comes from the nonlinearity of the stationary equation satisfied by the invariant measure, which makes the above reasoning much more cumbersome to apply.

As an example of application, we consider the Euler–Maryuama scheme (1.93). In this case we have, for any $\varphi \in C_c^\infty(\mathcal{X})$,

$$\forall x \in \mathcal{X}, \quad Q_{\Delta t} \varphi(x) = \mathbb{E}_G \left[\varphi \left(x + b(x) \Delta t + \sigma(x) \sqrt{\Delta t} G \right) \right],$$

where the expectation runs over the standard Gaussian variable G . Since a factor $\sqrt{\Delta t}$ is involved, a second order Taylor expansion shows that

$$Q_{\Delta t}\varphi = \varphi + \Delta t\mathcal{L}\varphi + O(\Delta t^2),$$

which is (1.99)-(1.100) with $p_i = 1$. As a result, we see that (1.100) provides a rather practical criterion for assessing the order of accuracy of a numerical scheme. This led to substantial developments in the past years (in particular in the context of molecular dynamics), see for instance [399, 324, 325, 110, 295]. Interestingly, efficient numerical methods are often built on integrators that have been known for a long time in the physics community, such as the Verlet scheme [418, 62, 294, 3], or modifications thereof.

Let us insist on the fact that in general it holds $p_s \leq p_w \leq p_i$, and these inequalities are rather often strict. In other words, it is possible to build Markov chains whose trajectories are quite different from the continuous ones, but whose *statistical properties* are close to that of the continuous process. This remark can be extended to the computation of dynamical quantities, see for instance [295, 304] and references therein. Note that, in Chapter 4, we will present a yet weaker notion of error, which does not exist for SDEs but is important when dealing with Feynman–Kac semigroups. We now give more details on the statistical estimation of the quantities we are considering.

1.4.2 Estimating rare events

We now present some numerical issues related to the estimation of the quantities of interest in Sections 1.1.2 and 1.3.1. After some preliminaries, we highlight two classes of methods used to overcome the high variance of naive estimators: those based on population dynamics, and the ones relying on optimal control arguments.

1.4.2.1 The variance issue

As we explained in Sections 1.1.2 and 1.3.1, the computation of the rate function I and the free energy λ is important to understand the system at hand, and to provide quantitative estimates on the probabilities of rare events. For the sake of this presentation we consider fluctuations of the empirical average

$$L_t(f) = \frac{1}{t} \int_0^t f(X_s) ds,$$

where $(X_t)_{t \geq 0}$ is solution to (1.1) and f is a measurable function such that a large deviations principle holds (see Section 1.3.1) with rate function $I_f : \mathbb{R} \rightarrow [0, +\infty]$. In order to estimate I_f , we may consider some small element $\Delta a \subset \mathbb{R}$ centered on $a \in \mathbb{R}$ and use the estimator

$$I_f(a) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}(L_t(f) \in \Delta a).$$

In words, I_f is the limiting normalized histogram of $L_t(f)$. However, $\mathbb{P}(L_t(f) \in \Delta a)$ decays like $e^{-tI_f(a)}$, meaning that an exponentially large (in t) number of samples is necessary to obtain a good estimate of the rate function. Since the large t limit must be reached, this is obviously an inefficient approach.

The Gärtner–Ellis theorem presented in Sections 1.2.2 and 1.3.1 is possibly a better alternative. Indeed, (1.80) provides the following alternative representation of the rate function,

$$I_f(a) = \sup_{k \in \mathbb{R}} \{ka - \lambda_f(k)\}, \quad (1.101)$$

where

$$\lambda_f(k) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right]. \quad (1.102)$$

Therefore, the computation of I_f can be deduced from that of λ_f (and possibly λ'_f to compute (1.101), provided the cumulant is sufficiently regular). Although this idea is alluring, exponential quantities like

$$\mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right] \quad (1.103)$$

typically have a variance scaling exponentially in time, and are therefore difficult to numerically estimate. In order to understand this phenomenon, we note that

$$\int_0^t f(X_s) ds \approx t\alpha + \sqrt{t}\sigma G,$$

where α is a real number and G is a standard Gaussian random variable. This motivates considering the simpler expression

$$\mathbb{E} \left[e^{k(t\alpha + \sqrt{t}\sigma G)} \right],$$

for $t > 0$. An easy computation shows that that relative variance reads

$$\frac{\text{Var} \left[e^{k(t\alpha + \sqrt{t}\sigma G)} \right]}{\left(\mathbb{E} \left[e^{k(t\alpha + \sqrt{t}\sigma G)} \right] \right)^2} = \frac{\mathbb{E} \left[e^{2k\sqrt{t}\sigma G} \right]}{\left(\mathbb{E} \left[e^{k\sqrt{t}\sigma G} \right] \right)^2} - 1 = e^{k^2\sigma^2 t} - 1. \quad (1.104)$$

As a result, the relative variance scales like $e^{k^2\sigma^2 t}$. Transferring this reasoning to the estimation of (1.103) motivates that the variance of a naive estimator of the free energy scales exponentially with the time t . In practice, the variance of (1.103) is high because the value of the expectation is dominated by a few rare trajectories taking very large values (a problem shared in the estimate of (1.35) for the small temperature systems described in Section 1.1.3). Therefore, not much seems to be gained from the Fenchel representation (1.101). It is however possible to design estimators of (1.102) with a reduced variance. We describe below two strategies: interacting particle systems and optimal control.

Remark 1.20. *As mentioned in Section 1.3.1, $\lambda_f(k)$ is also the largest eigenvalue of the Feynman–Kac generator $\mathcal{L} + kf$. Thus, Galerkin or spectral approximation methods can be used to estimate the cumulant. However, these strategies are often inefficient because the dimension d of the space is large. On the other hand, (1.102) naturally suggests to use a stochastic algorithm for computing a principal eigenvalue. Not surprisingly, this representation was first used in quantum physics [10, 77], see the discussion below and the comments in Chapter 5.*

1.4.2.2 Interacting particles systems

There is a wide variety of approaches relying on the principle of “divide and conquer” in order to compute (1.102) more efficiently. The basic idea is to run a set of replicas of the dynamics (1.1) with exponential weights attached to them. These weights correspond to the argument of the expectation in (1.103). When those weights become degenerate, the particles are resampled with some rule preserving the weighted expectation. In practice, particles with larger weights tend to duplicate, while those with smaller weights tend to be pruned. Doing so, the dynamics is conditioned on taking larger values of f , which reduces the variance and provides relatively well-behaved estimators [139, 113, 140, 301].

Let us make this idea slightly more precise by giving a basic version of a genetic algorithm to compute (1.102) (which will be detailed and used in Chapter 4). Consider a set $(X^m)_{m=1}^M$ of M replicas satisfying (1.1) and a time τ at which each replica $m \in \{1, \dots, M\}$ has weight

$$w^m = \exp \left(k \int_0^\tau f(X_s^m) ds \right).$$

We then define the probability distribution p over the set $\{1, \dots, M\}$ of replicas by

$$p^m = \frac{w^m}{\sum_{j=1}^M w^j}.$$

At time τ the replicas are then resampled according to the probability vector p (for instance by drawing M replicas from a multinomial distribution of probability p). The so-obtained distribution of replicas is an approximation of the distribution at time t of the Feynman–Kac semigroup $(\Phi_t)_{t \geq 0}$ introduced in Section 1.3.1, while the cumulant $\lambda_f(k)$ can be estimated *e.g.* by averaging the weights w over the replicas along the dynamics (this is made precise in Chapter 2). Typically, τ is deterministic and associated to the time step Δt of a numerical discretization, as in Section 1.4.1. It may also be random and correspond to a stopping time for a given level of degeneracy of the weights, see for instance [301, Chapter 6]. In comparison with (1.104), we see that the problem is reduced to smaller subproblems with smaller variance (because τ is typically small), while statistical properties are preserved if the resampling is correct [140].

To the best of our knowledge, this replica technique goes back to [213, 10] for estimating ground state energies of Schrödinger operators (based on papers by Donsker and Kac [258, 130]), followed by

a series of works in quantum physics [77, 320, 412, 328]. There has been after that a great interest in the mathematics community for these so-called Feynman–Kac models, see for instance [115, 139, 116, 117, 113, 377, 109]. Interestingly, these two communities, which used the same algorithms, did not seem to communicate much at that time. This is surely because of their different motivations. More recently, due to recent advances in nonequilibrium statistical mechanics [194, 280, 288, 202, 122], large deviations functions received a renewed interest, which generated some new algorithms, see *e.g.* [289, 337]. A first comprehensive history (and systematic treatment) of these techniques can be found in [12, 11].

1.4.2.3 Optimal control approaches

The strategy presented above relies on a “genetic” procedure, where a set of replicas are selected so as to enhance the probability of rare events. A complementary point of view is to introduce a bias in the dynamics to enforce the fluctuation in a single trajectory. The relation between this method and large deviations theory has been extensively studied by Dupuis, Ellis and collaborators (the so-called “weak convergence approach” to large deviations [68, 147, 149]), as well as by Chetrite and Touchette [98] in the context of nonequilibrium statistical mechanics. Rather than providing a long presentation of the theory, we prefer presenting the key idea leading to an alternative representation of the free energy and rate function, with the perspective of designing more efficient numerical methods, as illustrated¹³ in Chapter 5. This will also highlight the importance of the eigenproblem associated with the cumulant function, as mentioned in Section 1.3.1.

Let us recall that the cumulant function $\lambda_f(k)$ is associated to the eigenproblem

$$(\mathcal{L} + kf)h_k = \lambda_f(k)h_k, \quad (1.105)$$

where h_k is assumed to be a smooth function (this will be detailed in Chapter 3). We next introduce the tilted generator [97]

$$\mathcal{L}_k = h_k^{-1}(\mathcal{L} + kf)(h_k \cdot) - \lambda_f(k) = \mathcal{L} + D\nabla \log h_k, \quad (1.106)$$

where $D = \sigma\sigma^T$ is the diffusion matrix, which we assume constant and positive definite for simplicity. This “ h -transform” can be shown to condition the dynamics on having a fluctuation of $L_t(f)$ [136, 97] (note that we use this transform in Part II for proving convergence results). Actually, the operator \mathcal{L}_k defined in (1.106) is a Markovian generator associated to the tilted dynamics (also called driven process)

$$d\tilde{X}_t = b_k(\tilde{X}_t) dt + \sigma dB_t, \quad (1.107)$$

where

$$\forall x \in \mathcal{X}, \quad b_k(x) = b(x) + D\nabla \log h_k(x). \quad (1.108)$$

The idea of the modified drift (1.108) is that the additional term pushes the dynamics towards the regions where f is higher, hence reducing the variance.

A possible way to understand why the titled process $(\tilde{X}_t)_{t \geq 0}$ is in some sense optimal is to derive the zero-variance estimator for the SCGF $\lambda_f(k)$. For this, we denote by $u = \log h_k$, which solves (if h_k is positive and smooth)

$$\mathcal{L}u + \frac{|\sigma^T \nabla u|^2}{2} + kf = \lambda_f(k). \quad (1.109)$$

We next use the Girsanov theorem [262] to rewrite the exponential expectation in the definition of the free energy, which leads to

$$\mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right] = \mathbb{E} \left[e^{k \int_0^t f(\tilde{X}_s) ds - \frac{1}{2} \int_0^t |\sigma^T \nabla u|^2(\tilde{X}_s) ds - \int_0^t \nabla u(\tilde{X}_s) \cdot \sigma(\tilde{X}_s) dB_s} \right]. \quad (1.110)$$

On the other hand, Itô’s formula applied to u leads to

$$du(\tilde{X}_s) = (\mathcal{L}u(\tilde{X}_s) + (D\nabla u) \cdot \nabla u(\tilde{X}_s)) dt + \nabla u(\tilde{X}_s) \cdot \sigma(\tilde{X}_s) dB_t,$$

so that

$$-\frac{1}{2} \int_0^t |\sigma \nabla u|^2(\tilde{X}_s) ds - \int_0^t \nabla u(\tilde{X}_s) \cdot \sigma(\tilde{X}_s) dB_s = \int_0^t \mathcal{L}u(\tilde{X}_s) ds + \frac{1}{2} \int_0^t |\sigma \nabla u|^2(\tilde{X}_s) ds + u(\tilde{X}_0) - u(\tilde{X}_t).$$

¹³Similar ideas are made precise in Chapter 6 for the small temperature problem presented in Section 1.1.3. We however postpone the discussion on this aspect to this chapter for conciseness.

Plugging the above formula into (1.110) and using (1.109) shows that

$$\mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right] = \mathbb{E} \left[e^{\int_0^t (\mathcal{L}u + \frac{|\sigma \nabla u|^2}{2} + kf)(\tilde{X}_s) ds + u(\tilde{X}_0) - u(\tilde{X}_t)} \right] = e^{t\lambda_f(k)} \mathbb{E} \left[e^{u(\tilde{X}_0) - u(\tilde{X}_t)} \right].$$

In the situation where the state space \mathcal{X} is bounded, this shows that, for any $t > 0$,

$$\frac{1}{t} \log \mathbb{E} \left[e^{\int_0^t kf(X_s) ds} \right] = \lambda_f(k) + O\left(\frac{1}{t}\right).$$

We thus obtain an asymptotically zero-variance estimator from the tilted dynamics (1.107) and the Girsanov formula (1.110). This motivates introducing the tilted generator (1.106).

It is actually possible to obtain more information from the tilted process (1.107), as shown by manipulating the eigenproblem with appropriate changes of variable [98, 149]. For instance, (1.109) is a Hamilton–Jacobi equation with quadratic cost [352, 231]. This formally allows to derive the following representation of the free energy:

$$\lambda_f(k) = \sup_u \left[\lim_{t \rightarrow +\infty} \frac{1}{t} \int_0^t \left(kf(\tilde{X}_s^u) - \frac{|\sigma^T \nabla u(\tilde{X}_s^u)|^2}{2} \right) ds \right],$$

where the supremum runs over sufficiently smooth adapted controls $u : \mathcal{X} \rightarrow \mathbb{R}^d$, and $(\tilde{X}_t^u)_{t \geq 0}$ is the controlled process

$$d\tilde{X}_t^u = b(\tilde{X}_t^u) dt + D\nabla u(\tilde{X}_t^u) dt + \sigma dB_t.$$

The optimal u is then given by (1.108) through $u = \log h_k$, so that $\lambda_f(k)$ is solution to the control problem which consists in maximizing the fluctuation f under a quadratic cost constraint¹⁴ (and it can be shown that the optimal control is in gradient form). Quite interestingly, the rate function corresponds to the quadratic cost of the control [98], that is

$$I_f(a_k) = \lim_{t \rightarrow +\infty} \frac{1}{2t} \int_0^t \nabla \log h_k(\tilde{X}_s) \cdot D^{-1} \nabla \log h_k(\tilde{X}_s) ds, \quad (1.111)$$

where $a_k = \lambda'(k)$. In other words, the rate function is the cost to make a certain fluctuation of f likely: the larger the force needed to make f atypical, the rarer the event. We also mention that (1.111) can be written as

$$I_f(a_k) = \lim_{t \rightarrow +\infty} \frac{1}{2t} \int_0^t |D^{-1}(b(\tilde{X}_s) - b_k(\tilde{X}_s))|_D^2 ds,$$

where $|v|_D^2 = v \cdot D^{-1}v$ for $v \in \mathbb{R}^d$. This expression is very similar to that of the action functional presented in Section 1.1.3 in the context of systems at low temperature.

Coming back to our numerical problem, we may wonder to what extent the alternative representations provided by the optimal control framework are useful in practice compared to spectral or genetic approaches. There is actually an important literature on stochastic and risk-sensitive control [353, 58, 28, 57, 30], and we exploit in Chapter 5 some of these works to design an adaptive algorithm for estimating on the fly the optimal drift b_k and the large deviations functions (note that there is also an important literature on finite time control problems, see for instance [352, 231, 414, 232] and references therein). We will discuss there the advantages and problems of this technique.

It is however important to emphasize that, for large (and possibly low temperature) systems, the genetic and optimal control methods are both likely not to work, because of the important number of clones needed in the first case, and the difficulty to estimate the optimal control in the second. A good strategy seems to combine both methods: with an approximation of the optimal control, the rare event becomes more likely and a genetic method can be implemented more efficiently, see [337, 339] for recent accounts.

1.5 Contributions

After having presented the topics covered in this thesis, we make precise its contributions following the organization of the manuscript. Since precise results and bibliography are available later on, we present the main results in an informal way and postpone more detailed discussions to each individual chapter.

¹⁴This problem is well-known in finance under the name of *risk-sensitive control* [28, 57].

1.5.1 Feynman–Kac dynamics and large deviations

Part II is devoted to theoretical results on Feynman–Kac dynamics and large deviations in time for diffusion processes.

1.5.1.1 Long time behaviour of Feynman–Kac dynamics

As motivated in Section 1.3.1, the long time behaviour of Feynman–Kac models is closely related to large deviations. These dynamics are also interesting on their own (for instance in quantum physics or nonlinear filtering), and this motivates the study of Chapter 2. More precisely, for a Polish space \mathcal{X} , we study the following kind of dynamics on $\mathcal{P}(\mathcal{X})$: for $\nu \in \mathcal{P}(\mathcal{X})$ and $\varphi \in B^\infty(\mathcal{X})$,

$$\forall k \in \mathbb{N}, \quad \Phi_k(\nu)(\varphi) = \frac{\nu((Q^f)^k \varphi)}{\nu((Q^f)^k \mathbf{1})}. \quad (1.112)$$

In the above formula, Q^f is a transition kernel which is not probabilistic (*i.e.* $Q^f(x, \cdot)$ is a measure not normalized to 1, so $Q^f \mathbf{1} \neq \mathbf{1}$). The ratio in (1.112) ensures that Φ_k maps $\mathcal{P}(\mathcal{X})$ onto $\mathcal{P}(\mathcal{X})$ but makes the dynamics nonlinear. This abstract formulation encompasses many practical cases. For instance if $Q^f = e^f Q$ where Q is the evolution operator of a Markov chain $(x_n)_{n \in \mathbb{N}}$, (1.112) reads

$$\Phi_k(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(x_n) e^{\sum_{i=0}^{k-1} f(x_i)} \right]}{\mathbb{E}_\nu \left[e^{\sum_{i=0}^{k-1} f(x_i)} \right]},$$

which is a standard formula for Feynman–Kac dynamics [113].

The long time behaviour of the nonlinear dynamics (1.112) has been studied for quite some time in the context of nonlinear filtering, see for instance [114]. However, to the best of our knowledge, no result is available when the state space \mathcal{X} is unbounded. Our main contribution is to prove ergodicity for (1.112) under a Lyapunov condition [219]. More precisely, we assume that there exist $W : \mathcal{X} \rightarrow [1, +\infty)$, positive sequences $(\gamma_n)_{n \in \mathbb{N}}$, $(b_n)_{n \in \mathbb{N}}$ and an increasing sequence of compact sets $(K_n)_{n \in \mathbb{N}}$ such that

$$\forall x \in \mathcal{X}, \quad Q^f W \leq \gamma_n W + b_n \mathbf{1}_{K_n}, \quad (1.113)$$

with $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$. As discussed in Section 1.1.2, this is a confinement condition (although now stated for a discrete time dynamics). With a couple of additional assumptions (a minorization ensuring mixing and some regularization property of the kernel), we prove in Theorem 2.7 the existence of $\bar{\alpha} \in (0, 1)$ and a unique measure μ_f such that, for any initial measure $\nu \in \mathcal{P}_W(\mathcal{X})$, there is $C_\nu > 0$ for which

$$\sup_{\|\varphi\|_{B_W^\infty} \leq 1} |\Phi_k(\nu)(\varphi) - \mu_f(\varphi)| \leq C_\nu \bar{\alpha}^k. \quad (1.114)$$

We also obtain in Theorem 2.8 that the largest eigenvalue Λ of the operator Q^f over $B_W^\infty(\mathcal{X})$ is well-defined, real, simple, and satisfies (assuming that $Q^f = e^f Q$ as above for simplicity):

$$\log(\Lambda) = \lim_{k \rightarrow +\infty} \frac{1}{k} \log \mathbb{E}_\nu \left[e^{\sum_{i=0}^{k-1} f(x_i)} \right]. \quad (1.115)$$

This expression is reminiscent of the cumulant function presented in Section 1.3.1. Solving the spectral problem associated with Q^f is actually the crux of the proof of (1.114). In fact, we use the Lyapunov condition (1.113) together with a regularization property to prove compactness of Q^f , and ensure that the largest eigenvalue Λ is real, isolated, and such that

$$Q^f h = \Lambda h,$$

where $h \in B_W^\infty(\mathcal{X})$ is a properly normalized eigenvector. From this eigenproblem, we can introduce the h -transformed kernel operator

$$Q_h = \Lambda^{-1} h^{-1} Q^f (h \cdot).$$

It then suffices to note that Q_h induces a Markovian dynamics which can be studied for instance with [219], and that the non linear semigroup (1.112) can be written as the ratio of two linear semigroups as

$$\Phi_k(\nu)(\varphi) = \frac{\nu(h(Q_h)^k (h^{-1} \varphi))}{\nu(h(Q_h)^k h^{-1})}$$

to obtain the desired convergence.

In order to make more precise the connection with large deviations, we mention that the above results also apply in the continuous time context. In other words, we obtain similar results for the continuous time dynamics defined by

$$\Phi_t(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]}{\mathbb{E}_\nu \left[e^{\int_0^t f(X_s) ds} \right]}, \quad (1.116)$$

where $(X_t)_{t \geq 0}$ is solution to (1.1). In this case the Lyapunov condition (1.113) can be written: for all $a > 0$ there exists $C_a \in \mathbb{R}$ such that

$$(\mathcal{L} + f)W \leq -aW + C_a. \quad (1.117)$$

This condition should be natural, since this is a stronger version of the Lyapunov condition (1.25) where the generator of the diffusion \mathcal{L} is replaced by the Feynman–Kac operator $\mathcal{L} + f$. Under this condition (plus regularity and irreducibility), we can prove the continuous analogues of (1.114)–(1.115).

1.5.1.2 Large deviations for empirical measures of diffusions in fine topologies

As a sequel to the study of the long time behaviour of Feynman–Kac dynamics, Chapter 3 turns to a large deviations problem for the empirical measure of diffusions. The question is to know, given a diffusion $(X_t)_{t \geq 0}$ over $\mathcal{X} \subset \mathbb{R}^d$, for which functions f the empirical mean

$$L_t(f) = \frac{1}{t} \int_0^t f(X_s) ds$$

satisfies a large deviations principle (see Section 1.3.1 for definitions). In other words, we ask for the finest topology for a LDP to hold for the empirical average

$$L_t = \frac{1}{t} \int_0^t \delta_{X_s} ds. \quad (1.118)$$

In Section 1.3.1, we have presented results for a bounded state space \mathcal{X} , but the problem is difficult when \mathcal{X} is unbounded. We know from the works of Varadhan that a LDP holds for the weak topology under Lyapunov and Feller conditions [416]. However, to the best of our knowledge, satisfactory results are lacking for a topology associated to *unbounded* functions f (although [427, 275] bring interesting elements). Such results however exist for independent variables, as shown in Section 1.2.2.

The first contribution of Chapter 3 is to identify a class of unbounded functions f to consider the LDP over. This class can be read off the estimates used to prove the convergence of Feynman–Kac dynamics. Indeed, the Lyapunov condition (1.117) can be written, since $W \geq 1$,

$$f \leq -a \frac{\mathcal{L}W}{W} + C_a$$

for $a > 0$ arbitrary large. This identifies the class of appropriate functions f as those that are “strictly dominated” by (*i.e.* that grow strictly slower at infinity than)

$$\Psi = -\frac{\mathcal{L}W}{W}.$$

This condition makes sense provided Ψ diverges to infinity (which is a Lyapunov condition for the dynamics $(X_t)_{t \geq 0}$ with generator \mathcal{L} , as explained in Section 1.1.2). By adapting the proofs of Chapter 2 and under appropriate regularity and irreducibility conditions on the dynamics, we show in Lemma 3.32 that

$$\lambda(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right]$$

is well-defined for f in the identified class of functions. By the Gärtner–Ellis theorem presented in Section 1.2.2, we obtain in Theorem 3.10 a LDP for a precise topology associated to unbounded functions, which is a new result. By this we mean that the following scaling holds: for any $\Gamma \subset \mathcal{P}(\mathcal{X})$

$$\mathbb{P}(L_t \in \Gamma) \asymp e^{-t \inf_{\nu \in \Gamma} I(\nu)},$$

in a topology associated with Ψ , where $I : \mathcal{P}(\mathcal{X}) \rightarrow [0, +\infty]$ is a rate function given by an appropriate modification of (1.79). Interestingly, as mentioned in Section 1.1.2, Ψ can be related to a Witten Laplacian for reversible dynamics – it thus draws a link with spectral gap conditions and Poincaré inequalities [20].

The second contribution of Chapter 3 is to describe more precisely the structure of the rate function I , which is important for theoretical and computational purposes. First, we revisit the standard Donsker–Varadhan formula presented in Section 1.3.1 by showing in Proposition 3.13 that

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad I(\nu) = \sup \left\{ - \int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu, \quad u \in \mathcal{D}^+ \right\}, \quad (1.119)$$

where \mathcal{D}^+ is an appropriate domain defined in Section 3.3.

Maybe more importantly, we harvest [49] to understand better how irreversibility affects the rate function (and hence convergence to stationarity [363]). For this, we notice that the generator \mathcal{L} associated to (1.1) can be decomposed as $\mathcal{L} = \mathcal{L}_S + \mathcal{L}_A$ where \mathcal{L}_S and \mathcal{L}_A are respectively symmetric and antisymmetric in $L^2(\mu)$, with μ the invariant measure of the dynamics $(X_t)_{t \geq 0}$. We then prove in Theorem 3.16 that, for any $\nu \ll \mu$,

$$I(\nu) = \frac{1}{4} \left| \log \frac{d\nu}{d\mu} \right|_{\mathcal{H}^1(\nu)}^2 + \frac{1}{4} \left| \mathcal{L}_A \left(\log \frac{d\nu}{d\mu} \right) \right|_{\mathcal{H}^{-1}(\nu)}^2, \quad (1.120)$$

where $|\cdot|_{\mathcal{H}^1(\nu)}$ and $|\cdot|_{\mathcal{H}^{-1}(\nu)}$ refer to Sobolev seminorms defined in Section 3.2.1. The proof of this result is interesting since it uses a variant of the Witten transform inside (1.119), which we call *variational Witten transform*. Although the decomposition (1.120) looks quite abstract, we show in Chapter 3 that it can be used to obtain interesting information on the fluctuations of irreversible dynamics, in particular the underdamped Langevin dynamics presented in Section 1.1.

1.5.2 Numerical analysis and algorithms

Part III treats two complementary problems: the numerical analysis of quantities introduced in Part II when continuous dynamics are discretized with a finite time step, and the design of efficient numerical algorithms for estimating large deviations functions.

1.5.2.1 Error estimates on ergodic properties of Feynman–Kac semigroups

In Chapter 4, we consider the numerical discretization of the time continuous Feynman–Kac dynamics (1.116). We place ourselves in a setting where \mathcal{X} is bounded, hence we can easily prove the existence of an invariant measure μ_f such that¹⁵

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in C^\infty(\mathcal{X}), \quad \Phi_t^f(\nu)(\varphi) \xrightarrow{t \rightarrow \infty} \int_{\mathcal{X}} \varphi d\mu_f.$$

In order to perform numerical computations, we want to discretize in time the dynamics and obtain error estimates on its ergodic properties such as the average above. As should be clear from the presentation of Section 1.4, it is first natural to discretize the diffusion $(X_t)_{t \geq 0}$ into a Markov chain $(x_n)_{n \in \mathbb{N}}$ approximating the dynamics. An intuitive procedure to discretize (1.116) is for instance to consider the dynamics

$$\Phi_{\Delta t, n}^f(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(x_n) e^{\Delta t \sum_{i=0}^{n-1} f(x_i)} \right]}{\mathbb{E}_\nu \left[e^{\Delta t \sum_{i=0}^{n-1} f(x_i)} \right]}.$$

This is indeed a discrete time Feynman–Kac dynamics of the form (1.112) with $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$ (where $Q_{\Delta t}$ is the evolution operator of the Markov chain $(x_n)_{n \in \mathbb{N}}$ defined in (1.92)). The long time behaviour of $\Phi_{\Delta t, n}^f$ can be studied for example with the tools of Chapter 2, so we can prove the existence of an invariant measure $\mu_{f, \Delta t}$ such that

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in C^\infty(\mathcal{X}), \quad \Phi_{\Delta t, n}^f(\nu)(\varphi) \xrightarrow{n \rightarrow +\infty} \int_{\mathcal{X}} \varphi d\mu_{f, \Delta t}.$$

¹⁵In Chapter 4 we write Φ_t^f rather than Φ_t to emphasize the dependance on the weight function f .

The core of the chapter is then to make precise the difference between μ_f and $\mu_{f,\Delta t}$.

Our strategy is to adapt the works of Talay and Tubaro [400, 398] briefly presented in Section 1.4. This is however a difficult task because the dynamics is nonlinear (and hence its invariant measure satisfies a nonlinear equation, as proved in Chapter 2). A key element of the analysis is a fine understanding of the spectral problem associated to the Feynman–Kac evolution operator. We actually build an approximate spectral problem through a hierarchical decomposition in Poisson equations in the spirit of [110].

We then prove in Theorem 4.16 the existence of an integer $p \geq 1$ and a function ψ solution to a Poisson equation (depending on the numerical scheme at hand and the quadrature rule for the integral in (1.116)) such that, for all $\varphi \in C^\infty(\mathcal{X})$,

$$\int_{\mathcal{X}} \varphi d\mu_{f,\Delta t} = \int_{\mathcal{X}} \varphi d\mu_f + \Delta t^p \int_{\mathcal{X}} \varphi \psi d\mu_f + O(\Delta t^{p+1}). \quad (1.121)$$

This provides a generalization of the results known for the error on the invariant measure presented in Section 1.4 for this nonlinear dynamics. Moreover, we have seen at various places the importance of the principal eigenvalue λ of the operator $\mathcal{L}+f$ (which is also the cumulant function, see Section 1.3.1). We obtain in Theorem 4.21 the following error estimate for an estimator $\lambda_{\Delta t}$ of this quantity: there exists $p' > 0$ such that

$$\lambda_{\Delta t} = \frac{1}{\Delta t} \log \left[\int_{\mathcal{X}} Q_{\Delta t}^f \mathbf{1} d\mu_{f,\Delta t} \right] = \lambda + C\Delta t^{p'} + O(\Delta t^{p'+1}), \quad (1.122)$$

where $Q_{\Delta t}^f$ is the evolution operator of the discretized dynamics with weight function f . Typically it holds $p' \geq p$ because of Proposition 4.23, but we cannot state a general result. For concreteness, one can think of $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$ where $Q_{\Delta t}$ is the evolution operator of the discretization of $(X_t)_{t \geq 0}$. In this case $Q_{\Delta t}^f \mathbf{1} = e^{\Delta t f}$. In other words, (1.122) shows that the cumulant function is computed from the average weight along the discrete Feynman–Kac dynamics. This is a standard numerical procedure for which our results provide a precise error analysis.

1.5.2.2 Adaptive sampling of large deviations functions

Chapter 5 is concerned with the numerical estimation of the large deviations functions introduced in Section 1.3.1. More precisely, we focus on the cumulant function defined by

$$\forall k \in \mathbb{R}, \quad \lambda_f(k) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E} \left[e^{k \int_0^t f(X_s) ds} \right], \quad (1.123)$$

and on the rate function, which may be defined through

$$\forall a \in \mathbb{R}, \quad I(a) = \sup_{k \in \mathbb{R}} \{ka - \lambda_f(k)\}. \quad (1.124)$$

We recall that $(X_t)_{t \geq 0}$ is solution to (1.1). We have discussed in Section 1.4.2 that naive estimators of (1.123)–(1.124) show a large variance, but that this effect can be mitigated by replacing b by the modified drift

$$b_k = b + D\nabla \log h_k,$$

where $(\mathcal{L} + kf)h_k = \lambda(k)h_k$. However, the eigenvector h_k is typically difficult to estimate in practice. The goal of Chapter 5 is to propose an adaptive algorithm to estimate h_k along a single realization of the dynamics, together with $\lambda(k)$, $\lambda'(k)$ and $I(\lambda'(k))$. Let us mention that, in Chapter 5, we do not only consider the empirical average $L_t(f)$ but also empirical currents, which is useful when considering irreversible dynamics.

Our strategy builds on previous works of Borkar and collaborators [58, 5, 28], which we adapt to our situation and enhance with estimators arising from recent works in nonequilibrium statistical physics [97, 98]. The idea is quite simple, and relies on the multiplicative structure of the Feynman–Kac evolution operator. As we have seen above, for $k \in \mathbb{R}$, the evolution operator is typically approximated by, for any test function φ ,¹⁶

$$\forall x \in \mathcal{X}, \quad Q_{\Delta t}^k \varphi(x) = e^{\Delta t k f(x)} \mathbb{E}_x [\varphi(x_{n+1})], \quad (1.125)$$

¹⁶In Chapter 5 we use the notation $Q_{\Delta t}^k$ instead of $Q_{\Delta t}^{kf}$ to alleviate notation. This should not be mistaken with the iterates $(Q_{\Delta t})^k$.

where $(x_n)_{n \in \mathbb{N}}$ is a numerical discretization of $(X_t)_{t \geq 0}$. From the spectral structure of the dynamics (Chapter 2) and the numerical analysis of Chapter 4, we expect that, for any test function φ ,

$$(Q_{\Delta t}^k)^n \varphi(x) \underset{n \rightarrow +\infty}{\sim} C_\varphi h_k(x) e^{n\lambda(k)},$$

for some constant $C_\varphi > 0$, up to errors vanishing when the time step Δt goes to 0. Therefore, we approximate the eigenvector h_k by many applications of the evolution operator (this is nothing else than the power method for matrices). However, the expectation in (1.125) cannot be computed exactly during a simulation. To address this issue, a simple idea is to use a stochastic approximation, which builds on the fact that

$$\mathbb{E}_x [\varphi(x_{n+1})] \approx \varphi(x_{n+1})$$

where x_{n+1} is drawn with respect to $Q_{\Delta t}(x, \cdot)$. We thus obtain a practical implementation of the power method for this stochastic dynamics.

The last ingredient is to use at each step the current estimate at time n of the eigenvector h_k to bias on the fly the dynamics. The variance of the estimator is thus reduced as the dynamics evolves, which is crucial for the algorithm to be efficient. We illustrate the algorithm on one dimensional reversible and irreversible dynamics. This demonstrates that the method can be efficient by running only one trajectory (compared to the genetic algorithm presented in Section 1.4.2, which requires many replicas of a system).

1.5.2.3 Approximate controls for low temperature problems

Chapter 6 proposes a variance reduction technique for the slightly different problem of low temperature systems. Our goal is to compute efficiently the quantity

$$A_\varepsilon = \mathbb{E}_x \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right], \quad (1.126)$$

where $(X_t^\varepsilon)_{t \geq 0}$ is solution to

$$dX_t^\varepsilon = b(X_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma(X_t^\varepsilon) dB_t,$$

and $\varepsilon > 0$ is a small temperature parameter. As explained in Section 1.1.3, a naive estimator based on drawing independent trajectories has a very large variance, because the expectation (1.126) is dominated by a few rare trajectories taking very large values.

However, the modified dynamics

$$d\tilde{X}_t^\varepsilon = b(\tilde{X}_t^\varepsilon) dt + D(\tilde{X}_t^\varepsilon) \nabla g_\varepsilon(t, \tilde{X}_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma(\tilde{X}_t^\varepsilon) dB_t, \quad (1.127)$$

where¹⁷

$$g_\varepsilon(t, x) = \varepsilon \log \mathbb{E}_{t,x} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right], \quad (1.128)$$

leads to a zero variance estimator. This can be seen by a computation very similar to that performed in Section 1.4.2, see Proposition 6.2. However, from a practical point of view, this strategy is not so useful because computing the optimal control (1.128) is harder than computing (1.126).

The idea of Chapter 6 is to use the instanton theory presented in Section 1.1.3 to build an amenable approximation of g_ε , still useful in practical situations. In particular, we construct *offline* an approximate control, which can then be plugged in (1.127). This is an important improvement over previous works such as [414], since we do not need to estimate (1.128) at each time and position.

The first level of our approximation is the following choice for approximating the optimal control g_ε (see Section 1.1.3 for the definition of the instanton):

$$g(t, x) = \theta_t \cdot (x - \phi_t). \quad (1.129)$$

In this case $\nabla g(t, x) = \theta_t$, so the drift is position independent. Intuitively, θ_t is a momentum pushing the process along the instanton. We show formally that the resulting modified dynamics provides a finite variance estimator in the small temperature regime.

We next build a higher order control by setting

$$g(t, x) = \theta_t \cdot (x - \phi_t) + \frac{1}{2} (x - \phi_t) \cdot K_t (x - \phi_t), \quad (1.130)$$

¹⁷By $\mathbb{E}_{t,x}$ we mean that the dynamics is started at time t from position x .

where $(K_t)_{t \in [0, T]}$ is a time dependent matrix. However, contrarily to the control (1.129), there is no natural expression for K_t . We show in Section 6.3.2 that, in order to reduce the variance, K_t should solve the following Riccati equation:

$$\begin{cases} \partial_t K_t + (\nabla b)^T K_t + K_t \nabla b + (\nabla^2 b)^T \theta_t + K_t^T D K_t = 0, \\ K_T = \nabla^2 f(\phi_T). \end{cases}$$

Interestingly, we formally prove in Proposition 6.3 that our construction (1.130) is actually the Taylor expansion of the optimal drift (1.128) around the reaction path $(\phi_t)_{t \in [0, T]}$ as $\varepsilon \rightarrow 0$. We believe that proving rigorously this result is an interesting mathematical challenge. We illustrate the efficiency of the method for reducing the variance of (1.126) on low dimensional reversible and irreversible dynamics.

1.5.3 Coulomb gases and random matrices

The subject of Part IV somehow differs from the two previous parts, since it is concerned with *static* singular gas models arising in random matrix theory (as discussed in Section 1.1.4). We however draw a connection with the dynamical systems studied before, which are used as numerical tools to sample from the electrostatic Gibbs measure. This actually led us to studying Coulomb gases conditioned on satisfying an equality or inequality constraint, possibly with interesting applications to random matrix models.

1.5.3.1 Sampling Coulomb gases with Hamiltonian Monte Carlo

We have discussed in Sections 1.1.4 and 1.3.3 the model of a random gas $X_N = (X_{N,1}, \dots, X_{N,N})$ distributed according to the following Gibbs measure over $(\mathbb{R}^d)^N$:

$$P_N(dx) = Z_N^{-1} e^{-\beta N \left(\frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} K(x_i - x_j) \right)} dx. \quad (1.131)$$

In particular we have seen that, under appropriate growth conditions on the confinement potential V and regularity of the interaction kernel K , the empirical measure

$$\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$$

concentrates under P_N on a minimizer μ_* of the electrostatic energy \mathcal{E} defined by

$$\mathcal{E} : \mu \in \mathcal{P}(\mathbb{R}^d) \mapsto \int_{\mathbb{R}^d} V(x) \mu(dx) + \iint_{\mathbb{R}^d \times \mathbb{R}^d} K(x-y) \mu(dx) \mu(dy). \quad (1.132)$$

Even though the minimizer μ_* can often be proved to satisfy an integral equation, it is typically difficult to numerically estimate because the support of this measure is unknown. The goal of Chapter 7 is to design a method to compute numerically the distribution of μ_N for N finite but (hopefully) large.

When a random matrix model is available and identifies the random variable X_N distributed according to (1.131) with the spectrum of a random matrix, it is obviously quite efficient to draw random matrices to sample their spectrum. However, for many physical applications, such a model is not available [387], while the behaviour of μ_N as N becomes large is still of interest. Chapter 7 builds on the simple idea that (1.131) is a Gibbs measure, which can be sampled for instance with the overdamped or underdamped Langevin dynamics presented in Section 1.1.1. Although the numerical integration is typically made difficult by the singularity of the interaction kernel K , this problem can be addressed efficiently with a Hamiltonian Monte Carlo algorithm [142, 153, 60]. We apply the method to various models, from which we can recover well-known distributions and explore more original conjectures. Interestingly, our work [82] has already been used in [76] to numerically observe theoretical results, and it motivated the analysis performed in [316], see Remark 7.2.

1.5.3.2 Coulomb gases under constraint

Chapter 8 is a natural sequel of the previous chapter. Our motivation lies in the unlikely fluctuations of the spectrum of random matrices. Consider for instance a random matrix M of size N distributed

according to the Ginibre Ensemble described in Section 1.1.4. In this case, denoting by $(X_{N,i})_{i=1}^N$ the eigenvalues of M , the trace of the matrix

$$\mathrm{Tr}(M) = \sum_{i=1}^N X_{N,i}$$

has expectation zero. However, fluctuations may translate the trace of the spectrum by some amount, and our first aim was to numerically study this phenomenon. This can be motivated by applications, for instance in finance [65, 66], where the fluctuations of the spectrum of Wishart matrices can be important. Building on the analogy explained above, a certain class of random matrix models can be associated with a singular gas distributed according to (1.131). We therefore consider the measure P_N defined in (1.131) and our goal is to sample from the empirical distribution associated with the random variable

$$Y_N \sim \mathrm{Law}(X_N \mid \xi_N(X_N) \leq 0), \quad (1.133)$$

where $\xi_N : (\mathbb{R}^d)^N \rightarrow \mathbb{R}$ is a constraint function. In order to simplify the analysis, we assume that $K = g$ is the Coulomb kernel (see the definitions in Section 1.1.4). We however believe that other singular interactions are suitable for our analysis provided the empirical measure μ_N satisfies a large deviations principle under P_N .

In Chapter 7 we used a Hamiltonian Monte Carlo (HMC) algorithm to sample from the empirical distribution associated with X_N . Interestingly, in the computational physics literature, there has been a great interest in *constrained* versions of the HMC algorithm, in particular in the context of free energy computations [106, 302, 430]. Our first contribution is to use such an algorithm [303] to sample from the empirical measure

$$\frac{1}{N} \sum_{i=1}^N Y_{N,i}, \quad (1.134)$$

where Y_N is defined in (1.133). As a first case, we consider the Ginibre Ensemble with a trace constraint, *i.e.* $d = 2$ and

$$\forall x \in (\mathbb{R}^2)^N, \quad \xi_N(x) = c - \frac{1}{N} \sum_{i=1}^N x_i \cdot v, \quad (1.135)$$

with $v \in \mathbb{R}^2$, $|v| = 1$ and $c \in \mathbb{R}$. We numerically observe that, in this situation, the uniform distribution on the unit disk μ_* is simply translated by a factor c in the direction v . This surprising result motivates a more precise theoretical analysis of the conditioning phenomenon.

Our analysis follows two steps. We first consider the very particular case of Gaussian ensembles (when $V = |\cdot|^2$ and g is a translation invariant kernel) under trace constraint. We prove in Theorem 8.1 that these Gaussian ensembles have a spectacular factorization property: under P_N , the barycenter of the cloud of points is independent from the remainder of the distribution. As a result, the linear conditioning (1.135) only modifies the center of the distribution, without further modifications.

In order to understand more general constraints than (1.135) we resort to the large deviations machinery, in particular the Gibbs conditioning principle presented in Section 1.2.2. We first consider the case of *linear statistics constraints*, namely

$$\forall x \in (\mathbb{R}^d)^N, \quad \xi_N(x) = \frac{1}{N} \sum_{i=1}^N \varphi(x_i), \quad (1.136)$$

where $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$. The constraint (1.135) then corresponds to the special choice $\varphi(x) = c - x \cdot v$. In order to understand the limiting behaviour of (1.134) with Y_N defined by (1.133), we apply the Gibbs conditioning principle (Proposition 1.13 above) with a subset of $\mathcal{P}(\mathbb{R}^d)$ of the form

$$B = \{\nu \in \mathcal{P}(\mathbb{R}^d) \mid \nu(\varphi) \leq 0\}.$$

We are thus naturally led to the following minimization problem

$$\inf_B \mathcal{E},$$

where \mathcal{E} is defined for instance in (1.132). This turns out to be a tedious problem for two reasons:

1. The space of probability measures is not a vector space, hence we cannot rely on standard strategies for optimizing under constraint;
2. The electrostatic energy \mathcal{E} is very singular. For instance the domain $D_{\mathcal{E}} = \{\nu \in \mathcal{P}(\mathbb{R}^d) \mid \mathcal{E}(\nu) < +\infty\}$ has empty interior for any reasonable topology on $\mathcal{P}(\mathbb{R}^d)$!

In the end, we can however show under reasonable conditions on V and φ that Proposition 1.13 applies when $\mathcal{P}(\mathcal{X})$ is endowed with the p -Wasserstein topology for $p \geq 1$, see Theorem 8.9. This shows that the empirical measure (1.133) concentrates on the unique probability measure μ^φ satisfying

$$\{\mu^\varphi\} = \operatorname{argmin}_B \mathcal{E}.$$

Moreover, we are able to characterize μ^φ with an equation similar to (1.57) by showing that there exists $\alpha \geq 0$ such that¹⁸

$$2g * \mu^\varphi + V + \alpha\varphi = C_\varphi,$$

over the support of μ^φ , and where $C_\varphi = \mathcal{E}(\mu^\varphi)$, see Remark 8.13. This shows that conditioning with a linear statistics amounts to modifying the confinement potential V with the constraint function φ . We next perform the same analysis for quadratic statistics of the form

$$\forall x \in (\mathbb{R}^d)^N, \quad \xi_N(x) = \frac{1}{N^2} \sum_{i,j=1}^N \psi(x_i - x_j),$$

for which we refer to Theorem 8.16.

We finally come back to our simulation algorithm to numerically estimate the distribution of the empirical measure (1.134) for some choices of V , g , and constraint function. When the solution is explicit, our simulations confirm the theory. They also allow to estimate μ^φ when no analytical solution is available.

¹⁸Here, α has the natural interpretation of a Lagrange multiplier for the linear constraint $\nu \in \mathcal{P}(\mathbb{R}^d) \mapsto \nu(\varphi) \in \mathbb{R}$.

Bibliographical comments

I provide in this short section some comments on the bibliography used to complete this work. My aim is to highlight some papers and books that significantly contributed to my understanding of the problems along the way. This bibliography is purely personal and does not aim at being comprehensive by any mean.

Feynman–Kac dynamics and large deviations

Concerning the long time analysis of SDEs and Markov chains, the book chapter [362] contains a lot of useful tools, including control arguments and spectral theory of evolution operators. The paper [219] also provides a very nice proof of exponential ergodicity, while [304] is a great reference for both Lyapunov function techniques and hypocoercivity techniques à la Villani. I also enjoyed reading the course of M. Hairer [217] on ergodic theory. The references [324, 295] give interesting insights in the particular case of the Langevin equation (in particular the influence of the friction in the Langevin dynamics in [295]). Concerning ergodicity of Feynman–Kac dynamics, I liked [114] and the proofs in [377], but I also found interesting elements of understanding in [301], including for numerical applications. Although not directly related, the monograph [419] is a very good introduction to optimal transport problems and Wasserstein distances, where I found several useful elements on the topology associated with Wasserstein metrics. The classical books [361, 262] provided many useful stochastic analysis tools, while [112] has an excellent chapter on problems over cones (in particular about the Krein–Rutman theorem).

As far as large deviations are concerned, the book by A. Dembo and O. Zeitouni [119] has been a great source of inspiration, both for the intuition built in Chapter 1 and the powerful results of Chapter 4. I also found several technical elements in [124]. On the other hand, I am greatly indebted towards the papers of Touchette, Chetrite and collaborators [406, 408, 97, 98, 25, 410, 411, 342] for all their ideas and results, and their writing of rare quality. Concerning the particular problem of long time large deviations in weighted topology, the main source of ideas is the paper of L. Wu [427], where I found in particular the key martingale allowing for the spectral analysis developed in Chapter 3. Other interesting technical elements were found in [275, 424]. Finally, [49, 363] provided interesting insights on the influence of irreversibility on large deviations.

Numerical analysis and algorithms

For the numerical analysis of Chapter 4, I liked the very good presentation of [304]. The seminal works of Talay and Tubaro also remain a very instructive reading [400, 398, 399]. The construction of the adjusted invariant measure in [325] was very useful, and inspired the approximate eigenvector technique developed in Chapter 4, Section 4.3. There are other nice references on the error on the invariant measure of SDEs, see for instance [324, 295] and the other references in Chapter 4.

Concerning the algorithm designed in Chapter 5, I again refer to [98] for a nice background on the optimal control technique used there, while the works of Borkar and collaborators, in particular [58, 28], initiated most of our ideas. On the other hand, the important book [30] provides a clear analysis of the stochastic approximation techniques used in this chapter.

For the small temperature problems of Chapter 6, the recent review [211] provides a good entry point for the theory of instantons. Of course, the celebrated book of Freidlin and Wentzell [193] is a great source for the theory of small temperature dynamical systems, while [352] is a good introduction to stochastic control in general.

Coulomb gases

The study on random gases studied in Part IV was first motivated by an introductory course of D. Chafaï relating Coulomb and log-gases to random matrix problems. I then particularly enjoyed reading [84] as well as the seminal paper by Ben Arous and Guionnet [29]. There are very nice textbooks about singular gases and potential theory, such as [284, 380], which both helped me to build intuition and to find precise proofs of important results. I then liked reading the book by Tao [401] to discover other techniques of proofs for random matrices, but there are obviously other good textbooks on the subject such as [327, 9, 188]. Moreover, we mention that the book [26] is a great reference on abstract optimization, which helped me a lot for solving difficult constrained optimization problems on convex (non-vector) spaces.

Notational comments

Although the notation will be recalled in each chapter, I make here some general comments in order to facilitate the reader's task, starting with latin alphabet. In all the manuscript, \mathcal{X} stands for the state space of the system (in general a subset of \mathbb{R}^d , except in Chapter 2 where it denotes a Polish space), while \mathcal{Y} denotes the state in which a LDP takes place (which may vary quite a lot as seen in Section 1.3). The letter X is used for random variables and stochastic processes, while x (and small capital letters like y, z) typically denotes a point in \mathcal{X} . The letter d stands for the dimension, but is also used for infinitesimal elements (like dx for the Lebesgue measure on \mathbb{R}^d). We believe that this does not lead to any confusion. We generally write d for some distance on a metric space (for instance a Wasserstein distance). We denote by $\mathcal{P}(\mathcal{X})$ the set of probability measures on \mathcal{X} , and generally $\mathcal{M}(\mathcal{X})$ for the set of measures on \mathcal{X} (but \mathcal{M} is sometimes used to denote a submanifold). The letter M may be a constant, or a local martingale when indexed by time as M_t , just like \mathcal{M}_t . The letters W, \mathcal{W} are always used for Lyapunov functions, \mathcal{L} is the generator of an SDE with drift b and diffusion coefficient σ (although σ is used to denote the spectrum of an operator at a few places); \mathcal{D} denotes the domain of an operator \mathcal{L} (or $\mathcal{D}_{\mathcal{L}}$ to make the operator precise). Note that b sometimes denotes a real number, like a or c , but the notation is in general kept for the drift of a diffusion.

Capital letters A and B generally denote sets or matrices, C a positive constant, D the definite diffusion matrix $\sigma\sigma^T$ of an elliptic diffusion (we write S when it is symmetric but not necessarily definite, like in Chapter 3). Note that we also use the letter D_I to denote the domain of a functional I in the sense of convex analysis (*i.e.* the set of points where the functional is finite), while we use $C^0, C_b, C^k, C^\infty, C_c^\infty$ to denote respectively the space of continuous, continuous bounded, k -times differentiable with continuous derivatives, infinitely differentiable and infinitely differentiable with compact support functions. In general E is an abstract Banach space, with norm $\|\cdot\|_E$, and we use the notation $\mathcal{B}(E)$ for the space of linear operators on E . However, note that B is also used to define the space of bounded functions $B^\infty(\mathcal{X})$ and its weighted variants. The capital letter L denotes an empirical measure, except for the low temperature systems where it is sometimes used for the action functional, and H denotes a Hamiltonian. Note that the letter L is also used for the Lebesgue $L^p(dx)$ spaces, while H is used for standard Sobolev spaces such as $H^1(\mu)$ (while \mathcal{H} is reserved for the particular construction of Sobolev space in Chapter 3). The letters P and Q are generally used to denote continuous time and discrete time evolution operators respectively, while K denotes a compact set in Part II, a time dependent matrix in Chapter 6 and a kernel in Part IV.

Indices are generally indicated by i, j, k, ℓ, m, n , while p, q denote momenta and positions in the Langevin equation (for instance in Chapter 3) or powers for the Wasserstein distance (in particular in Part IV). Finally, f and g denote observable functions, while the notation h is reserved for a principal eigenvector of an operator (associated with the largest eigenvalue). The letter T is used for both a final time and matrix transposition (sometimes denoted by T), while s, t generally denote positive times.

Concerning greek letters, φ and ϕ often denote test functions, but they also used for paths in \mathcal{X} . In general, we keep the notation ψ for the solution to a PDE (generally a Poisson problem), while Φ is a (Feynman–Kac) semigroup, and Ψ, κ are Lyapunov potentials like in Chapter 3. We always use the letter λ (or Λ) for the principal eigenvalue (or cumulant function or free energy) of a Feynman–Kac operator, while α denotes an arbitrary real or complex variable, β is generally an inverse temperature, γ a friction and δ, ε small real numbers. Probability measures are always denoted by ν, μ, π and η (although η may be a small constant at a few places). The notation Π is used for projectors over a subspace or submanifold. The capital letter Γ usually refers to a subset of $\mathcal{P}(\mathcal{X})$.

Finally, we recall that we generally write \cdot for a scalar product and $:$ for a double contraction (*i.e.* $A : B = \text{Tr}(A^T B)$ for two matrices A, B). The symbol \otimes is used for tensorization at a few places. The standard norm on \mathbb{R}^d is written $|\cdot|$, but this notation is also used for the Lebesgue measure of sets with no ambiguity. The symbol Δ is the Laplacian operator, ∇ is the gradient and $\nabla \cdot$ is the divergence operator.

Part II

Feynman–Kac dynamics and large deviations: Theoretical analysis

CHAPTER 2

LONG TIME ANALYSIS OF FEYNMAN–KAC DYNAMICS

The material for this chapter has been released in [181] and is currently under review.

Contents

2.1	Motivation	76
2.2	Results	77
2.2.1	Framework	77
2.2.2	Results in discrete time	78
2.2.3	Results in continuous time	83
2.3	Applications	87
2.3.1	Examples in discrete time	87
2.3.2	Applications to diffusion processes	90
2.3.3	Convergence results uniform with respect to the time step	91
2.4	Discussion	94
2.5	Complementary elements	95
2.5.1	Stability of Markov chains	95
2.5.2	Useful theorems	95
2.5.3	Proof of Lemma 2.4	97
2.5.4	Proof of Lemma 2.5	98
2.5.5	Proof of Lemma 2.6	100
2.5.6	Proof of Lemma 2.20	101

Abstract. Feynman–Kac semigroups appear in various areas of mathematics: non-linear filtering, large deviations theory, spectral analysis of Schrödinger operators among others. Their long time behavior provides important information, for example in terms of ground state energy of Schrödinger operators, or scaled cumulant generating function in large deviations theory (see Section 1.3.1 in the Introduction). In this chapter, we propose a simple and natural extension of the stability of Markov chains for these non-linear evolutions. As other classical ergodicity results, it relies on two assumptions: a Lyapunov condition that induces some compactness, and a minorization condition ensuring some mixing. We show that these conditions are satisfied in a variety of situations. We also use our technique to provide uniform in the time step convergence estimates for discretizations of stochastic differential equations. As we will see in Chapter 3, the tools developed below are useful to obtain interesting large deviations results.

2.1 Motivation

Feynman–Kac semigroups have a long history in physics and mathematics. One of their traditional applications as a probabilistic representation of Schrödinger semigroups [262] is the computation of ground state energies through Diffusion Monte Carlo algorithms [213, 10, 77, 190]. It has then become a significant tool in non-linear filtering and genealogical models [115, 117, 113]. We have also motivated in the Introduction (Section 1.3.1) its relation with large deviations theory [133, 275, 202, 427] in the context of empirical measures of diffusions. In all these contexts, the dynamics is evolved and its paths are weighted depending on some cost function. This function is typically a potential energy, a likelihood, or a function whose fluctuations are of interest.

As for Markov chains, the long time behavior of such dynamics is important, in particular to define the cumulant function in the context of large deviations. However, the long-time analysis is made difficult by the non-linear character of the evolution, so the methods used for the stability of Markov chains [330, 219] cannot be straightforwardly adapted in this context. A series of papers [114, 116, 113] rely on the powerful Dobrushin ergodic coefficient [127, 128]. Although this tool enables to deal with the nonlinearity and to consider time-inhomogeneous processes, the conditions imposed on the dynamics are not realistic for unbounded domains.

The purpose of this chapter is to propose a new scheme of proof for the ergodicity of Feynman–Kac dynamics, suitable for cases where the state space is unbounded. It is based on the principal eigenvalue problem associated to a weighted evolution operator. It then relies on studying a h -transformed version of the dynamics [136] (also called Doob transform), where h is the eigenvector associated to the eigenproblem. This turns the non-linear dynamics into a linear Markov evolution, which can then be studied with standard techniques [330, 219]. However, the spectral properties of the generator fall out of the typical regime of self-adjoint operators, since the dynamics is in general non-reversible. A striking fact of our results is that, under Lyapunov and minorization conditions similar to those of [219] stated for non-probabilistic kernels, we perform a non self-adjoint spectral analysis that recasts the Feynman–Kac problem into the Markov chain framework studied in [219].

The works of Kontoyannis and Meyn [274, 275] provide elements of answer concerning the spectral properties of the evolution operator, and rely on a nonlinear Lyapunov condition and a regularity in terms of hitting times. If the latter Lyapunov condition is natural in terms of optimal stochastic control [186], we propose instead proofs based on linear conditions. Our generalized linear Lyapunov condition is inspired by [362], and comes together with a minorization condition and a local strong Feller assumption. We will see that these conditions apply to a variety of situations, with natural interpretations. From a broader perspective, it appears as a natural extension of previous works on the stability of Markov chains [219] for evolution kernels that do not conserve probability. To that extent, our work resonates with recent works on Quasi-Stationary Distributions (QSD) [205, 90, 89, 23]. However, our scope and assumptions being different, we leave the comparison for future studies. Let us also mention that our framework applies for both discrete and continuous time processes. This is interesting since one motivation for this work is to understand the behavior of time discretizations of continuous Feynman–Kac dynamics, see Chapter 4 below.

Let us outline our main results in an informal way. The quantities we are interested in typically correspond to Markov chains $(x_k)_{k \geq 0}$ over a state space \mathcal{X} , whose trajectories are weighted by a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$. The associated Feynman–Kac semigroups is of the form

$$\Phi_k(\nu)(\varphi) = \frac{\mathbb{E} \left[\varphi(x_k) e^{\sum_{i=0}^{k-1} f(x_i)} \mid x_0 \sim \nu \right]}{\mathbb{E} \left[e^{\sum_{i=0}^{k-1} f(x_i)} \mid x_0 \sim \nu \right]}, \quad (2.1)$$

where ν is an initial probability distribution, and φ is a test function. We show that, for more general semigroups and under some assumptions on $(x_k)_{k \geq 0}$ and f , there exists a measure μ_f such that for any initial probability measure ν and any φ belonging to a particular class of test functions,

$$\Phi_k(\nu)(\varphi) \xrightarrow[k \rightarrow +\infty]{} \mu_f(\varphi), \quad (2.2)$$

at an exponential rate. As a corollary of this result, we show that the principal eigenvalue Λ of the generator of the dynamics $(\Phi_k)_{k \geq 1}$ can be obtained as the following limit, for any suitable functions f and independently of the initial probability measure ν , it holds

$$\log(\Lambda) = \lim_{k \rightarrow +\infty} \frac{1}{k} \log \mathbb{E} \left[e^{\sum_{i=0}^{k-1} f(x_i)} \mid x_0 \sim \nu \right],$$

which is the scaled cumulant generating function in large deviations theory [119, 275], as explained in Section 1.3.1. Another natural situation corresponds to continuous semigroups of the form

$$\Phi_t(\nu)(\varphi) = \frac{\mathbb{E} \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \mid x_0 \sim \nu \right]}{\mathbb{E} \left[e^{\int_0^t f(X_s) ds} \mid x_0 \sim \nu \right]}, \quad (2.3)$$

where $(X_t)_{t \geq 0}$ is typically a diffusion process. Note that we use the same notation for the discrete and continuous processes and we believe there is no ambiguity here (but in Chapter 4 we explicitly write $\Phi_{\Delta t, k}$ for discretizations of continuous Feynman–Kac dynamics). Similar results are then derived for this continuous dynamics. We will see that ergodic properties such as (2.2) are proved under natural extensions of Lyapunov and minorization conditions, which should be reminiscent of the corresponding theory for Markov chains [219, 362], with additional regularity conditions. We emphasize again that the convergence results obtained here are useful to obtain large deviations principle in the next chapter.

The chapter is organized as follows. In Section 2.2, we present our main results on the stability of Feynman–Kac semigroups. Section 2.2.2 is devoted to discrete time results, while Section 2.2.3 is concerned with the continuous time case. Section 2.3 presents a number of natural applications of the method. In particular, Section 2.3.3 provides uniform in the time step convergence estimates. Section 2.4 discusses some links with related works and possible further directions.

2.2 Results

2.2.1 Framework

In this section, we present our main convergence results for generalizations of the dynamics (2.1). The state space \mathcal{X} is assumed to be a Polish space, and for a measurable set $A \subset \mathcal{X}$, we denote by A^c its complement, and $\mathbb{1}_A$ its indicator function. For a Banach space E , we denote by $\mathcal{B}(E)$ the space of bounded linear operators over E , with associated norm $\|T\|_{\mathcal{B}(E)} = \sup \{\|Tu\|_E, \|u\|_E \leq 1\}$. The Banach space of continuous functions is called $C^0(\mathcal{X})$, and the Banach space of measurable functions φ such that

$$\|\varphi\|_{B^\infty} = \sup_{x \in \mathcal{X}} |\varphi(x)| < +\infty$$

is referred to as $B^\infty(\mathcal{X})$. Given a measure ν over \mathcal{X} with finite mass, we use the notation $\nu(\varphi) = \int_{\mathcal{X}} \varphi(x) \nu(dx)$ for $\varphi \in B^\infty(\mathcal{X})$. The spaces of positive measures and probability measures over \mathcal{X} are denoted respectively by $\mathcal{M}(\mathcal{X})$ and $\mathcal{P}(\mathcal{X})$. When we consider Markov chains $(x_k)_{k \in \mathbb{N}}$ over \mathcal{X} , we write \mathbb{E}_ν for the expectation over all the realizations of the Markov chain with initial condition distributed according to the probability measure ν . Section 2.5.1 is devoted to reminders on the ergodicity of Markov chains extracted from [219], while Section 2.5.2 recalls some useful definitions and theorems used in the proofs of the results of this section.

We consider general kernel operators Q^f over \mathcal{X} , *i.e.* such that for any $x \in \mathcal{X}$, $Q^f(x, \cdot)$ is a positive measure with finite mass (*i.e.* $Q\mathbb{1}(x) < +\infty$), and for any measurable set $A \subset \mathcal{X}$, $Q^f(\cdot, A)$ is a measurable function. Such a kernel is referred to as Markov (also probabilistic or conserving) when $Q^f\mathbb{1} = \mathbb{1}$. The notation Q^f instead of Q emphasizes that in general $Q^f\mathbb{1} \neq \mathbb{1}$ depends on a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$. For $\varphi \in B^\infty(\mathcal{X})$, we denote by $Q^f\varphi = \int_{\mathcal{X}} \varphi(y) Q^f(\cdot, dy)$ the action of Q^f on test functions, and by $\nu Q^f = \int_{\mathcal{X}} \nu(dx) Q^f(x, \cdot)$ its action on finite measures ν . We call Feynman–Kac semigroups the dynamics $(\Phi_k)_{k \geq 1}$ defined as follows:

$$\forall k \geq 1, \quad \forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in B^\infty(\mathcal{X}), \quad \Phi_k(\nu)(\varphi) = \frac{\nu((Q^f)^k \varphi)}{\nu((Q^f)^k \mathbb{1})}. \quad (2.4)$$

Note that $\Phi_k = \Phi \circ \dots \circ \Phi$, where Φ is the one step evolution operator $\Phi : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$:

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in B^\infty(\mathcal{X}), \quad \Phi(\nu)(\varphi) = \frac{\nu(Q^f \varphi)}{\nu(Q^f \mathbb{1})}, \quad (2.5)$$

which is well-defined as soon as $\nu(Q^f \mathbb{1}) > 0$ for any $\nu \in \mathcal{P}(\mathcal{X})$. Lemma 2.4 below proves that (2.5) is indeed well-defined under the assumptions presented in Section 2.2.2.

Although Q^f is not probabilistic, the normalizing factor in (2.5) ensures that Φ evolves a positive measure of finite mass into a probability measure. An important motivation for studying the general

dynamics (2.5) is that (2.1) can be written in the form (2.4) with $Q^f = e^f Q$, where Q is the transition operator of the Markov chain $(x_k)_{k \in \mathbb{N}}$. In this typical setting, $Q^f \mathbf{1} = e^f$. Even when Q^f is not defined in this way (see for instance the continuous time situation (2.30) considered in Section 2.2.3), we keep the notation to emphasize that Q^f typically corresponds to a Markov dynamics whose trajectories are weighted by a function f .

2.2.2 Results in discrete time

We now introduce the assumptions ensuring the well-posedness and ergodicity of the semigroup (2.4), which should be reminiscent of the ones used in [219, 362] for showing the ergodicity of Markov chains. We start with discrete systems, before considering the continuous counterpart in Section 2.2.3. The first step of the proof is the existence of a principal eigenvector h for Q^f , as shown in Lemma 2.5. This eigenvector is used in Lemma 2.6 to study a h -transformed version of Q^f , which leads to our main result, Theorem 2.7. Note that, in practice, we have in mind the situation $\mathcal{X} = \mathbb{R}^d$ for $d \in \mathbb{N}^*$, but discrete spaces like $\mathcal{X} = \mathbb{Z}^d$ can also be considered, in which case the framework may be simplified.

The first assumption is that a generalized Lyapunov condition holds. We will see in Section 2.3 that it is satisfied for a large class of processes. In all this section, we consider an increasing sequence of compact sets $(K_n)_{n \geq 1}$ such that, for any compact $K \subset \mathcal{X}$, there exists $m \geq 1$ for which $K \subset K_m$.

Assumption 2.1 (Lyapunov condition). *There exist a function $W : \mathcal{X} \rightarrow [1, +\infty)$ bounded on compact sets, and positive sequences $(\gamma_n)_{n \geq 1}$, $(b_n)_{n \geq 1}$ with $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$ such that, for all $n \geq 1$,*

$$Q^f W \leq \gamma_n W + b_n \mathbf{1}_{K_n}. \quad (2.6)$$

Let us mention that, in many situations, the function W has compact level sets, so that a natural choice of compact sets is $K_n = \{x \in \mathcal{X} \mid W(x) \leq n\}$. When a Lyapunov function W exists, it is natural [219] to consider the following functional space

$$B_W^\infty(\mathcal{X}) = \left\{ \varphi \text{ measurable, } \left\| \frac{\varphi}{W} \right\|_{B^\infty} < +\infty \right\}, \quad (2.7)$$

as well as the space of measures integrating W , namely

$$\mathcal{P}_W(\mathcal{X}) = \left\{ \nu \in \mathcal{P}(\mathcal{X}) \mid \nu(W) < +\infty \right\}.$$

In particular, Assumption 2.1 implies that Q^f is a bounded operator on $B_W^\infty(\mathcal{X})$, since one easily shows that

$$\forall n \geq 1, \quad \|Q^f\|_{\mathcal{B}(B_W^\infty)} \leq \gamma_n + b_n.$$

We next assume that the following minorization condition holds.

Assumption 2.2 (Minorization and irreducibility). *For any $n \geq 1$, there exist $\eta_n \in \mathcal{P}(\mathcal{X})$ and $\alpha_n > 0$ such that*

$$\inf_{x \in K_n} Q^f(x, \cdot) \geq \alpha_n \eta_n(\cdot). \quad (2.8)$$

In addition, for any $n_0 \geq 1$ and any $\varphi \in B_W^\infty(\mathcal{X})$ with $\varphi \geq 0$,

$$\eta_n(\varphi) = 0, \forall n \geq n_0 \implies (Q^f \varphi)(x) = 0, \forall x \in \mathcal{X}. \quad (2.9)$$

Note that (2.9) expresses some form of irreducibility with respect to the minorizing measures. It can be reformulated in the following way: for any $n_0 \geq 1$ and any $x \in \mathcal{X}$, $Q^f(x, \cdot)$ is absolutely continuous with respect to the measure

$$\sum_{n \geq n_0} 2^{-n} \eta_n.$$

The typical situation for $\mathcal{X} = \mathbb{R}^d$ is to choose $\eta_n(dx) = \mathbf{1}_{K_n}(x) dx / |K_n|$, where $|K_n|$ denotes the Lebesgue measure of K_n . We also mention that, although we will consider the previous minorization measures η_n in our examples in Section 2.3, the first part of Assumption 2.2 can be obtained using irreducibility together with a strong Feller property, see [216], or through the Stroock–Varadhan support theorem [395] with some regularity property, see the discussion in [362]. In our context, we also need some local regularity for the operator Q^f .

Assumption 2.3 (Local regularity). *The operator Q^f is strong Feller on the compact sets K_n , i.e. for any $n \geq 1$ and any measurable function φ bounded on K_n , $Q^f(\varphi \mathbf{1}_{K_n})$ is continuous over K_n .*

From these assumptions we first state the following preliminary lemma, whose proof can be found in Section 2.5.3.

Lemma 2.4. *Let Q^f satisfy Assumptions 2.1 and 2.2. Then, for any $\nu \in \mathcal{P}_W(\mathcal{X})$, one has*

$$0 < \nu(Q^f \mathbf{1}) < +\infty. \quad (2.10)$$

Moreover, for any $n \geq 1$, it holds $1 \leq \eta_n(W) < +\infty$, and there exist infinitely many indices $\bar{n} \geq 1$ such that

$$\eta_{\bar{n}}(K_{\bar{n}}) > 0. \quad (2.11)$$

The lower bound in (2.10) implies in particular that the dynamics (2.4) is well-defined. The inequality (2.11) means that, for infinitely many minorization conditions, some mass of the minorizing measure remains in the associated compact set. It is used in the proof of Lemma 2.5 to show that Q^f has a positive spectral radius. Since (2.11) is satisfied for infinitely many indices, we could consider that it holds for any $n \geq 0$, upon extracting a subsequence and, in the situations considered in Section 2.3, we can actually check that $\eta_n(K_n) > 0$ for all $n \geq 1$.

We are now in position to state some spectral properties of the operator Q^f , which are a key ingredient for our analysis. Let us recall that the spectral radius of Q^f on $B_W^\infty(\mathcal{X})$, denoted by $\Lambda = \Lambda(Q^f)$, is given by the Gelfand formula [345]:

$$\Lambda = \lim_{k \rightarrow +\infty} \|(Q^f)^k\|_{B(B_W^\infty)}^{\frac{1}{k}}, \quad (2.12)$$

and that the essential spectral radius of Q^f , denoted by $\Theta(Q^f)$, reads (see Section 2.5.2):

$$\Theta(Q^f) = \lim_{k \rightarrow +\infty} \left(\inf \left\{ \|(Q^f)^k - T\|_{B(B_W^\infty)}, T \text{ compact} \right\} \right)^{\frac{1}{k}}.$$

Lemma 2.5. *Under Assumptions 2.1, 2.2 and 2.3, the operator Q^f considered on $B_W^\infty(\mathcal{X})$ satisfies $\Theta(Q^f) = 0$, admits its spectral radius $\Lambda > 0$ as a largest eigenvalue (in modulus), and has an associated eigenfunction $h \in B_W^\infty(\mathcal{X})$, normalized so that $\|h\|_{B_W^\infty} = 1$, and which satisfies*

$$\forall x \in \mathcal{X}, \quad 0 < h(x) < +\infty. \quad (2.13)$$

In particular, $0 < \eta_n(h) < +\infty$ for all $n \geq 1$.

We prove Lemma 2.5 in Section 2.5.4 by using arguments inspired by [362, Theorem 8.9] to show that the essential spectral radius of Q^f is zero, and then relying on the spectral theory of positive operators [112]. Some useful elements of operator theory are reminded in Section 2.5.2 for the reader's convenience. Note that the eigenspace associated to Λ is a priori not of dimension one. Our result is close to those obtained in [275], and the control of the essential spectral radius under Lyapunov and topological irreducibility conditions has already been studied in [428, 214]. However, our proof uses different techniques based on different assumptions.

Once such a principal eigenvector h is available, the geometric ergodicity of the Feynman–Kac dynamics (2.4) is derived from the one of a h -transformed kernel, as made clear in the proof of Theorem 2.7 below. This is the purpose of the next lemma whose proof is postponed to Section 2.5.5.

Lemma 2.6. *Suppose that Assumptions 2.1, 2.2 and 2.3 hold, and consider an eigenvector h associated with Λ as given by Lemma 2.5. Since $h > 0$ we can define the corresponding h -transformed operator Q_h as*

$$Q_h \phi = \Lambda^{-1} h^{-1} Q^f(h\phi). \quad (2.14)$$

Then Q_h is a Markov operator with Lyapunov function $Wh^{-1} : \mathcal{X} \rightarrow [1, +\infty)$. Moreover, there exist a unique $\mu_h \in \mathcal{P}(\mathcal{X})$, which satisfies $\mu_h(Wh^{-1}) < +\infty$, and constants $c > 0$, $\bar{\alpha} \in (0, 1)$ such that, for any $\phi \in B_{Wh^{-1}}^\infty(\mathcal{X})$ and any $k \geq 1$,

$$\|Q_h^k \phi - \mu_h(\phi)\|_{B_{Wh^{-1}}^\infty} \leq c \bar{\alpha}^k \|\phi - \mu_h(\phi)\|_{B_{Wh^{-1}}^\infty}. \quad (2.15)$$

Although this is not obvious at first glance, the operator Q_h is in fact independent of the choice of h in Lemma 2.5, and so is the invariant measure μ_h . Actually, Lemma 2.6 allows to show that the eigenspace associated with h has geometric dimension one, *i.e.* $\text{Ker}(Q^f - \Lambda I_d) = \text{Span}\{h\}$. Indeed,

if $\tilde{h} \in B_W^\infty(\mathcal{X})$ is another eigenvector associated with Λ (which may not be of constant sign), it holds, since $h(x) > 0$ for all $x \in \mathcal{X}$ by (2.13):

$$Q_h \left(\frac{\tilde{h}}{h} \right) = \Lambda^{-1} h^{-1} Q^f \tilde{h} = \frac{\tilde{h}}{h} \in B_{W h^{-1}}^\infty(\mathcal{X}).$$

From (2.15), we obtain

$$\frac{\tilde{h}}{h} = \mu_h \left(\frac{\tilde{h}}{h} \right),$$

hence \tilde{h} is proportional to h . It may be possible to directly obtain this uniqueness result from stronger Krein–Rutman theorems, like [112, Theorem 19.3], using the irreducibility condition (2.9) in Assumption 2.2. Under our assumptions, it is also possible to show that there cannot be another positive eigenvector (see Chapter 3, in particular the proof of Lemma 3.32).

We are now in position to state our main theorem.

Theorem 2.7. *Consider a kernel operator Q^f satisfying Assumptions 2.1, 2.2 and 2.3 and the associated dynamics (2.4) with one step evolution operator $\Phi : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$. Then Φ admits a unique fixed point $\mu_f \in \mathcal{P}(\mathcal{X})$, that is a probability measure such that*

$$\Phi(\mu_f) = \mu_f, \quad (2.16)$$

and this measure satisfies $\mu_f(W) < +\infty$. Moreover, there exists $\bar{\alpha} \in (0, 1)$ such that, for any $\nu \in \mathcal{P}_W(\mathcal{X})$, there is $C_\nu > 0$ for which

$$\forall \varphi \in B_W^\infty(\mathcal{X}), \quad \forall k \geq 1, \quad |\Phi_k(\nu)(\varphi) - \mu_f(\varphi)| \leq C_\nu \bar{\alpha}^k \|\varphi\|_{B_W^\infty}. \quad (2.17)$$

We call μ_f the invariant measure of Q^f , in analogy with Markov chains. However, in the situation of the non linear dynamics described by the evolution operator Φ , the invariant measure satisfies the non linear stationary equation (2.16). This remark is central for the numerical analysis performed in Chapter 4. Note that Theorem 2.7 also implies the convergence of $\Phi_k(\nu)$ towards μ_f in the weighted total variation distance [219] defined, for $\mu, \nu \in \mathcal{P}(\mathcal{X})$ with $\mu(W) < +\infty$, $\nu(W) < +\infty$, by

$$\rho_W(\mu, \nu) = \sup_{\|\varphi\|_{B_W^\infty} \leq 1} \int_{\mathcal{X}} \varphi(x) (\mu - \nu)(dx). \quad (2.18)$$

We now present the proof of Theorem 2.7.

Proof. The key idea of the proof is to reformulate the dynamics (2.4) using the h -transformed operator $Q_h = \Lambda^{-1} h^{-1} Q^f h$ of Lemma 2.6. Using the notation of Lemmas 2.5 and 2.6, we rewrite (2.4) as

$$\Phi_k(\nu)(\varphi) = \frac{\nu((Q^f)^k \varphi) \Lambda^{-k}}{\nu((Q^f)^k \mathbb{1}) \Lambda^{-k}} = \frac{\nu(h(\Lambda^{-1} h^{-1} Q^f h)^k (h^{-1} \varphi))}{\nu(h(\Lambda^{-1} h^{-1} Q^f h)^k h^{-1})} = \frac{\nu(h(Q_h)^k (h^{-1} \varphi))}{\nu(h(Q_h)^k h^{-1})}.$$

The dynamics (2.4) is therefore reformulated as the ratio of long time expectations of the Markov chains induced by Q_h , applied to the functions $h^{-1} \varphi$ and h^{-1} . It is then possible to resort to the convergence results given by Lemma 2.6.

We first construct a probability measure μ_f for which (2.17) is satisfied, namely

$$\mu_f(\varphi) = \frac{\mu_h(h^{-1} \varphi)}{\mu_h(h^{-1})}, \quad (2.19)$$

where μ_h is the probability measure introduced in Lemma 2.6. Note that μ_f is well-defined for $\varphi \in B_W^\infty(\mathcal{X})$. Indeed, for $\varphi \in B_W^\infty(\mathcal{X})$, it holds $h^{-1} \varphi \in B_{W h^{-1}}^\infty(\mathcal{X})$. Second, we show that

$$\mu_h(h^{-1}) > 0. \quad (2.20)$$

Indeed, since $\|h\|_{B_W^\infty} = 1$, it holds $h^{-1} \geq W^{-1}$, and since W is upper bounded on any compact set, W^{-1} is lower bounded by a positive constant on any compact set. As $\mu_h \in \mathcal{P}(\mathcal{X})$, we can use

Lemma 2.31 in Section 2.5.2 to conclude that $\mu_h(h^{-1}) > 0$. Moreover, μ_f does not depend on the choice of normalization for h . Finally, $\mu_f(W) < +\infty$ since $\mu_h(Wh^{-1}) < +\infty$.

From Lemma 2.6, for any $\varphi \in B_W^\infty(\mathcal{X})$, it holds $Q_h^k(h^{-1}\varphi) = \mu_h(h^{-1}\varphi) + a_k$ and $Q_h^k(h^{-1}) = \mu_h(h^{-1}) + b_k$ with $\|a_k\|_{B_{Wh^{-1}}^\infty} \leq c\bar{\alpha}^k \|h^{-1}\varphi - \mu_h(h^{-1}\varphi)\|_{B_{Wh^{-1}}^\infty}$ and $\|b_k\|_{B_{Wh^{-1}}^\infty} \leq c\bar{\alpha}^k \|h^{-1} - \mu_h(h^{-1})\|_{B_{Wh^{-1}}^\infty}$. Since $\varphi \in B_W^\infty(\mathcal{X})$, we have in particular (using also $\|h\|_{B_W^\infty} = 1$),

$$\|h^{-1}\varphi - \mu_h(h^{-1}\varphi)\|_{B_{Wh^{-1}}^\infty} \leq \|h^{-1}\varphi\|_{B_{Wh^{-1}}^\infty} + \mu_h(h^{-1}|\varphi|)\|h\|_{B_W^\infty} \leq (1 + \mu_h(Wh^{-1}))\|\varphi\|_{B_W^\infty} < +\infty.$$

Since $\mu_h(Wh^{-1}) < +\infty$, we can set $c' = 1 + \mu_h(Wh^{-1})$, so that

$$\|a_k\|_{B_{Wh^{-1}}^\infty} \leq c'\bar{\alpha}^k \|\varphi\|_{B_W^\infty}. \quad (2.21)$$

A similar estimate holds for the sequence $(b_k)_{k \geq 1}$ by taking $\varphi \equiv \mathbf{1}$. This leads to, for any $\varphi \in B_W^\infty(\mathcal{X})$,

$$\begin{aligned} |\Phi_k(\nu)(\varphi) - \mu_f(\varphi)| &= \left| \frac{\nu(h(Q_h)^k(h^{-1}\varphi))}{\nu(h(Q_h)^k h^{-1})} - \mu_f(\varphi) \right| = \left| \frac{\nu(h(\mu_h(h^{-1}\varphi) + a_k))}{\nu(h(\mu_h(h^{-1}) + b_k))} - \mu_f(\varphi) \right| \\ &= \left| \frac{\nu(h)\mu_h(h^{-1}\varphi) + \nu(ha_k)}{\nu(h)\mu_h(h^{-1}) + \nu(hb_k)} - \mu_f(\varphi) \right| = \left| \frac{\mu_f(\varphi) + c_{h,\nu}\nu(ha_k)}{1 + c_{h,\nu}\nu(hb_k)} - \mu_f(\varphi) \right|, \end{aligned}$$

where we used the definition (2.19) and introduced

$$c_{h,\nu} = \frac{1}{\nu(h)\mu_h(h^{-1})}. \quad (2.22)$$

It holds $0 < c_{h,\nu} < +\infty$ because:

- Lemma 2.5 shows that for any $\nu \in \mathcal{P}_W(\mathcal{X})$, it holds $0 < \nu(h) < +\infty$;
- we know that $\mu_h(h^{-1}) < +\infty$ from Lemma 2.6;
- $\mu_h(h^{-1}) > 0$ by (2.20).

Now, since $|b_k| \leq \|b_k\|_{B_{Wh^{-1}}^\infty} Wh^{-1}$ and (2.21) holds for b_k with $\varphi \equiv \mathbf{1}$, we have

$$1 + c_{h,\nu}\nu(hb_k) \geq 1 - c_{h,\nu}\nu(h|b_k|) \geq 1 - c_{h,\nu}\nu(W)\|b_k\|_{B_{Wh^{-1}}^\infty} \geq 1 - \bar{\alpha}^k c' c_{h,\nu}\nu(W).$$

Therefore, the choice

$$k \geq -\frac{\log(2c'c_{h,\nu}\nu(W))}{\log(\bar{\alpha})}$$

ensures that

$$1 + c_{h,\nu}\nu(hb_k) \geq \frac{1}{2}.$$

As a result, for k large enough, using $|a_k| \leq \|a_k\|_{B_{Wh^{-1}}^\infty} Wh^{-1}$ and recalling (2.21),

$$|\Phi_k(\nu)(\varphi) - \mu_f(\varphi)| \leq \frac{c_{h,\nu}(\nu(h|a_k|) + \mu_f(|\varphi|)\nu(h|b_k|))}{1 + c_{h,\nu}\nu(hb_k)} \leq C_\nu \|\varphi\|_{B_W^\infty} \bar{\alpha}^k, \quad (2.23)$$

with

$$C_\nu = 2c_{h,\nu}c'\nu(W)(1 + \mu_f(W)) = \frac{2}{\mu_h(h^{-1})}(1 + \mu_h(Wh^{-1}))(1 + \mu_f(W))\frac{\nu(W)}{\nu(h)}. \quad (2.24)$$

We therefore obtain (2.17) from (2.23) with the constant C_ν defined in (2.24). Note that C_ν depends on the initial measure ν only through the ratio $\nu(W)/\nu(h) \geq \|h\|_{B_W^\infty}^{-1}$.

Taking the supremum over $\varphi \in B_W^\infty(\mathcal{X})$ such that $\|\varphi\|_{B_W^\infty} \leq 1$, (2.23) rewrites, with (2.18):

$$\rho_W(\Phi_k(\nu), \mu_f) \leq C_\nu \bar{\alpha}^k.$$

Choosing $\nu = \Phi(\mu_f)$ and using the semigroup property we obtain

$$\rho_W(\Phi(\Phi_k(\mu_f)), \mu_f) \leq C_{\mu_f} \bar{\alpha}^k.$$

Taking the limit $k \rightarrow +\infty$ shows that $\Phi(\mu_f) = \mu_f$, so μ_f is a fixed point of Φ .

We have shown the existence of an invariant measure of the form (2.19), which is a fixed point of Φ and integrates W . We now turn to uniqueness, which follows by a standard fixed point argument. Assume that we have two probability measures ν_1 and ν_2 in $\mathcal{P}_W(\mathcal{X})$ satisfying (2.17), which are therefore fixed points of Φ . Then, there exists $\bar{\alpha} \in (0, 1)$ such that, for any measure $\nu \in \mathcal{P}(\mathcal{X})$ with $\nu(W) < +\infty$, there is a constant C_ν for which

$$\forall k \geq 1, \quad \rho_W(\Phi_k(\nu), \nu_1) \leq C_\nu \bar{\alpha}^k.$$

Choosing $\nu = \nu_2$ and using the invariance by Φ leads to

$$\rho_W(\nu_2, \nu_1) \leq C_{\nu_2} \bar{\alpha}^k.$$

Taking the limit $k \rightarrow +\infty$ shows that $\nu_1 = \nu_2$, so the invariant measure is unique. \square

Theorem 2.7 also leads to alternative representations of the spectral radius Λ as a scaled cumulant generating function [275] and as the average rate of creation of probability of the dynamics. This is the purpose of the following result.

Theorem 2.8. *Let Q^f be as in Theorem 2.7 and define $\lambda = \log(\Lambda)$. Then, for any $\nu \in \mathcal{P}_W(\mathcal{X})$,*

$$\lambda = \lim_{k \rightarrow +\infty} \frac{1}{k} \log \left(\nu[(Q^f)^k \mathbf{1}] \right). \quad (2.25)$$

Moreover,

$$\Lambda = \mu_f(Q^f \mathbf{1}). \quad (2.26)$$

Proof. Considering the operator Q_h introduced in Lemma 2.6, we have for any $\nu \in \mathcal{P}_W(\mathcal{X})$,

$$\nu[(Q^f)^k \mathbf{1}] = \nu(\Lambda^k h Q_h^k h^{-1}).$$

Taking the logarithm and dividing by k leads to

$$\frac{1}{k} \log \nu[(Q^f)^k \mathbf{1}] = \log(\Lambda) + \frac{1}{k} \log \nu(h Q_h^k h^{-1}).$$

Lemma 2.6 shows that $\nu(h Q_h^k h^{-1})$ converges to $c_{h,\nu}^{-1}$, where $0 < c_{h,\nu} < +\infty$ is defined in (2.22). Taking the limit $k \rightarrow +\infty$ then leads to (3.25).

In order to prove (2.26), we use that μ_f is a fixed point of Φ , *i.e.* for any $\varphi \in B_W^\infty(\mathcal{X})$,

$$\mu_f(\varphi) = \frac{\mu_f(Q^f \varphi)}{\mu_f(Q^f \mathbf{1})}.$$

Taking $\varphi = h \in B_W^\infty(\mathcal{X})$ and using $Q^f h = \Lambda h$ we obtain

$$\mu_f(h) = \frac{\mu_f(\Lambda h)}{\mu_f(Q^f \mathbf{1})},$$

so that $\Lambda = \mu_f(Q^f \mathbf{1})$, as claimed. \square

Although stated in an abstract setting, Theorem 2.8 has a natural interpretation. If $Q^f = e^f Q$ where Q is the evolution operator of a Markov chain $(x_k)_{k \in \mathbb{N}}$ with $x_0 \sim \nu$, then (3.25) rewrites

$$\lambda = \lim_{k \rightarrow +\infty} \frac{1}{k} \log \mathbb{E}_\nu \left[e^{\sum_{i=0}^{k-1} f(x_i)} \right],$$

which is a standard formula for the scaled cumulant generating function, as discussed in Section 1.3.1 (we will use this type of representation formula in Chapter 3). We remind that \mathbb{E}_ν stands for the expectation with respect to all trajectories with initial condition distributed according to ν . On the other hand, (2.26) means that this SCGF can be expressed as the average rate of creation of probability of the process under the invariant measure. In particular, if $Q^f = Q$ is the evolution operator of a Markov chain, $\Lambda = 1$ since there is no creation of probability. Formula (2.26) does not seem typical in the large deviations literature, but we use it in Chapter 4 to quantify the bias arising from discretizing a continuous Feynman–Kac dynamics.

Remark 2.9. *It should be clear from the proofs that Assumptions 1 to 3 can be adapted or relaxed depending on the context. In particular, we typically consider situations in which the state space \mathcal{X} is (a subset of) \mathbb{R}^d , and the transition kernel Q^f has a transition density $p^f(x, y) > 0$ jointly continuous in x, y . In this case, Assumptions 2.2 and 2.3 are immediately fulfilled by setting $\eta_n(dx) = \mathbb{1}_{K_n}(x)dx/|K_n|$ for each compact K_n , as we will see in Section 2.2.3. We will however use yet a slightly different setting in Chapter 3 more natural for continuous time processes. In this setting, the assumptions will appear more clearly as conditions on the coefficients of the diffusion.*

Similarly, the assumption that $W \geq 1$ can be weakened into: W is lower bounded by a positive constant on each compact set.

Another remark of interest is that the regularity condition (Assumption 2.3) is not satisfied by Metropolis type kernels [369], which are therefore not covered by our analysis. In Chapter 4, we show a different proof based on [114] which does not require regularization, but is limited to bounded state space.

Let us mention that, in Assumption 2.1, it seems sufficient to suppose that $\gamma_n < \Lambda$ for some $n \geq 1$ in order to obtain that $\Theta(Q^f) < \Lambda(Q^f)$ in the proof of Lemma 2.5. This is sufficient to use the Krein–Rutman theorem, and to obtain a Lyapunov condition for Q_h (see Remark 2.34 in Section 2.5.5).

It is also possible to keep track of the constants in the proofs of Lemma 2.6 and Theorem 2.7, like in [219], and observe that they depend on the assumptions through the coefficients γ_n , b_n , α_n , the measures η_n and the function W . More precisely, the constants deteriorate when γ_n , α_n and $\eta_n(h)$ become small, and b_n and $\sup_{K_n} W$ get large. Therefore, although the term $\eta_n(h)$ cannot be controlled more explicitly under our assumptions, it seems possible to optimize the final constants in Lemma 2.6 (and thus in Theorem 2.7) with respect to the choice of n .

2.2.3 Results in continuous time

Our analysis carries over to time continuous processes, in particular diffusions. In this case, it is possible to rephrase Assumption 2.1 in terms of the associated infinitesimal generator. Note that we do not consider the most realistic setting here and we use more general assumptions in Chapter 3. In order to avoid the technical difficulty of dealing with an infinite dimensional process, we consider a diffusion $(X_t)_{t \geq 0}$ over $\mathcal{X} = \mathbb{R}^d$ for some integer $d \geq 1$, satisfying the SDE

$$dX_t = b(X_t) dt + \sigma(X_t) dB_t, \quad (2.27)$$

where $b : \mathcal{X} \rightarrow \mathbb{R}^d$, $\sigma : \mathcal{X} \rightarrow \mathbb{R}^{d \times m}$ and $(B_t)_{t \geq 0}$ is an m -dimensional Brownian motion (for some integer $m \geq 1$). The associated infinitesimal generator is given by

$$\mathcal{L} = b \cdot \nabla + \frac{\sigma \sigma^T}{2} : \nabla^2 = \sum_{i=1}^d b_i \partial_{x_i} + \frac{1}{2} \sum_{i,j=1}^d (\sigma \sigma^T)_{ij} \partial_{x_i} \partial_{x_j}. \quad (2.28)$$

We also consider a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$ and the corresponding continuous Feynman–Kac semigroup that reads, for all $t > 0$ and all initial distribution $\nu \in \mathcal{P}(\mathcal{X})$,

$$\Phi_t(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]}{\mathbb{E}_\nu \left[e^{\int_0^t f(X_s) ds} \right]}. \quad (2.29)$$

We use the same notation for the continuous and discrete semigroups since we believe this does not lead to any confusion at this stage (we mention again that in Section 2.3.3 and Chapter 4 below we write more explicitly $\Phi_{\Delta t, n}$ for discretizations of Feynman–Kac dynamics with time step Δt).

In this setting, we define the operator

$$(P_t^f \varphi)(x) = \mathbb{E}_x \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right],$$

so that (2.29) is the natural continuous counterpart of (2.4) where, for a fixed time $t > 0$, we formally have

$$Q^f = P_t^f = e^{t(\mathcal{L}+f)}. \quad (2.30)$$

As a result, $(\Phi_t)_{t \geq 0}$ satisfies a semigroup property like the discrete time evolution through (2.5). In this case, we can show that the generator of the weighted evolution operator P_t^f is $\mathcal{L} + f$ by the

Feynman–Kac formula [361, Chapter VIII, Proposition 3.10]. As for the discrete semigroup (2.4), we are interested in the long time behavior of quantities such as (2.29). When $b = 0$ and $\sigma = \sqrt{2} I_d$, decay estimates of (2.29) in $L^2(\mathcal{X})$ towards a well-defined limit can be obtained by considering the spectral properties of the Schrödinger type operator $-\Delta - f$, as in [377]. When $\sigma = \sqrt{2} I_d$ and $b = -\nabla U$ is the gradient of a potential energy, the operator $\mathcal{L} + f$ is self-adjoint in $L^2(e^{-U})$ (see for instance [21]), and the unitary transform $\varphi \mapsto \varphi e^{-\frac{U}{2}}$ leads to an analysis similar to the Schrödinger case. More precisely, $\mathcal{L} + f$ is unitarily equivalent to

$$\Delta - \frac{1}{4}|\nabla U|^2 + \frac{1}{2}\Delta U + f,$$

which can be studied by the theory of symmetric operators [235]. In both cases, the operator $\mathcal{L} + f$ is self-adjoint on a suitable Hilbert space, so that the Rayleigh formula can be used. It is also possible to study the spectral properties of P_t^f when $b \neq -\nabla U$ and \mathcal{X} is bounded through the Krein–Rutman theorem (see for instance Proposition 4.1 in Chapter 4 below). To the best of our knowledge, the case $b \neq -\nabla U$ in an unbounded space \mathcal{X} remains open in general.

Our analysis provides a practical criterion to study the long time behavior of (2.29) through the Lyapunov function techniques developed in Section 2.2. The continuous counterpart of Assumption 2.1 can be stated in the following simple form.

Assumption 2.10. *Let $(X_t)_{t \geq 0}$ be the dynamics (2.27) with generator (2.28). There exists a $C^2(\mathcal{X})$ function $W : \mathcal{X} \rightarrow [1, +\infty)$ going to infinity at infinity such that*

$$-\frac{(\mathcal{L} + f)W}{W} \xrightarrow{|x| \rightarrow +\infty} +\infty. \quad (2.31)$$

In addition, there exist a $C^2(\mathcal{X})$ function $\mathscr{W} : \mathcal{X} \rightarrow [1, +\infty)$ and a constant $c \geq 0$ such that

$$\varepsilon(x) = \frac{\mathscr{W}(x)}{W(x)} \xrightarrow{|x| \rightarrow +\infty} 0, \quad \frac{(\mathcal{L} + f)\mathscr{W}}{\mathscr{W}} \leq c. \quad (2.32)$$

Condition (2.31) can be checked by direct computations, as shown on some examples in Section 2.3.2. Finding a function \mathscr{W} such that (2.32) holds is generally just a formality, since we build Lyapunov functions in an exponential form. More precisely, we consider in general $W(x) = e^{aU(x)}$ with $a > 0$ for some function $U : \mathcal{X} \rightarrow \mathbb{R}$ going to infinity at infinity, and $\mathscr{W}(x) = e^{a'U(x)}$ for $0 < a' < a$. Moreover, we mention that if a dynamics is given and possesses a Lyapunov function W such that

$$\Psi = -\frac{\mathcal{L}W}{W} \xrightarrow{+\infty} +\infty,$$

then the assumption (2.31) shows that the allowed weight function f should be dominated by Ψ . This is the central idea to obtain large deviations principles in weighted topologies in Chapter 3.

In the proof of Theorem 2.12, (2.31)–(2.32) are used to control P_t^f thanks to a Grönwall lemma. It is also important to remark that, in the case $f = 0$, we are exactly back to typical conditions for the ergodicity of SDEs and compactness of the evolution operator P_t , see [362, Theorem 8.9]. As in Section 2.2.2, some regularity of the transition kernel is required. A natural condition in the context of diffusions reads as follows [362, Section 7].

Assumption 2.11. *The functions f and σ are continuous and, for any $t > 0$, the transition kernel P_t^f has a continuous density p_t^f with respect to the Lebesgue measure, that is*

$$\forall x, y \in \mathcal{X}, \quad P_t^f(x, dy) = p_t^f(x, y) dy.$$

Moreover, it holds

$$\forall x, y \in \mathcal{X}, \quad p_t^f(x, y) > 0.$$

This assumption is natural for diffusion processes and, as shown in the proof of Theorem 2.12, it implies Assumptions 2.2 and 2.3 in Section 2.2.2. It holds true in particular for elliptic diffusions with regular coefficients and additive noise ($b \in C^\infty(\mathcal{X})$ and $\sigma = I_d$). For degenerate diffusions, possibly with multiplicative noise, this result can be obtained through hypoelliptic conditions and controllability. We will extensively come back on these problems in Chapter 3 and refer for now to [395, 362, 427] for more details.

We now state the continuous version of Theorem 2.7.

Theorem 2.12. *Consider the dynamics (2.29) induced by the SDE (2.27) and suppose that Assumptions 2.10 and 2.11 hold. Then, there exist a unique invariant measure μ_f and $\kappa > 0$ such that, for any initial measure $\nu \in \mathcal{P}_W(\mathcal{X})$ there is $C_\nu > 0$ for which*

$$\forall \varphi \in B_W^\infty(\mathcal{X}), \quad \forall t > 0, \quad |\Phi_t(\nu)(\varphi) - \mu_f(\varphi)| \leq C_\nu e^{-\kappa t} \|\varphi\|_{B_W^\infty}. \quad (2.33)$$

Moreover, the invariant measure satisfies $\mu_f(W) < +\infty$.

Proof. The idea of the proof is to show that, for any $t > 0$, the evolution operator

$$(P_t^f \varphi)(x) = \mathbb{E}_x \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]$$

satisfies the assumptions of Theorem 2.7.

Step 1: Minorization and regularity. We first show that, by Assumption 2.11 and for any $t > 0$, P_t^f satisfies Assumptions 2.2 and 2.3. A first remark is that, since P_t^f is assumed to have a continuous density with respect to the Lebesgue measure, Assumption 2.3 immediately holds.

It is enough to prove the minorization condition (Assumption 2.2) for measurable subsets of $\mathcal{X} = \mathbb{R}^d$. Consider the compact sets $K_n = B(0, n)$, i.e. the balls centered at 0 with radius $n \geq 1$. For a measurable set $A \subset \mathbb{R}^d$ and $n \geq 1$, we have, for all $x \in K_n$,

$$(P_t^f \mathbf{1}_A)(x) = \int_A p_t^f(x, y) dy \geq \int_{A \cap K_n} p_t^f(x, y) dy \geq \left(\inf_{x, y \in K_n} p_t^f(x, y) \right) |A \cap K_n|, \quad (2.34)$$

where we denote by $|A|$ the Lebesgue measure of $A \subset \mathbb{R}^d$. As a result, (2.8) holds for all $n \geq 1$ with

$$\eta_n(A) = \frac{|A \cap K_n|}{|K_n|}, \quad \alpha_n = |K_n| \left(\inf_{x, y \in K_n} p_t^f(x, y) \right) > 0.$$

Finally, let us check that (2.9) is satisfied. Take $\varphi \in B_W^\infty(\mathcal{X})$ with $\varphi \geq 0$ such that

$$\eta_n(\varphi) = \frac{1}{|K_n|} \int_{K_n} \varphi(x) dx = 0,$$

for any $n \geq n_0$ for an arbitrary $n_0 \geq 1$. Since for any compact set $K \subset \mathcal{X}$ there exists $m \geq 1$ such that $K \subset K_m$, this implies that $\varphi = 0$ almost everywhere, so $Q^f \varphi = 0$ everywhere since Q^f has a continuous density with respect to the Lebesgue measure. Therefore, Assumption 2.2 is satisfied.

Step 2: Lyapunov condition. Let us now show that Assumption 2.1 holds for P_t^f with $t > 0$ fixed. First, Assumption 2.10 is equivalent to the existence of positive sequences $(a_n)_{n \in \mathbb{N}}$, $(b_n)_{n \in \mathbb{N}}$ such that (recall that W is bounded on compact sets)

$$(\mathcal{L} + f)W \leq -a_n W + b_n, \quad (2.35)$$

with $a_n \rightarrow +\infty$ as $n \rightarrow +\infty$. We then compute, for any $t > 0$ and $n \in \mathbb{N}$,

$$\frac{d}{dt} \left(e^{a_n t} P_t^f W \right) = e^{a_n t} P_t^f (a_n W + (\mathcal{L} + f)W) \leq b_n e^{a_n t} P_t^f \mathbf{1}. \quad (2.36)$$

We can now bound the right hand side of the latter expression using (2.32). Since $\mathscr{W} \geq 1$,

$$(P_t^f \mathbf{1})(x) = \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right] \leq \mathbb{E}_x \left[\mathscr{W}(X_t) e^{\int_0^t f(X_s) ds} \right]. \quad (2.37)$$

From the second condition in (2.32), (2.37) becomes

$$(P_t^f \mathbf{1})(x) \leq e^{ct} \mathbb{E}_x \left[\mathscr{W}(X_t) e^{-\int_0^t \frac{c\mathscr{W}}{\mathscr{W}}(X_s) ds} \right].$$

Inspired by a similar calculation in [427], we see that the right hand side of the above equation is a supermartingale. Indeed, introducing

$$M_t = \mathscr{W}(X_t) e^{-\int_0^t \frac{c\mathscr{W}}{\mathscr{W}}(X_s) ds},$$

Itô formula shows that

$$dM_t = e^{-\int_0^t \frac{c\mathscr{W}}{\mathscr{W}}(X_s) ds} \nabla \mathscr{W}^T(X_t) \sigma(X_t) dB_t,$$

so that M_t is a local martingale since σ is continuous (see [262, Proposition 2.24]). Since M_t is nonnegative, it is a supermartingale by Fatou's lemma. As a result, $\mathbb{E}_x[M_t] \leq M_0 = \mathscr{W}(x)$. The inequality (2.37) then becomes

$$(P_t^f \mathbf{1})(x) \leq e^{ct} \mathbb{E}_x[M_t] \leq e^{ct} \mathscr{W}(x).$$

Coming back to (2.36), we obtain

$$\frac{d}{dt} \left(e^{a_n t} P_t^f W \right) \leq b_n e^{(a_n+c)t} \mathscr{W}.$$

Integrating in time shows that

$$(e^{a_n t} P_t^f W - W)(x) \leq b_n \frac{e^{(a_n+c)t}}{a_n + c} \mathscr{W}(x).$$

As a result

$$P_t^f W(x) \leq \tilde{\gamma}_n W(x) + c_n \mathscr{W}(x), \quad (2.38)$$

with

$$\tilde{\gamma}_n = e^{-a_n t}, \quad c_n = \frac{b_n e^{ct}}{a_n + c} \geq 0.$$

At this stage, (2.6) holds with the indicator function replaced by the function \mathscr{W} . However, using the first condition in (2.32), we can find a compact set K_n such that $c_n \varepsilon(x) \leq \tilde{\gamma}_n$ outside K_n . Using this set and $\mathscr{W} = \varepsilon W$, (2.38) becomes

$$\begin{aligned} P_t^f W(x) &\leq \tilde{\gamma}_n W(x) + c_n \mathbf{1}_{K_n}(x) \mathscr{W}(x) + c_n \varepsilon(x) W(x) \mathbf{1}_{K_n^c}(x) \\ &\leq 2\tilde{\gamma}_n W(x) + c_n \left(\sup_{K_n} \mathscr{W} \right) \mathbf{1}_{K_n}(x). \end{aligned}$$

Setting $\gamma_n = 2\tilde{\gamma}_n$ and $b_n = c_n \sup_{K_n} \mathscr{W}$, we see that

$$P_t^f W \leq \gamma_n W + b_n \mathbf{1}_{K_n}, \quad (2.39)$$

with $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$. This means that P_t^f satisfies Assumption 2.1, and hence fulfils all the assumptions of Theorem 2.7.

Step 3: using Theorem 2.7. We now use that P_t^f satisfies the assumptions of Theorem 2.7 to conclude the proof. Fix $t_0 > 0$. There exist a unique measure μ_{f,t_0} and a constant $\kappa_{t_0} > 0$ such that, for any $\nu \in \mathcal{P}_W(\mathcal{X})$, it holds (with the constant $C_\nu > 0$ defined in (2.24))

$$\forall \varphi \in B_W^\infty(\mathcal{X}), \quad \forall k \geq 1, \quad \left| \frac{\nu((P_{t_0}^f)^k \varphi)}{\nu((P_{t_0}^f)^k \mathbf{1})} - \mu_{f,t_0}(\varphi) \right| \leq C_\nu e^{-k\kappa_{t_0}} \|\varphi\|_{B_W^\infty}.$$

We next show that (2.33) can be obtained for any $t > 0$ (and not only multiples of t_0) and that the invariant measure μ_{f,t_0} actually does not depend on t_0 . This follows by a standard time decomposition argument [304, 239]. Indeed, for any $t > 0$, we set $t = kt_0 + r$ with $r \in [0, t_0)$, and we use the semigroup property to obtain

$$\Phi_t(\nu)(\varphi) = \Phi_{kt_0} \left(\frac{\nu P_r^f}{\nu(P_r^f \mathbf{1})} \right) (\varphi) = \frac{\nu_r((P_{t_0}^f)^k \varphi)}{\nu_r((P_{t_0}^f)^k \mathbf{1})},$$

where we defined ν_r as

$$\nu_r(\varphi) = \frac{\nu(P_r^f \varphi)}{\nu(P_r^f \mathbf{1})}.$$

We then only need to control the family of initial distributions $(\nu_r)_{r \in [0, t_0)}$. Step 1 in the proof shows that $\nu(P_r^f \mathbf{1}) > 0$ (using (2.34)). Then, in view of (2.39), the evolution operator P_r^f maps $B_W^\infty(\mathcal{X})$ to $B_W^\infty(\mathcal{X})$ for any $r > 0$, so $\nu_r(W) < +\infty$ and thus ν_r defines an admissible initial condition in Theorem 2.7. This leads to:

$$\forall \varphi \in B_W^\infty(\mathcal{X}), \quad \forall t > 0, \quad |\Phi_t(\nu)(\varphi) - \mu_{f,t_0}(\varphi)| \leq \left(\sup_{r \in [0, t_0)} C_{\nu_r} \right) e^{-\kappa_{t_0} \frac{t}{t_0}} \|\varphi\|_{B_W^\infty}, \quad (2.40)$$

where the constant C_{ν_r} is given in (2.24). In view of (2.24), it remains to bound

$$\sup_{r \in [0, t_0]} \frac{\nu_r(W)}{\nu_r(h_{t_0})} = \sup_{r \in [0, t_0]} \frac{\nu(P_r^f W)}{\nu(P_r^f h_{t_0})}, \quad (2.41)$$

where h_{t_0} is the principal eigenvector associated to $P_{t_0}^f$ with eigenvalue Λ_{t_0} (using Lemma 2.5). The numerator in the latter expression is easily bounded uniformly in r using (2.39). Standard semigroup analysis shows that $h_{t_0} = h$ does not depend on t_0 and $\Lambda_{t_0} = e^{t_0 \alpha}$ for some $\alpha \in \mathbb{R}$. Therefore, for any $r \in [0, t_0]$, $P_r^f h_{t_0} = e^{r \alpha} h$, and the denominator in (2.41) is bounded away from 0 independently on r .

We finally prove that the invariant measure μ_{f, t_0} does not depend on t_0 . Following the same procedure for another time $t_1 > 0$ shows that (2.40) holds with an invariant measure μ_{f, t_1} . Then, for any $\varphi \in B_W^\infty(\mathcal{X})$, $\nu \in \mathcal{P}_W(\mathcal{X})$ and $t > 0$ we have

$$\begin{aligned} |\mu_{f, t_0}(\varphi) - \mu_{f, t_1}(\varphi)| &\leq |\Phi_t(\nu)(\varphi) - \mu_{f, t_0}(\varphi)| + |\Phi_t(\nu)(\varphi) - \mu_{f, t_1}(\varphi)| \\ &\leq \left(\sup_{r \in [0, t_0]} C_{\nu_r} \right) e^{-\kappa t_0 \frac{t}{t_0}} \|\varphi\|_{B_W^\infty} + \left(\sup_{r \in [0, t_1]} C_{\nu_r} \right) e^{-\kappa t_1 \frac{t}{t_1}} \|\varphi\|_{B_W^\infty}. \end{aligned}$$

Taking the limit $t \rightarrow +\infty$ on the right hand side shows that $\mu_{f, t_0} = \mu_{f, t_1}$, so the invariant measure is independent of the arbitrary time t_0 . This concludes the proof of Theorem 2.12. \square

We close this section by mentioning that, under the assumptions of Theorem 2.12, it is also possible to define the logarithmic spectral radius of the dynamics as in Theorem 2.8, which reads in this case

$$\lambda = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_\nu \left[e^{\int_0^t f(X_s) ds} \right],$$

for any initial measure $\nu \in \mathcal{P}_W(\mathcal{X})$. We do not reproduce the proof of this result which is similar to that of Theorem 2.8 and will be detailed in Chapter 3 (see in particular Lemma 3.32).

2.3 Applications

Since our study was first motivated by practical situations, we provide in this section a number of finite dimensional examples where our framework provides simple criteria for proving convergence of the Feynman–Kac semigroup towards an invariant measure. Sections 2.3.1 and 2.3.2 are concerned with discrete and continuous time applications respectively. Section 2.3.3 presents a convergence result for numerical discretizations of (2.29) (with some anticipation on Chapter 4), where convergence rates are uniform in the time step.

2.3.1 Examples in discrete time

In this section, we provide two typical examples of Markov chains for which our results apply. First of all, let us consider the Diffusion Monte Carlo case where $f = -V$ and V stands for a Schrödinger potential.

Proposition 2.13. *Consider a weighted evolution operator $Q^V = e^{-V} Q$ in $\mathcal{X} = \mathbb{R}^d$ with Gaussian increments $Q(x, dy) = (2\pi\sigma^2)^{-\frac{d}{2}} e^{-\frac{(x-y)^2}{2\sigma^2}} dy$, and where V is a continuous function. Then, if $V(x) \rightarrow +\infty$ when $|x| \rightarrow +\infty$, $W(x) = \mathbf{1}$ is a Lyapunov function for Q^V in the sense of Assumption 2.1. Moreover, if there exist constants $a > 0$ and $c \in \mathbb{R}$ such that*

$$V(x) \geq a|x|^2 - c, \quad (2.42)$$

then $W(x) = e^{\beta x^2}$ is a Lyapunov function for

$$0 < \beta < \frac{a}{2} \left(\sqrt{1 + \frac{2}{a\sigma^2}} - 1 \right).$$

Finally, Assumptions 2.2 and 2.3 hold true, so that Theorem 2.7 applies for these choices of Lyapunov function.

The interpretation of this result is the following. In the Diffusion Monte Carlo setting, the confinement cannot be provided by the dynamics, since it is a Gaussian random walk over \mathbb{R}^d . However, the external potential V gives a small weight to the trajectories going to infinity, which makes the dynamics stable. If more information is available on the growth of V , we obtain better integrability results for the invariant measure μ_V^* through Lyapunov functions growing faster at infinity.

Proof. Let us first check that $W = \mathbf{1}$ is a Lyapunov function when V goes to infinity at infinity. Note that, for any compact set $K \subset \mathbb{R}^d$,

$$(Q^V \mathbf{1})(x) = e^{-V(x)} = \mathbf{1}_{K^c}(x) e^{-V(x)} + \mathbf{1}_K(x) e^{-V(x)}.$$

Taking an increasing sequence of compact sets K_n (in the sense of inclusion) and setting $\gamma_n = \sup_{K_n^c} e^{-V}$, $b_n = \sup_{K_n} e^{-V} < +\infty$, we obtain

$$Q^V \mathbf{1} \leq \gamma_n \mathbf{1} + b_n \mathbf{1}_{K_n},$$

which proves the first assertion since $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$.

Let us now assume that (2.42) holds. Setting $W(x) = e^{\beta x^2}$, under the condition

$$\beta < \frac{1}{2\sigma^2}, \quad (2.43)$$

an easy computation shows that

$$QW(x) = \frac{e^{\frac{\beta}{1-2\beta\sigma^2}x^2}}{(1-2\beta\sigma^2)^{\frac{d}{2}}}.$$

We remark that W is not a Lyapunov function for Q since $1 - 2\beta\sigma^2 < 1$. However, setting $C_d = (1 - 2\beta\sigma^2)^{-\frac{d}{2}}$,

$$Q^V W(x) = C_d e^{-V(x) + \frac{\beta}{1-2\beta\sigma^2}x^2} \leq C_d e^{c-ax^2 + \frac{\beta}{1-2\beta\sigma^2}x^2 - \beta x^2} W(x) = C'_d e^{-ax^2 + \frac{2\beta^2\sigma^2}{1-2\beta\sigma^2}x^2} W(x),$$

with $C'_d = C_d e^c$. One can then check that the choice

$$0 < \beta < \frac{a}{2} \left(\sqrt{1 + \frac{2}{a\sigma^2}} - 1 \right) \quad (2.44)$$

leads to

$$-a + \frac{2\beta^2\sigma^2}{1-2\beta\sigma^2} < 0.$$

Note that, since

$$\frac{a}{2} \left(\sqrt{1 + \frac{2}{a\sigma^2}} - 1 \right) < \frac{1}{2\sigma^2},$$

the condition (2.43) is automatically satisfied when β is chosen according to (2.44). Next, when β satisfies (2.44), the function

$$\varepsilon(x) = e^{-ax^2 + \frac{\beta^2\sigma^2}{1-2\beta\sigma^2}x^2}$$

goes to zero at infinity. Therefore, taking increasing compact sets K_n (such as balls of increasing radii),

$$(Q^V W)(x) = \mathbf{1}_{K_n^c}(x) \varepsilon(x) W(x) + \mathbf{1}_{K_n}(x) \varepsilon(x) W(x) \leq \gamma_n W(x) + b_n \mathbf{1}_{K_n}(x),$$

with $\gamma_n = \sup_{K_n^c} \varepsilon \rightarrow 0$ as $n \rightarrow +\infty$ and $b_n = \sup_{K_n} \varepsilon W < +\infty$. Hence W is a Lyapunov function for Q^V for this choice of β , *i.e.* Assumption 2.1 is satisfied.

Assumption 2.3 is easily seen to hold. It therefore suffices to prove the minorization condition (Assumption 2.2). Take a compact set K with non zero Lebesgue measure, and let us first show that the condition of Assumption 2.2 holds for Q . It is enough to prove the condition for the indicator function of any borel set $A \subset \mathcal{X}$. Denoting by $D_K = \sup\{|x - y|, x \in K, y \in K\}$ the diameter of K , we compute for any $x \in K$

$$\begin{aligned} (Q\mathbf{1}_A)(x) &= Q(x, A) = \int_A e^{-\frac{(x-y)^2}{2\sigma^2}} dy \geq \int_{A \cap K} e^{-\frac{(x-y)^2}{2\sigma^2}} dy \geq \inf_{x \in K} \int_{A \cap K} e^{-\frac{(x-y)^2}{2\sigma^2}} dy \\ &\geq e^{-\frac{D_K^2}{2\sigma^2}} \int_{A \cap K} dy \geq e^{-\frac{D_K^2}{2\sigma^2}} |A \cap K|, \end{aligned}$$

where we denote again by $|A|$ the Lebesgue measure of $A \subset \mathbb{R}^d$. This motivates defining

$$\alpha_K = e^{-\frac{D_K^2}{2\sigma^2}} |K| > 0, \quad \eta_K(A) = \frac{|A \cap K|}{|K|}.$$

Note also that, since $|K| \in (0, +\infty)$, η_K is a probability measure. Finally, since V is continuous,

$$\forall x \in K, \quad Q^V(x, \cdot) \geq \alpha_V \eta_K(\cdot),$$

with $\alpha_V = \alpha_K e^{-\sup_K V} > 0$. Choosing for instance $K_n = B(0, n)$ the centered balls of radius n , we see that (2.9) holds using arguments similar to the ones used for the proof of Theorem 2.12, hence Q^V satisfies Assumption 2.2. \square

We now provide an example where the dynamics Q admits a Lyapunov function W in the sense of the condition (2.58) recalled in Section 2.5.1, and this function is also a Lyapunov function for Q^f when f does not grow too fast.

Proposition 2.14. *Consider the dynamics corresponding to a discrete Ornstein–Uhlenbeck process in \mathbb{R}^d , namely*

$$x_{k+1} = \rho x_k + \sigma G_k,$$

where $\rho \in (-1, 1)$, $\sigma \in \mathbb{R}$ and $(G_k)_{k \geq 1}$ is a family of independent standard d -dimensional Gaussian random variables. Define the operator $Q^f = e^f Q$ with f a continuous function such that there exist constants $a > 0$, $c \geq 0$, $0 \leq p < 2$ for which $f(x) \leq a|x|^p + c$.

Then, the Feynman–Kac dynamics associated to Q^f satisfies the assumptions of Theorem 2.7 with Lyapunov function $W(x) = e^{\beta x^2}$ when

$$0 < \beta < \frac{1 - \rho^2}{2\sigma^2}.$$

The interpretation of this result is quite different from the interpretation of Proposition 2.13. Here, the confinement is provided by the dynamics itself, and the weight f has to be controlled by the Lyapunov function of the dynamics. In that case it is important to find a «strong enough» Lyapunov function in order for this control to be possible. Quite typically, if f is unbounded, $W(x) = x^2$ is a Lyapunov function for Q , but not for Q^f . On the other hand, if f is bounded above, the result is straightforward.

Proof. We set $W(x) = e^{\beta x^2}$ and first compute

$$QW(x) = \mathbb{E} [W(x_{k+1}) \mid x_k = x] = \mathbb{E}_G [e^{\beta|\rho x + \sigma G|^2}] = e^{\beta\rho^2 x^2} \mathbb{E} [e^{\beta(2\sigma\rho xG + \sigma^2 G^2)}].$$

For $\beta < 1/(2\sigma^2)$, an easy computation similar to that of Proposition 2.13 leads to

$$QW(x) = \frac{1}{(1 - 2\beta\sigma^2)^{\frac{d}{2}}} e^{\frac{\rho^2}{1 - 2\beta\sigma^2} \beta x^2}.$$

Define now

$$\delta_\beta = \frac{\rho^2}{1 - 2\beta\sigma^2}.$$

Then $\delta_\beta \in (0, 1)$ and $1 - 2\beta\sigma^2 > 0$ when

$$\beta \in \left(0, \frac{1 - \rho^2}{2\sigma^2}\right).$$

This leads to

$$e^{f(x)} QW(x) = \frac{1}{(1 - 2\beta\sigma^2)^{\frac{d}{2}}} e^{f(x) + (\delta_\beta - 1)x^2} W(x) \leq \frac{1}{(1 - 2\beta\sigma^2)^{\frac{d}{2}}} e^{a|x|^p + c + (\delta_\beta - 1)x^2} W(x) = \varepsilon(x) W(x),$$

with $\varepsilon(x) \rightarrow 0$ as $|x| \rightarrow +\infty$. Therefore, by considering again $K_n = B(0, n)$, we see that

$$(Q^f W)(x) = \mathbf{1}_{K_n^c}(x) \varepsilon(x) W(x) + \mathbf{1}_{K_n}(x) \varepsilon(x) W(x) \leq \gamma_n W(x) + \mathbf{1}_{K_n}(x) b_n,$$

where $\gamma_n = \sup_{K_n^c} \varepsilon \rightarrow 0$ as $n \rightarrow +\infty$, and $b_n = \sup_{K_n} \varepsilon W < +\infty$. This shows that Assumption 2.1 is satisfied. Assumptions 2.2 and 2.3 follow by arguments similar to those used in the proof of Proposition 2.13. \square

The latter examples do not intend to form a complete overview of the possible practical cases. However, they seem characteristic of two typical situations: one where the confinement arises from the dynamics, and another where it comes from the potential $V = -f$. These two strategies correspond respectively to a large deviations context [275] and a diffusion Monte Carlo context [221]. Interestingly, they are both encoded in the condition (2.6).

2.3.2 Applications to diffusion processes

We now provide some examples where the conditions of Section 2.2.3 are met. Our main concern is the Lyapunov condition, Assumption 2.10, so we assume f and the coefficients of the SDE (2.27) to be regular enough for Assumption 2.11 to be satisfied (the issue of regularity will be treated at length in Chapter 3). Let us start with a reversible diffusion.

Proposition 2.15. *Consider a diffusion process $(X_t)_{t \geq 0}$ over \mathbb{R}^d satisfying (2.27) with $\sigma = \sqrt{2} I_d$, and assume that the drift is given by $b = -\nabla U$, where $U : \mathcal{X} \rightarrow \mathbb{R}$ is a smooth potential such that $U(x) \rightarrow +\infty$ as $|x| \rightarrow +\infty$. Assume moreover that U satisfies*

$$\lim_{|x| \rightarrow +\infty} \frac{|\nabla U(x)|^2}{|\Delta U(x)|} = +\infty, \quad (2.45)$$

and there exists $1/2 < \beta < 1$ such that

$$\lim_{|x| \rightarrow +\infty} \left(-\beta(1 - \beta)|\nabla U|^2 + \beta\Delta U + f \right) = -\infty. \quad (2.46)$$

Then Assumption 2.10 holds for the Lyapunov function $W(x) = e^{\beta U(x)}$.

Proof. The proof follows from simple computations. Indeed, it holds

$$\mathcal{L}W = -\beta\nabla U \cdot (\nabla U)W + \beta\nabla \cdot [(\nabla U)W] = -\beta|\nabla U|^2W + \beta W\Delta U + \beta^2|\nabla U|^2W,$$

so that

$$(\mathcal{L} + f)W = (-\beta(1 - \beta)|\nabla U|^2 + \beta\Delta U + f)W, \quad (2.47)$$

hence (2.31) in Assumption 2.10 is satisfied. The conditions in (2.32) are obtained by setting

$$\mathscr{W}(x) = e^{\theta U(x)},$$

for some $\theta \in (1/2, \beta)$. It is clear that \mathscr{W}/W goes to zero at infinity, so the first condition in (2.32) holds true. The key remark is then to note that for our choice of θ, β , we have

$$\beta(1 - \beta) \leq \theta(1 - \theta).$$

Therefore, (2.45) and (2.46) show that there exist $c, c' \geq 0$ such that

$$f \leq \beta(1 - \beta)|\nabla U|^2 - \beta\Delta U + c \leq \theta(1 - \theta)|\nabla U|^2 - \theta\Delta U + c' = -\frac{\mathcal{L}\mathscr{W}}{\mathscr{W}} + c'.$$

This proves that the second condition in (2.32) holds, which concludes the proof. \square

Let us mention that the conditions in Proposition 2.15 are similar to conditions appearing in works on Poincaré inequalities (see [20] and references therein), and corresponds to the case where the confinement comes from the potential U , f being a perturbation that should not go too fast to $+\infty$ with respect to U . We will again comment on this point in Chapter 3.

Remark 2.16. *Proposition 2.15 is also related to confinement conditions for Schrödinger operators. Indeed, using the parameters of Proposition 2.15, the dynamics is reversible with respect to the measure e^{-U} and, as noted in Section 2.2.3, it is possible to turn the diffusion operator \mathcal{L} into a Schrödinger operator using the unitary transform:*

$$\mathcal{L} \rightarrow e^{-\frac{U}{2}} \mathcal{L} e^{\frac{U}{2}}.$$

Using this transformation, $\mathcal{L} + f$ is unitarily equivalent [304] to the following Schrödinger operator:

$$\Delta - \frac{1}{4}|\nabla U|^2 + \frac{1}{2}\Delta U + f.$$

We then notice that the confinement condition for this Schrödinger operator is precisely (2.46) for the limit value $\beta = 1/2$. This shows that our Lyapunov condition (2.31) is a natural extension of this condition for non-reversible dynamics. As a side product, it shows that a slightly modified confinement condition for a Schrödinger operator does not only provide convergence in L^2 -norm, but also in a weighted uniform norm, which does not seem to be a standard result.

In the non-reversible setting one cannot hope for a Schrödinger representation, and the Lyapunov function framework shows its usefulness. Let us present such an application, drawn from [160], where the drift provides polynomial confinement at infinity.

Proposition 2.17. *Let $(X_t)_{t \geq 0}$ satisfy the SDE (2.27) with $\sigma = \sqrt{2}I_d$ and a drift b such that there exist $q > 1$, $\delta > 0$, $R > 0$ for which*

$$\forall |x| \geq R, \quad b(x) \cdot x \leq -\delta|x|^q. \quad (2.48)$$

Assume also that f is smooth and satisfies $f(x) \leq a|x|^p$ for $|x| \geq R$ and some $p < 2q - 2$. Then, Assumption 2.10 holds for the Lyapunov function

$$W(x) = e^{\beta|x|^q}, \quad \text{with } 0 < \beta < \frac{\delta}{q}. \quad (2.49)$$

Proof. Setting $W(x) = e^{\beta|x|^q}$, a simple computation shows that

$$\begin{aligned} \mathcal{L}W(x) &= \beta qb(x) \cdot x|x|^{q-2}W(x) + \beta q \nabla \cdot (x|x|^{q-2}W(x)) \\ &= \beta qb(x) \cdot x|x|^{q-2}W(x) + \beta q d|x|^{q-2}W(x) + \beta q(q-2)|x|^{q-2}W(x) + \beta^2 q^2|x|^{2q-2}W(x), \end{aligned} \quad (2.50)$$

so

$$\frac{\mathcal{L}W}{W}(x) = \beta qb(x) \cdot x|x|^{q-2} + \beta q(q+d-2)|x|^{q-2} + \beta^2 q^2|x|^{2q-2}.$$

Using (2.48) and the bound on f leads to, for $|x| \geq R$,

$$\frac{\mathcal{L}W}{W}(x) + f(x) \leq -\beta q(\delta - \beta q)|x|^{2q-2} + \beta q(q+d-2)|x|^{q-2} + a|x|^p. \quad (2.51)$$

Since $p < 2q - 2$, (2.31) is readily satisfied when $0 < \beta < \delta/q$.

We end the proof by showing that (2.32) holds. Similarly to the proof of Proposition 2.15, we consider

$$\mathscr{W}(x) = e^{\theta|x|^q}, \quad \text{with } 0 < \theta < \beta,$$

which satisfies the first condition in (2.32). Repeating the calculations leading to (2.51), since $\theta < \delta/q$ and $p < 2q - 2$, we obtain the existence of a constant $c \geq 0$ such that

$$\frac{\mathcal{L}\mathscr{W}}{\mathscr{W}}(x) + f(x) \leq -\theta q(\delta - \theta q)|x|^{2q-2} + \theta q(q-1)|x|^{q-2} + a|x|^p \leq c,$$

so the second condition in (2.32) holds true, and Assumption 2.10 is satisfied. \square

2.3.3 Convergence results uniform with respect to the time step

When one considers continuous semigroups as in Section 2.2.3, it is natural in practical applications to discretize (2.29) for example with

$$\Phi_{\Delta t, k}(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(x_k) e^{\Delta t \sum_{i=0}^{k-1} f(x_i)} \right]}{\mathbb{E}_\nu \left[e^{\Delta t \sum_{i=0}^{k-1} f(x_i)} \right]}, \quad (2.52)$$

where $(x_k)_{k \in \mathbb{N}}$ is a discretization of the SDE (2.27) with time step $\Delta t > 0$, *i.e.* x_k is an approximation of $X_{k\Delta t}$. The framework developed in the present chapter covers this situation, as shown by the examples provided in Section 2.3.1 (we recall that the issue of discretization of Feynman–Kac semigroups on a bounded state space is treated in detail in Chapter 4).

Another interesting consequence of our analysis is that we are able to obtain convergence estimates uniform in the time step Δt , in the sense that the decay rate on fact depends on $k\Delta t$, the physical

time of the system, with a prefactor independent of Δt . It has been the purpose of several works to develop such uniform in Δt estimates for long time convergence, in particular in the context of Metropolized discretizations of overdamped Langevin dynamics [61, 177], discretization of the Langevin dynamics [304, 295], and other discretizations of SDEs [110, 276, 277]. Our goal is to show that similar results can be obtained for Feynman–Kac semigroups, and we slightly anticipate the results of Chapter 4 (we therefore suggest to skip this section at first reading). For the remainder of this section, we assume that

$$\mathcal{X} = \mathbb{T}^d$$

is the d -dimensional torus, the function σ in (2.27) is constant, and we denote by $[a]$ the upper integer part of a for $a \in \mathbb{R}$. Considering an unbounded state space \mathcal{X} is also possible but, as noted in Chapter 4, this leads to serious technical difficulties – we therefore postpone this case to future works.

We consider here a simplified version of the framework extensively developed in Chapter 4, Section 4.3. We say that a kernel operator $Q_{\Delta t}^f$ defines a consistent discretization of the semigroup (2.29) if it satisfies Assumption 2.3 and there exist $\Delta t^* > 0$, $C > 0$, $p \in \mathbb{N}$, and an operator $\mathcal{R}_{\Delta t} : C^\infty(\mathcal{X}) \rightarrow C^\infty(\mathcal{X})$ (which encodes remainder terms) such that, for any $\varphi \in C^\infty(\mathcal{X})$,

$$Q_{\Delta t}^f \varphi = \varphi + \Delta t(\mathcal{L} + f)\varphi + \Delta t^2 \mathcal{R}_{\Delta t} \varphi,$$

where, for all $\Delta t \in (0, \Delta t^*]$,

$$\|\mathcal{R}_{\Delta t} \varphi\|_{B^\infty} \leq C \sup_{\substack{m \in \mathbb{N}^d \\ |m| \leq p}} \|\partial^m \varphi\|_{B^\infty},$$

using the notation $\partial^m = \partial_{x_1}^{m_1} \dots \partial_{x_d}^{m_d}$ for $m = (m_1, \dots, m_d) \in \mathbb{N}^d$ with $|m| = m_1 + \dots + m_d$. The dynamics (2.29) is then approximated by the discrete semigroup

$$\forall k \geq 1, \quad \forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in B^\infty(\mathcal{X}), \quad \Phi_{\Delta t, k}(\nu)(\varphi) = \frac{\nu \left((Q_{\Delta t}^f)^k \varphi \right)}{\nu \left((Q_{\Delta t}^f)^k \mathbf{1} \right)}. \quad (2.53)$$

The latter definition encompasses many numerical schemes, as presented in Chapter 4 together with the subsequent numerical analysis. In order to obtain uniform in the time step estimates, we now assume a uniform minorization and boundedness condition of the following form.

Assumption 2.18. *Fix a time $T > 0$. There exist $\Delta t^* > 0$, $\eta \in \mathcal{P}(\mathcal{X})$ and $\alpha \in (0, 1)$ such that, for any $\Delta t \in (0, \Delta t^*]$, the operator $Q_{\Delta t}^f$ is strong Feller and for any $\varphi \in B^\infty(\mathcal{X})$ with $\varphi \geq 0$,*

$$\forall x \in \mathcal{X}, \quad \alpha \eta(\varphi) \leq \left(\left(Q_{\Delta t}^f \right)^{\lceil \frac{T}{\Delta t} \rceil} \varphi \right) (x) \leq \frac{1}{\alpha} \eta(\varphi). \quad (2.54)$$

The lower bound in (2.54) corresponds to a minorization condition with respect to a physical time $T > 0$, see [304, Section 3]. The upper bound is a standard ingredient for studying Feynman–Kac semigroups, see for instance [114, 113] and the proof of Theorem 4.9 presented in Section 4.2.2 below.

Remark 2.19. *Although Assumption 2.18 holds in many situations when \mathcal{X} is compact, the requirement that the upper bound in (2.54) holds may not seem natural in view of the results of Section 2.2.2. Indeed, our framework shows that this upper bound is not necessary to prove the ergodicity of Feynman–Kac semigroups, as opposed to previous works [114, 116, 113, 182]. A careful look at the proof of Theorem 2.21 shows that this upper bound is only used to show the uniform boundedness of $h_{\Delta t}$ in (2.56). However, controlling $h_{\Delta t}$ as $\Delta t \rightarrow 0$ does not seem to be an easy task in general. We therefore stick to this assumption here.*

Before stating our uniform in Δt convergence result, we need the following estimate deduced from Lemma 4.19 in Chapter 4, whose proof can be found in Section 2.5.6.

Lemma 2.20. *Consider the process $(X_t)_{t \geq 0}$ solution to (2.27) with $\sigma = I_d$, $b \in C^\infty(\mathcal{X})$, and a function $f \in C^\infty(\mathcal{X})$. Then the operator $\mathcal{L} + f$ admits a real isolated largest (in modulus) eigenvalue λ with eigenvector $h \in C^\infty(\mathcal{X})$ and associated eigenspace of dimension one, which satisfies*

$$(\mathcal{L} + f)h = \lambda h, \quad \text{and} \quad P_t^f h = e^{t\lambda} h, \quad \forall t \geq 0.$$

If $Q_{\Delta t}^f$ is a consistent discretization of (2.29) satisfying Assumption 2.18, then for any $\Delta t > 0$, the operator $Q_{\Delta t}^f$ has a largest (in modulus) eigenvalue $\Lambda_{\Delta t} \in \mathbb{R}$, which is non-degenerate. The associated eigenvector $h_{\Delta t}$, such that

$$Q_{\Delta t}^f h_{\Delta t} = \Lambda_{\Delta t} h_{\Delta t},$$

is normalized as $\eta(h_{\Delta t}) = 1$. Finally, there exist $\Delta t^* > 0$, $C > 0$, $\varepsilon > 0$ such that for all $\Delta t \in (0, \Delta t^*]$, there is $c_{\Delta t} \in \mathbb{R}$ for which

$$\Lambda_{\Delta t} = e^{\Delta t \lambda + \Delta t^2 c_{\Delta t}}, \quad (2.55)$$

with $|c_{\Delta t}| \leq C$ and

$$\forall x \in \mathcal{X}, \quad \forall \Delta t \in (0, \Delta t^*], \quad \varepsilon \leq h_{\Delta t}(x) \leq \varepsilon^{-1}. \quad (2.56)$$

This result means that the evolution operator associated with a consistent discretization has a principal eigenvalue approximating the principal eigenvalue of the continuous dynamics, and that its associated principal eigenvector remains uniformly bounded from below and above if Δt is sufficiently small. We will see in Proposition 2.22 that Assumption 2.18 is naturally satisfied if a similar condition holds for $Q_{\Delta t}$ and the evolution operator reads $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$ (which corresponds to the discretization (2.52)). Let us now state the uniform in Δt version of Theorem 2.7.

Theorem 2.21. *Consider a consistent discretization $Q_{\Delta t}^f$ of the dynamics (2.29) satisfying Assumption 2.18. Then, there exists $\Delta t^* > 0$ such that, for any $\Delta t \in (0, \Delta t^*]$, the dynamics (2.53) admits a unique invariant measure $\mu_{f, \Delta t} \in \mathcal{P}(\mathcal{X})$. Moreover, there exist $\kappa > 0$, $C > 0$ such that for any $\varphi \in B^\infty(\mathcal{X})$, $\nu \in \mathcal{P}(\mathcal{X})$, and $\Delta t \in (0, \Delta t^*]$,*

$$\forall k \geq 0, \quad |\Phi_{\Delta t, k}(\nu)(\varphi) - \mu_{f, \Delta t}(\varphi)| \leq C e^{-\kappa k \Delta t} \|\varphi\|_{B^\infty}.$$

Let us note that the uniformity of the prefactor C in the initial condition is a consequence of the boundedness of \mathcal{X} . Indeed, in this case, we can choose $W \equiv 1$ as a Lyapunov function, so the constant C_ν in (2.24) can be uniformly bounded using (2.56). Such a uniformity does not hold for Theorem 2.7 since in that case \mathcal{X} was not assumed to be bounded. The important part of the theorem is the control of C and κ in the time step, which provides convergence with respect to the physical time $k\Delta t$.

Proof. The proof essentially relies on the fact that if $Q_{\Delta t}^f$ satisfies Assumption 2.18, then $Q_{h, \Delta t}$ defined as in Lemma 2.6 satisfies a uniform minorization condition. For controlling the dependencies in the time step, we rely on Lemma 2.20, and use the same notation.

We want to prove a uniform minorization condition (in the sense of [304, Lemma 3.4]) for the operator defined by

$$Q_{h, \Delta t} = \Lambda_{\Delta t}^{-1} h_{\Delta t}^{-1} Q_{\Delta t}^f h_{\Delta t},$$

and apply [304, Corollary 3.5]. Fix $T > 0$. From (2.54) and (2.56) we have, for any $\varphi \geq 0$ and $x \in \mathcal{X}$,

$$Q_{h, \Delta t}^{\lceil \frac{T}{\Delta t} \rceil} \varphi(x) = \Lambda_{\Delta t}^{-\lceil \frac{T}{\Delta t} \rceil} h_{\Delta t}^{-1} \left(Q_{\Delta t}^f \right)^{\lceil \frac{T}{\Delta t} \rceil} (h_{\Delta t} \varphi)(x) \geq \Lambda_{\Delta t}^{-\lceil \frac{T}{\Delta t} \rceil} \varepsilon^2 \alpha \eta(\varphi). \quad (2.57)$$

Moreover, from (4.23),

$$\Lambda_{\Delta t}^{-\lceil \frac{T}{\Delta t} \rceil} = e^{-\Delta t (\lambda + \Delta t c_{\Delta t}) \lceil \frac{T}{\Delta t} \rceil} \geq e^{-2|\lambda|T} > 0,$$

upon possibly reducing Δt^* . Then, (2.57) becomes

$$\forall x \in \mathcal{X}, \quad Q_{h, \Delta t}^{\lceil \frac{T}{\Delta t} \rceil}(x, \cdot) \geq \alpha \varepsilon^2 e^{-2|\lambda|T} \eta(\cdot).$$

As a result, $Q_{h, \Delta t}$ satisfies the assumptions of [304, Corollary 3.5]: there exist a unique measure $\mu_{h, \Delta t} \in \mathcal{P}(\mathcal{X})$, $C > 0$, $\kappa > 0$ such that, for any $\phi \in B^\infty(\mathcal{X})$, $k \in \mathbb{N}$ and $\Delta t \in (0, \Delta t^*]$,

$$\|Q_{h, \Delta t}^k \phi - \mu_{h, \Delta t}(\phi)\|_{B^\infty} \leq C e^{-\kappa k \Delta t} \|\phi\|_{B^\infty}.$$

This is a version of Lemma 2.6 uniform with respect to Δt . The result then follows by rewriting the proof of Theorem 2.7, with $\bar{\alpha}^k$ replaced by $e^{-\kappa k \Delta t}$.

It only remains to study the constant $C_{\nu, \Delta t}$ arising in Theorem 2.7 (see (2.24)), which now also depends on Δt through the eigenvector $h_{\Delta t}$ and the invariant measure $\mu_{h, \Delta t}$. Since \mathcal{X} is bounded, we

can actually choose a constant Lyapunov function, *i.e.* $W = \mathbb{1}$. Next, using (2.56) we obtain that for any $\Delta t \in (0, \Delta t^*]$ and any $\nu \in \mathcal{P}(\mathcal{X})$, it holds

$$C_{\nu, \Delta t} = \frac{4}{\nu_{h, \Delta t}(h_{\Delta t}^{-1})} (1 + \nu_{h, \Delta t}(h_{\Delta t}^{-1})) \frac{1}{\nu(h_{\Delta t})} \leq 4\varepsilon^{-2}(1 + \varepsilon^{-1}).$$

This provides a uniform bound on $C_{\nu, \Delta t}$, which concludes the proof. \square

We now show that the setting of Theorem 2.21 is natural, since Assumption 2.18 can be deduced from a similar assumption on the Markov dynamics $Q_{\Delta t}$ when the evolution operator is $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$, which corresponds to the discretization (2.52). For proving the condition on $Q_{\Delta t}$, we refer to [304] and the references therein.

Proposition 2.22. *Assume that \mathcal{X} is bounded, $f \in C^0(\mathcal{X})$, and the SDE (2.27) is discretized for a given time step $\Delta t > 0$ with a Markov chain $(x_k)_{k \in \mathbb{N}}$ whose evolution operator $Q_{\Delta t}$ is strong Feller and satisfies the following uniform minorization and boundedness condition: for a fixed $T > 0$, there exist $\Delta t^* > 0$, $\eta \in \mathcal{P}(\mathcal{X})$ and $\alpha \in (0, 1)$ such that, for any $\Delta t \in (0, \Delta t^*]$ and $\varphi \in B^\infty(\mathcal{X})$ with $\varphi \geq 0$,*

$$\forall x \in \mathcal{X}, \quad \alpha \eta(\varphi) \leq (Q_{\Delta t})^{\lceil \frac{T}{\Delta t} \rceil} \varphi(x) \leq \frac{1}{\alpha} \eta(\varphi).$$

Then, the transition operator $Q_{\Delta t}^f$ defined as $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$ satisfies Assumption 2.18.

Proof. Since $Q_{\Delta t}$ is strong Feller and f is continuous, $Q_{\Delta t}^f$ is strong Feller. Then, for any $k \in \mathbb{N}$ and $\varphi \in B^\infty(\mathcal{X})$,

$$(Q_{\Delta t}^f)^k \varphi(x) = \mathbb{E}_x \left[\varphi(x_k) e^{\Delta t \sum_{i=0}^{k-1} f(x_i)} \right] \geq e^{-k \Delta t \|f\|_{B^\infty}} \mathbb{E}_x [\varphi(x_k)] = e^{-k \Delta t \|f\|_{B^\infty}} ((Q_{\Delta t})^k \varphi)(x).$$

Taking $k = \lceil T/\Delta t \rceil$ with $0 < \Delta t \leq \Delta t^*$ then shows that

$$(Q_{\Delta t}^f)^{\lceil \frac{T}{\Delta t} \rceil} \varphi(x) \geq e^{-2T \|f\|_{B^\infty}} (Q_{\Delta t}^{\lceil \frac{T}{\Delta t} \rceil} \varphi)(x) \geq e^{-2T \|f\|_{B^\infty}} \alpha \eta(\varphi).$$

A similar computation for the upper bound allows to conclude the proof. \square

2.4 Discussion

The ideas developed in this chapter concerning the ergodicity of Feynman–Kac semigroups solve several problems for which, to the best of our knowledge, no solution was available. They are closely related to previous works and we want to highlight two important connections.

First, as we mentioned in the introduction, our framework can be considered as an extension of ergodic theory for Markov chains [330], when the evolution operator of the dynamics does not conserve probability. For this reason, we tried to formulate our assumptions in the flavour of [219]. However, the spectral theory on which we crucially rely in our study requires stronger conditions. This leaves open a few questions, as the converge of Feynman–Kac dynamics based on Metropolis type kernels, which lack regularity, or the case of non-Polish spaces, which may arise for stochastic partial differential equations. Finally, another interesting feature of our framework is that we can prove ergodicity for Feynman–Kac dynamics for which the underlying Markov chain is not ergodic – a case we called Diffusion Monte Carlo (DMC) in analogy with quantum physics models (see Proposition 2.13).

The other clear connection concerns large deviations theory. Indeed, one motivation for studying Feynman–Kac dynamics is to prove large deviations principles for additive functionals of Markov chains [133, 119, 427, 275], which can be achieved by proving the existence of formulas such as (3.25). It is then no surprise that the spectral theory we develop, although based on [362], is reminiscent of works such as [275], and requires stronger assumptions than the ones needed for proving ergodicity in [219] (see the discussion on the Gärtner–Ellis theorem in Section 1.2.1). However, the tools we use seem new in this context, and more adapted to the situation at hand, for instance the Krein–Rutman theorem based on the minorization condition. In particular, [275] (like [178]) makes use of nonlinear generators related to an optimal control problem. In the next chapter, we use instead linear spectral analysis to obtain large deviations principles for the empirical average of SDEs in topologies on measures corresponding to convergence with respect to unbounded functions.

Acknowledgements

The authors are grateful to Jonathan C. Mattingly for interesting discussions at a preliminary stage of this work. The authors also warmly thank Nicolas Champagnat and Denis Villemonais for pointing out a gap in one argument in the first version of the manuscript.

2.5 Complementary elements

2.5.1 Stability of Markov chains

In this short section, we recall the results presented in [219]. We consider a measurable space \mathcal{X} and Markov chain $(x_k)_{k \geq 0}$ with transition kernel Q on \mathcal{X} . By transition kernel, we mean that (i) for all $x \in \mathcal{X}$, $Q(x, \cdot)$ is a positive measure on \mathcal{X} , (ii) for any measurable set $A \subset \mathcal{X}$, $Q(\cdot, A)$ is measurable, and (iii) $Q\mathbf{1} = \mathbf{1}$. In the notation of Section 2.2, Q is a kernel operator (*i.e.* (i) and (ii) are satisfied) such that $Q\mathbf{1} = \mathbf{1}$.

The stability of Markov dynamics can be obtained from minorization and Lyapunov conditions [330, 362, 219].

Assumption 2.23. *There exist a function $\mathcal{W} : \mathcal{X} \rightarrow [1, +\infty)$ and constants $C \geq 0$, $\gamma \in (0, 1)$ such that*

$$\forall x \in \mathcal{X}, \quad (Q\mathcal{W})(x) \leq \gamma\mathcal{W}(x) + C. \quad (2.58)$$

Given such a Lyapunov function, we consider the associated functional space as in (2.7).¹ A second key ingredient in the ergodicity of Q is the following minorization condition.

Assumption 2.24. *There exist $\alpha \in (0, 1)$ and $\eta \in \mathcal{P}(\mathcal{X})$ such that*

$$\inf_{x \in \mathcal{C}} Q(x, \cdot) \geq \alpha\eta(\cdot), \quad (2.59)$$

where $\mathcal{C} = \{x \in \mathcal{X} \mid \mathcal{W}(x) \leq R + 1\}$ for some $R > 2C/(1 - \gamma)$, and γ, C are the constants from Assumption 2.23.

Under these conditions, we have the following result (see [219, Theorem 1.2]).

Theorem 2.25. *Let Assumptions 2.23 and 2.24 hold. Then, Q has a unique invariant measure μ_* , which is such that $\mu_*(\mathcal{W}) < +\infty$. Moreover, there exist $C > 0$ and $\bar{\alpha} \in (0, 1)$ such that, for any $\varphi \in B_{\mathcal{W}}^{\infty}(\mathcal{X})$,*

$$\forall k \geq 0, \quad \|Q^k\varphi - \mu_*(\varphi)\|_{B_{\mathcal{W}}^{\infty}} \leq C\bar{\alpha}^k \|\varphi - \mu_*(\varphi)\|_{B_{\mathcal{W}}^{\infty}}.$$

2.5.2 Useful theorems

We remind here some definitions and results around the Krein–Rutman theorem, as well as some basic results from analysis. Let us start with some operator-theoretic definitions from [344, 359, 112, 362].

Definition 2.26. *For a Banach space E and an operator $T \in \mathcal{B}(E)$, we denote by $\Lambda(T)$ its spectral radius defined by:*

$$\Lambda(T) = \lim_{k \rightarrow +\infty} \|T^k\|_{\mathcal{B}(E)}^{\frac{1}{k}} = \inf_{k \geq 1} \|T^k\|_{\mathcal{B}(E)}^{\frac{1}{k}}.$$

We denote by $\Theta(T)$ the essential spectral radius of T defined by (see [345, Eq. (1.14)] and [344, Theorem 1]):

$$\Theta(T) = \lim_{k \rightarrow +\infty} \left(\inf \left\{ \|T^k - Q\|_{\mathcal{B}(E)}, Q \text{ compact} \right\} \right)^{\frac{1}{k}} = \inf_{k \geq 1} \left(\inf \left\{ \|T^k - Q\|_{\mathcal{B}(E)}, Q \text{ compact} \right\} \right)^{\frac{1}{k}}.$$

An operator $T \in \mathcal{B}(E)$ is said to be compact if it maps bounded sets into precompact sets. In other words, T is compact if, for any bounded sequence $(u_n)_{n \in \mathbb{N}}$ in E , there is a subsequence $(n_k)_{k \in \mathbb{N}}$ such that $(Tu_{n_k})_{k \in \mathbb{N}}$ is convergent in E , see [359].

In order to recall the Krein–Rutman theorem, let us first give some definitions for cones in Banach spaces.

¹Compared to [219], we replace \mathcal{W} by $1 + \mathcal{W}$; this is for notational convenience only.

Definition 2.27. Let E be a Banach space. A closed convex set $\mathbb{K} \subset E$ is said to be a cone if $\mathbb{K} \cap -\mathbb{K} = \{0\}$ and for all $u \in \mathbb{K}$ and $\alpha \in \mathbb{R}_+$, it holds $\alpha u \in \mathbb{K}$. A cone is total if the norm closure of $\mathbb{K} - \mathbb{K}$ is equal to E .

We now recall a weak version of the Krein–Rutman theorem, which can be found in [345, Theorem 1.1]. Interesting remarks and comments are also available in [112, Section 19.8].

Theorem 2.28. Let E be a Banach space, $\mathbb{K} \subset E$ a total cone, and $T \in \mathcal{B}(E)$ be such that $\Theta(T) < \Lambda(T)$ and $T\mathbb{K} \subset \mathbb{K}$. Then $\Lambda(T)$ is an eigenvalue of T with an eigenvector in \mathbb{K} .

In Theorem 2.28, there is no uniqueness of the eigenvector. The non degeneracy can be obtained under stronger positivity conditions on the operator T , as made precise in [112, Theorems 19.3 and 19.5]. In order to control the essential spectral radius and apply the Krein–Rutman theorem, we will need the following classical results, see [379, Theorem 11.28] and [384, Theorem 2.7.19].

Theorem 2.29 (Ascoli). Let $(\mathcal{Y}, d_{\mathcal{Y}})$ be a compact metric space and $C^0(\mathcal{Y})$ be the space of continuous functions over \mathcal{Y} endowed with the uniform norm $\|f\|_{C^0} = \sup_{y \in \mathcal{Y}} |f(y)|$. Consider a uniformly bounded and equicontinuous sequence $(f_n)_{n \in \mathbb{N}}$, i.e. a sequence for which there exists $M > 0$ such that $\|f_n\|_{C^0} \leq M$ for all $n \geq 1$, and for any $\varepsilon > 0$ there exists $\delta > 0$ such that $d_{\mathcal{Y}}(x, y) \leq \delta$ implies $|f(x) - f(y)| \leq \varepsilon$. Then $(f_n)_{n \in \mathbb{N}}$ converges in the uniform norm to some limit f up to extraction.

Theorem 2.30 (Heine–Cantor). Let $f : E \rightarrow F$ where (E, d_E) and (F, d_F) are two metric spaces and E is compact. Then, if f is continuous, it is uniformly continuous: for any $\varepsilon > 0$, there is $\delta > 0$ such that for any $x, x' \in E$ with $d_E(x, x') \leq \delta$, it holds $d_F(f(x), f(x')) \leq \varepsilon$.

We close this section with some results in probability theory. The next lemma can be found in [217, Lemma 4.14].

Lemma 2.31. If \mathcal{X} is a Polish space and $\nu \in \mathcal{P}(\mathcal{X})$, then the family constituted of the single measure ν is tight, i.e. for any $\varepsilon > 0$, there exists a compact set $K \subset \mathcal{X}$ such that $\nu(K) \geq 1 - \varepsilon$.

We finally present results concerning ultra-Feller operators, extending the ones of [218, Appendix A]. Recall that the total variation distance between two positive measures $\mu, \nu \in \mathcal{M}(\mathcal{X})$ is defined by the norm:

$$\|\mu - \nu\|_{\text{TV}} = \sup_{\substack{\varphi \in B^\infty(\mathcal{X}) \\ \|\varphi\|_{B^\infty} \leq 1}} \int_{\mathcal{X}} \varphi d\mu - \int_{\mathcal{X}} \varphi d\nu. \quad (2.60)$$

Definition 2.32 (Ultra-Feller). A kernel operator Q is ultra-Feller if the mapping $x \mapsto Q(x, \cdot) \in \mathcal{M}(\mathcal{X})$ is continuous in the total variation distance (2.60).

The next lemma, used to show that an operator is ultra-Feller, is adapted from [218, Appendix A].

Lemma 2.33. Suppose that P and Q are two kernel operators over a Polish space \mathcal{X} that satisfy the following properties:

- for all $\varphi \in B^\infty(\mathcal{X})$, $Q\varphi$ is continuous and finite;
- for all ψ such that $|\psi| \leq Q\mathbf{1}$, $P\psi$ is continuous and finite.

Then PQ is ultra-Feller.

We remind some elements of the proof from [218, Theorem 1.6.6], which is based on the Banach–Alaoglu theorem. The details are left to the reader.

Proof. A first element to prove Lemma 2.33 is to show that, if Q is strong Feller, then there exists a reference probability measure $\zeta \in \mathcal{P}(\mathcal{X})$ such that for any $x \in \mathcal{X}$, $Q(x, \cdot)$ is absolutely continuous with respect to ζ . This is shown in [218, Lemma 1.6.4] for operators Q such that $Q\mathbf{1} = 1$. Even for a non-probabilistic Q , we can consider the normalized probabilities

$$\frac{Q(x, \cdot)}{Q\mathbf{1}(x)},$$

for x in the open set $\tilde{X} = \{x \in \mathcal{X} \mid Q\mathbf{1}(x) > 0\}$. We can apply [218, Lemma 1.6.4] to these probabilities defined over the set \tilde{X} , so there exists a measure ζ such that, for any $x \in \tilde{X}$, $Q(x, \cdot)$ is absolutely

continuous with respect to ζ . If $x \in \mathcal{X} \setminus \widetilde{X}$, $Q(x, \cdot) = 0$, which is also absolutely continuous with respect to ζ , so that $Q(x, \cdot)$ is absolutely continuous with respect to ζ for any $x \in \mathcal{X}$.

Once this is done, one can write the kernel Q as $Q(y, dz) = k(y, z) dz$ with $k(y, \cdot) \in L^1(\mathcal{X}, \zeta)$ for all $x \in \mathcal{X}$. If one supposes by contradiction that PQ is not ultra-Feller, then Definition 2.32 shows that there exist a sequence of functions $(g_n)_{n \in \mathbb{N}}$ with $\|g_n\|_{B^\infty} \leq 1$ and a sequence $(x_n)_{n \in \mathbb{N}}$ converging to an element $x \in \mathcal{X}$ such that for some $\delta > 0$ it holds

$$\forall n \in \mathbb{N}, \quad PQg_n(x_n) - PQg_n(x) > \delta. \quad (2.61)$$

Since the sequence $(g_n)_{n \in \mathbb{N}}$ is bounded, it possesses a weak-* converging subsequence in $L^\infty(\mathcal{X}, \zeta)$ (the space of ζ -essentially bounded functions) to an element $g \in B^\infty(\mathcal{X}, \zeta)$. In particular it holds (upon extracting a subsequence), for any $y \in \mathcal{X}$,

$$\lim_{n \rightarrow +\infty} Qg_n(y) = \lim_{n \rightarrow +\infty} \int_{\mathcal{X}} k(y, z) g_n(z) \zeta(dz) = \int_{\mathcal{X}} k(y, z) g(z) \zeta(dz) = Qg(y).$$

Defining, $f_n = Qg_n$, the latter limit shows that f_n converges pointwise to $f = Qg$. Since $(g_n)_{n \in \mathbb{N}}$ is bounded in $B^\infty(\mathcal{X})$, the second condition in Lemma 2.33 ensures that $Pf_n(x) \rightarrow Pf(x)$ for all $x \in \mathcal{X}$, by the dominated convergence theorem. This is the main difference compared to the proof in [218, Theorem 1.6.6]. The contradiction follows similarly. Indeed, defining the positive decreasing function $h_n = \sup_{m \geq n} |f_m - f|$ we have, for any $m \in \mathbb{N}$,

$$\lim_{n \rightarrow +\infty} Ph_n(x_n) \leq \lim_{n \rightarrow +\infty} Ph_m(x_n) = Ph_m(x),$$

so that $Ph_n(x_n) \rightarrow 0$ as $n \rightarrow +\infty$. In the end,

$$\lim_{n \rightarrow +\infty} Pf_n(x_n) - Pf(x) \leq \lim_{n \rightarrow +\infty} |Pf_n(x_n) - Pf(x_n)| + \lim_{n \rightarrow +\infty} |Pf(x_n) - Pf(x)| = 0,$$

which comes in contradiction with (2.61) and concludes the proof. \square

2.5.3 Proof of Lemma 2.4

Let us show that $\nu(Q^f \mathbf{1}) > 0$ for any $\nu \in \mathcal{P}(\mathcal{X})$. First, Lemma 2.31 in Section 2.5.2 ensures that, for any $\varepsilon > 0$, there exists a compact $K \subset \mathcal{X}$ such that $\nu(K) \geq 1 - \varepsilon$. Consider next a compact set K_n of Assumption 2.1 such that $K \subset K_n$. Then, with the corresponding $\alpha_n > 0$ and $\eta_n \in \mathcal{P}(\mathcal{X})$ defined in Assumption 2.2, we have

$$\forall x \in K_n, \quad (Q^f \mathbf{1})(x) \geq \alpha_n \eta_n(\mathbf{1}) \geq \alpha_n > 0.$$

Integrating with respect to ν leads to

$$\int_{\mathcal{X}} (Q^f \mathbf{1})(x) \nu(dx) \geq \int_{K_n} (Q^f \mathbf{1})(x) \nu(dx) \geq \alpha_n \int_{K_n} \nu(dx) = \alpha_n \nu(K_n) \geq \alpha_n (1 - \varepsilon) > 0,$$

since $K \subset K_n$, which proves the statement. Moreover $W \geq 1$, so (2.6) implies that $\nu(Q^f \mathbf{1}) \leq \nu(Q^f W) < +\infty$ if $\nu(W) < +\infty$.

Since $W \geq 1$, we immediately have that $\eta_n(W) \geq 1 > 0$ for any $n \geq 1$. Now, for any $n \geq 1$ and $x \in K_n$, Assumptions 2.1 and 2.2 lead to

$$\alpha_n \eta_n(W) \leq Q^f W(x) \leq \gamma_n W(x) + b_n \mathbf{1}_{K_n}(x) < +\infty,$$

since W is finite. Moreover, $\alpha_n > 0$, so that $\eta_n(W) < +\infty$ for any $n \geq 1$.

Let us conclude with the proof of (2.11). We proceed by contradiction and assume that, for any $n \geq n_0$, we have $\eta_n(K_n) = 0$, with $n_0 \geq 1$ an arbitrary integer. Consider $m \geq n_0$ and $\varphi_m = \mathbf{1}_{K_m} \geq 0$. Then, using (2.8) with $n = n_0$,

$$\forall x \in K_{n_0}, \quad (Q^f \varphi_m)(x) \geq \alpha_{n_0} \eta_{n_0}(K_m).$$

Using again Lemma 2.31 in Section 2.5.2, we see that for m large enough, $(Q^f \varphi_m)(x) > 0$ for $x \in K_{n_0}$ and so $Q^f \varphi_m \neq 0$. However, for $n \geq m$, we have, using that $K_m \subset K_n$ (since the sets are increasing):

$$0 \leq \eta_n(\varphi_m) = \eta_n(K_m) \leq \eta_n(K_n) = 0,$$

since we assumed $\eta_n(K_n) = 0$ for $n \geq n_0$. The contradiction with (2.9) shows that there exists $\bar{n} \geq n_0$ such that $\eta_{\bar{n}}(K_{\bar{n}}) > 0$. Since n_0 is arbitrary, \bar{n} can be chosen arbitrarily large, and this concludes the proof of Lemma 2.4.

2.5.4 Proof of Lemma 2.5

The proof is decomposed into three steps. First we show that the essential spectral radius of the operator Q^f considered over $B_W^\infty(\mathcal{X})$ is zero. We next prove that the spectral radius Λ of Q^f is positive. Finally, we use the Krein–Rutman theorem to obtain that Λ is a eigenvalue of Q^f with largest modulus, and that the associated eigenvector is positive.

Step 1: Q^f has zero essential spectral radius

We first perform the following decomposition, for any $n \geq 1$:

$$(Q^f)^2 = \mathbb{1}_{K_n} Q^f \mathbb{1}_{K_n} Q^f + \mathbb{1}_{K_n^c} (Q^f)^2 + \mathbb{1}_{K_n} Q^f \mathbb{1}_{K_n^c} Q^f,$$

where $K_n \subset \mathcal{X}$ are the compact sets from Section 2.2.2. Applying again Q^f leads to

$$(Q^f)^3 = (\mathbb{1}_{K_n} Q^f \mathbb{1}_{K_n})^2 Q^f + \mathbb{1}_{K_n^c} Q^f (\mathbb{1}_{K_n} Q^f)^2 + Q^f \mathbb{1}_{K_n^c} (Q^f)^2 + Q^f \mathbb{1}_{K_n} Q^f \mathbb{1}_{K_n^c} Q^f. \quad (2.62)$$

We will show that $Q_n^f = \mathbb{1}_{K_n} Q^f \mathbb{1}_{K_n}$ is such that $(Q_n^f)^2$ is compact on $B_W^\infty(\mathcal{X})$, while $\mathbb{1}_{K_n^c} Q^f$ tends to zero in operator norm. This will prove that $(Q^f)^3$ is compact as limit of compact operators in operator norm, so the essential spectral radius of Q^f in $B_W^\infty(\mathcal{X})$, denoted by $\Theta(Q^f)$, is equal to zero.

Let us first prove that $(Q_n^f)^2$ is compact on $B_W^\infty(\mathcal{X})$ for any $n \in \mathbb{N}$. For this, we use the ultra-Feller property proved in Lemma 2.33 (see Section 2.5.2) to apply the Ascoli theorem. Consider a sequence $(\varphi_k)_{k \in \mathbb{N}}$ in $B_W^\infty(\mathcal{X})$ such that $\|\varphi_k\|_{B_W^\infty} \leq M$ for some $M \geq 0$. By Assumption 2.3, the operator Q_n^f is strong Feller over the compact set K_n . In particular, for $\varphi \in B_W^\infty(\mathcal{X})$, $\varphi \mathbb{1}_{K_n} \in B^\infty(\mathcal{X})$, so $Q_n^f \varphi$ is continuous over K_n and finite, so that Lemma 2.33 in Section 2.5.2 applies. Indeed, the second condition in the lemma is easy to check since $Q_n \mathbb{1}$ is equal to zero outside the compact K_n . Therefore, $(Q_n^f)^2$ is ultra-Feller by Lemma 2.33. By Definition 2.32, the application $x \in K_n \mapsto (Q_n^f)^2(x, \cdot) \in \mathcal{M}(\mathcal{X})$ is continuous in total variation norm. Since K_n is compact in the metric space \mathcal{X} and $\mathcal{P}(\mathcal{X})$ is a metric space, the Heine–Cantor theorem (Theorem 2.30 in Section 2.5.2) ensures that this application is continuous over K_n . This means that, for any $\varepsilon > 0$, there exists $\delta > 0$ such that, for any $x, x' \in K_n$ with $|x - x'| \leq \delta$, it holds

$$\sup_{\|\varphi\|_{B^\infty} \leq 1} \left| ((Q_n^f)^2 \varphi)(x) - ((Q_n^f)^2 \varphi)(x') \right| \leq \varepsilon. \quad (2.63)$$

Noting that Assumption 2.1 implies that $1 \leq \sup_{K_n} W < +\infty$, it holds $M_n = (\sup_{K_n} W)^{-1} \in (0, 1]$ for any $n \geq 1$, so

$$\{\varphi \text{ measurable} \mid \|\mathbb{1}_{K_n} \varphi\|_{B^\infty} \leq 1\} \supset \{\varphi \text{ measurable} \mid \|\mathbb{1}_{K_n} \varphi\|_{B_W^\infty} \leq M_n\}. \quad (2.64)$$

Since $Q_n^f = \mathbb{1}_{K_n} Q^f \mathbb{1}_{K_n}$, (2.64) shows that (2.63) becomes

$$\sup_{\|\varphi\|_{B_W^\infty} \leq M_n} \left| ((Q_n^f)^2 \varphi)(x) - ((Q_n^f)^2 \varphi)(x') \right| \leq \varepsilon.$$

As a consequence, if $(\varphi_k)_{k \in \mathbb{N}}$ is such that $\|\varphi_k\|_{B_W^\infty} \leq M$, we see that $((Q_n^f)^2 \varphi_k)_{k \in \mathbb{N}}$ is equicontinuous. By the Ascoli theorem, it therefore converges uniformly to a continuous limit on K_n (since the function is supported on K_n , we extend it by 0 on \mathcal{X} outside K_n). Since $W \geq 1$, it also converges as a function in $B_W^\infty(\mathcal{X})$, showing that $(Q_n^f)^2$ is a compact operator on $B_W^\infty(\mathcal{X})$. Since Q^f is bounded over $B_W^\infty(\mathcal{X})$ and the space of compact operators is stable by composition with bounded operators [359], $(Q_n^f)^2 Q^f$ is also compact.

We now show that the second, third and fourth operators on the right hand side of (2.62) tend to 0 in the operator norm of $B_W^\infty(\mathcal{X})$. For any $\varphi \in B_W^\infty(\mathcal{X})$,

$$\|\mathbb{1}_{K_n^c} Q^f \varphi\|_{B_W^\infty} = \left\| \frac{\mathbb{1}_{K_n^c} Q^f \varphi}{W} \right\|_{B^\infty} \leq \|\varphi\|_{B_W^\infty} \left\| \mathbb{1}_{K_n^c} \frac{Q^f W}{W} \right\|_{B^\infty} \leq \gamma_n \|\varphi\|_{B_W^\infty}.$$

Taking the supremum over $\varphi \in B_W^\infty(\mathcal{X})$ and using $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$, we obtain:

$$\|\mathbb{1}_{K_n^c} Q^f\|_{B(B_W^\infty)} \xrightarrow{n \rightarrow +\infty} 0. \quad (2.65)$$

Since Q^f is bounded on $\mathcal{B}(B_W^\infty)$, the second, third and fourth operators on the right hand side of (2.62) vanish in norm as $n \rightarrow +\infty$. As a result, $(Q^f)^3$ is the norm-limit of the compact operators $(Q_n^f)^2 Q^f$ as $n \rightarrow +\infty$ in $\mathcal{B}(B_W^\infty)$. Since the set of compact operators over $B_W^\infty(\mathcal{X})$ is closed in the Banach space $\mathcal{B}(B_W^\infty)$, $(Q^f)^3$ is compact, see *e.g.* [359, Theorem VI.12]. Using Definition 2.26, we conclude that $\Theta(Q^f) = 0$. In this procedure, we see that working in the weighted space $B_W^\infty(\mathcal{X})$ as opposed to $B^\infty(\mathcal{X})$ is crucial in order to obtain the compactness of $(Q^f)^3$ from the control (2.65) provided by the Lyapunov condition (2.6).

Step 2: The spectral radius is positive

We now show that the spectral radius Λ of Q^f defined in (2.12) is positive, in order to use Theorem 2.28. Given the definition of the operator norm, choosing some arbitrary non negative function $\phi \in B_W^\infty(\mathcal{X})$ with $\|\phi\|_{B_W^\infty} \leq 1$ leads to

$$\|Q^f\|_{\mathcal{B}(B_W^\infty)} \geq \left\| \frac{Q^f \phi}{W} \right\|_{B^\infty} \geq \frac{(Q^f \phi)(x_0)}{W(x_0)},$$

where $x_0 \in \mathcal{X}$ is arbitrary. We now consider a compact set corresponding to some $n = \bar{n}$ as defined in Lemma 2.4, which satisfies $\eta_{\bar{n}}(K_{\bar{n}}) > 0$, and take $x_0 \in K_{\bar{n}}$. For any non negative function $\phi \in B_W^\infty(\mathcal{X})$ with $\|\phi\|_{B_W^\infty} \leq 1$,

$$\eta_{\bar{n}}(Q^f \phi) = \left(\int_{K_{\bar{n}}} (Q^f \phi)(x) \eta_{\bar{n}}(dx) + \int_{\mathcal{X} \setminus K_{\bar{n}}} (Q^f \phi)(x) \eta_{\bar{n}}(dx) \right) \geq \int_{K_{\bar{n}}} \alpha_{\bar{n}} \eta_{\bar{n}}(\phi) \eta_{\bar{n}}(dx) \geq \alpha_{\bar{n}} \eta_{\bar{n}}(\phi) \eta_{\bar{n}}(K_{\bar{n}}), \quad (2.66)$$

where we used (2.8) with $n = \bar{n}$. Iterating the inequality shows that

$$\forall k \geq 1, \quad \eta_{\bar{n}}((Q^f)^k \phi) \geq \alpha_{\bar{n}}^k \eta_{\bar{n}}(K_{\bar{n}})^k \eta_{\bar{n}}(\phi).$$

This leads to the following lower bound on the operator norm of $(Q^f)^k$:

$$\begin{aligned} \|(Q^f)^k\|_{\mathcal{B}(B_W^\infty)} &\geq \frac{((Q^f)^k \phi)(x_0)}{W(x_0)} = \frac{(Q^f((Q^f)^{k-1} \phi))(x_0)}{W(x_0)} \\ &\geq \alpha_{\bar{n}} \frac{\eta_{\bar{n}}((Q^f)^{k-1} \phi)}{W(x_0)} \geq \frac{\alpha_{\bar{n}}^k \eta_{\bar{n}}(K_{\bar{n}})^{k-1}}{W(x_0)} \eta_{\bar{n}}(\phi). \end{aligned}$$

Taking the power $1/k$ and the limit $k \rightarrow +\infty$, together with the choice $\phi = \mathbf{1} \in B_W^\infty(\mathcal{X})$, leads to

$$\Lambda \geq \alpha_{\bar{n}} \eta_{\bar{n}}(K_{\bar{n}}).$$

From Lemma 2.4, it holds $\eta_{\bar{n}}(K_{\bar{n}}) > 0$, hence $\Lambda > 0$ and Q^f has a positive spectral radius. Note that the existence of $\bar{n} \geq 1$ such that $\eta_{\bar{n}}(K_{\bar{n}}) > 0$ is crucial for this step.

Step 3: Existence of a principal eigenvector

In order to use Theorem 2.28, we introduce the closed cone:

$$\mathbb{K}_W = \{u \in B_W^\infty(\mathcal{X}) \mid u \geq 0\}.$$

This cone is total, and the positiveness of $Q^f \in B_W^\infty(\mathcal{X})$ shows that $Q^f \mathbb{K} \subset \mathbb{K}$. At this stage, Theorem 2.28 in Section 2.5.2 ensures that the spectral radius Λ is an eigenvalue of Q^f of largest modulus with an associated eigenvector $h \in \mathbb{K}_W \setminus \{0\}$.

Step 4: Positivity

We now use the irreducibility condition (2.9) to show that, for the eigenvector h obtained in Step 3, it holds $h(x) > 0$ for all $x \in \mathcal{X}$ and hence $\eta_n(h) > 0$ all $n \geq 1$.

Let us show the first property by contradiction. Assume that there exists $x_0 \in \mathcal{X}$ such that $h(x_0) = 0$. Since the sets K_n are increasing, there exists n_0 such that for all $n \geq n_0$ it holds $x_0 \in K_n$ so that, by (2.8),

$$\forall n \geq n_0, \quad (Q^f h)(x_0) \geq \alpha_n \eta_n(h).$$

Since $Q^f h = \Lambda h$ with $\Lambda > 0$, this leads to

$$0 \geq \eta_n(h),$$

and so $\eta_n(h) = 0$ for $n \geq n_0$. By the irreducibility assumption (2.9), we therefore have $(Q^f h)(x) = 0$ for all $x \in \mathcal{X}$. Using again $Q^f h = \Lambda h$, this shows that $h = 0$, which is in contradiction with the fact that h is an eigenvector associated with Λ .

The second property follows from $h(x) > 0$ for all $x \in \mathcal{X}$ and $\eta_n \in \mathcal{P}(\mathcal{X})$ for all $n \geq 1$. Indeed,

$$\mathcal{X} = \bigcup_{k \geq 1} h^{-1} \left[\frac{1}{k}, +\infty \right), \quad (2.67)$$

where h^{-1} denotes here the pre-image of h . Therefore, for a given $n \geq 1$,

$$\eta_n(\mathcal{X}) = \eta_n(h^{-1}[1, +\infty)) + \sum_{k \geq 1} \eta_n \left(h^{-1} \left[\frac{1}{k+1}, \frac{1}{k} \right) \right) = 1.$$

Thus, there exists $N \geq 1$ such that

$$\eta_n \left(h^{-1} \left[\frac{1}{N}, +\infty \right) \right) \geq \frac{1}{2},$$

so

$$\eta_n(h) \geq \eta_n(h \mathbb{1}_{h \geq \frac{1}{N}}) \geq \frac{1}{N} \eta_n(\mathbb{1}_{h \geq \frac{1}{N}}) \geq \frac{1}{2N}.$$

Since $n \geq 1$ is arbitrary, this shows that $\eta_n(h) > 0$ for all $n \geq 1$.

2.5.5 Proof of Lemma 2.6

A first important remark is that Q_h is a Markov operator. Indeed, it is a well-defined kernel operator (since $0 < h(x) < +\infty$ for all $x \in \mathcal{X}$), and $Q_h \mathbb{1} = \Lambda^{-1} h^{-1} Q^f h = \Lambda^{-1} h^{-1} \Lambda h = \mathbb{1}$. Our goal is therefore to show that the Markov operator Q_h fits the framework reminded in Section 2.5.1, in particular that it satisfies Assumptions 2.23 and 2.24.

Let us show that this operator satisfies Assumption 2.23 in Section 2.5.1 with Lyapunov function Wh^{-1} . We first note that the normalization $\|h\|_{B_W^\infty} = 1$ implies that $Wh^{-1} \geq 1$. Using Assumption 2.1, we obtain

$$Q_h(Wh^{-1}) = \Lambda^{-1} h^{-1} Q^f W \leq \Lambda^{-1} h^{-1} (\gamma_n W + b_n \mathbb{1}_{K_n}) \leq \frac{\gamma_n}{\Lambda} Wh^{-1} + \frac{b_n}{\Lambda h} \mathbb{1}_{K_n}.$$

Noting that, for all $x \in K_n$,

$$\Lambda h(x) = (Q^f h)(x) \geq \alpha_n \eta_n(h),$$

with $\eta_n(h) > 0$ from Lemma 2.5, the above inequality becomes

$$Q_h(Wh^{-1}) \leq \frac{\gamma_n}{\Lambda} Wh^{-1} + \frac{b_n}{\alpha_n \eta_n(h)} \mathbb{1}_{K_n}. \quad (2.68)$$

Since γ_n can be taken arbitrarily small and $\eta_n(h) > 0$ for any $n \geq 1$, we deduce that Wh^{-1} is a Lyapunov function for Q_h in the sense of Assumption 2.23 in Section 2.5.1.

Remark 2.34. *Let us mention that, in order for (2.68) to define a Lyapunov condition in the sense of Assumption 2.23, it is not necessary to have $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$. The existence of $n \geq 1$ such that $\gamma_n < \Lambda$ is sufficient.*

We will now prove that: (i) Wh^{-1} has compact level sets, and (ii) Q_h satisfies Assumption 2.23 in Section 2.5.1 on any compact set K_n , that is $\inf_{K_n} Q_h$ is lower bounded by some probability measure. First, choosing $x_n \notin K_n$ in Assumption 2.1 leads to

$$\Lambda h(x_n) = (Q^f h)(x_n) \leq \gamma_n W(x_n),$$

so that

$$\frac{W(x_n)}{h(x_n)} \geq \frac{\Lambda}{\gamma_n}. \quad (2.69)$$

Since $\gamma_n \rightarrow 0$ as $n \rightarrow +\infty$, the function Wh^{-1} diverges outside the compact sets K_n defined in Assumption 2.1. In other words, Wh^{-1} has compact level sets, which shows (i).

Next, for $n \geq 1$, consider $\alpha_n > 0$ and $\eta_n \in \mathcal{P}(\mathcal{X})$ as in Assumption 2.2, so that, for any bounded measurable function $\varphi \geq 0$ and $x \in K_n$,

$$Q_h \varphi(x) = \Lambda^{-1} \frac{Q^f(h\varphi)(x)}{h(x)} \geq \frac{1}{\Lambda \sup_{K_n} h} \alpha_n \eta_n(h\varphi) \geq \tilde{\alpha}_n \tilde{\eta}_n(\varphi),$$

with

$$\tilde{\alpha}_n = \alpha_n \frac{\eta_n(h)}{\Lambda \sup_{K_n} h} > 0, \quad \tilde{\eta}_n(\varphi) = \frac{\eta_n(h\varphi)}{\eta_n(h)} \in \mathcal{P}(\mathcal{X}).$$

The latter expression is well-defined because, from Lemma 2.5, we know that $0 < \eta_n(h) < +\infty$ for any $n \geq 1$. Moreover, $0 < \sup_{K_n} h < +\infty$ (since $h \in B_{W^\infty}(\mathcal{X})$ and $\sup_{K_n} W < +\infty$ by Assumption 2.1), and this yields precisely (ii). Finally, (i) and (ii) show that Q_h satisfies Assumption 2.24, so that Q_h satisfies the assumptions of Theorem 2.25. As a result there exist a unique $\mu_h \in \mathcal{P}(\mathcal{X})$ and constants $c > 0$, $\bar{\alpha} \in (0, 1)$ such that for any $\phi \in B_{Wh^{-1}}^\infty(\mathcal{X})$,

$$\forall k \geq 0, \quad \|Q_h^k \phi - \mu_h(\phi)\|_{B_{Wh^{-1}}^\infty} \leq c \bar{\alpha}^k \|\phi - \mu_h(\phi)\|_{B_{Wh^{-1}}^\infty}.$$

Moreover, the measure μ_h satisfies $\mu_h(Wh^{-1}) < +\infty$.

2.5.6 Proof of Lemma 2.20

From Proposition 4.1 in Chapter 4, we obtain that $\mathcal{L} + f$ has a largest (in modulus) eigenvalue λ with associated smooth eigenvector h . Similarly, Lemma 2.5 shows that for any $\Delta t \in (0, \Delta t^*]$ the operator $Q_{\Delta t}^f$ has a largest (in modulus) eigenvalue $\Lambda_{\Delta t}$ with continuous eigenvector $h_{\Delta t}$ (since $Q_{\Delta t}^f$ is assumed to be strong Feller). Moreover, there is no restriction of generality in normalizing $h_{\Delta t}$ so that $\eta(h_{\Delta t}) = 1$.

We now turn to the estimate (4.23) on the spectral radius. In the notation of Chapter 4, we have $\Lambda_{\Delta t} = e^{\Delta t \lambda_{\Delta t}}$. A direct application of Theorem 4.21 below then shows that there exist $\Delta t^* > 0$ and $C > 0$ such that $\lambda_{\Delta t} = \lambda + \Delta t c_{\Delta t}$ with $|c_{\Delta t}| \leq C$ for $\Delta t \in (0, \Delta t^*]$, which is the desired result.

Finally, since $Q_{\Delta t}^f h_{\Delta t} = \Lambda_{\Delta t} h_{\Delta t}$, the lower bound (2.54) applied to $\varphi = h_{\Delta t} \geq 0$ leads to

$$\forall x \in \mathcal{X}, \quad \left(Q_{\Delta t}^f\right)^{\lceil \frac{T}{\Delta t} \rceil} h_{\Delta t}(x) = \Lambda_{\Delta t}^{\lceil \frac{T}{\Delta t} \rceil} h_{\Delta t}(x) \geq \alpha \eta(h_{\Delta t}).$$

Using the estimate on $\Lambda_{\Delta t}$ and the normalization $\eta(h_{\Delta t}) = 1$ we obtain, for $\Delta t \in (0, \Delta t^*]$ (possibly upon decreasing Δt^*) and $x \in \mathcal{X}$,

$$h_{\Delta t}(x) \geq \Lambda_{\Delta t}^{-\lceil \frac{T}{\Delta t} \rceil} \alpha \eta(h_{\Delta t}) \geq \alpha e^{-\Delta t(\lambda + \Delta t c_{\Delta t}) \lceil \frac{T}{\Delta t} \rceil} \geq \alpha e^{-2T|\lambda|}.$$

A similar computation leads to an analogous upper bound, which shows (2.56).

CHAPTER 3

LARGE DEVIATIONS OF EMPIRICAL MEASURES OF DIFFUSIONS IN FINE TOPOLOGIES

The material for this chapter has been released on arXiv in [183].

Contents

3.1	Presentation	104
3.2	Large deviations principle	106
3.2.1	Setting	106
3.2.2	Statement of the main results	109
3.3	Decomposition of the rate function	113
3.3.1	Donsker–Varadhan variational formula	113
3.3.2	Entropy decomposition: symmetry and antisymmetry	114
3.4	Applications	116
3.4.1	Overdamped Langevin dynamics	116
3.4.2	Underdamped Langevin dynamics	118
3.5	Conclusion and perspectives	122
3.6	Proofs	123
3.6.1	Proof of the large deviations principle	123
3.6.2	Proofs of Section 3.3	135
3.6.3	Proof of Proposition 3.8	139
3.6.4	Proof of Lemma 3.23	139

Abstract. We consider large deviations of the empirical measure of diffusion processes, as presented in the Introduction of the thesis. In a first part, we present conditions to obtain a large deviations principle (LDP) in a Wasserstein-like topology. In particular, we derive a precise class of unbounded functions for which the LDP holds. This provides an analogue to the standard Cramér condition (as presented in Theorem 1.9 in Section 1.2.2) in the context of diffusion processes. This condition turns out to be related to a spectral gap condition for a Witten–Schrödinger operator. Secondly, we study more precisely the properties of the Donsker–Varadhan rate functional associated to the LDP. We revisit and generalize some standard duality results as well as a more original decomposition of the rate functional with respect to the symmetric and antisymmetric parts of the dynamics, which is motivated by physical and practical considerations presented in Section 1.1. Finally, we apply our results to overdamped and underdamped Langevin dynamics, showing the applicability of our framework in both unbounded and degenerate situations which are relevant for applications.

3.1 Presentation

Empirical averages of diffusion processes and their convergence are commonly studied in statistical mechanics, probability theory and machine learning, as motivated in Part I. In statistical physics, an observable averaged along the trajectory of a diffusion typically converges to the macroscopic average representing the stationary distribution of the system [266, 301]. For reversible dynamics, this convergence is known to be characterized by an entropy functional [406, 25], which generalizes results for small fluctuations such as the central limit theorem [273] or Berry-Esseen type inequalities [343], see the discussion in Section 1.1.2. It has been shown for some time that the approach can be extended to nonequilibrium systems by considering generalized entropy and free energy functionals, as provided by the large deviations theory [119, 166, 406]. In a more computational perspective, studying the convergence of empirical averages is an important problem for the efficiency of Monte Carlo Markov Chain methods [15, 368, 363].

Since its initiation by Cramér in the 1930s [103], large deviations theory has been given many extensions. The theory takes its origin in the study of fluctuations for sums of independent variables, leading to the celebrated Sanov theorem [121]. Interestingly, the necessity of Cramér’s exponential moment condition for the Sanov theorem to hold in a Wasserstein topology has been only recently proved [424], see the presentation in Section 1.2.2 for more intuition.

Due to the above mentioned applications, it is natural to try to apply such a theory to diffusions, or more generally Markovian dynamics. This is useful for instance in statistical physics, when considering Gallavotti–Cohen fluctuation relations for irreversible systems [194, 288, 280], as well as for characterizing dynamical phase transitions in physical systems [198, 18, 339, 346]. From a more computational perspective, studying the rate function associated with a given dynamics is interesting for designing better sampling strategies [363, 364, 146], which is important for instance in a Bayesian framework [70, 47]. The approach can also be used for deriving concentration results such as Bernstein-type inequalities [195, 43] and uncertainty quantification bounds [264, 206].

However, proving a large deviations principle for correlated processes turns out to be a difficult task. A milestone in the theory is the series of papers by Donsker and Varadhan [131, 132, 134, 135] and the dual approach followed by Gärtner and Ellis [199, 165]. The strategy of the first is to build explicitly the lower and upper large deviations bounds from the Tchebychev inequality and the Girsanov theorem [416]. On the other hand, the Gärtner–Ellis theorem relies on the existence and regularity of a free energy function (we refer again to the Introduction for a more throughout discussion). This technique has been later related to optimal control problems through the so-called weak convergence approach [147, 149].

Whatever strategy is chosen, proving large deviations principles for empirical measures of unbounded diffusions remains difficult. Indeed, proving the stability of unbounded Markov processes is already challenging, and often relies on Lyapunov function techniques [330, 324, 362, 219]. Such a Lyapunov function can be interpreted as an energy associated with the system, which decreases in average and provides a control on the excursions of the process far away from the origin. This technique can be used for proving LDPs, see for instance [416, Section 9] and [124, 427, 149]. However, the above mentioned works consider LDP in the so-called strong (resp. weak) topology, *i.e.* with respect to the topology on measures associated with the convergence of measurable bounded (resp. continuous bounded) functions. To the best of our knowledge, Wasserstein topologies (*i.e.* associated with unbounded functions) for diffusions has only been addressed in [275], and [427, Section 2.2]. Unfortunately, the nonlinear approach of [275] does not allow to characterize precisely the set of functions for which the LDP holds, while [427] considers a particular system (Langevin dynamics). In both cases, the rate function is not related to the standard Donsker–Varadhan theory [133]. Our first result is to derive the LDP in a Wasserstein topology under very natural conditions, and to express the rate function in duality with a free energy. From a practical point of view, this allows to compute the rate function from the free energy, a standard procedure [202, 406, 97, 337, 184].

Once a large deviations principle has been derived, providing alternative expressions of the rate function is an important problem. Our first contribution in this direction is to derive a variational representation of the rate function similar to the Donsker–Varadhan formula [133]. This provides a variational representation of the principal eigenvalue for any non-symmetric linear second order differential operator associated with a diffusion, under confinement and regularity conditions. To the best of our knowledge, there is no such formula in an unbounded setting, a fortiori for unbounded functions and Wasserstein topologies on probability measures. Finally, it has been shown in a pioneering work [49] that (in the case of a Langevin dynamics) the above mentioned duality allows to

decompose the rate function into two parts: one corresponding to a «reversible» part and the other to an «irreversible part» of the dynamics. We extend these results to general diffusions by using Sobolev seminorms, a feature inspired by the small fluctuations framework developed in [273]. This decomposition turns out to be useful for various purposes. We apply it to study the rate function of the Langevin dynamics, in particular its dependence on the friction.

We now sketch the main results of this chapter, the precise setting being presented in Section 3.2.1. Consider a diffusion process $(X_t)_{t \geq 0}$ over a state space $\mathcal{X} \subset \mathbb{R}^d$ with generator \mathcal{L} and invariant measure μ , and the empirical measure

$$L_t = \frac{1}{t} \int_0^t \delta_{X_s} ds, \quad (3.1)$$

where δ_x is the Dirac mass at $x \in \mathcal{X}$. Our first contribution is to prove a large deviations principle with respect to μ for the empirical measure $(L_t)_{t \geq 0}$, in a Wasserstein topology associated with an unbounded function $\kappa : \mathcal{X} \rightarrow [1, +\infty)$. That is, we prove the following type of long time scaling: for $\Gamma \subset \mathcal{P}(\mathcal{X})$,

$$\mathbb{P}(L_t \in \Gamma) \asymp e^{-t \inf_{\nu \in \Gamma} I(\nu)}, \quad (3.2)$$

where I is a rate function. Here, $\mathcal{P}(\mathcal{X})$ denotes the set of probability measures on \mathcal{X} , and the above scaling holds for the topology on measures associated with the weak convergence against functions f satisfying

$$\|f\|_{B_\kappa^\infty} = \sup_{x \in \mathcal{X}} \frac{|f(x)|}{\kappa(x)} < +\infty. \quad (3.3)$$

As is standard for a LDP on unbounded state spaces [416, 427], our result relies on the existence of a Lyapunov function $W : \mathcal{X} \rightarrow [1, +\infty)$ twice differentiable and such that

$$\Psi = -\frac{\mathcal{L}W}{W} \quad (3.4)$$

has compact level sets (in words, it goes to infinity at infinity). Contrarily to previous works where this condition implies (3.2) in a topology of bounded test functions [416, 149, 427], we show in Section 3.2 that the LDP holds for the Wasserstein topology associated with any cost function κ controlled by Ψ . Moreover, the associated rate function $I : \mathcal{P}(\mathcal{X}) \rightarrow [0, +\infty]$, also called entropy, reads

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad I(\nu) = \sup_{\|f\|_{B_\kappa^\infty} < +\infty} \{\nu(f) - \lambda(f)\},$$

where

$$\lambda(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E} \left[e^{\int_0^t f(X_s) ds} \right] \quad (3.5)$$

is the cumulant or free energy function.

We mention that our strategy relies on the Gärtner–Ellis theorem, according to which the existence and regularity of (3.5) implies the large deviations principle. We actually show that (3.5) is well-defined because it matches the principal eigenvalue of the Feynman–Kac operator

$$\varphi \mapsto \mathbb{E} \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]. \quad (3.6)$$

A key remark for defining the above operator is that

$$M_t = W(X_t) e^{-\int_0^t \frac{\mathcal{L}W}{W}(X_s) ds} \quad (3.7)$$

is a local martingale, as noted by Wu in [427]. This allows to define (3.6) for functions φ such that $\|\varphi\|_{B_W^\infty} < +\infty$, as soon as f is dominated by the function Ψ defined in (3.4). As a result, for any such f , the operator (3.6) can be shown to be compact over the space of functions controlled by W (see [199, 181]). The functional (3.5) is then obtained as the largest eigenvalue of the operator (3.6) through a generalized Perron–Frobenius theorem (the Krein–Rutman theorem [112]). This is actually an adaption of the strategy developed in Chapter 2 in this large deviations context.

The second part of our work consists in rewriting the large deviations functional I . For this, we revisit [133] by showing that

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad I(\nu) = \sup \left\{ -\int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu, \quad u \in \mathcal{D}^+ \right\}, \quad (3.8)$$

where \mathcal{D}^+ is an appropriate domain defined in Section 3.3, which differs from that of [133] through an additional growth condition. This result leads to a variational formula for the largest eigenvalue $\lambda(f)$ of the operator $\mathcal{L} + f$ defined on a suitable functional space through

$$\lambda(f) = \sup_{\nu \in \mathcal{P}(\mathcal{X})} \{ \nu(f) - I(\nu) \}.$$

We mention that proving (3.8) is easily realized thanks to the spectral problem associated with the Feynman–Kac operator (3.6), relying on the recent work [181].

Finally, the variational representation (3.8) allows to generalize the results of [49] by splitting I into two parts. This is motivated by considerations on nonequilibrium systems, as explained in Section 1.1. More specifically, denoting by $\mathcal{L} = \mathcal{L}_S + \mathcal{L}_A$ the decomposition into symmetric and antisymmetric parts of the generator with respect to the invariant measure μ , we obtain, for any $\nu \ll \mu$:

$$I(\nu) = \frac{1}{4} \left| \log \frac{d\nu}{d\mu} \right|_{\mathcal{H}^1(\nu)}^2 + \frac{1}{4} \left| \mathcal{L}_A \left(\log \frac{d\nu}{d\mu} \right) \right|_{\mathcal{H}^{-1}(\nu)}^2,$$

where $|\cdot|_{\mathcal{H}^1(\nu)}$ and $|\cdot|_{\mathcal{H}^{-1}(\nu)}$ refer to Sobolev seminorms defined in Section 3.2.1. Interestingly, the proof is derived from a generalized Witten transform performed in the variational representation (3.8), which we may therefore call variational Witten transform. This shows that, for a given invariant measure, an irreversible dynamics ($\mathcal{L}_A \neq 0$) produces more entropy than a reversible one, in accordance with the second law of thermodynamics. This decomposition is useful for instance to study the entropy production of the Langevin dynamics, which is irreversible but has a particular structure which allows to naturally separate its reversible and irreversible parts (see Section 1.1.1).

The chapter is organized as follows. Section 3.2 presents a large deviations principle in a Wasserstein topology for the empirical measure of diffusions, under Lyapunov and regularity conditions. Section 3.3 provides a rewriting of the rate functional, and its decomposition into symmetric and antisymmetric parts. Some examples of application are given in Section 3.4, in particular for overdamped and underdamped Langevin dynamics. Section 3.5 discusses possible extensions and connections with related works. Finally, most of the proofs are postponed to Section 3.6.

3.2 Large deviations principle

3.2.1 Setting

This section introduces the main notation used throughout the Chapter, which matches those of the Introduction. We consider a diffusion process $(X_t)_{t \geq 0}$ evolving in $\mathcal{X} \subset \mathbb{R}^d$ with $d \in \mathbb{N}^*$, and satisfying the following stochastic differential equation (SDE):

$$dX_t = b(X_t) dt + \sigma(X_t) dB_t, \tag{3.9}$$

where $b : \mathcal{X} \rightarrow \mathbb{R}^d$, $\sigma : \mathcal{X} \rightarrow \mathbb{R}^{d \times m}$ and $(B_t)_{t \geq 0}$ is a m -dimensional Brownian motion for some $m \in \mathbb{N}^*$. In order to avoid the issue with boundaries of \mathcal{X} , we assume that either $\mathcal{X} = \mathbb{R}^d$, $\mathcal{X} = \mathbb{T}^d$ or $\mathcal{X} = \mathbb{T}^d \times \mathbb{R}^d$, where \mathbb{T}^d is the d -dimensional torus (applications to other subdomains of \mathbb{R}^d can be treated upon appropriate modifications). The last case is motivated by applications to the Langevin equation, where \mathbb{T}^d would be a bounded position space and \mathbb{R}^d the unbounded momentum space (see Section 3.4.2). The generator of the dynamics (3.9), denoted by \mathcal{L} , reads

$$\mathcal{L} = b \cdot \nabla + S : \nabla^2, \quad \text{with } S = \frac{\sigma \sigma^T}{2}, \tag{3.10}$$

where σ^T denotes the transpose of the matrix σ and \cdot is the scalar product on \mathbb{R}^d . Moreover, ∇^2 denotes the Hessian matrix, while for two matrices A, B belonging to $\mathbb{R}^{d \times d}$, we write $A : B = \text{Tr}(A^T B)$. The domain of \mathcal{L} and the conditions on b and σ will be made precise in Section 3.2.2. The function S takes values in the set of symmetric positive matrices (non necessarily definite). We also introduce the *carré du champs* operator [21] associated with \mathcal{L} defined by, for two regular functions φ, ψ :

$$\mathcal{C}(\varphi, \psi) = \frac{1}{2} (\mathcal{L}(\varphi\psi) - \varphi\mathcal{L}\psi - \psi\mathcal{L}\varphi) = \nabla\varphi \cdot S\nabla\psi. \tag{3.11}$$

We will use the space of smooth functions with at most polynomial growth:

$$\mathcal{S} = \left\{ \varphi \in C^\infty(\mathcal{X}) \mid \forall \alpha \in \mathbb{N}^d, \exists N > 0 \text{ such that } \sup_{x \in \mathcal{X}} \frac{|\partial^\alpha \varphi(x)|}{(1 + |x|^2)^N} < +\infty \right\},$$

as well as $C_c^\infty(\mathcal{X})$ (resp. $C_b(\mathcal{X})$) the space of smooth functions with compact support (resp. continuous and bounded).

The space of bounded measurable functions, denoted by $B^\infty(\mathcal{X})$, is endowed with the norm

$$\|\varphi\|_{B^\infty} = \sup_{x \in \mathcal{X}} |\varphi(x)|.$$

Moreover, like in Chapter 2, we will need weighted function spaces and the corresponding probability measure spaces, which commonly appear in Markov chain theory [330, 275, 219]. For any measurable function $W : \mathcal{X} \rightarrow [1, +\infty)$ we define

$$B_W^\infty(\mathcal{X}) = \left\{ \varphi : \mathcal{X} \rightarrow \mathbb{R} \mid \|\varphi\|_{B_W^\infty} = \sup_{x \in \mathcal{X}} \frac{|\varphi(x)|}{W(x)} < +\infty \right\}, \quad (3.12)$$

and the corresponding space of probability measures (see [379, Chapter 2] for duality results on measure spaces):

$$\mathcal{P}_W(\mathcal{X}) = \left\{ \nu \in \mathcal{P}(\mathcal{X}) \mid \nu(W) < +\infty \right\}. \quad (3.13)$$

The associated weighted total variation distance is (see for instance [219]):

$$\forall \nu, \eta \in \mathcal{P}_W(\mathcal{X}), \quad d_W(\nu, \eta) = \sup_{\|\varphi\|_{B_W^\infty} \leq 1} \left\{ \int_{\mathcal{X}} \varphi d\nu - \int_{\mathcal{X}} \varphi d\eta \right\} = \int_{\mathcal{X}} W(x) |\nu - \eta|(dx). \quad (3.14)$$

Note that the spaces (3.12) and (3.13) are defined for an arbitrary measurable function $W \geq 1$. It is possible to weaken the assumption $W \geq 1$ but we do not need these refinements in this work. We denote by τ -topology the weak topology associated with the convergence of functions belonging to $B^\infty(\mathcal{X})$. This means that for a sequence $(\nu_n)_{n \in \mathbb{N}}$ in $\mathcal{P}(\mathcal{X})$, $\nu_n \rightarrow \nu$ in the τ -topology if $\nu_n(\varphi) \rightarrow \nu(\varphi)$ for any $\varphi \in B^\infty(\mathcal{X})$. When considering functions $\varphi \in B_W^\infty(\mathcal{X})$, we denote by τ^W the associated topology [427, 275], see also [124, Lemma 3.3.8] for details. When $W(x) = 1 + |x|^\alpha$ for $\alpha > 0$, this is nothing else than the topology induced by the Wasserstein distance associated with the cost $x \mapsto |x|^\alpha$, see [419, Theorem 7.12]. This is why we sometimes call τ^W a Wasserstein-like topology.

Remark 3.1. *One may wonder why we aim at using such a Wasserstein-like topology to prove the LDP over. Our first motivation was to understand the fluctuations of the empirical average*

$$L_t(f) = \frac{1}{t} \int_0^t f(X_s) ds$$

when f is unbounded. However, if a LDP on \mathbb{R} can be reached for $L_t(f)$, we cannot say much about the rate function in general. This is an important motivation for considering the empirical average (3.1) instead, as Section 3.3 will prove.

Then, the coarsest topology which makes the application

$$L_t \mapsto L_t(f)$$

continuous for the functions $f \in B_W^\infty(\mathcal{X})$ is precisely the τ^W -topology. As a result, although other topologies could be considered and in view of the contraction principle presented in Section 1.2.2, the τ^W -topology appears as the natural framework to consider unbounded functions f while studying large deviations principle for the empirical average (3.1).

We associate to the dynamics $(X_t)_{t \geq 0}$ the semigroup $(P_t)_{t \geq 0}$ defined through

$$\forall \varphi \in B^\infty(\mathcal{X}), \quad (P_t \varphi)(x) = \mathbb{E}_x[\varphi(X_t)], \quad (3.15)$$

where \mathbb{E}_x stands for the expectation with respect to all realizations of the Brownian motion in (3.9), for dynamics starting at $x \in \mathcal{X}$. We say that $\mu \in \mathcal{P}(\mathcal{X})$ is invariant with respect to the dynamics $(X_t)_{t \geq 0}$ if $(\mu P_t)(\varphi) = \mu(\varphi)$ for any $\varphi \in C_b(\mathcal{X})$, with the notation

$$(\mu P_t)(\varphi) = \mu(P_t \varphi) = \int_{\mathcal{X}} \mathbb{E}_x[\varphi(X_t)] \mu(dx).$$

An equivalent condition is that $\mu(\mathcal{L}\varphi) = 0$ for $\varphi \in C_c^\infty(\mathcal{X})$.

We now follow the path of [273, Chapter 2] for defining other useful functional spaces. For any probability measure $\mu \in \mathcal{P}(\mathcal{X})$, let

$$L^2(\mu) = \left\{ \varphi \text{ measurable} \mid \int_{\mathcal{X}} |\varphi|^2 d\mu < +\infty \right\}. \quad (3.16)$$

For $\varphi \in C_c^\infty(\mathcal{X})$, we introduce the seminorm

$$|\varphi|_{\mathcal{H}^1(\mu)}^2 = \int_{\mathcal{X}} \mathcal{C}(\varphi, \varphi) d\mu, \quad (3.17)$$

and the equivalence relation \sim_1 through $\varphi \sim_1 \psi$ if and only if $|\varphi - \psi|_{\mathcal{H}^1(\mu)} = 0$. We denote by $\mathcal{H}^1(\mu)$ the adherence of $C_c^\infty(\mathcal{X})$ quotiented by \sim_1 for the norm $|\cdot|_{\mathcal{H}^1(\mu)}$. Note that $\mathcal{H}^1(\mu)$ and $L^2(\mu)$ are not subspaces of each other in general, but $L^2(\mu) \subset \mathcal{H}^1(\mu)$ if μ satisfies a Poincaré inequality and S is positive definite. The difference between $L^2(\mu)$ and $\mathcal{H}^1(\mu)$ is however important for degenerate dynamics, see the application in Section 3.4.2. We now construct a space dual to $\mathcal{H}^1(\mu)$ with the same density argument by introducing the seminorm, for $\varphi \in C_c^\infty(\mathcal{X})$,

$$|\varphi|_{\mathcal{H}^{-1}(\mu)}^2 = \sup_{\psi \in C_c^\infty(\mathcal{X})} \left\{ 2 \int_{\mathcal{X}} \varphi \psi d\mu - |\psi|_{\mathcal{H}^1(\mu)}^2 \right\}. \quad (3.18)$$

We define the equivalence relation \sim_{-1} on $C_c^\infty(\mathcal{X})$ by $\varphi \sim_{-1} \psi$ if and only if $|\varphi - \psi|_{\mathcal{H}^{-1}(\mu)} = 0$. The dual space $\mathcal{H}^{-1}(\mu)$ is then the adherence of $C_c^\infty(\mathcal{X})$ quotiented by \sim_{-1} for the $\mathcal{H}^{-1}(\mu)$ -norm.

Let us relate $\mathcal{H}^{-1}(\mu)$ to the more standard $H^1(\mu)$ Sobolev space [304]. If μ is invariant with respect to \mathcal{L} , it holds for $\varphi \in C_c^\infty(\mathcal{X})$ (using $\mathcal{L}(\varphi^2) = 2\varphi\mathcal{L}\varphi + 2\mathcal{C}(\varphi, \varphi)$)

$$|\varphi|_{\mathcal{H}^1(\mu)}^2 = 2 \int_{\mathcal{X}} \varphi(-\mathcal{L}\varphi) d\mu.$$

In particular, when $S = I_d$ we have

$$|\varphi|_{\mathcal{H}^1(\mu)}^2 = \int_{\mathcal{X}} |\nabla\varphi|^2 d\mu.$$

In this case, $|\cdot|_{\mathcal{H}^1(\mu)}$ is the standard $H^1(\mu)$ Sobolev seminorm [304]. For precisions on $\mathcal{H}^1(\mu)$ and its relation to the central limit theorem for Markov processes, we refer to [273, Chapter 2].

Remark 3.2. *The space $\mathcal{H}^{-1}(\mu)$ can be thought of as a weaker version of the space $L_0^2(\mu)$ of functions in $L^2(\mu)$ with average zero with respect to μ . Indeed, assume for instance that $\varphi \in L^2(\mu)$ (so $\varphi \in L^1(\mu)$), $\varphi \geq 0$ and $|\varphi|_{\mathcal{H}^{-1}} < +\infty$. We may choose $\psi \in C_c^\infty(\mathcal{X})$ such that*

$$\psi(x) = \begin{cases} 1, & \text{if } |x| \leq 1, \\ 0, & \text{if } |x| \geq 2, \end{cases}$$

and set $\psi_n(x) = n\psi(x/n)$, so $|\psi_n|_{\mathcal{H}^1(\mu)} \leq C$ for some constant $C > 0$ independent of n . The definition (3.18) of the $\mathcal{H}^{-1}(\mu)$ -norm shows that

$$|\varphi|_{\mathcal{H}^{-1}(\mu)} \geq 2n \int_{|x| \leq n} \varphi d\mu - C.$$

Since $|\varphi|_{\mathcal{H}^{-1}(\mu)} < +\infty$, we obtain by letting $n \rightarrow +\infty$ that $\mu(\varphi) = 0$. Since the functions of $\mathcal{H}^{-1}(\mu)$ may not belong to $L^2(\mu)$, this dual space generalizes $L_0^2(\mu)$.

We also introduce some notation concerning the growth of functions. A function $f = \mathcal{X} \rightarrow \mathbb{R}$ is said to have *compact level sets* if for any $M \in \mathbb{R}$, the set

$$\{x \in \mathcal{X} \mid f(x) \leq M\}$$

is compact (with the convention that \emptyset is compact). A function g is said to be negligible with respect to f (denoted by $g \ll f$) if f/g has compact level sets, and g is said to be equivalent to f (denoted by $g \sim f$) if there exist constants $c, c' > 0$ and $R, R' \in \mathbb{R}$ such that

$$\forall x \in \mathcal{X}, \quad c'g(x) - R' \leq f(x) \leq cg(x) + R.$$

Remark 3.3. *The above definitions are useful when the state space \mathcal{X} is unbounded. A sufficient condition for f to have compact level sets in this case is for this function to be lower semicontinuous and to go to infinity when $|x| \rightarrow +\infty$. If \mathcal{X} is bounded, all these criteria are automatically met for smooth functions.*

Finally, we again denote by $\underline{\lim}$ and $\overline{\lim}$ the inferior and superior limits respectively, while for a subset $A \subset \mathcal{Y}$ of a Polish space \mathcal{Y} , $\overset{\circ}{A}$ and \bar{A} denote the interior and closure of A for the chosen topology on \mathcal{Y} . The function $\mathbb{1}_A$ denotes the indicator function of the set A , *i.e.* $\mathbb{1}_A(x) = 1$ if $x \in A$ and $\mathbb{1}_A(x) = 0$ otherwise. For a Banach space E , $\mathcal{B}(E)$ refers to the Banach space of bounded linear operators over E with the usual norm. Some elements of large deviations theory are reminded in Section 2.5.1.

3.2.2 Statement of the main results

The large deviations principle relies on three standard assumptions: hypoellipticity of the generator, irreducibility of the dynamics, and a Lyapunov condition. We start with our hypoellipticity assumption (which could certainly be relaxed for particular applications, see for instance [427]). It will be useful for proving regularization properties of the Feynman–Kac semigroup in Lemma 3.30.

Assumption 3.4 (Hypoellipticity). *The functions b and σ in (3.9) belong to \mathcal{S}^d and $\mathcal{S}^{d \times m}$ respectively, and the generator \mathcal{L} defined in (3.10) satisfies the hypoelliptic Hörmander condition. More precisely, \mathcal{L} can be written as*

$$\mathcal{L} = \sum_{i=1}^d A_i^\dagger A_i + A_0,$$

where $(A_i)_{i=0}^d$ are first order differential operators with coefficients belonging to \mathcal{S} and such that the family

$$\{A_i\}_{i=1}^d \cup \{[A_i, A_j]\}_{i,j=0}^d \cup \{[[A_i, A_j], A_k]\}_{i,j,k=0}^d \cdots$$

spans \mathbb{R}^d at any $x \in \mathcal{X}$ for a finite number of commutators $n_x \in \mathbb{N}$.

This assumption is natural in practical situations, as illustrated in the applications of Section 3.4 covering elliptic and hypoelliptic diffusions, see [244, 161, 362] for details. Note that excluding the operator A_0 from the first family means that, if \mathcal{L} satisfies Assumption 3.4, $\partial_t + \mathcal{L}$ is hypoelliptic and the transition kernel of $(X_t)_{t \geq 0}$ has a smooth density for any $t > 0$. This regularity requirement comes together with a controllability condition (recall that σ takes values in $\mathbb{R}^{d \times m}$).

Assumption 3.5 (Controllability). *For any $x, y \in \mathcal{X}$ and $T > 0$, there exists a C^1 -control $(u_t)_{t \in [0, T]}$ in \mathbb{R}^m such that the path $(\phi_t)_{t \in [0, T]}$ in \mathcal{X} defined as*

$$\begin{cases} \phi_0 = x, \\ \dot{\phi}_t = b(\phi_t) + \sigma(\phi_t) \dot{u}_t, \end{cases} \quad (3.19)$$

is well-defined and satisfies $\phi_T = y$.

Assumption 3.5 implies that the process is irreducible, *i.e.* that the transition density of $(X_t)_{t \geq 0}$ is everywhere positive, which will be used in Lemma 3.31. Note that constructing a control $(u_t)_{t \in [0, T]}$ may be difficult in general [257]. However, for the overdamped and underdamped Langevin dynamics we are interested in, building such a control turns out to be genuinely feasible, see [324, 362, 304, 316] and the references therein. Let us mention that the above two assumptions are not specific to our problem of large deviations [362].

A typical idea for dealing with Markov chain stability and large deviations on an unbounded state space is to reduce the analysis to a compact set and to control the excursions of the dynamics out of this set with a Lyapunov function [330, 427]. Our Witten–Lyapunov condition for the dynamics reads as follows (for the terminology, see Remark 3.11 below).

Assumption 3.6 (Witten–Lyapunov condition). *There exists a function $W : \mathcal{X} \rightarrow [1, +\infty)$ of class $C^2(\mathcal{X})$, with compact level sets, and such that*

$$\Psi = -\frac{\mathcal{L}W}{W} \quad (3.20)$$

has compact level sets. Moreover, there exists a $C^2(\mathcal{X})$ function $\mathcal{W} : \mathcal{X} \rightarrow [1, +\infty)$ with $\mathcal{W} \ll W$ and such that, for some constants $C_1 > 0$, $C_2 \in \mathbb{R}$,

$$\mathcal{W}^2 \leq C_1 W, \quad \Psi \sim -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}, \quad -2\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}} \leq -\frac{\mathcal{L}W}{W} + C_2. \quad (3.21)$$

In all what follows, we consider an arbitrary continuous function $\kappa : \mathcal{X} \rightarrow [1; +\infty)$ such that:

- $\kappa \ll \Psi$;
- κ is either bounded or has compact level sets.

Remark 3.7. Note that, since $\kappa \ll \Psi$ and $\Psi \sim -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}$, it holds $\kappa \ll -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}$. This fact will be frequently used in the proofs. Moreover the conditions (3.21) are not restrictive for exponential-like Lyapunov function as shown in Proposition 3.8 below – the idea being that \mathcal{W} can be set to \sqrt{W} . In practice, the auxiliary function \mathcal{W} is used to obtain some control in the proofs of Lemmas 3.28 and 3.30 (in particular to apply a Grönwall lemma), and it could certainly be phrased differently. The continuity condition on κ may be relaxed for instance by assuming this function to be lower semicontinuous and bounded on compact sets.

Although we stated Assumption 3.6 in order to fit standard conditions when considering large deviations on unbounded state spaces [416, 427], in practice it can be obtained from a non-linear Lyapunov condition in the spirit of [275] and [149, Condition 2.2]. This is the purpose of the next proposition, whose proof is postponed to Section 3.6.3.

Proposition 3.8. Assume that there exists $V \in \mathcal{S}$ such that:

- V has compact level sets;
- $|\sigma^T \nabla V|$ has compact level sets;
- for any $\theta \in (0, 1)$ it holds

$$-\mathcal{L}V - \frac{\theta}{2} |\sigma^T \nabla V|^2 \sim |\sigma^T \nabla V|^2. \quad (3.22)$$

Then Assumption 3.6 is satisfied with

$$W(x) = e^{\theta V(x)}, \quad \mathcal{W}(x) = e^{\varepsilon V(x)},$$

for $\theta \in (0, 1)$ and $\varepsilon < \theta/2$ small enough. Moreover it holds

$$\Psi \sim |\sigma^T \nabla V|^2.$$

Note that (3.22) means that the term $-\mathcal{L}V$ coming from the dynamics must compensate the quadratic loss proportional to $|\sigma^T \nabla V|^2$. A first consequence of Assumptions 3.4-3.6 is the ergodicity of the dynamics.

Proposition 3.9. Under Assumptions 3.4, 3.5 and 3.6, (3.9) has a global strong solution, and the process $(X_t)_{t \geq 0}$ admits a unique invariant measure $\mu \in \mathcal{P}_W(\mathcal{X})$. This measure has a positive $C^\infty(\mathcal{X})$ -density with respect to the Lebesgue measure: there exists $\rho^\mu \in C^\infty(\mathcal{X})$ with $\rho^\mu > 0$ such that $\mu(dx) = \rho^\mu(x) dx$. Moreover, the dynamics is ergodic with respect to μ : there is $C, c > 0$ such that

$$\forall t \geq 0, \quad \forall \varphi \in B_W^\infty(\mathcal{X}), \quad \|P_t \varphi - \mu(\varphi)\|_{B_W^\infty} \leq C e^{-ct} \|\varphi - \mu(\varphi)\|_{B_W^\infty}.$$

Equivalently,

$$\forall t \geq 0, \quad \forall \nu \in \mathcal{P}_W(\mathcal{X}), \quad d_W(\nu P_t, \mu) \leq C e^{-ct} d_W(\nu, \mu).$$

Proof. The existence of a unique local strong solution is standard when Assumption 3.4 holds, see [361, Chapter IX, Exercise (2.10)]. Assumption 3.6 then implies the existence of $a > 0$, $b \in \mathbb{R}$ such that

$$\mathcal{L}W \leq -aW + b,$$

and global existence can be deduced from the above Lyapunov inequality [362]. The end of the proof is a direct application of [362, Theorem 8.9] since Assumption 3.5 together with Assumption 3.4 ensures irreducibility. \square

We can now present the large deviations principle associated with the empirical measure of the process $(X_t)_{t \geq 0}$ with respect to its invariant measure μ . Recall that the empirical measure of the process is defined by

$$L_t = \frac{1}{t} \int_0^t \delta_{X_s} ds, \quad (3.23)$$

where δ_x denotes the Dirac mass at $x \in \mathcal{X}$. When one considers large deviations principles for empirical averages of the form (3.23), the topology on probability measures has to be specified. As mentioned in the introduction, most of the proofs of LDPs consider topologies associated with bounded measurable functions (resp. continuous bounded), the so-called strong topology or τ -topology (resp. weak topology). We now prove that, in our setting, a LDP holds in the τ^κ -topology defined in Section 3.2.1, for any function κ satisfying Assumption 3.6. The proof of Theorem 3.10 is presented in Section 3.6.1.

Theorem 3.10. *Suppose that Assumptions 3.4, 3.5 and 3.6 hold true, and consider a function κ as in Assumption 3.6. Then,*

$$L_t \xrightarrow[t \rightarrow +\infty]{} \mu, \quad (3.24)$$

almost surely in the τ^κ -topology. Moreover, the functional

$$f \in B_\kappa^\infty(\mathcal{X}) \mapsto \lambda(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right] \quad (3.25)$$

is well-defined, convex and finite, and $(L_t)_{t \geq 0}$ satisfies a LDP in the τ^κ -topology with the good rate function defined by:

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad I(\nu) = \begin{cases} \sup_{f \in B_\kappa^\infty} \{ \nu(f) - \lambda(f) \}, & \text{if } \nu \in \mathcal{P}_\kappa(\mathcal{X}) \text{ and } \nu \ll \mu, \\ +\infty, & \text{otherwise.} \end{cases} \quad (3.26)$$

More precisely, for any τ^κ -measurable set $\Gamma \subset \mathcal{P}(\mathcal{X})$, it holds

$$-\inf_{\nu \in \Gamma} I(\nu) \leq \liminf_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}_x(L_t \in \Gamma) \leq \overline{\lim}_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}_x(L_t \in \Gamma) \leq -\inf_{\nu \in \Gamma} I(\nu), \quad (3.27)$$

where the interior and closure of Γ are taken with respect to the τ^κ -topology. Finally, for any $\nu \in \mathcal{P}(\mathcal{X})$, it holds $I(\nu) = 0$ if and only if $\nu = \mu$.

Our conclusion is in essence close to that of [275], but the conditions to reach it seem more natural to us and correspond to usual conditions for proving large deviations principles in an unbounded state space, see [427, 149] and [416, Section 9]. In particular, they allow to derive the duality representation (3.26), and we do not need to consider non-linear operators. Our strategy (presented in Section 3.6.1) relies on the Gärtner–Ellis theorem [199, 165, 166, 119], for which the existence of the free-energy (3.25) is a key element (see Section 1.2.2 in Part I). The originality of our work is to make use of the local martingale (3.7) introduced by Wu [427] in order to solve the spectral problem associated with the Feynman–Kac operator $\mathcal{L} + f$, which proves the existence of (3.25). This directly provides the LDP in the τ^κ -topology by duality. However, there may be cases in which a LDP holds although the conditions of the Gärtner–Ellis theorem are not satisfied, for instance in the framework of the Sanov theorem [424], so our condition may not be necessary.

Another advantage of our approach is to characterize precisely the set of functions for which a LDP holds from the standard condition on Ψ defined in (3.20), like in [131, 416]. This condition is also used in [427, Corollary 2.3] for proving a level 1 LDP for the Langevin dynamics. We present below a clear connection with a spectral gap condition for the Witten–Schrödinger operator in the reversible case. The comparison with Cramér’s condition for independent variables highlights the effect of correlations on fluctuations.

Remark 3.11 (Reversible processes, Witten Laplacian and Cramér’s condition). *Consider the following reversible diffusion*

$$dX_t = -\nabla V(X_t) dt + \sqrt{2} dB_t,$$

where $V : \mathcal{X} \rightarrow \mathbb{R}$ is a smooth potential with compact level sets. The generator of this dynamics is $\mathcal{L} = -\nabla V \cdot \nabla + \Delta$ and its invariant measure reads $\mu(dx) = Z^{-1} e^{-V(x)} dx$. Define

$$W_\theta(x) = e^{\theta V(x)},$$

for some $\theta \in (0, 1)$. This is a standard choice for obtaining compactness of the evolution operator [362, Section 8], and optimal control representations of rate functions [149], see also Proposition 3.8. An easy computation shows that

$$\Psi_\theta = -\frac{\mathcal{L}W_\theta}{W_\theta} = \theta(1 - \theta)|\nabla V|^2 - \theta\Delta V. \quad (3.28)$$

However, we also know [425] that the generator \mathcal{L} considered on $L^2(\mu)$ is unitarily equivalent to the operator

$$\tilde{\mathcal{L}} = e^{-\frac{V}{2}} \mathcal{L}(e^{\frac{V}{2}} \cdot),$$

defined on $L^2(dx)$ (a procedure also called symmetrization [404, Section 4.3]), which is actually the opposite of the Witten Laplacian [425, 234]:

$$\tilde{\mathcal{L}} = \Delta - \frac{1}{4}|\nabla V|^2 + \frac{1}{2}\Delta V = -\left(-\Delta + \Psi_{\frac{1}{2}}\right). \quad (3.29)$$

In this case, the condition for (3.28) to have compact level sets when $\theta = 1/2$ is actually equivalent to a confinement condition (or spectral gap condition [235]) for the Witten–Schrödinger operator $\tilde{\mathcal{L}}$ defined in (3.29). In that sense, Assumption 3.6 is a natural generalization of a spectral gap condition for the Witten Laplacian in the case of possibly non-reversible dynamics. This is why we call Assumption 3.6 a Witten–Lyapunov condition.

We now compare this Witten–Lyapunov condition to Cramér’s exponential moment condition in the case of independent variables of law μ . Consider the case when $V(x) = |x|^q$ for $q > 1$, for which Assumption 3.6 is satisfied by application of Proposition 3.8. The standard Cramér’s condition in the case of independent variables states that the empirical measure

$$\frac{1}{n} \sum_{i=1}^n \delta_{X_i}$$

satisfies a large deviations principle in the τ^κ -topology if and only if (see [424, Theorem 1.1] recalled in Section 1.2.2):

$$\forall \theta \in \mathbb{R}, \quad \int_{\mathcal{X}} e^{\theta \kappa} d\mu < +\infty.$$

If $\mu(dx) \propto e^{-|x|^q} dx$, a sufficient condition for the above condition to hold is to choose $\kappa(x) = 1 + |x|^\alpha$ with $0 \leq \alpha < q$. On the other hand, the Witten–Lyapunov potential (3.28) reads in this case

$$\Psi_\theta(x) = \theta(1 - \theta)q^2|x|^{2(q-1)} - \theta q(q-1)|x|^{q-2},$$

so that we may choose $\kappa(x) = 1 + |x|^\alpha$ for $0 \leq \alpha < 2(q-1)$. When comparing the two conditions, we obtain the following different situations depending on q :

- $q > 2$ (super-Gaussian case): $2(q-1) > q$, the Witten–Lyapunov condition is less restrictive than Cramér’s condition;
- $q = 2$ (Gaussian case): $2(q-1) = q$, the two conditions are equivalent;
- $q \in (1, 2)$ (sub-Gaussian case): $2(q-1) < q$, the Witten–Lyapunov condition is more restrictive than Cramér’s condition.

This simple example shows that considering a correlated system instead of independent variables has a non-trivial effect on the stability of the system. Depending on the confinement potential, the Witten–Lyapunov condition for (3.28) to have compact level sets can be more or less restrictive than Cramér’s condition for independent variables distributed according to the invariant measure μ . Finally, we remark that for $q \in (1, 3/2)$, the process is heavy-tailed in the sense that $2(q-1) < 1$ and the observable $f(x) = x$ (assuming $d = 1$) does not satisfy a LDP. In other words, the average position of the process defined by

$$\frac{1}{t} \int_0^t X_s ds$$

a priori does not satisfy a large deviations principle at speed t .

We finally mention that, in the case where the observable f grows faster than the potential Ψ , it seems possible to derive a level 1 large deviations principle at a speed smaller than t . We refer to [342] for a recent account dealing with the case of an Ornstein–Uhlenbeck process, and to [54, 17] for related issues.

We close this section with a practical corollary of Theorem 3.10 which generalizes the level 1 LDP proved in [427, Corollary 2.3]. It is a classical application of the contraction principle presented in Section 1.2.2.

Corollary 3.12 (Level 1 large deviations principle). *Suppose that Assumptions 3.4, 3.5 and 3.6 hold true and consider a function $f \in B_\kappa^\infty(\mathcal{X})$. Then, the function*

$$\theta \in \mathbb{R} \mapsto \lambda_f(\theta) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\theta \int_0^t f(X_s) ds} \right] \quad (3.30)$$

is well-defined and differentiable. Moreover,

$$L_t(f) = \frac{1}{t} \int_0^t f(X_s) ds \xrightarrow{t \rightarrow +\infty} \int_{\mathcal{X}} f d\mu$$

almost surely, and $L_t(f)$ satisfies a large deviations principle in \mathbb{R} at speed t with good rate function given by

$$\forall a \in \mathbb{R}, \quad I_f(a) = \inf \{ I(\nu), \nu \in \mathcal{P}(\mathcal{X}) \mid \nu(f) = a \}, \quad (3.31)$$

where I is defined in (3.26). Moreover, it holds

$$I_f(a) = \sup_{\theta \in \mathbb{R}} \{ \theta a - \lambda_f(\theta) \}. \quad (3.32)$$

As already noted in Section 1.3.1, Corollary 3.12 is useful for practical applications, since (3.32) is a natural way to estimate the rate function I_f associated with an observable f , see also [202, 377, 397, 97, 184] and the numerical method proposed in Chapter 5.

Proof. For $f \in B_\kappa^\infty(\mathcal{X})$, the application $L_t \in \mathcal{P}_\kappa(\mathcal{X}) \mapsto L_t(f) \in \mathbb{R}$ is continuous in the τ^κ -topology [124, Lemma 3.3.8]. Therefore, $L_t(f)$ obeys a large deviations principle in \mathbb{R} by the contraction principle [119, Theorem 4.2.1], with good rate function given by (3.31). Moreover, one can redo the proofs leading to Theorem 3.10 and show that λ_f defined in (3.30) is smooth and well-defined on \mathbb{R} . This implies that a LDP with good rate function (3.32) holds through the Gärtner–Ellis theorem applied in \mathbb{R} . Since the rate function is unique, the expressions (3.31) and (3.32) coincide. \square

3.3 Decomposition of the rate function

Our goal in this section is to rewrite I in various ways, which is useful for theoretical understanding and practical purposes. In Section 3.3.1, we first show an extension of the standard Donsker–Varadhan formulation for I . This result is easily obtained by making use of the spectral analysis of the operator $\mathcal{L} + f$ for $f \in B_\kappa^\infty(\mathcal{X})$, which is presented in Section 3.6.1. We then apply this result to obtain a variational representation for the principal eigenvalue $\lambda(f)$ of the operator $\mathcal{L} + f$. Next, in Section 3.3.2, we split the expression of the rate function into symmetric and antisymmetric parts of the dynamics, extending the work [49] to general diffusions. Such a decomposition will prove useful in Section 3.4 to compare the entropy of overdamped and underdamped Langevin dynamics. Most of the proofs of this section are postponed to Section 3.6.2.

3.3.1 Donsker–Varadhan variational formula

We start with the variational representation of the entropy. Our proof, which can be found in Section 3.6.2.1, is an adaption of [124, Lemma 4.2.35] relying on the Feynman–Kac operator and its spectral elements.

Proposition 3.13. *The rate function defined in (3.26) admits the following representation:*

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad I(\nu) = \sup \left\{ - \int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu, u \in \mathcal{D}^+ \right\}, \quad (3.33)$$

where

$$\mathcal{D}^+ = \left\{ u \in B_W^\infty(\mathcal{X}) \mid u > 0, -\frac{\mathcal{L}u}{u} \in B_\kappa^\infty(\mathcal{X}) \right\}. \quad (3.34)$$

In particular, the functional defined in (3.33) is equal to $+\infty$ if $\nu \notin \mathcal{P}_\kappa(\mathcal{X})$ or ν is not absolutely continuous with respect to μ .

This result is standard when \mathcal{X} is compact [133], but does not seem to be known for an unbounded space \mathcal{X} and for the τ^κ -topology we consider. In this situation the space \mathcal{D}^+ has to be designed with caution (in particular, \mathcal{D}^+ is not empty since it contains the functions of the form $u = e^\psi$ for $\psi \in C_c^\infty(\mathcal{X})$). Note also that the last statement of Proposition 3.13 is consistent with the Fenchel definition (3.26) of the rate function. In order to get some intuition on the formula (3.33), let us mention that the proof relies on replacing the maximum over functions $u \in \mathcal{D}^+$ by the supremum over eigenfunctions h_f satisfying

$$(\mathcal{L} + f)h_f = \lambda(f)h_f,$$

for $f \in B_\kappa^\infty(\mathcal{X})$. The above equation rewrites, since $h_f > 0$,

$$-\frac{\mathcal{L}h_f}{h_f} = f - \lambda(f).$$

By integrating with respect to a measure $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ we find (3.33) on the left hand side, and the Fenchel transform (3.26) on the right hand side. The functional spaces associated with f and h_f motivate the choice of \mathcal{D}^+ .

A natural consequence of Proposition 3.13 is the following variational representation for the principal eigenvalue. The proof, postponed to Section 3.6.2.2, relies on the convexity of the cumulant function to invert the Fenchel transform (3.26).

Corollary 3.14. *For $f \in B_\kappa^\infty(\mathcal{X})$, the principal eigenvalue $\lambda(f)$ associated with the operator $\mathcal{L} + f$ over $B_W^\infty(\mathcal{X})$ is isolated and admits the following representation:*

$$\lambda(f) = \sup_{\nu \in \mathcal{P}_\kappa} \{ \nu(f) - I(\nu) \}, \tag{3.35}$$

where I is defined in (3.33).

Remark 3.15. *Corollary 3.14 may seem anecdotal, but it provides a variational representation for the principal eigenvalue of non-symmetric diffusion operators, as pioneered by Donsker and Varadhan in their seminal paper [133] for a compact space \mathcal{X} . To the best of our knowledge, this formula had not been shown in an unbounded setting, for which we need to introduce the «generalized domain» \mathcal{D}^+ defined in (3.34). However, our set of assumptions implies that the largest eigenvalue $\lambda(f)$ is isolated for any f (because of the compactness of the resolvent provided by Lemma 3.32), whereas in [133], (3.35) may be the supremum of the essential spectrum of the operator. This suggests that (3.35) holds under weaker assumptions. A possible approach for such a generalization may be to consider different methods for studying the long time behaviour of unnormalized semigroups, see for instance [89, 24, 91] or to resort to more subtle spectral analysis tools [426, 428, 195, 43].*

3.3.2 Entropy decomposition: symmetry and antisymmetry

Our goal is now to provide refined expressions for the rate function I in terms of symmetric and antisymmetric parts of the dynamics, inspired in particular by [49]. In the following, for any closed operator T , we denote by T^* its adjoint on $L^2(\mu)$, where μ is the invariant measure of the process, as obtained in Proposition 3.9. Considering the generator \mathcal{L} of the diffusion (3.9), we can always decompose it into symmetric and antisymmetric parts with respect to μ through

$$\mathcal{L} = \mathcal{L}_S + \mathcal{L}_A, \quad \mathcal{L}_S = \frac{\mathcal{L} + \mathcal{L}^*}{2}, \quad \mathcal{L}_A = \frac{\mathcal{L} - \mathcal{L}^*}{2}. \tag{3.36}$$

It is important to note that \mathcal{L}_A is a first order differential operator (and therefore obeys the chain rule of first order differentiation). The decomposition (3.36) then allows to separate the rate function (3.33) into two parts. This is the purpose of the next key result, whose proof can be read in Section 3.6.2.3. It is inspired by the computations in [49, Proposition 2], which we simplify and generalize here through a variational Witten transform and the use of the Sobolev spaces introduced in Section 3.2.1. The algebra of the proof also suggests to consider $I(\nu)$ for probability measures ν of the form $d\nu = e^v d\mu$.

Theorem 3.16. *Suppose that Assumptions 3.4, 3.5 and 3.6 hold true, consider a measure $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ such that $\nu \ll \mu$ and $d\nu = e^v d\mu$ with $v \in \mathcal{H}^1(\nu)$ and $\mathcal{L}_A v \in \mathcal{H}^{-1}(\nu)$. Then, the rate function I defined in (3.33) admits the following decomposition:*

$$I(\nu) = I_S(\nu) + I_A(\nu), \tag{3.37}$$

where

$$I_S(\nu) = \frac{1}{4} |v|_{\mathcal{H}^1(\nu)}^2 \quad (3.38)$$

and

$$I_A(\nu) = \frac{1}{4} |\mathcal{L}_A v|_{\mathcal{H}^{-1}(\nu)}^2. \quad (3.39)$$

Theorem 3.16 expresses the rate function as the sum of dual norms of the symmetric and anti-symmetric parts of the dynamics. Note also that we consider a measure of the form $d\nu = e^v d\mu$, that is the Radon–Nikodym derivative of ν with respect to μ is positive. However, we believe that we can consider more general measures ν , see Remark 3.36 in the proof. Since the measure ν at hand appears both inside the norms and in the definition of the norms themselves, a possibly clearer rewriting is the following:

$$I(\nu) = \frac{1}{4} \left| \log \frac{d\nu}{d\mu} \right|_{\mathcal{H}^1(\nu)}^2 + \frac{1}{4} \left| \mathcal{L}_A \left(\log \frac{d\nu}{d\mu} \right) \right|_{\mathcal{H}^{-1}(\nu)}^2.$$

Moreover, the symmetric part of the rate function (3.38) can be written as a Fisher information for the invariant measure μ , a standard result [199]: denoting by $\rho = d\nu/d\mu$, it holds

$$I_S(\nu) = \frac{1}{4} \int_{\mathcal{X}} \frac{\nabla \rho \cdot S \nabla \rho}{\rho} d\mu.$$

The next corollary builds upon (3.39) by rewriting I_A using a Poisson equation, which can be manipulated more easily. The proof can be found in Section 3.6.2.4.

Corollary 3.17. *Suppose that Assumptions 3.4, 3.5 and 3.6 hold true, and consider κ as in Assumption 3.6. Consider a measure $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ such that $d\nu = e^v d\mu$ with $v \in \mathcal{H}^1(\nu)$ and $\mathcal{L}_A v \in \mathcal{H}^{-1}(\nu)$. Then, the antisymmetric part of the rate function (3.39) reads*

$$I_A(\nu) = \frac{1}{4} \int_{\mathcal{X}} \mathcal{C}(\psi_v, \psi_v) d\nu, \quad (3.40)$$

where ψ_v is the unique solution in $\mathcal{H}^1(\nu)$ to the Poisson equation

$$\tilde{\nabla}(S \nabla \psi_v) = \mathcal{L}_A v, \quad (3.41)$$

the symmetric matrix S being defined in (3.10) and $\tilde{\nabla}$ denoting the adjoint of the gradient operator in $L^2(\nu)$.

It has been known for a long time [133] that the rate function of a reversible process is a Fisher information as in (3.38). The antisymmetric part of the rate function has been less investigated, although an expression like (3.40) already appears in [199] (see also [363, 49]). However, our setting provides natural well-posedness conditions for both parts of the rate function to be finite. Interestingly, the solution ψ_v of (3.41) can be formally represented through [304]

$$\psi_v = \int_0^{+\infty} e^{t\mathcal{L}_\nu} (\mathcal{L}_A v) dt,$$

where $\mathcal{L}_\nu = -\tilde{\nabla}(S \nabla \cdot)$. In particular, the uniqueness of ψ_v is a consequence of the definition of $\mathcal{H}^1(\nu)$ through equivalence classes, see Section 3.2.1. Then, the stochastic process $(X_t^\nu)_{t \geq 0}$ associated with \mathcal{L}_ν is reversible with respect to ν . Denoting by e^{-V_ν} the density of ν with respect to the Lebesgue measure, $(X_t^\nu)_{t \geq 0}$ is solution to the following SDE:

$$dX_t^\nu = -S \nabla V_\nu(X_t^\nu) dt + \nabla \cdot S(X_t^\nu) dt + \sigma(X_t^\nu) dB_t.$$

Finally (3.40) takes the form

$$I_A(\nu) = \frac{1}{4} \int_0^{+\infty} \mathbb{E}_\nu \left[(\mathcal{L}_A v)(X_0^\nu) (\mathcal{L}_A v)(X_t^\nu) \right] dt. \quad (3.42)$$

The antisymmetric part of the entropy is then the autocorrelation of $\mathcal{L}_A v$ along a reversible process that realizes the fluctuation corresponding to the measure ν , with some optimality in the sense that it minimizes the non-reversible part (3.39) of the entropy. From a mathematical point of view, it seems

interesting to relate (3.42) to the so-called level 2.5 of large deviations [25, 98], since this approach consists in considering joint fluctuations of the empirical measure and the associated empirical current. In this case, the large deviations function is explicit: this reflects the fact that a Markov process is characterized *entirely* by its density and current. Exploring further the connection between (3.42) and level 2.5 large deviations is an interesting direction for future works.

Remark 3.18. *It is also possible to consider the adjoint \mathcal{L}^* not with respect to the invariant measure μ , but instead a reference measure μ_{ref} such that*

$$\mathcal{L}^* = \mathcal{L}_S - \mathcal{L}_A + \xi,$$

for a measurable function ξ . The measure μ_{ref} may be an equilibrium measure for systems subject to a small external nonequilibrium forcing. In [49], the technique is used to study atom chains in contact with an inhomogeneous heat bath, μ_{ref} being the Gibbs measure associated with a fixed temperature profile. This leads to an additional term $-\int_{\mathcal{X}} \xi d\nu$ in the expression of the rate function (3.37), as can be readily checked by a straightforward adaptation of the proof.

3.4 Applications

In order to illustrate the interest of the results presented above, we come back to the two dynamics introduced in Section 1.1.1: the overdamped and underdamped Langevin dynamics.

3.4.1 Overdamped Langevin dynamics

In this section, we come back to the setting of Remark 3.11 by considering a diffusion process over $\mathcal{X} = \mathbb{R}^d$ subject to

$$dX_t = b(X_t) dt + \sqrt{2} dB_t, \tag{3.43}$$

where $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is a smooth function and $(B_t)_{t \geq 0}$ is a d -dimensional Brownian motion. This corresponds to (3.9) with $\sigma = \sqrt{2}$, in which case the generator reads

$$\mathcal{L} = b \cdot \nabla + \Delta.$$

We will treat the reversible case where $b = -\nabla V$ for a smooth potential V , and $b = -\nabla V + F$ for a smooth function F such that $\nabla \cdot (Fe^{-V}) = 0$. In both cases, the invariant probability measure μ of the process is (assuming $e^{-V} \in L^1(\mathcal{X})$)

$$\mu(dx) = Z^{-1} e^{-V(x)} dx, \quad Z = \int_{\mathcal{X}} e^{-V} < +\infty. \tag{3.44}$$

When $b = -\nabla V$, the dynamics (3.43) is reversible (*i.e.* $\mathcal{L}^* = \mathcal{L}$, where \mathcal{L}^* denotes the adjoint of \mathcal{L} in $L^2(\mu)$). We now give a standard condition on V under which the framework developed in Sections 3.2 and 3.3 applies.

Assumption 3.19. *The potential $V \in \mathcal{S}$ has compact level sets, satisfies $e^{-V} \in L^1(\mathcal{X})$ and, for any $\theta \in (0, 1)$, it holds*

$$(1 - \theta) |\nabla V|^2 - \Delta V \xrightarrow{|x| \rightarrow +\infty} +\infty. \tag{3.45}$$

This assumption is satisfied for smooth potentials growing like $|x|^q$ for $q > 1$ at infinity, and it also implies that the invariant probability measure μ satisfies a Poincaré inequality [20]. Similar conditions are derived in [275] in the context of large deviations. The next proposition is a direct application of Propositions 3.8 and 3.9, Theorem 3.10 and Corollary 3.17 (recall the definition (3.14) of the weighted total variation distance d_W).

Proposition 3.20. *Under Assumption 3.19, the process (3.43) with $b = -\nabla V$ admits the function*

$$W(x) = e^{\theta V(x)}$$

for any $\theta \in (0, 1)$ as a Lyapunov function in the sense of Assumption 3.6. For any fixed $\theta \in (0, 1)$, there exist $C, c > 0$ such that for any initial measure $\nu \in \mathcal{P}_W(\mathcal{X})$,

$$d_W(\nu P_t, \mu) \leq C e^{-ct} d_W(\nu, \mu).$$

Moreover, the function

$$\Psi = -\frac{\mathcal{L}W}{W} = \theta((1-\theta)|\nabla V|^2 - \Delta V) \quad (3.46)$$

has compact level sets and, for any $\kappa : \mathcal{X} \rightarrow [1, +\infty)$ bounded or with compact level sets and such that

$$\frac{\Psi(x)}{\kappa(x)} \xrightarrow{|x| \rightarrow +\infty} +\infty,$$

the empirical measure

$$L_t = \frac{1}{t} \int_0^t \delta_{X_s} ds$$

satisfies a large deviation principle in the τ^κ -topology. The good rate function is defined by: for all $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ with $d\nu = e^v d\mu = \rho d\mu$,

$$I(\nu) = \frac{1}{4} \int_{\mathcal{X}} |\nabla v|^2 d\nu = \frac{1}{4} \int_{\mathcal{X}} \frac{|\nabla \rho|^2}{\rho} d\mu, \quad (3.47)$$

and $I(\nu) = +\infty$ otherwise.

In this reversible example, we see that the rate function is only defined through its symmetric part (3.38), as shown in Theorem 3.16. We now consider a modification of this dynamics when a divergence-free drift is added. The next proposition is an extension of the examples proposed in [363] to the unbounded state space case.

Proposition 3.21. *Suppose that Assumption 3.19 holds and consider the diffusion process:*

$$dX_t = (-\nabla V(X_t) + F(X_t))dt + \sqrt{2} dB_t,$$

with F a smooth vector field such that $\nabla \cdot (Fe^{-V}) = 0$ and

$$\frac{F \cdot \nabla V}{\Psi} \xrightarrow{|x| \rightarrow +\infty} 0,$$

where Ψ is defined in (3.46). Then $\mathcal{L}_S = -\nabla V \cdot \nabla + \Delta$ and $\mathcal{L}_A = F \cdot \nabla$. Moreover

$$\Psi_F = -\frac{(\mathcal{L} + F \cdot \nabla)W}{W} = \theta((1-\theta)|\nabla V|^2 - \Delta V - F \cdot \nabla V) \sim \Psi, \quad (3.48)$$

and $(X_t)_{t \geq 0}$ satisfies an LDP in the τ^κ -topology for any function κ bounded, or with compact level sets and such that

$$\frac{\Psi(x)}{\kappa(x)} \xrightarrow{|x| \rightarrow +\infty} +\infty.$$

The associated rate function I_F reads: for any ν with $d\nu = e^v d\mu$ and $v \in \mathcal{H}^1(\nu)$ together with $F \cdot \nabla v \in \mathcal{H}^{-1}(\nu)$,

$$I_F(\nu) = \frac{1}{4} \int_{\mathcal{X}} |\nabla v|^2 d\nu + \frac{1}{4} \int_{\mathcal{X}} |\nabla \psi_v|^2 d\nu,$$

where ψ_v is the unique $\mathcal{H}^1(\nu)$ solution to

$$-\Delta \psi_v + \nabla(V - v) \cdot \nabla \psi = F \cdot \nabla v.$$

Proposition 3.21 shows that in this simple case, the equilibrium and nonequilibrium dynamics admit a LDP for the same class of functions but with different rate functions, the irreversible dynamics producing more entropy. It is therefore an extension of the case treated in [363, Theorem 2.2]. As for this result, Proposition 3.21 can be used to design algorithms with accelerated convergence to equilibrium, see also [248, 249, 146]. A setting in which Proposition 3.21 typically applies is when $V(x) = |x|^q$ for some $q > 1$ and $F = A\nabla V$ with $A \in \mathbb{R}^{d \times d}$ such that $A^T = -A$ (see [363]).

3.4.2 Underdamped Langevin dynamics

We now apply our framework to the underdamped Langevin dynamics. A first nice feature of our results is that, compared to [427], we obtain a stronger result with similar assumptions – that is our LDP for the empirical measure holds for a finer topology than the one associated with bounded measurable functions. Note however that [427, Corollary 2.3] obtains results similar to ours for a contraction of the rate function (level 1 LDP). In addition, Theorem 3.16 and Corollary 3.17 allow to obtain fine results on the dependency of the rate function on the friction parameter γ .

We start describing the Langevin equation in Section 3.4.2.1, before stating the large deviations principle in Section 3.4.2.2. Finally Section 3.4.2.3 provides asymptotics on the rate function depending on the friction.

3.4.2.1 Description of the dynamics

The dynamics is set on $\mathcal{X} = \mathbb{R}^d \times \mathbb{R}^d$, with $(X_t)_{t \geq 0} = (q_t, p_t)_{t \geq 0} \in \mathbb{R}^d \times \mathbb{R}^d$ evolving as

$$\begin{cases} dq_t = p_t dt, \\ dp_t = -\nabla V(q_t) dt - \gamma p_t dt + \sqrt{2\gamma} dB_t, \end{cases} \quad (3.49)$$

where $\gamma > 0$ is a friction parameter, $V : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth potential, and $(B_t)_{t \geq 0}$ is a d -dimensional Brownian motion. We could also consider the easier case where the position space is bounded ($q \in \mathbb{T}^d$) but leave this easy modification to the reader. The generator of the dynamics is

$$\mathcal{L}_\gamma = \mathcal{L}_{\text{ham}} + \gamma \mathcal{L}_{\text{FD}}, \quad (3.50)$$

where

$$\mathcal{L}_{\text{ham}} = p \cdot \nabla_q - \nabla V \cdot \nabla_p, \quad \mathcal{L}_{\text{FD}} = -p \cdot \nabla_p + \Delta_p.$$

The operator \mathcal{L}_γ leaves invariant the measure

$$\mu(dx) = \mu(dq dp) = \bar{\mu}(dq)\omega(dp), \quad \bar{\mu}(dq) = Z_q^{-1} e^{-V(q)} dq, \quad \omega(dp) = (2\pi)^{-d/2} e^{-\frac{p^2}{2}} dp. \quad (3.51)$$

The invariant measure (3.51) can also be written

$$\mu(dq dp) = Z^{-1} e^{-H(q,p)} dq dp, \quad (3.52)$$

where

$$H(q, p) = V(q) + \frac{p^2}{2} \quad (3.53)$$

is the Hamiltonian of the system, and we assume that the normalization constant Z in (3.52) is finite. In (3.50), the Liouville operator \mathcal{L}_{ham} corresponding to the Hamiltonian part of the dynamics is antisymmetric in $L^2(\mu)$. On the other hand, the fluctuation-dissipation part with generator \mathcal{L}_{FD} is symmetric in $L^2(\mu)$, so that $\mathcal{L}_A = \mathcal{L}_{\text{ham}}$ and $\mathcal{L}_S = \gamma \mathcal{L}_{\text{FD}}$ with the notation of Section 3.3.2.

Before turning to the LDP associated with the Langevin dynamics (3.49), we give some intuition on the behaviour of the process as γ varies. First, it is clear that in the small γ limit, (3.49) becomes the Hamiltonian dynamics

$$\begin{cases} dq_t = p_t dt, \\ dp_t = -\nabla V(q_t) dt. \end{cases}$$

To be more precise, we introduce the process $(Q_t^\gamma, P_t^\gamma) = (q_{t/\gamma}, p_{t/\gamma})$ where $(q_t, p_t)_{t \geq 0}$ is solution to (3.49). It can then be shown that, in the limit $\gamma \rightarrow 0$, the Hamiltonian $H(Q_t^\gamma, P_t^\gamma)$ converges to an effective diffusion on a graph [192, 193, 191, 220]. In particular the relevant time scale in the underdamped limit is $\gamma^{-1}t$.

On the other hand, in the limit $\gamma \rightarrow +\infty$ and under appropriate rescaling, we recover the overdamped dynamics studied in Section 3.4.1. To see this, we integrate the second line in (3.49) to obtain

$$p_t - p_0 = - \int_0^t \nabla V(q_s) ds - \gamma(q_t - q_0) + \sqrt{2\gamma} B_t.$$

By introducing now $Q_t^\gamma = q_{\gamma t}$ and $P_t^\gamma = p_{\gamma t}$, this becomes

$$Q_t^\gamma - Q_0^\gamma = \frac{P_0^\gamma - P_t^\gamma}{\gamma} - \int_0^t \nabla V(Q_s^\gamma) ds + \sqrt{2} B_t.$$

When $\gamma \rightarrow +\infty$, we observe that Q_t^∞ converges formally at speed $1/\gamma$ towards the solution of (3.43), see [350, Section 6.5]. The relevant time scale in the overdamped limit is therefore γt . These remarks will be of interest when studying the rate function below.

3.4.2.2 Large deviations

In order to obtain a large deviations principle for (3.49), let us make the following classical assumption on the growth of the potential [324, 277, 304, 427].

Assumption 3.22. *The potential $V \in \mathcal{S}$ has compact level sets, satisfies $e^{-V} \in L^1(\mathcal{X})$ and there exist $c_V > 0$, $C_V \in \mathbb{R}$ such that*

$$q \cdot \nabla V(q) \geq c_V |q|^2 - C_V.$$

We can now find a Lyapunov function for (3.49) by following *e.g.* [427, 399, 324], as made precise in Section 3.6.4. Recall that the Hamiltonian H is defined in (3.53).

Lemma 3.23. *Suppose that $(X_t)_{t \geq 0} = (q_t, p_t)_{t \geq 0}$ solves (3.49) where V satisfies Assumption 3.22. Then for any $\gamma > 0$ and $\theta \in (0, 1)$, there exists $\varepsilon > 0$ such that*

$$W(q, p) = e^{\theta H(q, p) + \varepsilon q \cdot p} \quad (3.54)$$

is a Lyapunov function in the sense of Assumption 3.6. More precisely, for any $\gamma > 0$ and $\theta \in (0, 1)$, there exist $\varepsilon > 0$ and $a, b, C > 0$ such that

$$-\frac{\mathcal{L}_\gamma W}{W} \geq a|q|^2 + b|p|^2 - C.$$

The Lyapunov function (3.54) can be adapted to cases where V has singularities, see [239, 316]. We can now deduce our main theorem on the Langevin dynamics, since Assumptions 3.4 and 3.5 are readily satisfied, see for instance [324].

Theorem 3.24. *Assume that $(X_t)_{t \geq 0} = (q_t, p_t)_{t \geq 0}$ solves (3.49) where V satisfies Assumption 3.22, and consider $\kappa(q, p) = 1 + |q|^\alpha + |p|^\beta$ with $\alpha \in [0, 2)$, $\beta \in [0, 2)$. Then $(X_t)_{t \geq 0}$ is ergodic with respect to the measure μ in the sense of Proposition 3.9, with Lyapunov function defined in (3.54). Moreover, the empirical measure*

$$L_t = \frac{1}{t} \int_0^t \delta_{(q_s, p_s)} ds$$

satisfies a LDP in the τ^κ -topology. Finally, for any $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ such that $d\nu = e^v d\mu$ with $v \in \mathcal{H}^1(\nu)$ and $\mathcal{L}_{\text{ham}} v \in \mathcal{H}^{-1}(\nu)$, the rate function reads

$$I_\gamma(\nu) = \frac{\gamma}{4} \int_{\mathcal{X}} |\nabla_p v|^2 d\nu + \frac{1}{4\gamma} \int_{\mathcal{X}} |\nabla_p \psi|^2 d\nu, \quad (3.55)$$

where ψ is the unique solution in $\mathcal{H}^1(\nu)$ to the Poisson problem:

$$-\Delta_p \psi + (p - \nabla_p v) \cdot \nabla_p \psi = \mathcal{L}_{\text{ham}} v. \quad (3.56)$$

The proof of Theorem 3.24 is a direct application of the results of Sections 3.2 and 3.3. For the expression of the rate function, we use (3.41) and (3.50) together with the fact that in this case, the matrix S defined in Section 3.2.1 reads

$$S = \gamma \begin{pmatrix} 0 & 0 \\ 0 & I_{d \times d} \end{pmatrix} \in \mathbb{R}^{2d \times 2d}.$$

While κ can be chosen independently of the friction γ , it is interesting to note the dependency of the rate function (3.55) with respect to this parameter. We discuss more precisely the scaling of the rate function with respect to γ in the next section, depending on the form of ν .

3.4.2.3 Low and large friction asymptotics

The next corollary shows how the decomposition (3.55) allows to identify the most likely fluctuations in the overdamped and underdamped limits. By this we mean that, when $\gamma \rightarrow 0$ or $\gamma \rightarrow +\infty$, most fluctuations become exponentially rare in γ , but some of them are associated with rate functions that vanish as $\gamma \rightarrow 0$ and $\gamma \rightarrow +\infty$. The expression of these typical fluctuations is motivated by the discussion on the overdamped and underdamped limits in Section 3.4.2.1, from which the scalings of the rate function appear natural. Recall the definition of the marginal in position $\bar{\mu}$ in (3.51).

Corollary 3.25. *Suppose the assumptions of Theorem 3.24 hold true.*

- Overdamped limit $\gamma \rightarrow +\infty$. Take a measure $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ that has equilibrated in speed, i.e. such that $v(q, p) = v(q)$ with $v \in \mathcal{H}^1(\nu)$ and $p \cdot \nabla_q v \in \mathcal{H}^{-1}(\nu)$. Then, for any $\gamma > 0$,

$$I_\gamma(\nu) = \frac{1}{4\gamma} \int_{\mathbb{R}^d} |\nabla v(q)|^2 \bar{\nu}(dq), \quad (3.57)$$

where $\bar{\nu} = e^v \bar{\mu}$.

- Hamiltonian limit $\gamma \rightarrow 0$. Consider a Hamiltonian fluctuation, i.e. $d\nu = e^v d\mu$ with $v(q, p) = g(H(q, p)) \in \mathcal{H}^1(\nu)$ for $g \in C^1(\mathbb{R})$, where H is defined in (3.54). Then, for any $\gamma > 0$,

$$I_\gamma(\nu) = \frac{\gamma}{4} \int_{\mathcal{X}} |pg'(H(q, p))|^2 \nu(dq dp). \quad (3.58)$$

The proof is an immediate consequence of (3.55).

Proof. Consider first the case where $d\nu = e^v d\mu$ with $v(q, p) = v(q)$. We have

$$\frac{\gamma}{4} \int_{\mathcal{X}} |\nabla_p v|^2 d\nu = 0.$$

Next, (3.56) becomes

$$-(\Delta_p - p \cdot \nabla_p) \psi(q, p) = p \cdot \nabla_q v(q).$$

The solution to this equation is $\psi(q, p) = -p \cdot \nabla_q v(q)$ which indeed belongs to $\mathcal{H}^1(\nu)$ since $\mathcal{L}_{\text{ham}} v \in \mathcal{H}^{-1}(\nu)$ (in fact we may add to ψ any function depending on q only but the solutions would be equivalent by definition of the space $\mathcal{H}^1(\nu)$ in Section 3.2.1). Plugging this solution into (3.55) leads to (3.57).

Assume now that $v(q, p) = g(H(q, p))$ belongs to $\mathcal{H}^1(\nu)$ with $g \in C^1(\mathbb{R})$. It holds

$$\mathcal{L}_{\text{ham}} v(q, p) = g'(H(q, p)) \mathcal{L}_{\text{ham}} H(q, p) = 0.$$

As a result, the solution ψ to (3.56) is $\psi = 0$ (again, up to a function of q only), from which (3.58) follows since $v \in \mathcal{H}^1(\nu)$. \square

Corollary 3.25 characterizes the dominant fluctuations in the small and large friction regimes. In the overdamped limit $\gamma \rightarrow +\infty$ the dominant fluctuations are in position only, and the rate function is actually that of the overdamped limit (3.47) up to a time rescaling in $t \mapsto \gamma t$, which is consistent with the discussion on the overdamped limit in Section 3.4.2.1. On the other hand, in the Hamiltonian limit $\gamma \rightarrow 0$, the dominant fluctuations have a Hamiltonian form, with the inverse time rescaling $t \rightarrow \gamma^{-1} t$. This is consistent with the small temperature limit of Hamiltonian systems [193].

Although Corollary 3.25 provides interesting information, its structure is quite rigid. For instance, in the overdamped limit, we consider only position-dependent perturbation, which is not realistic. We now refine the asymptotics by considering the next order correction in γ for the perturbation in both regimes, which shows the robustness of the analysis.

Corollary 3.26. *Suppose the assumptions of Theorem 3.24 hold true.*

- Overdamped limit $\gamma \rightarrow +\infty$. Consider the measure $\nu_\gamma \in \mathcal{P}_\kappa(\mathcal{X})$ defined by $\nu_\gamma = e^{v_\gamma} d\mu$ with $v_\gamma(q, p) = v(q) + \gamma^{-1} \tilde{v}(q, p)$ where $\mathcal{L}_{\text{ham}} v \in \mathcal{H}^{-1}(\nu)$, and $\tilde{v} \in \mathcal{H}^1(\nu)$ is bounded and satisfies $\nabla_q v \cdot \nabla_p \tilde{v} \in \mathcal{H}^{-1}(\nu)$ and $\mathcal{L}_{\text{ham}} \tilde{v} \in \mathcal{H}^{-1}(\nu)$. Then

$$\forall \gamma \geq 1, \quad I_\gamma(\nu_\gamma) = \frac{1}{4\gamma} \left[\int_{\mathcal{X}} |\nabla_p \tilde{v}|^2 d\nu + \int_{\mathbb{R}^d} |\nabla_q v|^2 d\bar{\nu} \right] + O\left(\frac{1}{\gamma^2}\right), \quad (3.59)$$

where $\bar{\nu} = e^v \bar{\mu}$.

- Hamiltonian limit $\gamma \rightarrow 0$. Consider $\nu_\gamma = e^{v_\gamma} d\mu$ with $v_\gamma(q, p) = g(H(q, p)) + \gamma \tilde{v}(q, p)$, where $g \in C^1(\mathbb{R})$, $g(H) \in \mathcal{H}^1(\nu)$, and $\tilde{v} \in \mathcal{H}^1(\nu)$ is bounded and satisfies $\mathcal{L}_{\text{ham}} \tilde{v} \in \mathcal{H}^{-1}(\nu)$. Then

$$\forall \gamma \leq 1, \quad I_\gamma(\nu_\gamma) = \frac{\gamma}{4} \left[\int_{\mathcal{X}} |pg'(H(q, p))|^2 \nu(dq dp) + \int_{\mathcal{X}} |\nabla_p \tilde{\psi}|^2 d\nu \right] + O(\gamma^2), \quad (3.60)$$

where $\tilde{\psi}$ is the unique $\mathcal{H}^1(\nu)$ -solution to

$$-\Delta_p \tilde{\psi} - (1 - g'(H(q, p)))p \cdot \nabla_p \tilde{\psi} = \mathcal{L}_{\text{ham}} \tilde{v}. \quad (3.61)$$

We believe it is also instructive to mention the relation between the rate function (3.55) and the asymptotic variance of the Langevin dynamics. Indeed, when considering small perturbations of the invariant measure, Corollary 3.26 shows that

$$I_\gamma \sim \min\left(\gamma, \frac{1}{\gamma}\right). \quad (3.62)$$

However, the resolvent estimates in [295, Section 2.1] show that the asymptotic variance σ_γ^2 scales like

$$\sigma_\gamma^2 \sim \max\left(\gamma, \frac{1}{\gamma}\right). \quad (3.63)$$

Since we expect the asymptotic variance to be the inverse of the rate function around the invariant measure [363], the scalings (3.62) and (3.63) are consistent. However, as (3.55) suggests, this scaling is no longer true for general fluctuations. We now present the proof of Corollary 3.26.

Proof. We first consider the overdamped limit $\gamma \rightarrow +\infty$. By boundedness of \tilde{v} we have, for any $\gamma \geq 1$ and $\psi \in \mathcal{H}^1(\nu_\gamma)$,

$$e^{\frac{\inf \tilde{v}}{\gamma}} |\psi|_{\mathcal{H}^1(\nu)}^2 \leq |\psi|_{\mathcal{H}^1(\nu_\gamma)}^2 \leq e^{\frac{\sup \tilde{v}}{\gamma}} |\psi|_{\mathcal{H}^1(\nu)}^2. \quad (3.64)$$

Thus, the norms $\mathcal{H}^1(\nu_\gamma)$ and $\mathcal{H}^1(\nu)$ are equivalent for any fixed $\gamma \geq 1$, and the functions of $\mathcal{H}^1(\nu_\gamma)$ and $\mathcal{H}^1(\nu)$ coincide (we repeatedly use this fact below, and we will use a similar argument when $\gamma \leq 1$). A similar conclusion holds for the corresponding dual norms. This consequence of the boundedness of \tilde{v} makes the analysis simpler.

Recall that we consider $v_\gamma = v + \gamma^{-1} \tilde{v}$ in the overdamped limit. The symmetric part of the rate function is easily computed since v only depends on the position variable, namely

$$I_S(\nu_\gamma) = \frac{\gamma}{4} \int_{\mathcal{X}} |\nabla_p(v + \gamma^{-1} \tilde{v})|^2 e^{v + \frac{\tilde{v}}{\gamma}} d\mu = \frac{1}{4\gamma} \int_{\mathcal{X}} |\nabla_p \tilde{v}|^2 d\nu + O\left(\frac{1}{\gamma^2}\right),$$

where we used that \tilde{v} belongs to $\mathcal{H}^1(\nu)$ and is bounded to expand the exponential. For the antisymmetric part, by (3.56), we have to consider the solution $\psi_\gamma \in \mathcal{H}^1(\nu_\gamma)$ to

$$-\Delta_p \psi_\gamma + \left(p - \frac{1}{\gamma} \nabla_p \tilde{v}\right) \cdot \nabla_p \psi_\gamma = \mathcal{L}_{\text{ham}} v_\gamma.$$

Corollary 3.25 suggests that at leading order in γ it holds $\psi_\gamma = \psi + O(\gamma^{-1})$ where $\psi(q, p) = p \cdot \nabla_q v(q)$. In order to make this idea more precise we compute

$$\left(-\Delta_p + \left(p - \frac{1}{\gamma} \nabla_p \tilde{v}\right) \cdot \nabla_p\right) (\psi_\gamma - \psi) = \frac{1}{\gamma} (\mathcal{L}_{\text{ham}} \tilde{v} + \nabla_q v \cdot \nabla_p \tilde{v}).$$

In what follows, we denote by $u = \mathcal{L}_{\text{ham}} \tilde{v} + \nabla_q v \cdot \nabla_p \tilde{v}$ the right hand side of the above equation. Since $\nabla_q v \cdot \nabla_p \tilde{v} \in \mathcal{H}^{-1}(\nu_\gamma)$ and $\mathcal{L}_{\text{ham}} \tilde{v} \in \mathcal{H}^{-1}(\nu_\gamma)$ by assumption, it holds $u \in \mathcal{H}^{-1}(\nu_\gamma)$. Thus, multiplying by $\psi_\gamma - \psi$ and integrating with respect to ν_γ we obtain

$$\int_{\mathcal{X}} |\nabla_p(\psi_\gamma - \psi)|^2 d\nu_\gamma = -\frac{1}{\gamma} \int_{\mathcal{X}} (\psi_\gamma - \psi) u d\nu_\gamma.$$

Using the duality between $\mathcal{H}^1(\nu_\gamma)$ and $\mathcal{H}^{-1}(\nu_\gamma)$ (see [273, Section 2.2 Claim F]) and (3.64) we obtain

$$\forall \gamma \geq 1, \quad |\psi_\gamma - \psi|_{\mathcal{H}^1(\nu)} \leq \frac{C}{\gamma} |u|_{\mathcal{H}^{-1}(\nu)},$$

where C is some constant independent of γ . This shows that $\psi_\gamma = \psi + \gamma^{-1}\tilde{\psi}_\gamma$ with $|\tilde{\psi}_\gamma|_{\mathcal{H}^1(\nu)} \leq C'$ for a constant $C' > 0$ and all $\gamma \geq 1$. Plugging this estimate into (3.55) and using that $\nabla_p \psi = \nabla_q v$, we obtain the second term on the right hand side of (3.59).

The arguments to prove the limit $\gamma \rightarrow 0$ follow a similar path, so we only sketch the proof. First, the boundedness of \tilde{v} allows again to compare the Sobolev norms associated with ν and ν_γ for any $\gamma \leq 1$ (by writing the counterpart of (3.64) in this regime). The first term on the right hand side of (3.60) is easily obtained as in Corollary 3.25 using that $g(H) \in \mathcal{H}^1(\nu)$ and \tilde{v} is bounded. Concerning the antisymmetric part, (3.56) now reads

$$(-\Delta_p + (p - \nabla_p v_\gamma) \cdot \nabla_p) \psi_\gamma = \gamma \mathcal{L}_{\text{ham}} \tilde{v},$$

since $\mathcal{L}_{\text{ham}} g(H(q, p)) = 0$. Because of the scaling in γ on the right hand side of the above equation, the solution ψ_γ can be expanded as $\psi_\gamma = \gamma \tilde{\psi} + O(\gamma^2)$ in $\mathcal{H}^1(\nu)$, where $\tilde{\psi}$ is solution to

$$-\Delta_p \tilde{\psi} + (1 - g'(H(q, p))) p \cdot \nabla_p \tilde{\psi} = \mathcal{L}_{\text{ham}} \tilde{v}.$$

This reasoning can be made rigorous by a precise asymptotic analysis as above. Plugging this expansion into (3.55) provides the second term on the right hand side of (3.60). \square

3.5 Conclusion and perspectives

The goal of the work presented in this chapter was twofold. Our first aim was to provide, given a diffusion process, a precise class of unbounded functions for which a large deviations principle holds. This question is answered in Section 3.2 where we prove a LDP for the empirical measure in a topology associated with unbounded functions, in relation with a Witten–Lyapunov condition. In particular, a comparison with Cramér’s condition for independent variables shows the effect of correlations on the stability of the SDE at hand. These results extend in several directions and refine previous works [427, 275]. However, the necessity of our Lyapunov condition for a LDP to hold (as is known for the Sanov theorem [424]) is still an open problem. Our second concern was to provide finer expressions of the rate function governing the LDP, in particular in order to study the Langevin equation. We answer to this question in two ways in Section 3.3. We first provide an alternative variational formula for the rate function in Section 3.3.1, which gives as a by product a very general representation formula for the principle eigenvalue of second order differential operators, without symmetry assumption. This extends the important work of Donsker and Varadhan [133] in an unbounded setting. Next, in Section 3.3.2, we show a general decomposition of the rate function into symmetric and antisymmetric parts of the dynamics based on the computations in [49]. Interestingly, the proof of the result relies on a Witten-like transform in the above mentioned variational representation of the rate function. These results allow us to describe precisely the rate function of an irreversible overdamped Langevin dynamics in Section 3.4.1, revisiting results from [363] in an unbounded setting. More interestingly in Section 3.4.2 we provide, for the Langevin dynamics, asymptotics of the rate function for the overdamped and the underdamped limits. We thus characterize the most likely fluctuations in both regimes with a natural physical interpretation. Considering piecewise deterministic processes [39, 151, 152] (which lacks regularity) instead of the Langevin dynamics is also an interesting problem.

We would like to mention several interesting directions for future works. A first natural issue is to rephrase our results in the optimal control framework developed *e.g.* in [68, 147, 149]. This is particularly interesting for numerical purposes, since the optimal control representations can be learnt on the fly with stochastic approximation methods [30, 58, 28, 184]. We believe that such results can be obtained by harvesting the contraction principle provided by Corollary 3.12.

On a more theoretical ground, dual Sobolev norms have recently attracted attention in the optimal transport community due to the so-called optimal matching problem, see for instance [291, 292] and the references therein. With these works in mind, the dual Sobolev norm in the antisymmetric part of the rate function described in Section 3.3.2 could be interpreted as an infinitesimal transport cost related to the antisymmetric part of the dynamics, which is an alluring interpretation of irreversibility. Note that the relations between optimal transport and large deviations theory have a fruitful history, see *e.g.* [207].

Then, it has been known for some time in the physics literature that the empirical density of a diffusion may not contain enough information to describe its fluctuations in an irreversible regime. It is actually more relevant to consider the fluctuations of both the empirical density and current, a

procedure sometimes called level 2.5 large deviations [98, 25]. This framework can be used to provide a clear description of the rate function of irreversible dynamics. As shown in [25], such large deviations results can be derived by Krein–Rutman arguments like those used in the present paper. Therefore, we believe that our results can be extended to prove level 2.5 large deviations principles and characterize precisely the admissible currents.

Finally, it is important to understand the behaviour of observables which are not covered by our analysis. It has been recently shown [342] in the case of the Ornstein–Uhlenbeck process that observables growing too fast at infinity with respect to the confinement are characterized by a *heavy tail* behaviour. This leads in particular to a localization in time of the fluctuating trajectories, and the Krein–Rutman strategy developed in the present paper does not apply. We therefore believe there are several interesting open questions in this direction.

Acknowledgements

We warmly thank Hugo Touchette for reading an early version of the manuscript and providing useful comments. We are grateful to O. Zeitouni for an interesting discussion about scalings in large deviations theory as well as Jianfeng Lu for mentioning the work [49]. We also thank Julien Reygner for general discussions on large deviations.

3.6 Proofs

3.6.1 Proof of the large deviations principle

As mentioned after Theorem 3.10, our proof relies on the Gärtner–Ellis theorem [119], for which we need several preliminary results. The key object is the functional

$$f \in B_\kappa^\infty(\mathcal{X}) \mapsto \lambda(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right].$$

As mentioned in Section 1.2.2, the Gärtner–Ellis theorem (Theorem 1.10 in Section 1.2.2) states that if this functional is finite and Gateau-differentiable over $B_\kappa^\infty(\mathcal{X})$ and $(L_t)_{t \geq 0}$ defined in (3.2) is exponentially tight for the τ^κ -topology, then $(L_t)_{t \geq 0}$ satisfies a LDP in the dual space of $B_\kappa^\infty(\mathcal{X})$. We refer to Section 1.2 for more details on large deviations theory.

However, studying the range of functions f for which the functional λ is finite and Gateau-differentiable is not an easy task. Our strategy is to prove that $r(f)$, the largest eigenvalue in modulus of the Feynman–Kac operator $\mathcal{L} + f$, is real for any $f \in B_\kappa^\infty(\mathcal{X})$, and to show that it is actually equal to the cumulant function $\lambda(f)$ defined in (3.25). This amounts to showing the well-posedness and regularity of a family of spectral problems. For this, we use several ideas from Chapter 2, where we show that under a Lyapunov and an irreducibility conditions, the eigenvalue problem to which λ is associated is well defined. The seminal paper by Gärtner [199, Section 3] provides useful technical tools, as well as [165, 427].

In all this section, we suppose that Assumptions 3.4, 3.5 and 3.6 hold true and consider a function $\kappa : \mathcal{X} \rightarrow [1, +\infty)$ as in Assumption 3.6, *i.e.* such that $\kappa \ll \Psi$ and either κ has compact level sets or is bounded. We repeatedly use that $\kappa \ll -\frac{\mathcal{L}\Psi}{\Psi}$ in view of (3.21). We start with important properties of key martingales that repeatedly appear in the proofs of the required technical results.

Lemma 3.27. *If $(X_t)_{t \geq 0}$ is solution to (3.9), then the stochastic processes defined by*

$$M_t = W(X_t) e^{-\int_0^t \frac{\mathcal{L}W}{W}(X_s) ds} \quad \text{and} \quad \mathcal{M}_t = \mathcal{W}(X_t) e^{-\int_0^t \frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}(X_s) ds} \quad (3.65)$$

are continuous non-negative local martingales, hence supermartingales. Moreover, it holds almost surely

$$\mathcal{M}_t^2 \leq C_1 e^{tC_2} M_t, \quad (3.66)$$

where $C_1 > 0$ and $C_2 \in \mathbb{R}$ are the constants from Assumption 3.6.

Proof. First, Itô formula shows that

$$dM_t = \sigma(X_t) \cdot \nabla W(X_t) e^{-\int_0^t \frac{\mathcal{L}W}{W}(X_s) ds} dB_t.$$

Since W is $C^2(\mathcal{X})$ and σ is continuous, M_t is a continuous local martingale [262]. Since it is non-negative, it is a supermartingale by Fatou's lemma, and the same conclusion holds for \mathcal{M}_t . On the other hand, (3.21) shows that for any continuous path $(X_t)_{t \geq 0}$ we have

$$\mathcal{M}_t^2 = \mathcal{W}(X_t)^2 e^{\int_0^t -2 \frac{\mathcal{L}W}{W}(X_s) ds} \leq C_1 W(X_t) \exp \left[\int_0^t \left(-\frac{\mathcal{L}W}{W}(X_s) + C_2 \right) ds \right] \leq C_1 e^{C_2 t} M_t,$$

which concludes the proof. \square

The use of the martingale M_t is inspired by [427] where it is considered to control return times to compact sets. Here, it allows to define the Feynman–Kac semigroup associated with the dynamics $(X_t)_{t \geq 0}$ with weight function $f \in B_\kappa^\infty(\mathcal{X})$.

Lemma 3.28. *For $f \in B_\kappa^\infty(\mathcal{X})$, consider the Feynman–Kac operator defined as*

$$(P_t^f \varphi)(x) = \mathbb{E}_x \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]. \quad (3.67)$$

Then for any $f \in B_\kappa^\infty(\mathcal{X})$, any $t > 0$ and any $a > 0$, there exist $c_{a,t} \geq 0$ and a compact subset $K_{a,t} \subset \mathcal{X}$ such that

$$\forall x \in \mathcal{X}, \quad (P_t^f W)(x) \leq e^{-at} W(x) + c_{a,t} \mathbb{1}_{K_{a,t}}(x). \quad (3.68)$$

Moreover, for any $f \in B_\kappa^\infty(\mathcal{X})$ and $t > 0$ it holds $P_t^f \in \mathcal{B}(B_W^\infty(\mathcal{X}))$, and $(P_t^f)_{t \geq 0}$ is a semigroup of bounded operators over $B_W^\infty(\mathcal{X})$. It has generator $\mathcal{L} + f$ defined on the domain

$$\mathcal{D}_{\mathcal{L},f} = \left\{ \varphi \in B_W^\infty(\mathcal{X}) \mid (\mathcal{L} + f)\varphi \in B_W^\infty(\mathcal{X}) \right\}. \quad (3.69)$$

Note that for any $f \in B_\kappa^\infty(\mathcal{X})$ the domain $\mathcal{D}_{\mathcal{L},f}$ contains $C_c^\infty(\mathcal{X})$, but $C_c^\infty(\mathcal{X})$ is not dense in $B_W^\infty(\mathcal{X})$ for the associated norm. However $(P_t^f)_{t \geq 0}$ generates a semigroup of bounded operators [351], so $\mathcal{L} + f$ can be defined on $B_W^\infty(\mathcal{X})$ with domain $\mathcal{D}_{\mathcal{L},f}$.

Proof. We first show that for any $f \in B_\kappa^\infty(\mathcal{X})$, $(P_t^f)_{t \geq 0}$ is a semigroup of bounded operators on $B_W^\infty(\mathcal{X})$, before turning to the proof of (3.68). For a fixed $f \in B_\kappa^\infty(\mathcal{X})$, since $\kappa \ll \Psi$ there is $c > 0$ such that for any $t > 0$ it holds

$$P_t^f W(x) = \mathbb{E}_x \left[W(X_t) e^{\int_0^t f(X_s) ds} \right] \leq e^{ct} \mathbb{E}_x \left[W(X_t) e^{-\int_0^t \frac{\mathcal{L}W}{W}(X_s) ds} \right].$$

Using Lemma 3.27, the supermartingale property leads to

$$P_t^f W(x) \leq e^{ct} \mathbb{E}_x [M_t] \leq e^{ct} W(x).$$

Thus $(P_t^f)_{t \geq 0}$ is a semigroup of bounded operators over $B_W^\infty(\mathcal{X})$. This semigroup has generator $\mathcal{L} + f$, which can be shown (although f is not regular) by using the Itô formula on smooth test functions, see [361, Chapter VIII, Proposition 3.10]. The domain of this generator is then defined by (3.69) through standard semigroup definitions [351].

We next prove (3.68) for a fix $f \in B_\kappa^\infty(\mathcal{X})$, which we assume non-zero without loss of generality. Note that

$$\frac{\mathcal{L}W}{W} + f \leq -\Psi + \|f\|_{B_\kappa^\infty} \kappa = -\left(1 - \|f\|_{B_\kappa^\infty} \frac{\kappa}{\Psi}\right) \Psi.$$

Since Ψ has compact level sets and $\kappa \ll \Psi$, for any $a > 0$ there exists a compact set $K_a \subset \mathcal{X}$ and a constant $b_{0,a}$ such that

$$\frac{\mathcal{L}W}{W} + f \leq -a + b_{0,a} \mathbb{1}_{K_a},$$

which implies

$$(\mathcal{L} + f)W \leq -aW + b_a \mathbb{1}_{K_a},$$

with $b_a = b_{0,a} \sup_K W < +\infty$ since $W \in C^2(\mathcal{X})$. Therefore,

$$\frac{d}{dt} \left(e^{at} P_t^f W \right) = e^{at} P_t^f (aW + (\mathcal{L} + f)W) \leq b_a e^{at} P_t^f \mathbb{1}_{K_a} \leq b_a e^{at} P_t^f \mathbb{1}. \quad (3.70)$$

We can now bound the right hand side of the above equation with a technique similar to the one used in [181, Section 2.3]. Indeed, for any $x \in \mathcal{X}$,

$$(P_t^f \mathbf{1})(x) = \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right] \leq \mathbb{E}_x \left[e^{\|f\|_{B_\kappa^\infty} \int_0^t \kappa(X_s) ds} \right]. \quad (3.71)$$

Since $\kappa \ll -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}$, there exists a constant $c \geq 0$ depending on f such that

$$\kappa \leq \frac{1}{\|f\|_{B_\kappa^\infty}} \left(-\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}} \right) + c.$$

Plugging this estimate into (3.71) and using that $\mathcal{W} \geq 1$ leads to

$$(P_t^f \mathbf{1})(x) \leq e^{ct} \mathbb{E}_x \left[\mathcal{W}(X_t) e^{\int_0^t -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}(X_s) ds} \right] = e^{ct} \mathbb{E}_x[\mathcal{M}_t] \leq e^{ct} \mathcal{W}(x),$$

where the last bound comes from Lemma 3.27. Therefore, using this estimate to bound the right hand side of (3.70), we end up with

$$\frac{d}{dt} \left(e^{at} P_t^f W \right) \leq b_a e^{(a+c)t} \mathcal{W}.$$

Integrating with respect to time leads to

$$(P_t^f W)(x) \leq e^{-at} W(x) + \tilde{b}_a \mathcal{W}(x), \quad \tilde{b}_a = \frac{b_a}{a+c} e^{ct}.$$

Since $\mathcal{W} \ll W$, there exists a compact set $K_{a,t} \subset \mathcal{X}$ such that $\tilde{b}_a \mathcal{W} \leq e^{-at} W$ outside $K_{a,t}$, so that we have

$$\forall x \in \mathcal{X}, \quad (P_t^f W)(x) \leq 2e^{-at} W(x) + \left(\tilde{b}_a \sup_{K_{a,t}} \mathcal{W} \right) \mathbf{1}_{K_{a,t}}(x).$$

Consider finally $a' \leq a$ such that $2e^{-at} \leq e^{-a't}$ (for instance by defining $a' = a - \log(2)/t$ for a large enough), and set $K_{a',t} = K_{a,t}$ and $c_{a',t} = \tilde{b}_a \sup_{K_{a,t}} \mathcal{W}$. It then holds

$$\forall x \in \mathcal{X}, \quad (P_t^f W)(x) \leq e^{-a't} W(x) + c_{a',t} \mathbf{1}_{K_{a',t}}(x),$$

which proves (3.68). \square

Lemma 3.28 proves crucial to obtain the compactness of the evolution operator P_t^f over $B_W^\infty(\mathcal{X})$, as noted in Chapter 2 (a result inspired by [362, Theorem 8.9]). Another key ingredient is the regularization property of the evolution. The following bound on the Feynman–Kac semigroup depending on the weight function f is one element in this direction.

Lemma 3.29. *Suppose that Assumptions 3.4, 3.5 and 3.6 hold true, and fix $f, g \in B_\kappa^\infty(\mathcal{X})$. Then, for any $t > 0$, there exists $C_t > 0$ such that for all $\varphi \in B_W^\infty(\mathcal{X})$, and all $x \in \mathcal{X}$,*

$$|P_t^f \varphi(x) - P_t^g \varphi(x)| \leq \|\varphi\|_{B_W^\infty} \mathbb{E}_x \left[W(X_t) \left(\int_0^t |f(X_s) - g(X_s)| ds \right) e^{(\|f\|_{B_\kappa^\infty} + \|g\|_{B_\kappa^\infty}) \int_0^t \kappa(X_s) ds} \right] \quad (3.72)$$

Proof. Using the inequality $|e^a - e^b| \leq |a - b| e^{|a|+|b|}$ for $a, b \in \mathbb{R}$, we first write, for $x \in \mathcal{X}$,

$$\begin{aligned} |P_t^f \varphi(x) - P_t^g \varphi(x)| &\leq \mathbb{E}_x \left[|\varphi(X_t)| \left| e^{\int_0^t f(X_s) ds} - e^{\int_0^t g(X_s) ds} \right| \right] \\ &\leq \|\varphi\|_{B_W^\infty} \mathbb{E}_x \left[W(X_t) \left| \int_0^t f(X_s) ds - \int_0^t g(X_s) ds \right| e^{\int_0^t |f(X_s)| ds + \int_0^t |g(X_s)| ds} \right], \\ &\leq \|\varphi\|_{B_W^\infty} \mathbb{E}_x \left[W(X_t) \left(\int_0^t |f(X_s) - g(X_s)| ds \right) e^{(\|f\|_{B_\kappa^\infty} + \|g\|_{B_\kappa^\infty}) \int_0^t \kappa(X_s) ds} \right], \end{aligned}$$

which is the desired conclusion. \square

We can now use Lemma 3.29 to show an important regularization property of the Feynman–Kac semigroup.

Lemma 3.30. *For any $f \in B_\kappa^\infty(\mathcal{X})$, $\varphi \in B_W^\infty(\mathcal{X})$, any $t > 0$ and any compact $K \subset \mathcal{X}$, the function $P_t^f(\varphi \mathbf{1}_K)$ is continuous.*

Let us insist on the fact that the statement of Lemma 3.30 is a consequence of Hörmander's theorem [161, Theorem 4.1] when f has polynomial growth and is smooth. However, the result is more difficult to obtain when f is irregular. The idea of the proof is to use the local martingales introduced in Lemma 3.27 to show that the regularization property of Hörmander's theorem is preserved when f does not grow too fast. This gives a clear answer to the somehow non-optimal framework of Section 2.2.3 in Chapter 2.

Proof. We use Assumption 3.4 to revisit [199, pages 34-35] in an unbounded setting and with a hypoelliptic flavour. First, we note that for $f \in C_c^\infty(\mathcal{X})$, the result is a direct application of Assumption 3.4 combined with Hörmander's theorem. Indeed, for any $\varphi \in B_W^\infty(\mathcal{X})$ and compact set $K \subset \mathcal{X}$, the function $\varphi \mathbf{1}_K$ is bounded, so that

$$u(t, x) = P_t^f(\varphi \mathbf{1}_K)(x)$$

is solution to the PDE

$$\partial_t u = \mathcal{L}u + fu, \quad u(0, x) = \varphi \mathbf{1}_K(x).$$

Therefore, since $\partial_t - \mathcal{L} - f$ is hypoelliptic with regular coefficients of polynomial growth (recall that $b, \sigma \in \mathcal{S}$), [161, Theorem 4.1] ensures that $u(t, \cdot)$ is $C^\infty(\mathcal{X})$ for any $t > 0$ (see also [244]). In particular, $P_t^f(\varphi \mathbf{1}_K)$ is continuous.

We now use an approximation argument inspired by [199, Section 3] for a generic function $f \in B_\kappa^\infty(\mathcal{X})$. Consider a sequence $(f_n)_{n \in \mathbb{N}}$ of functions belonging to $C_c^\infty(\mathcal{X})$ with $\|f_n\|_{B_\kappa^\infty} \leq \|f\|_{B_\kappa^\infty}$ for any $n \in \mathbb{N}$, and such that $f \rightarrow f_n$ almost everywhere (such a sequence exists by Lusin's theorem, see [379, Chapter 2]). Using Lemma 3.29 with an additional indicator function, and since $\|f_n\|_{B_\kappa^\infty} \leq \|f\|_{B_\kappa^\infty}$, we have for any $\varphi \in B_W^\infty(\mathcal{X})$, $n \in \mathbb{N}$ and $x \in \mathcal{X}$,

$$\left| P_t^f(\varphi \mathbf{1}_K)(x) - P_t^{f_n}(\varphi \mathbf{1}_K)(x) \right| \leq \|\varphi\|_{B_W^\infty} \mathbb{E}_x \left[\mathbf{1}_K(X_t) W(X_t) \left(\int_0^t |f(X_s) - f_n(X_s)| ds \right) e^{\delta \int_0^t \kappa(X_s) ds} \right], \quad (3.73)$$

with $\delta = 2\|f\|_{B_\kappa^\infty}$.

Our goal is now to show that $P_t^{f_n}(\varphi \mathbf{1}_K)$ converges uniformly over any compact K' to $P_t^f(\varphi \mathbf{1}_K)$, by proving that the right hand side of (3.73) goes uniformly to 0 over K' . This will conclude the proof since a uniform limit of continuous functions is continuous. We introduce to this end the events

$$\forall m \geq 1, \quad \mathcal{E}_m = \left\{ \frac{1}{t} \int_0^t \Psi(X_s) ds \leq m \right\}, \quad (3.74)$$

and fix a compact set $K' \subset K$. The right hand side of (3.73) can then be split into two terms

$$\begin{aligned} (A) &= \mathbb{E}_x \left[\mathbf{1}_K(X_t) \mathbf{1}_{\mathcal{E}_m^c} W(X_t) \left(\int_0^t |f(X_s) - f_n(X_s)| ds \right) e^{\delta \int_0^t \kappa(X_s) ds} \right], \\ (B) &= \mathbb{E}_x \left[\mathbf{1}_K(X_t) \mathbf{1}_{\mathcal{E}_m} W(X_t) \left(\int_0^t |f(X_s) - f_n(X_s)| ds \right) e^{\delta \int_0^t \kappa(X_s) ds} \right], \end{aligned}$$

for which we show convergence to 0, uniformly for $x \in K'$, starting with (A). Since $\kappa \ll -\mathcal{L}\mathcal{W}/\mathcal{W}$, there exists $c > 0$ such that

$$2\delta\kappa \leq -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}} + c.$$

Moreover, $\|f_n\|_{B_\kappa^\infty} \leq \|f\|_{B_\kappa^\infty}$, $a \leq e^a$, and $\mathcal{W} \geq 1$, so that

$$(A) \leq e^{ct} \left(\sup_K W \right) \mathbb{E}_x \left[\mathbf{1}_K(X_t) \mathbf{1}_{\mathcal{E}_m^c} \mathcal{W}(X_t) e^{\int_0^t -\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}(X_s) ds} \right].$$

By definition of \mathcal{M}_t in (3.65) we have

$$(A) \leq e^{ct} \left(\sup_K W \right) \mathbb{E}_x \left[\mathbf{1}_{\mathcal{E}_m^c} \mathcal{M}_t \right].$$

The Cauchy-Schwarz inequality then shows that

$$(A) \leq e^{ct} \left(\sup_K W \right) W \sqrt{\mathbb{E}_x[\mathcal{M}_t^2]} \left(\mathbb{P}_x \left(\int_0^t \Psi(X_s) ds > mt \right) \right)^{\frac{1}{2}}.$$

By (3.66) it holds $\sqrt{\mathbb{E}_x[\mathcal{M}_t^2]} \leq \sqrt{C_1} e^{C_2 t/2} \sqrt{W(x)}$. Next, by Tchebychev's inequality and since $W \geq 1$,

$$\mathbb{P}_x \left(\int_0^t \Psi(X_s) ds > mt \right) \leq e^{-mt} \mathbb{E}_x \left[e^{\int_0^t \Psi(X_s) ds} \right] \leq e^{-mt} \mathbb{E}_x \left[W(X_t) e^{-\int_0^t \frac{\kappa W}{W}(X_s) ds} \right] \leq e^{-mt} W(x).$$

As a result, we obtain

$$(A) \leq e^{-\frac{mt}{2}} \left(\sup_K W \right) \left(\sup_{K'} W \right) \sqrt{C_1} e^{ct+C_2 t/2}.$$

Therefore, for any $\varepsilon > 0$, we can choose $m \geq 0$ such that $(A) \leq \varepsilon$.

Let us now control (B) , introducing $g_n = |f - f_n|$. Since $\kappa \ll \Psi$, it holds for some $c' \geq 0$,

$$\delta \kappa \leq \Psi + c'.$$

Using the definition (3.74) we have

$$\begin{aligned} (B) &\leq e^{(m+c')t} \left(\sup_K W \right) \mathbb{E}_x \left[\mathbb{1}_{\mathcal{E}_m} \int_0^t g_n(X_s) ds \right] \\ &\leq \underbrace{e^{(m+c')t} \left(\sup_K W \right) \mathbb{E}_x \left[\mathbb{1}_{\mathcal{E}_m} \int_0^t g_n(X_s) \mathbb{1}_{B_R^c}(X_s) ds \right]}_{(B')} + \underbrace{e^{(m+c')t} \left(\sup_K W \right) \mathbb{E}_x \left[\mathbb{1}_{\mathcal{E}_m} \int_0^t g_n(X_s) \mathbb{1}_{B_R}(X_s) ds \right]}_{(B'')}, \end{aligned}$$

where B_R is the ball of center 0 and radius $R > 0$. Let us first bound (B') , which retains only the parts of the trajectories performing excursions out of B_R . Using $\kappa \ll \Psi$, for $\varepsilon > 0$ and $m \geq 0$ as fixed above, there exist $R > 0$, $C_R > 0$ such that

$$\kappa \leq \varepsilon \frac{e^{-(m+c')t}}{tm \sup_K W \|f\|_{B_\infty}} \Psi + C_R \mathbb{1}_{B_R}.$$

We fix $R > 0$ and $C_R > 0$ such that the above inequality holds. Using again $g_n \leq \|f\|_{B_\infty} \kappa$, we are led to

$$\begin{aligned} (B') &\leq e^{(m+c')t} \left(\sup_K W \right) \|f\|_{B_\infty} \mathbb{E}_x \left[\mathbb{1}_{\mathcal{E}_m} \int_0^t \kappa(X_s) \mathbb{1}_{B_R^c}(X_s) ds \right] \\ &\leq \mathbb{E}_x \left[\mathbb{1}_{\mathcal{E}_m} \int_0^t \frac{\varepsilon}{tm} \Psi(X_s) \mathbb{1}_{B_R^c}(X_s) ds \right] \leq \frac{\varepsilon}{tm} \mathbb{E}_x \left[\mathbb{1}_{\mathcal{E}_m} \int_0^t \Psi(X_s) ds \right] \leq \varepsilon, \end{aligned}$$

where the last line follows from the definition (3.74) of \mathcal{E}_m . Therefore, once m is fixed, there exists $R > 0$ such that for any $n \geq 1$ and $x \in K'$, it holds $(B') \leq \varepsilon$. It remains to control (B'') in order to obtain the uniform convergence to zero of (3.73) over K' as $n \rightarrow +\infty$. In fact,

$$(B'') \leq e^{(m+c')t} \left(\sup_K W \right) \int_0^t \mathbb{E}_x [g_n(X_s) \mathbb{1}_{B_R}(X_s)] ds = e^{(m+c')t} \sup_K W \int_0^t P_s(g_n \mathbb{1}_{B_R})(x) ds,$$

where $(P_s)_{s \geq 0}$ is the evolution semigroup defined in (3.15). Since $(\mathbb{1}_{B_R} g_n)_{n \geq 1}$ is a sequence of bounded functions converging almost everywhere to zero and the transition kernel P_s has a smooth density for $s > 0$, it follows that $(P_s(g_n \mathbb{1}_{B_R}))_{n \geq 1}$ goes uniformly to zero over compact sets for any $s > 0$, see e.g. [199, 362]. Moreover, it can be shown that

$$\delta \mapsto \int_0^\delta P_s(g_n \mathbb{1}_{B_R}) ds$$

goes to zero when $\delta \rightarrow 0$, uniformly in $x \in K'$ and $n \in \mathbb{N}$. Therefore, for $\varepsilon > 0$, $R > 0$ and $m \geq 0$ fixed as above, there exist $\delta > 0$, and $n' \in \mathbb{N}$ such that for all $n \geq n'$ and $x \in K'$,

$$0 \leq \int_0^t P_s(g_n \mathbb{1}_{B_R})(x) ds = \int_0^\delta P_s(g_n \mathbb{1}_{B_R})(x) ds + \int_\delta^t P_s(g_n \mathbb{1}_{B_R})(x) ds \leq \varepsilon \frac{e^{-(m+c')t}}{\sup_K W}. \quad (3.75)$$

Then, for some constant $C > 0$ and any $n \geq n'$, $x \in K'$,

$$(B'') \leq \varepsilon,$$

so (B'') can be made arbitrarily small uniformly in $x \in K'$ as $n \rightarrow +\infty$.

Let us summarize the various approximations: for any $\varepsilon > 0$, we first fix $m \geq 0$ so that $(A) \leq \varepsilon$. Then, we choose $R > 0$ large enough so that $(B') \leq \varepsilon$. Finally, we take δ small enough and n large enough in (3.75) so that $(B'') \leq \varepsilon$ for $n \geq n'$. As a result, for any $\varepsilon > 0$ there is $n' \geq 0$ such that for $n \geq n'$ and $x \in K'$, it holds $(A) + (B) \leq 3\varepsilon$.

In conclusion, the right hand side of (3.73) goes to zero uniformly as $n \rightarrow +\infty$ over any compact set K' . Therefore $P_t^{f_n}(\varphi \mathbf{1}_K)$ is continuous and converges uniformly over K' to $P_t^f(\varphi \mathbf{1}_K)$, which is therefore continuous over K' . Since the compact $K' \subset \mathcal{X}$ is arbitrary, $P_t^f(\varphi \mathbf{1}_K)$ is continuous over \mathcal{X} , which concludes the proof. \square

Before presenting the main result concerning the spectral properties of the semigroup P_t^f and its consequences on the definition of the cumulant function $\lambda(f)$, we need the following «irreducibility» lemma, which relies on Assumption 3.5 (this replaces the minorization condition Assumption 2.2 in Chapter 2).

Lemma 3.31. *For any time $t > 0$, $x \in \mathcal{X}$ and any set $A \subset \mathcal{X}$ with non-empty interior, it holds*

$$(P_t^f \mathbf{1}_A)(x) > 0. \tag{3.76}$$

Proof. Take $x \in \mathcal{X}$ and $y \in \overset{\circ}{A}$ (which is possible since A has non-empty interior). By Assumption 3.5, there exists a C^1 -path $(\phi_s)_{s \in [0, t]}$ solving (3.19) such that $\phi_0 = x$ and $\phi_t = y$. We can then use the proof of the Stroock–Varadhan support theorem, see [362, Theorem 6.1] for an overview. In particular, Assumption 3.5 implies that [395, eq. (5.5)] is satisfied. Therefore, [395, eq. (5.6)] ensures that, for any $\varepsilon > 0$,

$$\mathbb{P}_x \left(\sup_{0 \leq s \leq t} |X_s - \phi_s| \leq \varepsilon \right) > 0. \tag{3.77}$$

Moreover, since $\phi_t = y \in \overset{\circ}{A}$ and upon reducing $\varepsilon > 0$ we may assume that $B(y, \varepsilon) \subset A$, where $B(y, \varepsilon)$ denotes the ball of center y and radius $\varepsilon > 0$. Recalling that $f \in B_\kappa^\infty(\mathcal{X})$, we then obtain

$$\begin{aligned} (P_t^f \mathbf{1}_A)(x) &= \mathbb{E}_x \left[\mathbf{1}_{\{X_t \in A\}} e^{\int_0^t f(X_u) du} \right] \geq \mathbb{E}_x \left[\mathbf{1}_{\{\sup_{0 \leq s \leq t} |X_s - \phi_s| \leq \varepsilon\}} e^{-\|f\|_{B_\kappa^\infty} \int_0^t \kappa(X_u) du} \right] \\ &\geq \exp\left(-t\|f\|_{B_\kappa^\infty} \sup_{S_{\phi, \varepsilon}} \kappa\right) \mathbb{P}_x \left(\sup_{0 \leq s \leq t} |X_s - \phi_s| \leq \varepsilon \right), \end{aligned} \tag{3.78}$$

where we denote by $S_{\phi, \varepsilon}$ the ε -tube around the path $(\phi_s)_{s \in [0, t]}$, namely

$$S_{\phi, \varepsilon} = \{x \in \mathcal{X} \mid \exists s \in [0, t] \text{ with } |\phi_s - x| \leq \varepsilon\}.$$

Since $S_{\phi, \varepsilon}$ is a bounded set and κ is continuous over \mathcal{X} , it holds

$$\sup_{S_{\phi, \varepsilon}} \kappa < +\infty.$$

The combination of (3.77) and (3.78) leads to the desired result (3.76). \square

At this stage, we can adapt the spectral analysis developed in Chapter 2 to our situation. However, we trade the minorization condition made there (Assumption 2.2) for the irreducibility granted by Lemma 3.31.

Lemma 3.32. *For any $f \in B_\kappa^\infty(\mathcal{X})$ the operator $\mathcal{L} + f$ considered over $B_W^\infty(\mathcal{X})$ has a real largest eigenvalue $r(f)$ with eigenspace of dimension one, and an associated continuous eigenvector $h_f \in \mathcal{D}_{\mathcal{L}, f}$ such that $h_f(x) > 0$ for any $x \in \mathcal{X}$. Moreover, it holds*

$$(\mathcal{L} + f)h_f = r(f)h_f \quad \text{and} \quad -\frac{\mathcal{L}h_f}{h_f} \in B_\kappa^\infty(\mathcal{X}), \tag{3.79}$$

and h_f is the only positive eigenvector of $\mathcal{L} + f$ (up to multiplication by a positive constant). Finally, $r(f)$ is equal to the cumulant function defined in (3.25):

$$r(f) = \lambda(f). \tag{3.80}$$

The result of Lemma 3.32 is twofold: it entails the well-posedness of the principal eigenproblem associated with $\mathcal{L} + f$ for any $f \in B_\kappa^\infty(\mathcal{X})$, and then identifies this principal eigenvalue with the free energy function (3.25).

Proof. We closely follow the path of Chapter 2 and split the proof into several steps.

Step 1: Compactness of the evolution operator. We first show that, for given $t > 0$ and $f \in B_\kappa^\infty(\mathcal{X})$, the operator P_t^f defined in Lemma 3.28 is compact when considered on $B_W^\infty(\mathcal{X})$. For any $\varepsilon < t/2$ and any compact set $K \subset \mathcal{X}$ we have the decomposition

$$P_t^f = P_{t-2\varepsilon}^f \mathbb{1}_K P_\varepsilon^f \mathbb{1}_K P_\varepsilon^f + P_{t-2\varepsilon}^f \mathbb{1}_{K^c} P_{2\varepsilon}^f + P_{t-2\varepsilon}^f \mathbb{1}_K P_\varepsilon^f \mathbb{1}_{K^c} P_\varepsilon^f. \quad (3.81)$$

We first consider the compact sets K_a from (3.68) for $a > 0$ (omitting the dependence on t in the notation since the time is fixed here) and show that $\mathbb{1}_{K_a^c} P_t^f$ converges to 0 in operator norm as $a \rightarrow +\infty$. Indeed, for any $\varphi \in B_W^\infty(\mathcal{X})$, (3.68) leads to

$$\|\mathbb{1}_{K_a^c} P_t^f \varphi\|_{B_W^\infty} \leq \|\varphi\|_{B_W^\infty} e^{-at}. \quad (3.82)$$

Since for any $a > 0$, there is a compact set K_a such that (3.68) holds, we have, taking the supremum over φ , that $\mathbb{1}_{K_a^c} P_t^f$ tends to 0 in operator norm as $a \rightarrow +\infty$.

We next show that $P_{t-2\varepsilon}^f \mathbb{1}_K P_t^f \mathbb{1}_K$ is compact over $B_W^\infty(\mathcal{X})$ for any compact set $K \subset \mathcal{X}$. Consider a sequence $(\varphi_k)_{k \in \mathbb{N}}$ bounded in $B_W^\infty(\mathcal{X})$. Following the first step of the proof of Lemma 2.5 and using our strong Feller result, Lemma 3.30, we see that $P_{t-2\varepsilon}^f \mathbb{1}_K$ and $P_t^f \mathbb{1}_K$ are strong Feller operators, so $P_{t-2\varepsilon}^f \mathbb{1}_K P_t^f \mathbb{1}_K$ is ultra-Feller (see Lemma 2.33 in Chapter 2). This means that the operator $P_{t-2\varepsilon}^f \mathbb{1}_K P_t^f \mathbb{1}_K$ is continuous in total variation norm, so that the family $(P_{t-2\varepsilon}^f \mathbb{1}_K P_t^f \mathbb{1}_K \varphi_k)_{k \in \mathbb{N}}$ is uniformly equicontinuous. We used here that since $\varphi \in B_W^\infty(\mathcal{X})$ and W is continuous, it holds $\mathbb{1}_K \varphi \in B^\infty(\mathcal{X})$. The sequence $(P_{t-2\varepsilon}^f \mathbb{1}_K P_t^f \mathbb{1}_K \varphi_k)_{k \in \mathbb{N}}$ therefore converges in $B^\infty(\mathcal{X})$ up to extraction by the Ascoli theorem [379, Theorem 11.28], and in $B_W^\infty(\mathcal{X})$ since $W \geq 1$. Therefore, the operator $P_{t-2\varepsilon}^f \mathbb{1}_K P_t^f \mathbb{1}_K$ sends a bounded sequence into a convergent one (up to extraction), so it is compact in $B_W^\infty(\mathcal{X})$ [359]. The decomposition (3.81) and the bound (3.82) then show that P_t^f is the limit in operator norm of the compact operators $P_{t-2\varepsilon}^f \mathbb{1}_{K_a} P_t^f \mathbb{1}_{K_a}$ as $a \rightarrow +\infty$, so it is compact in $B_W^\infty(\mathcal{X})$ (see *e.g.* [359, Theorem VI.12]).

Step 2: Existence of the principal eigenvalue. We can now use the Krein–Rutman theorem on the (closed) total cone $\mathbb{K}_W = \{\varphi \in B_W^\infty \mid \varphi \geq 0\}$ (see [112, 181] for definitions). It is clear that P_t^f leaves this cone invariant. We next show that P_t^f has a non-zero spectral radius

$$R_t(f) = \lim_{n \rightarrow +\infty} \left\| (P_t^f)^n \right\|_{\mathcal{B}(B_W^\infty)}^{\frac{1}{n}}.$$

To this end, fix a compact set K with non-empty interior. We have shown in Lemma 3.31 that

$$\forall x \in K, \quad (P_t^f \mathbb{1}_K)(x) > 0.$$

Since $P_t^f \mathbb{1}_K$ is continuous by Lemma 3.30, this shows that

$$\alpha_K = \min_{x \in K} (P_t^f \mathbb{1}_K)(x) > 0. \quad (3.83)$$

Therefore, for any $x \in K$,

$$\begin{aligned} \left[(P_t^f)^2 \mathbb{1}_K \right] (x) &= \mathbb{E}_x \left[(P_t^f \mathbb{1}_K)(X_t) e^{\int_0^t f(X_s) ds} \right] \geq \mathbb{E}_x \left[\mathbb{1}_K(X_t) (P_t^f \mathbb{1}_K)(X_t) e^{\int_0^t f(X_s) ds} \right] \\ &\geq \alpha_K \mathbb{E}_x \left[\mathbb{1}_K(X_t) e^{\int_0^t f(X_s) ds} \right] = \alpha_K (P_t^f \mathbb{1}_K)(x) \geq \alpha_K^2, \end{aligned}$$

i.e. $\mathbb{1}_K(x) (P_t^f)^2 \mathbb{1}_K(x) \geq \alpha_K^2 \mathbb{1}_K(x)$ for $x \in \mathcal{X}$. Iterating the procedure for any $n \geq 1$ we get

$$\left\| (P_t^f)^n \right\|_{\mathcal{B}(B_W^\infty)} \geq \left\| \mathbb{1}_K (P_t^f)^n \mathbb{1}_K \right\|_{B_W^\infty} \geq \frac{\alpha_K^n}{\sup_K W}.$$

As a result, since $1 \leq \sup_K W < +\infty$, we obtain in the large n limit the following lower bound for the spectral radius:

$$R_t(f) \geq \alpha_K > 0,$$

which shows that $R_t(f)$ is positive. Since P_t^f is compact, [112, Theorem 19.2] ensures that $R_t(f)$ is a real eigenvalue of P_t^f with associated eigenvector $h_f \in \mathbb{K}_W$ (in particular, $h_f \geq 0$). Using the semigroup property of P_t^f and standard arguments (see [351, Theorem 2.4]), we can show that $R_t(f) = e^{r(f)t}$ where $r(f)$ is the largest eigenvalue of $\mathcal{L} + f$, and $h_f \in \mathcal{D}_{\mathcal{L},f}$ satisfies

$$(\mathcal{L} + f)h_f = r(f)h_f,$$

as well as

$$P_t^f h_f = e^{r(f)t} h_f. \quad (3.84)$$

Step 3: Properties of h_f . For the remainder of the proof, we write for simplicity $r = r(f)$ and $h = h_f$ (the function f being fixed). We show here that h is continuous and positive. For any compact $K \subset \mathcal{X}$, (3.84) leads to

$$\begin{aligned} \left| P_t^f(\mathbb{1}_K h) - e^{rt} h \right| &= \left| P_t^f(\mathbb{1}_K h) - P_t^f h \right| = \left| P_t^f(\mathbb{1}_{K^c} h) \right| = \left| P_t^f(\mathbb{1}_{K^c} e^{-rt} P_t^f h) \right| \\ &\leq e^{-rt} \|h\|_{B_W^\infty} \|P_t^f\|_{B(B_W^\infty)} \left| \mathbb{1}_{K^c} P_t^f W \right|. \end{aligned}$$

Using Lemma 3.28 we obtain that, for any $a > 0$, there exists a compact set K_a such that

$$\|e^{-rt} P_t^f(\mathbb{1}_{K_a} h) - h\|_{B_W^\infty} \leq C e^{-at} \quad \text{with} \quad C = e^{-2rt} \|h\|_{B_W^\infty} \|P_t^f\|_{B(B_W^\infty)},$$

so that h is continuous as uniform limit of continuous functions (since $P_t^f(\mathbb{1}_K h)$ is continuous by Lemma 3.30). Finally, since $h \geq 0$ and h is not identically equal to 0, there exists $x_0 \in \mathcal{X}$ such that $h(x_0) > 0$. Moreover h is continuous, so there is $\varepsilon > 0$ for which $h > 0$ on $B(x_0, \varepsilon)$. By (3.84) it holds, for any $x \in \mathcal{X}$,

$$e^{rt} h(x) = (P_t^f h)(x) \geq P_t^f(h \mathbb{1}_{B(x_0, \varepsilon)})(x) \geq \left(\inf_{B(x_0, \varepsilon)} h \right) (P_t^f \mathbb{1}_{B(x_0, \varepsilon)})(x).$$

Since $h > 0$ on $B(x_0, \varepsilon)$ and h is continuous, $\inf_{B(x_0, \varepsilon)} h > 0$. Thus, Lemma 3.31 combined with the previous lower bound shows that $(P_t^f \mathbb{1}_{B(x_0, \varepsilon)})(x) > 0$ for any $x \in \mathcal{X}$, which allows to conclude that $h(x) > 0$ for all $x \in \mathcal{X}$.

Step 4: Properties of eigenspaces and eigenfunctions. We now show that the eigenspace associated with h is of dimension one, and that any other eigenvector vanishes somewhere in \mathcal{X} . For this, we introduce the so called h -transform [275, 377, 97, 181], just like in Chapter 2. A key element here is the fact that $h(x) > 0$ for all $x \in \mathcal{X}$, which allows to define the following Markov operator, for an arbitrary time $t > 0$:

$$Q_h \varphi = e^{-rt} h^{-1} P_t^f(h\varphi),$$

where h, h^{-1} refer here to the multiplication operators by the functions h and h^{-1} respectively. We now prove that Q_h is ergodic by first noting that Q_h admits Wh^{-1} as a Lyapunov function (using (3.68) and the normalization $\|h\|_{B_W^\infty} = 1$ which implies that $Wh^{-1} \geq 1$). Moreover, we can prove that Q_h satisfies a minorization condition on any compact set. For this, consider $K \subset \mathcal{X}$ compact with non-empty interior and denote by η_K the uniform Lebesgue measure on K . By (3.83), for any $t > 0$ there is $\alpha_K > 0$ such that, for any measurable set $A \subset \mathcal{X}$,

$$\forall x \in K, \quad (P_t^f \mathbb{1}_A)(x) \geq (P_t^f \mathbb{1}_{K \cap A})(x) \geq \alpha_K \eta_K(A).$$

Since h is continuous, this implies that, for any measurable $\varphi \geq 0$,

$$\forall x \in K, \quad (Q_h \varphi)(x) \geq \frac{\alpha_K \min_K h}{\max_K h} \eta_K(\varphi),$$

where both the minimum and maximum above are finite and non-zero. This shows that Q_h satisfies a minorization condition [219] over any compact set. Using Assumption 3.6, we can also show that Wh^{-1}

has compact level sets, see [181, Appendix E] for details. Therefore, the Markovian dynamics with kernel Q_h admits a unique invariant probability measure μ_h , with respect to which it is ergodic in $B_{W^h}^\infty(\mathcal{X})$. By this we mean that [219, Theorem 1.2] there exist $\bar{\alpha} > 0$ and $C > 0$ such that for any $\varphi \in B_{W^h}^\infty(\mathcal{X})$,

$$\forall n \geq 1, \quad \|(Q_h)^n \varphi - \mu_h(\varphi)\|_{B_{W^h}^\infty} \leq C e^{-\bar{\alpha}n} \|\varphi - \mu_h(\varphi)\|_{B_{W^h}^\infty}, \quad (3.85)$$

and it holds $\mu_h(W/h) < +\infty$.

We can now use this ergodic behaviour to show that the eigenspace associated with r has dimension one and that P_t^f cannot have another positive eigenvector with norm 1 in $B_W^\infty(\mathcal{X})$. Indeed, if there were another eigenvector $\tilde{h} \in B_W^\infty(\mathcal{X})$ associated with r , then $\tilde{h}/h \in B_{W^h}^\infty(\mathcal{X})$, and (3.85) ensures that

$$(Q_h)^n \begin{pmatrix} \tilde{h} \\ h \end{pmatrix} = \frac{\tilde{h}}{h} \xrightarrow{n \rightarrow +\infty} \mu_h \begin{pmatrix} \tilde{h} \\ h \end{pmatrix}.$$

This shows that h and \tilde{h} would be proportional, and answers the claim that the eigenspace associated with r has dimension 1. Assume now that there is another eigenvalue $\tilde{r} < r$ with eigenvector $\tilde{h} \in B_W^\infty(\mathcal{X})$ such that $\tilde{h}(x) > 0$ for all $x \in \mathcal{X}$. Noting again that $\tilde{h}/h \in B_{W^h}^\infty(\mathcal{X})$ and since $\tilde{h} > 0$, (3.85) shows that

$$(Q_h)^n \begin{pmatrix} \tilde{h} \\ h \end{pmatrix} \xrightarrow{n \rightarrow +\infty} \mu_h \begin{pmatrix} \tilde{h} \\ h \end{pmatrix} > 0. \quad (3.86)$$

However it now holds, for any $x \in \mathcal{X}$,

$$(Q_h)^n \begin{pmatrix} \tilde{h} \\ h \end{pmatrix} (x) = e^{(\tilde{r}-r)tn} \frac{\tilde{h}}{h}(x) \xrightarrow{n \rightarrow +\infty} 0,$$

where we used that $h > 0$ and $\tilde{r} < r$. Combining the two above equations shows that

$$\mu_h \begin{pmatrix} \tilde{h} \\ h \end{pmatrix} = 0,$$

which contradicts (3.86). As a result, there cannot be another eigenvalue with a positive eigenvector.

Step 5: The principal eigenvalue is the cumulant function. Proving (3.80) now follows by a simple rewriting. For $x \in \mathcal{X}$ and $t_0 > 0$ fixed, it holds for any $n \in \mathbb{N}^*$,

$$\mathbb{E}_x \left[e^{\int_0^{nt_0} f(X_s) ds} \right] = \left[(P_{t_0}^f)^n \mathbf{1} \right] (x) = e^{rnt_0} [h(Q_h)^n h^{-1}](x),$$

so that

$$\frac{1}{nt_0} \log \mathbb{E}_x \left[e^{\int_0^{nt_0} f(X_s) ds} \right] = \frac{1}{nt_0} \log [e^{rnt_0} h(Q_h)^n h^{-1}(x)] = r + \frac{1}{nt_0} \log [h(Q_h)^n h^{-1}(x)].$$

By (3.85) (since $h^{-1} \in B_{W^h}^\infty(\mathcal{X})$), we see that $h(Q_h)^n h^{-1}(x)$ converges to a constant, so that

$$r(f) = \lim_{n \rightarrow +\infty} \frac{1}{nt_0} \log \mathbb{E}_x \left[e^{\int_0^{nt_0} f(X_s) ds} \right].$$

We have chosen to work with an arbitrary time $t_0 > 0$ for convenience, so a priori the above limit depends on t_0 . Showing that the limit actually does not depend on this t_0 and that

$$r(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right]$$

follows by standard arguments not reproduced here (see *e.g.* [239, 181]), which concludes the proof. \square

An important ingredient for the lower bound of the LDP is the Gateau-differentiability of the cumulant functional, which we prove below.

Lemma 3.33. *The functional*

$$f \in B_\kappa^\infty(\mathcal{X}) \mapsto \lambda(f) = \lim_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right] \quad (3.87)$$

is convex and Gateau-differentiable.

Proof. The convexity of λ is a standard consequence of Hölder's inequality. Concerning Gateau-differentiability, we follow the strategy of [199, Section 3] for a compact state space, relying on results of Kato [263]. For this, we interpret the cumulant function (3.87) as the largest eigenvalue of the tilted generator $r(f)$ as defined in Lemma 3.32. More precisely, for $f, g \in B_\kappa^\infty(\mathcal{X})$ and $\alpha \in \mathbb{R}$, $\lambda(f + \alpha g)$ is associated with the largest eigenvalue of the operator $P_t^{f+\alpha g}$ in $B_W^\infty(\mathcal{X})$ through

$$P_t^{f+\alpha g} h_{f+\alpha g} = e^{t\lambda(f+\alpha g)} h_{f+\alpha g},$$

so that derivability in α can be shown through the differentiability of the spectrum of a bounded operator. We thus show that the operator $P_t^{f+\alpha g}$ is differentiable in operator norm. To this end, we fix $C > 0$ and prove that for $|\alpha| \leq C$ it holds

$$P_t^{f+\alpha g} \varphi = P_t^f \varphi + \alpha Q_t^{f,g,(1)} \varphi + \alpha^2 Q_t^{f,g,(2),\alpha_0}, \quad (3.88)$$

where $Q_t^{f,g,(1)}$ and $Q_t^{f,g,(2),\alpha_0}$ are bounded operators on $B_W^\infty(\mathcal{X})$ and $|\alpha_0| \leq C$.

For this we first define

$$Q_t^{f,g,(1)} : \varphi \in B_W^\infty(\mathcal{X}) \mapsto \mathbb{E}_x \left[\varphi(X_t) \left(\int_0^t g(X_s) ds \right) e^{\int_0^t f(X_s) ds} \right].$$

This operator is bounded in $B_W^\infty(\mathcal{X})$ by the same martingale estimate used to prove Lemma 3.29. In the same way, the second order operator reads

$$Q_t^{f,g,(2),\alpha} : \varphi \in B_W^\infty(\mathcal{X}) \mapsto \frac{1}{2} \mathbb{E}_x \left[\varphi(X_t) \left(\int_0^t g(X_s) ds \right)^2 e^{\int_0^t (f(X_s) + \alpha g(X_s)) ds} \right].$$

This operator is also bounded in $B_W^\infty(\mathcal{X})$ since, for $a \geq 0$ it holds $a^2/2 \leq e^a$ so that, for $\varphi \in B_W^\infty(\mathcal{X})$ and $x \in \mathcal{X}$ we obtain

$$\begin{aligned} |Q_t^{f,g,(2),\alpha} \varphi(x)| &\leq \|\varphi\|_{B_W^\infty} \|g\|_{B_\kappa^\infty}^2 \mathbb{E}_x \left[W(X_t) \frac{1}{2} \left(\int_0^t \kappa(X_s) ds \right)^2 e^{(\|f\|_{B_\kappa^\infty} + |\alpha| \|g\|_{B_\kappa^\infty}) \int_0^t \kappa(X_s) ds} \right] \\ &\leq \|\varphi\|_{B_W^\infty} \|g\|_{B_\kappa^\infty}^2 \mathbb{E}_x \left[W(X_t) e^{(1+\|f\|_{B_\kappa^\infty} + |\alpha| \|g\|_{B_\kappa^\infty}) \int_0^t \kappa(X_s) ds} \right] \\ &\leq \|\varphi\|_{B_W^\infty} \|g\|_{B_\kappa^\infty}^2 e^{ct} W(x), \end{aligned}$$

for some constant $c > 0$ depending on $\|f\|_{B_\kappa^\infty}$, $\|g\|_{B_\kappa^\infty}$ and α . Next, it suffices to note that

$$Q_t^{f,g,(1)} = \frac{d}{d\alpha} P_t^{f+\alpha g}, \quad Q_t^{f,g,(2),\alpha} = \frac{1}{2} \frac{d^2}{d\alpha^2} P_t^{f+\alpha g},$$

to obtain (3.88) through a Taylor expansion, where $|\alpha_0| \leq C$.

This shows that $\alpha \mapsto P_t^{f+\alpha g}$ is differentiable in operator norm. Thus, the principal eigenvalue $\lambda(f + \alpha g)$, which is always isolated, is differentiable, see [263, Chapter II, Theorem 5.4] and [263, Chapter IV, Theorem 3.5]. This concludes the proof of Gateau-differentiability. \square

Remark 3.34. *By pursuing further the Taylor expansion (3.88) in the proof of Lemma 3.33, we can actually show that, for any $f, g \in B_\kappa^\infty(\mathcal{X})$, the function*

$$\alpha \in \mathbb{C} \mapsto \lambda(f + \alpha g)$$

is analytic. This relies on the simple inequality $a^n/n! \leq e^a$ for any $a \geq 0$ together with the series expansion of the exponential and martingale estimates as in the proof of Lemma 3.33. This analyticity was already proven in [275] using a different argument that can be simplified with our tools. Indeed, our proof, based on martingales, shows that for any $t > 0$, the function

$$\alpha \mapsto \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t (f(X_s) + \alpha g(X_s)) ds} \right]$$

is analytic. Moreover, it is finite on \mathbb{R} and converges pointwise to a finite valued function as $t \rightarrow +\infty$, as shown in Lemma 3.32. Therefore, the convergence holds uniformly on any compact as $t \rightarrow +\infty$ (see [166, Theorem VI.3.3]). Now a locally uniform limit of analytic functions is analytic (see [379, Theorem 10.28]). Therefore, $\alpha \mapsto \lambda(f + \alpha g)$ is analytic.

The last step before proving the large deviations principle itself is an exponential tightness result, see [119, Section 1.2]. At this stage, the finiteness of $\lambda(f)$ together with the Gateau-differentiability of $f \in B_\kappa^\infty(\mathcal{X}) \mapsto \lambda(f)$ already provides the upper bound over compact sets and the lower bound in (3.27). In order to extend the upper bound to all closed sets, we prove exponential tightness in the τ^κ topology, see Section 1.2.1 for definitions (note that this exponential tightness is not explicitly stated in [275]).

Lemma 3.35. *The family of probability measures $t \mapsto \mathbb{P}_x(L_t \in \cdot)$ over $(\mathcal{P}_\kappa(\mathcal{X}), \tau^\kappa)$ is exponentially tight.*

Proof. We adapt the strategy of [427, Corollary 2.3] and [424, Section 2.2] by introducing the family of sets

$$\Gamma_N = \{\nu \in \mathcal{P}(\mathcal{X}) \mid \nu(\Psi) \leq N\}, \quad N > 0.$$

For $N > 0$, the sets Γ_N are subsets of $\mathcal{P}_\kappa(\mathcal{X})$ since $\kappa \ll \Psi$. We show that they are actually precompact in the τ^κ -topology.

Let us first show that Γ_N is precompact in the usual weak topology for any $N > 0$. Consider for this the compact sets $K_\beta = \{x \in \mathcal{X} \mid \Psi(x) \leq \beta\} \subset \mathcal{X}$ for $\beta > 0$ (recall that Ψ has compact level sets). Then, for any $\nu \in \Gamma_N$, we have

$$\beta\nu(K_\beta^c) + \nu(\Psi\mathbf{1}_{K_\beta}) \leq \nu(\Psi\mathbf{1}_{K_\beta^c}) + \nu(\Psi\mathbf{1}_{K_\beta}) = \nu(\Psi) \leq N.$$

This shows that for any $\beta > 0$ and any $\nu \in \Gamma_N$,

$$\nu(K_\beta^c) \leq \frac{N}{\beta},$$

hence for any $N > 0$ the family of measures Γ_N is tight, so it is precompact for the weak topology by the Prohorov theorem [42]. Now, if κ is bounded, Γ_N is tight for the τ^κ -topology and the theorem is shown, so we may assume that κ has compact level sets (see Assumption 3.6). For proving compactness in our finer topology, we show that κ is uniformly integrable over Γ_N in order to use [419, Theorem 7.12]. Since $\kappa \ll \Psi$, the set

$$A_n = \left\{ x \in \mathcal{X} \mid \frac{\Psi(x)}{\kappa(x)} \leq n \right\}$$

is compact for any $n \geq 1$. Moreover, since we assume κ to be continuous with compact level sets, for any $n \geq 1$ there exists $m_n \geq n$ such that

$$\left\{ \frac{\Psi}{\kappa} \leq n \right\} \subset \{\kappa \leq m_n\},$$

with $m_n \rightarrow +\infty$ when $n \rightarrow +\infty$. Therefore, for any $\nu \in \Gamma_N$ and $n \geq 1$,

$$\int_{\{\kappa > m_n\}} \kappa \, d\nu \leq \int_{A_n^c} \kappa \, d\nu = \frac{1}{n} \int_{A_n^c} n\kappa \, d\nu \leq \frac{1}{n} \int_{\mathcal{X}} \Psi \, d\nu = \frac{1}{n} \nu(\Psi) \leq \frac{N}{n}.$$

Taking the supremum over $\nu \in \Gamma_N$ in the above equation and recalling that $m_n \rightarrow +\infty$ when $n \rightarrow +\infty$ we obtain

$$\lim_{m \rightarrow +\infty} \sup_{\nu \in \Gamma_N} \int_{\{\kappa > m\}} \kappa \, d\nu = 0. \quad (3.89)$$

We can then conclude that Γ_N is compact for the τ^κ -topology. Consider indeed a sequence $(\nu_n)_{n \in \mathbb{N}} \subset \Gamma_N$. By Prohorov's theorem, $(\nu_n)_{n \in \mathbb{N}}$ has a subsequence weakly converging towards a measure ν , i.e. $\nu_n(\varphi) \rightarrow \nu(\varphi)$ for any $\varphi \in C_b(\mathcal{X})$. Then, by [419, Theorem 7.12], (3.89) ensures that $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ and for any $f \in B_\kappa^\infty(\mathcal{X})$, $\nu_n(f) \rightarrow \nu(f)$ as $n \rightarrow +\infty$. In other words, Γ_N is precompact for the τ^κ -topology.

We can now prove the τ^κ -exponential tightness of the empirical distribution $(L_t)_{t \geq 0}$ in $\mathcal{P}(\mathcal{X})$. Indeed, for any $N, t > 0$, Techebytchev's inequality leads to

$$\begin{aligned} \mathbb{P}_x(L_t \in \Gamma_N^c) &= \mathbb{P}_x\left(\int_0^t \Psi(X_s) ds > Nt\right) = \mathbb{P}_x\left(\int_0^t \Psi(X_s) ds > Nt\right) \\ &\leq e^{-Nt} \mathbb{E}_x \left[e^{\int_0^t \Psi(X_s) ds} \right] = e^{-Nt} P_t^\Psi \mathbf{1}(x). \end{aligned}$$

Renormalizing at log scale and using (3.25) leads to

$$\overline{\lim}_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}_x(L_t \in \Gamma_N^c) \leq -N + \lambda(\Psi). \quad (3.90)$$

The right hand side of the above quantity may look infinite since Ψ grows faster than κ . However, using again the martingale M_t defined in Lemma 3.27 we obtain, for any $t > 0$,

$$\mathbb{E}_x \left[e^{\int_0^t \Psi(X_s) ds} \right] \leq \mathbb{E}_x \left[W(X_t) e^{-\int_0^t \frac{\kappa W}{W}(X_s) ds} \right] = \mathbb{E}_x(M_t) \leq W(x).$$

Thus it holds

$$\lambda(\Psi) \leq \overline{\lim}_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[W(X_t) e^{\int_0^t \Psi(X_s) ds} \right] \leq \overline{\lim}_{t \rightarrow +\infty} \frac{1}{t} \log W(x) \leq 0.$$

As a result, (3.90) becomes

$$\overline{\lim}_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{P}_x(L_t \in \Gamma_N^c) \leq -N.$$

Since Γ_N is precompact in the τ^κ -topology for any $N > 0$, and N can be chosen arbitrarily large, this proves the exponential tightness of the family of empirical distributions in the τ^κ -topology. \square

We are now in position to prove Theorem 3.10.

Proof of Theorem 3.10. We assemble the previous lemmas to check that the assumptions of the Gärtner–Ellis theorem reminded in Section 1.2.2 are fulfilled. For this, we note that the cumulant function

$$\lambda : f \in B_\kappa^\infty(\mathcal{X}) \mapsto \overline{\lim}_{t \rightarrow +\infty} \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f(X_s) ds} \right]$$

can be identified to the function Λ in Theorem 1.10. The topological dual of $(\mathcal{M}_\kappa(\mathcal{X}), \tau^\kappa)$ is $B_\kappa^\infty(\mathcal{X})$, where $\mathcal{M}_\kappa(\mathcal{X})$ is the set of measures over \mathcal{X} integrating κ (see [379, 275] and [124, Lemma 3.3.8] for details). We have proved that λ is well defined, Gateau differentiable, and that the family of measures

$$t \rightarrow \pi_t(\cdot) = \mathbb{P}(L_t \in \cdot),$$

is exponentially tight in the τ^κ -topology. Therefore, π_t satisfies a large deviations principle in the τ^κ -topology with good rate function given by

$$\forall \nu \in \mathcal{M}(\mathcal{X}), \quad I(\nu) = \sup_{f \in B_\kappa^\infty} \{ \nu(f) - \lambda(f) \}. \quad (3.91)$$

We observe that $I(\nu) = +\infty$ if ν is not normalized to 1 (take f to be constant in the supremum (3.91)), so we may consider I over $\mathcal{P}(\mathcal{X})$. Moreover, choosing $f = \kappa$ in (3.91) and noting that $\lambda(\kappa) < +\infty$, we get $I(\nu) = +\infty$ if $\nu \notin \mathcal{P}_\kappa(\mathcal{X})$. If ν is not absolutely continuous with respect to μ , there exists $A \subset \mathcal{X}$ such that $\mu(A) = 0$ and $\nu(A) > 0$. Since μ has a positive Lebesgue density, this means that A has zero Lebesgue measure. Consider then $f_a = a \mathbf{1}_A \in B_\kappa^\infty(\mathcal{X})$ for $a \in \mathbb{R}$. Since A has zero Lebesgue measure and $(X_t)_{t \geq 0}$ has a smooth density for all $t > 0$ (as a consequence of Assumption 3.4) it holds, for all $t > 0$,

$$\mathbb{E}_x[f_a(X_t)] = a \mathbb{P}_x(X_t \in A) = 0.$$

Therefore, the process

$$Z_t = \int_0^t f_a(X_s) ds,$$

satisfies $\mathbb{E}_x[Z_t] = 0$ for all $t > 0$. Since $Z_t \geq 0$, it holds $Z_t = 0$ almost surely, for any $t > 0$. As a consequence we obtain

$$\forall t > 0, \quad \frac{1}{t} \log \mathbb{E}_x \left[e^{\int_0^t f_a(X_s) ds} \right] = \frac{1}{t} \log \mathbb{E}_x \left[e^{Z_t} \right] = 0.$$

This shows that $\lambda(f_a) = 0$, so that from (3.91) we obtain

$$I(\nu) \geq a\nu(A),$$

with $\nu(A) > 0$. By letting $a \rightarrow +\infty$ we are led to $I(\nu) = +\infty$.

Finally, we show that $I(\nu) = 0$ if and only if $\nu = \mu$, and that $(L_t)_{t \geq 0}$ converges almost surely to μ in the τ^κ -topology (see [119, Appendix B] for definitions). For this, we introduce the set of minimizers

$$\mathcal{I} = \left\{ \nu \in \mathcal{P}(\mathcal{X}) \mid I(\nu) = \inf_{\mathcal{P}(\mathcal{X})} I \right\}.$$

Since I has compact level sets, \mathcal{I} is a non-empty closed subset of $\mathcal{P}(\mathcal{X})$ for the τ^κ -topology. Moreover, in order for the LDP upper bound to make sense, it holds $\inf_{\mathcal{P}(\mathcal{X})} I = 0$. If \mathcal{I}_δ denotes an open neighborhood of \mathcal{I} in the τ^κ -topology, the lower semicontinuity of I implies that

$$\inf_{\mathcal{I}_\delta} I > 0.$$

Therefore, by the large deviations upper bound,

$$\mathbb{P}_x(L_t \notin \mathcal{I}_\delta) = \mathbb{P}_x(L_t \in \mathcal{I}_\delta^c) \leq C \exp\left(-t \inf_{\mathcal{I}_\delta^c} I\right), \quad (3.92)$$

for some constant $C > 0$. We next introduce the random subset of \mathbb{R}_+ of indices $t \geq 0$ for which L_t does not belong to \mathcal{I}_δ , namely $T = \{t \geq 0 \mid L_t \notin \mathcal{I}_\delta\}$. Since

$$\mathbb{P}_x(L_t \notin \mathcal{I}_\delta) = \mathbb{E}_x[\mathbb{1}_{\{L_t \notin \mathcal{I}_\delta\}}],$$

we have, by the Fubini theorem, for any $t > 0$,

$$\int_0^t \mathbb{P}_x(L_s \notin \mathcal{I}_\delta) ds = \mathbb{E}_x \left[\int_0^t \mathbb{1}_{\{L_s \notin \mathcal{I}_\delta\}} ds \right] = \mathbb{E}_x[|\min(T, t)|].$$

By using (3.92) and the dominated convergence theorem, we obtain

$$\mathbb{E}_x[|T|] = \int_0^{+\infty} \mathbb{P}_x(L_t \notin \mathcal{I}_\delta) dt < +\infty.$$

As a result, $|T| < +\infty$ almost surely. This means that, for any neighborhood \mathcal{I}_δ of \mathcal{I} in the τ^κ -topology, the empirical measure $(L_t)_{t \geq 0}$ almost surely spends a finite Lebesgue measure time outside \mathcal{I}_δ . In other words, $(L_t)_{t \geq 0}$ converges almost surely in the τ^κ -topology to the set \mathcal{I} . However, we know by Proposition 3.9 that the only possible limit for $(L_t)_{t \geq 0}$ is μ , hence $\mathcal{I} = \{\mu\}$ and $(L_t)_{t \geq 0}$ almost surely converges to μ in the τ^κ -topology, which concludes the proof. \square

3.6.2 Proofs of Section 3.3

3.6.2.1 Proof of Proposition 3.13

For the proof, which is partly inspired by [124, Lemma 4.1.36], we denote by I_F the rate function given by the Fenchel transform in (3.26) and I_V for the Varadhan functional on the right hand side of (3.33). We repeatedly use the results of Lemma 3.32 in what follows.

We first show that $I_V(\nu) = +\infty$ if ν is not absolutely continuous with respect to μ or does not belong to $\mathcal{P}_\kappa(\mathcal{X})$. Assume first that $\nu \ll \mu$ does not hold: there exists a set $A \subset \mathcal{X}$ such that $\nu(A) > 0$ and $\mu(A) = 0$. For any $a \in \mathbb{R}$ we introduce $f_a = a\mathbb{1}_A$ and denote by h_a the eigenvector associated to the principal eigenvalue $\lambda(f_a)$ (which is continuous and positive by Lemma 3.32). As shown in the proof of Theorem 3.10, it holds $\lambda(f_a) = 0$, so

$$\forall a \in \mathbb{R}, \quad (\mathcal{L} + f_a)h_a = 0.$$

This leads to

$$-\frac{\mathcal{L}h_a}{h_a} = a\mathbb{1}_A.$$

Therefore, since $h_a \in B_W^\infty(\mathcal{X})$, it holds $h_a \in \mathcal{D}^+$ and

$$I_V(\nu) \geq \int_{\mathcal{X}} -\frac{\mathcal{L}h_a}{h_a} d\nu = a\nu(A) > 0.$$

By letting $a \rightarrow +\infty$, we conclude that $I_V(\nu) = +\infty$ when ν is not absolutely continuous with respect to μ . Next, if $\nu \notin \mathcal{P}_\kappa(\mathcal{X})$, since $\kappa \geq 1$ it holds $\nu(\kappa) = +\infty$. We may then choose $f = \kappa \in B_\kappa^\infty(\mathcal{X})$. By Lemma 3.32 the principal eigenvector h_κ of $\mathcal{L} + \kappa$ belongs to \mathcal{D}^+ with $\lambda(\kappa) < +\infty$, so we have

$$I_V(\nu) \geq \int_{\mathcal{X}} -\frac{\mathcal{L}h_\kappa}{h_\kappa} d\nu = \int_{\mathcal{X}} \kappa d\nu - \lambda(\kappa) = +\infty,$$

i.e. $I_V(\nu) = +\infty$ if $\nu \notin \mathcal{P}_\kappa(\mathcal{X})$. This shows that $I_F(\nu) = I_V(\nu)$ when ν is not absolutely continuous with respect to μ or $\nu \notin \mathcal{P}_\kappa(\mathcal{X})$. We next show that $I_F = I_V$ when $\nu \ll \mu$ and $\nu \in \mathcal{P}_\kappa(\mathcal{X})$, which we assume until the end of the proof.

Let us first show that $I_F \geq I_V$. For this, we consider $u \in \mathcal{D}^+$ and introduce

$$f_u = -\frac{\mathcal{L}u}{u}.$$

Because of the definition (3.34) of \mathcal{D}^+ , we know that $f_u \in B_\kappa^\infty(\mathcal{X})$, so we can write, since $\nu \in \mathcal{P}_\kappa(\mathcal{X})$,

$$I_F(\nu) \geq \nu(f_u) - \lambda(f_u). \quad (3.93)$$

Moreover,

$$(\mathcal{L} + f_u)u = \mathcal{L}u - \left(\frac{\mathcal{L}u}{u}\right)u = 0. \quad (3.94)$$

As a result, $u > 0$ is an eigenvector of $\mathcal{L} + f_u$ associated with the eigenvalue 0 (and hence it is an eigenvector of $P_t^{f_u}$ with eigenvalue 1). But we know from Lemma 3.32 that a positive eigenvector can only be associated with the principal eigenvalue $\lambda(f_u)$, so that $\lambda(f_u) = 0$ by (3.94). Therefore, (3.93) leads to

$$I_F(\nu) \geq \nu(f_u) - \lambda(f_u) = \int_{\mathcal{X}} -\frac{\mathcal{L}u}{u} d\nu.$$

Since $u \in \mathcal{D}^+$ is arbitrary, taking the supremum shows that $I_F(\nu) \geq I_V(\nu)$ for any $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ with $\nu \ll \mu$.

We now turn to the inequality $I_F \leq I_V$. We again draw elements from [124, Lemma 4.1.36], but we use simpler arguments based on the spectral analysis of the operator $\mathcal{L} + f$. Consider for any arbitrary $f \in B_\kappa^\infty(\mathcal{X})$ the associated eigenvector $h_f \in B_W^\infty(\mathcal{X})$ defined in Lemma 3.32. It then holds:

- $h_f \in B_W^\infty(\mathcal{X})$;
- $h_f > 0$;
- by Lemma 3.32, $-\frac{\mathcal{L}h_f}{h_f} = f - \lambda(f) \in B_\kappa^\infty(\mathcal{X})$;

Thus $h_f \in \mathcal{D}^+$. As a result we have, since $\nu \in \mathcal{P}_\kappa(\mathcal{X})$,

$$I_V(\nu) \geq \int_{\mathcal{X}} -\frac{\mathcal{L}h_f}{h_f} d\nu = \nu(f) - \lambda(f).$$

Given that, in the above equation, f is an arbitrary function belonging to $B_\kappa^\infty(\mathcal{X})$, taking the supremum leads to

$$I_V(\nu) \geq \sup_{f \in B_\kappa^\infty} \{\nu(f) - \lambda(f)\}.$$

This finally shows that $I_F(\nu) = I_V(\nu)$ for all $\nu \in \mathcal{P}_\kappa(\mathcal{X})$ with $\nu \ll \mu$, which concludes the proof.

3.6.2.2 Proof of Corollary 3.14

Since I is the Fenchel transform of λ , the result follows if we can show that λ defined on $B_\kappa^\infty(\mathcal{X})$ is stable by bi-Fenchel conjugacy. Moreover, the convexity and finiteness of λ shows that a (necessary and) sufficient condition for λ to be bi-Fenchel stable is for the functional $f \mapsto \lambda(f)$ to be lower-semicontinuous (see [26, Theorem 2.22]). We show below that it is actually continuous: for any sequence $(f_n)_{n \geq 0}$ in $B_\kappa^\infty(\mathcal{X})$ such that $\|f_n - f\|_{B_\kappa^\infty} \rightarrow 0$ for some $f \in B_\kappa^\infty(\mathcal{X})$, it holds $\lambda(f_n) \rightarrow \lambda(f)$. We shall use for this a stability result from [92].

Consider a sequence $(f_n)_{n \geq 0}$ converging to f in $B_\kappa^\infty(\mathcal{X})$. Using Lemma 3.29, for any $\varphi \in B_W^\infty(\mathcal{X})$, $t > 0$, $x \in \mathcal{X}$ and $n \in \mathbb{N}$, it holds (using again the inequality $a \leq e^a$ for $a \geq 0$)

$$\begin{aligned} \left| (P_t^f \varphi)(x) - (P_t^{f_n} \varphi)(x) \right| &\leq \|\varphi\|_{B_W^\infty} \mathbb{E}_x \left[W(X_t) \int_0^t |f(X_s) - f_n(X_s)| ds e^{(\|f\|_{B_\kappa^\infty} + \|f_n\|_{B_\kappa^\infty}) \int_0^t \kappa(X_s) ds} \right] \\ &\leq \|\varphi\|_{B_W^\infty} \|f - f_n\|_{B_\kappa^\infty} \mathbb{E}_x \left[W(X_t) e^{(1 + \|f\|_{B_\kappa^\infty} + \|f_n\|_{B_\kappa^\infty}) \int_0^t \kappa(X_s) ds} \right] \\ &\leq C \|\varphi\|_{B_W^\infty} \|f - f_n\|_{B_\kappa^\infty} \mathbb{E}_x [M_t] \\ &\leq C \|\varphi\|_{B_W^\infty} \|f - f_n\|_{B_\kappa^\infty} W(x), \end{aligned}$$

for some constant $C > 0$ depending on $\|f\|_{B_\kappa^\infty}$ and $t > 0$. We used Lemma 3.27 and the supermartingale property of M_t to obtain the last line. This leads to

$$\|P_t^f - P_t^{f_n}\|_{\mathcal{B}(B_W^\infty)} \leq C \|f - f_n\|_{B_\kappa^\infty} \xrightarrow{n \rightarrow +\infty} 0. \quad (3.95)$$

We know by Lemma 3.32 that $\lambda(f)$ and $\lambda(f_n)$ are associated with the isolated largest eigenvalue of the operators P_t^f and $P_t^{f_n}$ respectively. Therefore, (3.95) shows that the approximation is strongly stable (we refer to [92], in particular the definition in Section 2.2 and Proposition 2.11), and [92, Proposition 2.2] ensures that $\lambda(f_n) \rightarrow \lambda(f)$ as $n \rightarrow +\infty$. This shows that the function $\lambda : B_\kappa^\infty(\mathcal{X}) \rightarrow \mathbb{R}$ is continuous and concludes the proof.

3.6.2.3 Proof of Theorem 3.16

The proof, inspired by [49], relies on two ideas: performing a Witten transform inside the variational representation (3.33) and separating the symmetric and antisymmetric parts of the generator \mathcal{L} . We write $d\nu = \rho d\mu = e^v d\mu$ and consider first that $v \in C_c^\infty(\mathcal{X})$ instead of $\mathcal{H}^1(\nu)$. Starting from (3.33), we consider a function u of the form

$$u = e^{\frac{v}{2}} \sqrt{\rho}, \quad \psi \in C_c^\infty(\mathcal{X}). \quad (3.96)$$

We call this choice ‘‘variational Witten transform’’ for its similarity with the standard Witten transform [425, 234, 304] and its use in the variational formula defining I in (3.33). Since $u = e^{\frac{v}{2} + \frac{v}{2}}$ with $v, \psi \in C_c^\infty(\mathcal{X})$ it is clear that $u \in \mathcal{D}^+$. This follows by noting that, using the shorthand notation $w = \psi/2 + v/2 \in C_c^\infty(\mathcal{X})$, we have

$$-\frac{\mathcal{L}u}{u} = -e^{-w} \mathcal{L}e^w = -\mathcal{L}w - \frac{1}{2} |\sigma^T \nabla w|^2 \in C_c^\infty(\mathcal{X}) \subset B_\kappa^\infty(\mathcal{X}).$$

Moreover, it holds $u = e^w > 0$ and u is constant outside a compact set, so $u \in B_W^\infty(\mathcal{X})$ and it holds $u \in \mathcal{D}^+$.

We now rewrite the expression in (3.33) for u given by (3.96), using again the notation $w = \psi/2 + v/2$:

$$-\int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu = -\int_{\mathcal{X}} \mathcal{L}w d\nu - \frac{1}{2} \int_{\mathcal{X}} |\sigma^T \nabla w|^2 d\nu.$$

Recalling that $S = \sigma\sigma^T/2$ and expanding $w = \psi/2 + v/2$, we obtain

$$-\int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu = -\frac{1}{2} \int_{\mathcal{X}} \mathcal{L}\psi d\nu - \frac{1}{2} \int_{\mathcal{X}} \mathcal{L}v d\nu - \frac{1}{4} \int_{\mathcal{X}} \nabla\psi \cdot S \nabla\psi d\nu - \frac{1}{2} \int_{\mathcal{X}} \nabla v \cdot S \nabla\psi d\nu - \frac{1}{4} \int_{\mathcal{X}} \nabla v \cdot S \nabla v d\nu. \quad (3.97)$$

We now decompose \mathcal{L} into symmetric and antisymmetric parts. First, it holds

$$-\frac{1}{2} \int_{\mathcal{X}} \mathcal{L}\psi d\nu = -\frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_S \psi) e^v d\mu - \frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_A \psi) d\nu = \frac{1}{2} \int_{\mathcal{X}} \nabla\psi \cdot S \nabla v d\nu - \frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_A \psi) d\nu. \quad (3.98)$$

On the other hand, using that \mathcal{L}_A is a first order differential operator satisfying $\mathcal{L}_A^* \mathbf{1} = 0$, we obtain

$$\int_{\mathcal{X}} (\mathcal{L}_A v) e^v d\mu = \int_{\mathcal{X}} (\mathcal{L}_A e^v) d\mu = \int_{\mathcal{X}} (\mathcal{L}_A^* \mathbf{1}) e^v d\mu = 0.$$

As a result

$$-\frac{1}{2} \int_{\mathcal{X}} \mathcal{L}v d\nu = -\frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_S v) e^v d\mu - \frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_A v) e^v d\mu = \frac{1}{2} \int_{\mathcal{X}} \nabla v \cdot S \nabla v d\nu. \quad (3.99)$$

By plugging (3.98)-(3.99) into (3.97), we obtain

$$-\int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu = \frac{1}{4} \int_{\mathcal{X}} \nabla v \cdot S \nabla v d\nu - \frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_A \psi) d\nu - \frac{1}{4} \int_{\mathcal{X}} \nabla \psi \cdot S \nabla \psi d\nu. \quad (3.100)$$

The first term in the above equation reads (recalling that $\rho = e^v$)

$$\frac{1}{4} \int_{\mathcal{X}} \nabla v \cdot S \nabla v d\nu = \int_{\mathcal{X}} \nabla(\sqrt{\rho}) \cdot S \nabla(\sqrt{\rho}) d\mu.$$

By density of $C_c^\infty(\mathcal{X})$ in $\mathcal{H}^1(\mu)$, the above expression is valid for any ρ such that $\sqrt{\rho} \in \mathcal{H}^1(\mu)$. The above computation shows that this condition is equivalent to $v \in \mathcal{H}^1(\nu)$, and

$$\frac{1}{4} \int_{\mathcal{X}} \nabla v \cdot S \nabla v d\nu = \frac{1}{4} |v|_{\mathcal{H}^1(\nu)}^2,$$

which does not involve the function $\psi \in C_c^\infty(\mathcal{X})$. Moreover, since \mathcal{L}_A is a first order differential operator, antisymmetric in $L^2(\mu)$, it holds

$$\int_{\mathcal{X}} (\mathcal{L}_A \psi) d\nu = - \int_{\mathcal{X}} (\mathcal{L}_A e^v) \psi d\mu = - \int_{\mathcal{X}} (\mathcal{L}_A v) \psi d\nu.$$

As a result, (3.100) rewrites

$$-\int_{\mathcal{X}} \frac{\mathcal{L}u}{u} d\nu = \frac{1}{4} |v|_{\mathcal{H}^1(\nu)}^2 + \frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_A v) \psi d\nu - \frac{1}{4} |\psi|_{\mathcal{H}^1(\nu)}^2, \quad (3.101)$$

and this expression is finite for any $\psi \in C_c^\infty(\mathcal{X})$.

Our goal is now to take the supremum over functions $\psi \in C_c^\infty(\mathcal{X})$ in (3.101), and prove that this is enough to obtain the supremum over \mathcal{D}^+ . We consider for this the terms depending on ψ in (3.101) and, using the duality between $\mathcal{H}^1(\nu)$ and $\mathcal{H}^{-1}(\nu)$ (see [273, Section 2, Claim F]) we obtain

$$\begin{aligned} \frac{1}{2} \int_{\mathcal{X}} (\mathcal{L}_A v) \psi d\nu - \frac{1}{4} \int_{\mathcal{X}} \nabla \psi \cdot S \nabla \psi d\nu &\leq \frac{1}{2} |\mathcal{L}_A v|_{\mathcal{H}^{-1}(\nu)} |\psi|_{\mathcal{H}^1(\nu)} - \frac{1}{4} |\psi|_{\mathcal{H}^1(\nu)}^2 \\ &\leq \frac{1}{4\varepsilon} |\mathcal{L}_A v|_{\mathcal{H}^{-1}(\nu)}^2 - \frac{1}{4} (1 - \varepsilon) |\psi|_{\mathcal{H}^1(\nu)}^2, \end{aligned} \quad (3.102)$$

where we used Young's inequality with $\varepsilon < 1$ to obtain the second line. Since $\mathcal{L}_A v \in \mathcal{H}^{-1}(\nu)$, the supremum over the functions $\psi \in C_c^\infty(\mathcal{X})$ takes the value $-\infty$ when $\psi \notin \mathcal{H}^1(\nu)$. Therefore, by density of $C_c^\infty(\mathcal{X})$ in $\mathcal{H}^1(\nu)$, the supremum over the functions of the form (3.96) for $\psi \in C_c^\infty(\mathcal{X})$ recovers the supremum over \mathcal{D}^+ and it holds

$$I(\nu) = \frac{1}{4} |v|_{\mathcal{H}^1(\nu)}^2 + \frac{1}{4} |\mathcal{L}_A v|_{\mathcal{H}^{-1}(\nu)}^2, \quad (3.103)$$

by definition of the $\mathcal{H}^{-1}(\nu)$ -norm in Section 6.2, which concludes the proof.

Remark 3.36. *We have proved our result for measures of the form $d\nu = e^v d\mu$. Considering more general measures $\nu \ll \mu$ is made difficult because the Radon–Nikodym derivative $\rho = d\nu/d\mu$ may vanish on some region of \mathcal{X} , hence the definition of $\mathcal{L}_A(\log \rho)$ is not clear. Given (3.102), we see that we can give a sense to our computations provided $\mathcal{L}_A(\log \rho)$ defines a linear form on $\mathcal{H}^1(\nu)$, namely: there exists $C > 0$ such that*

$$\forall \psi \in \mathcal{H}^1(\nu), \quad \left| \int_{\mathcal{X}} \psi \mathcal{L}_A(\log \rho) d\nu \right| \leq C \|\psi\|_{\mathcal{H}^1(\nu)}.$$

We find it however clearer to work directly with exponential perturbations of the invariant measure μ .

3.6.2.4 Proof of Corollary 3.17

The proof follows from the variational formulation of Theorem 3.16. Indeed, let us rewrite (3.39) as

$$I_A(\nu) = -\frac{1}{2} \inf_{\psi \in \mathcal{H}^1(\nu)} \mathcal{I}_\nu(\psi), \quad (3.104)$$

where ν is fixed and satisfies the assumptions of the theorem, and

$$\mathcal{I}_\nu(\psi) = \frac{1}{2} \int_{\mathcal{X}} \mathcal{C}(\psi, \psi) d\nu - \int_{\mathcal{X}} \psi(\mathcal{L}_A v) d\nu.$$

By [273, Section 2, Claim F], we can identify $\mathcal{H}^{-1}(\nu)$ with the dual of $\mathcal{H}^1(\nu)$, so that \mathcal{I}_ν reads

$$\forall \psi \in \mathcal{H}^1(\nu), \quad \mathcal{I}_\nu(\psi) = \frac{1}{2} |\psi|_{\mathcal{H}^1(\nu)}^2 - \langle \mathcal{L}_A v, \psi \rangle_{\mathcal{H}^{-1}(\nu), \mathcal{H}^1(\nu)}.$$

Denoting by $\tilde{\nabla}$ the adjoint of the gradient operator in $L^2(\nu)$, the Lax–Milgram theorem [72, Corollary V.8], whose conditions are readily fulfilled, shows that the minimum is attained at a unique $\psi_v \in \mathcal{H}^1(\nu)$ solution to

$$\tilde{\nabla}(S\nabla\psi_v) = \mathcal{L}_A v. \quad (3.105)$$

Inserting ψ_v solution to (3.105) in (3.104) leads to

$$I_A(\nu) = \frac{1}{4} \int_{\mathcal{X}} \mathcal{C}(\psi_v, \psi_v) d\nu, \quad (3.106)$$

which concludes the proof.

3.6.3 Proof of Proposition 3.8

The proposition is a consequence of the equality

$$\Psi = -\frac{\mathcal{L}W}{W} = \theta \left(-\mathcal{L}V - \frac{\theta}{2} |\sigma^T \nabla V|^2 \right).$$

Since $|\sigma^T \nabla V|$ has compact level sets and $\Psi \sim |\sigma^T \nabla V|^2$ by (3.22), Ψ has compact level sets. Since V has compact level sets, for $\varepsilon < \theta/2$ it holds $\mathcal{W} \ll W$ and $\mathcal{W}^2 \leq C_1 W$ for some constant $C_1 > 0$. Moreover, outside a compact set, the function

$$\frac{\Psi}{-\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}}} = \frac{\theta(-\mathcal{L}V - \frac{\theta}{2} |\sigma^T \nabla V|^2)}{\varepsilon(-\mathcal{L}V - \frac{\varepsilon}{2} |\sigma^T \nabla V|^2)}$$

is bounded above and below since the numerator and denominator are both equivalent to $|\sigma^T \nabla V|^2$, so the second condition in (3.21) holds. Finally,

$$\begin{aligned} -2\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}} &= 2\varepsilon \left(-\mathcal{L}V - \frac{\varepsilon}{2} |\sigma^T \nabla V|^2 \right) = 2\frac{\varepsilon}{\theta} \theta \left(-\mathcal{L}V - \frac{\theta}{2} |\sigma^T \nabla V|^2 \right) + \varepsilon(\theta - \varepsilon) |\sigma^T \nabla V|^2 \\ &= 2\frac{\varepsilon}{\theta} \Psi + \varepsilon(\theta - \varepsilon) |\sigma^T \nabla V|^2. \end{aligned}$$

Since $\Psi \sim |\sigma^T \nabla V|^2$, we may choose ε small enough so as to obtain

$$-2\frac{\mathcal{L}\mathcal{W}}{\mathcal{W}} \leq \Psi + C_2,$$

for some constant $C_2 \in \mathbb{R}$. This proves the third item of (3.21).

3.6.4 Proof of Lemma 3.23

The proof relies on manipulations similar to those of [324]. A simple computation shows that

$$-\frac{\mathcal{L}_\gamma W}{W}(q, p) = \varepsilon q \cdot \nabla V - \gamma \varepsilon^2 |q|^2 + \gamma \varepsilon (1 - 2\theta) p \cdot q + \theta \gamma (1 - \theta) |p|^2 - \varepsilon |p|^2 - \theta \gamma d. \quad (3.107)$$

For any $\eta > 0$ it holds

$$p \cdot q \geq -\eta \frac{|q|^2}{2} - \frac{|p|^2}{2\eta}.$$

As a result, Assumption 3.22 leads to

$$-\frac{\mathcal{L}_\gamma W}{W}(q, p) \geq |q|^2 \left(c_V \varepsilon - \gamma \varepsilon^2 - \frac{\eta \gamma \varepsilon}{2} (1 - 2\theta) \right) + |p|^2 \left(\theta \gamma - \theta^2 \gamma - \varepsilon - \frac{\gamma \varepsilon}{2\eta} (1 - 2\theta) \right) - \theta \gamma d - C_V.$$

Since $\theta > 0$, it holds

$$-\frac{\mathcal{L}_\gamma W}{W}(q, p) \geq a|q|^2 + b|p|^2 - C,$$

where

$$a = \varepsilon \left(c - \frac{\eta \gamma}{2} \right) - \gamma \varepsilon^2, \quad b = \theta \gamma (1 - \theta) - \varepsilon - \frac{\gamma \varepsilon}{2\eta}, \quad C = \theta \gamma d + C_V.$$

The claim follows for $\theta \in (0, 1)$ by choosing $\eta < 2c/\gamma$ and ε sufficiently small.

Part III

Numerical analysis and algorithms

CHAPTER 4

ERROR ESTIMATES FOR ERGODIC PROPERTIES OF DISCRETIZED FEYNMAN–KAC SEMIGROUPS

The material for this chapter has been released in [182] and is accepted for publication in Numerische Mathematik.

Contents

4.1	Presentation	144
4.2	Convergence properties of Feynman–Kac semigroups	146
4.2.1	Continuous dynamics	146
4.2.2	Discretization	148
4.3	Numerical analysis of the discretization	151
4.3.1	Error estimates on the invariant probability measure	151
4.3.2	Alternative error estimate for the principal eigenvalue	156
4.3.3	TU Lemma	157
4.3.4	Second order schemes	158
4.4	Numerical application	161
4.4.1	Monte Carlo discretization	161
4.4.2	Galerkin discretization	162
4.4.3	Numerical results	163
4.5	Discussion	166
4.6	Proofs	167
4.6.1	Proof of the results of Section 4.2.1	167
4.6.2	Proof of Theorem 4.9	169
4.6.3	Proofs related to Theorem 4.16	171
4.6.4	Proof of Proposition 4.23	176
4.7	Markov contractions and Dobrushin coefficients	177

Abstract. We consider the numerical analysis of the time discretization of Feynman–Kac semigroups associated with diffusion processes. These semigroups naturally appear in large deviation theory (as motivated in Sections 1.3.1 and 2.1) but also in Diffusion Monte Carlo methods or non-linear filtering. We present in this chapter error estimates à la Talay–Tubaro on their invariant measures when the underlying continuous stochastic differential equation is discretized; as well as on the leading eigenvalue of the generator of the dynamics, which corresponds to the rate of creation of probability. This provides criteria to construct efficient integration schemes of Feynman–Kac dynamics, as well as a mathematical justification of numerical results already observed in the Diffusion Monte Carlo community. Our analysis is illustrated by numerical simulations.

4.1 Presentation

The study of Feynman–Kac semigroups for stochastic differential equations (SDEs) has been a topic of growing importance in the past two decades, since these dynamics are related to several theoretical and applied areas of mathematics. They can be seen as standard SDEs whose paths are reweighted according to the exponential of the time integral of some weight function.

Feynman–Kac semigroups naturally appear in large deviation theory, where they can be used to enhance the likelihood of observing rare fluctuations and henceforth computing cumulant generating functions [416, 121]. This is motivated at length in Section 1.3.1. They also have important practical applications, such as in the Diffusion Monte Carlo (DMC) method [190], which is a probabilistic way of estimating the ground state energy of Schrödinger operators; or in computational statistics, in particular in (non-linear) filtering [113, 137], where relevant trajectories are selected from observations. We refer to Chapter 2 for further motivations and references.

We focus in this chapter on the bias arising from the time discretization of the underlying continuous stochastic dynamics and of the time integrated weight. Our interest resides in the ergodic properties of the discretization, namely the invariant measure as well as the average rate of creation of probability (*i.e.* the cumulant generating function). We refer to Section 1.4 in Part I for general notions of numerical analysis for discretization of SDEs. Let us briefly present our setting and results.

We study a system $X_t \in \mathcal{X}$ evolving in a d -dimensional space. We assume in this chapter that

$$\mathcal{X} = \mathbb{T}^d$$

is a d -dimensional torus (where $\mathbb{T} = \mathbb{R} \setminus \mathbb{Z}$, the extension to unbounded spaces poses non-trivial issues, as discussed at various places later on). For convenience, we consider that the evolution is dictated by a stochastic differential equation with additive noise:

$$dX_t = b(X_t) dt + \sigma dB_t, \quad (4.1)$$

where $b : \mathcal{X} \rightarrow \mathbb{R}^d$ is a $C^\infty(\mathcal{X})$ vector field, B_t is a standard d -dimensional Brownian motion. Note that the dynamics (4.1) may be non-reversible. The diffusion coefficient $\sigma > 0$ is assumed to be *constant*, but our results could be extended to dynamics with multiplicative noise upon appropriate modifications. The infinitesimal generator of the dynamics (4.1) reads

$$\mathcal{L} = b \cdot \nabla + \frac{\sigma^2}{2} \Delta, \quad (4.2)$$

and we denote by \mathcal{L}^\dagger the adjoint of \mathcal{L} on $L^2(\mathcal{X})$ endowed with the Lebesgue measure. Since \mathcal{X} is compact and b is smooth, (4.1) admits a unique invariant distribution, denoted by μ , which is solution to the Fokker–Planck equation

$$\mathcal{L}^\dagger \mu = 0,$$

see *e.g.* [362, 262, 304]. In all this chapter, we use the shorthand notation $\mathcal{C} = C^\infty(\mathcal{X})$ for the space of smooth functions over \mathcal{X} , which is a core for the generator \mathcal{L} in $L^2(\mu)$.

Denoting by $\mathcal{P}(\mathcal{X})$ the set of probability measures over \mathcal{X} , we use the definition of Chapter 2 for Feynman–Kac type semigroups. That is, for a measurable function $f : \mathcal{X} \rightarrow \mathbb{R}$ and initial probability measure $\nu \in \mathcal{P}(\mathcal{X})$ we have,

$$\forall t \geq 0, \quad \forall \varphi \in \mathcal{C}, \quad \Phi_t^f(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right]}{\mathbb{E}_\nu \left[e^{\int_0^t f(X_s) ds} \right]}, \quad (4.3)$$

where the expectations run over initial conditions X_0 distributed according to ν and all realizations of (4.1). The family of mappings $\{\Phi_t^f\}_{t \geq 0}$ is a measure-valued non-linear semigroup in the sense that $\Phi_t^f : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$ depends non-linearly on the initial condition and, for all $\nu \in \mathcal{P}(\mathcal{X})$ and $t, s \in \mathbb{R}_+$, $\Phi_t^f(\Phi_s^f(\nu)) = \Phi_{t+s}^f(\nu)$, see Chapter 2. Such semigroups have been studied for a long time in the context of Diffusion Monte Carlo (DMC) [213, 10, 77, 412, 190] in order to estimate the principal eigenvalue of Schrödinger type operators $-\Delta + f$, which correspond in our case to $b \equiv 0$. They also appear in the large deviations community [202, 289, 397, 406, 337, 338] where they are related to the principal eigenvalue of $\mathcal{L} + f$, as discussed at length in Chapter 3, see also [133, 416, 121, 406, 119]. Other fields such as non-linear filtering, Hidden Markov Models [253, 137, 138] and free energy computation [252, 251, 378, 301] also motivate the study of such semigroups.

As we have studied in Chapter 2, the semigroup (4.3) converges in general to the average of φ with respect to a tilted measure μ_f (recall that \mathcal{X} is bounded here). More precisely, the operator $\mathcal{L}^\dagger + f$ has a largest eigenvalue λ which is isolated from the remainder of the spectrum and non-degenerate, with associated eigenfunction μ_f , and

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in \mathcal{C}, \quad \Phi_t^f(\nu)(\varphi) \xrightarrow[t \rightarrow \infty]{} \int_{\mathcal{X}} \varphi d\mu_f,$$

exponentially fast. We address in this work the time discretization of the semigroup (4.3) using a finite time step Δt . The underlying continuous evolution (4.1) is discretized by a Markov chain $(x_n)_{n \in \mathbb{N}}$ and (4.3) is approximated by (using a simple quadrature rule for the time integral)

$$\Phi_{\Delta t, n}^f(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(x_n) e^{\Delta t \sum_{i=0}^{n-1} f(x_i)} \right]}{\mathbb{E}_\nu \left[e^{\Delta t \sum_{i=0}^{n-1} f(x_i)} \right]}. \quad (4.4)$$

Under mild assumptions on the discretization scheme (made precise in Section 4.2.2), the discrete semigroup (4.4) converges to an invariant measure $\mu_{f, \Delta t}$ in the following sense,

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in \mathcal{C}, \quad \Phi_{\Delta t, n}^f(\nu)(\varphi) \xrightarrow[n \rightarrow +\infty]{} \int_{\mathcal{X}} \varphi d\mu_{f, \Delta t}.$$

The core of this chapter consists in making precise the difference between μ_f and $\mu_{f, \Delta t}$. We aim in particular at designing numerical schemes leading to the smallest possible biases. Although a series of papers study the statistical error of estimators such as (4.4) (see [114, 117, 113, 376, 377]), there are, to our knowledge, no available estimates on the bias of the limiting measure with respect to Δt . However, in the context of DMC (where we recall $b = 0$), it was numerically observed that some discretizations provide first, second or fourth order of convergence in Δt for the largest eigenvalue λ of $\mathcal{L} + f$, see for example [10, 412, 328, 381], and [164] for the numerical analysis in a simple case (note that we write λ instead of $\lambda(f)$ to shorten notation). The results presented in this chapter provide a mathematical justification of such convergences, while extending them to the case $b \neq 0$. Let us also mention that Hairer and Weare have studied in [221, 222] the convergence with respect to the time step of discretized dynamics similar to the one we consider, over a finite time and for a finite population of replicas. They obtain in the limit $\Delta t \rightarrow 0$ a limiting process, the so-called Brownian fan.

We rely on the techniques developed since the works of Talay and Tubaro [400, 398], taking advantage of the analytical tools developed in a series of papers [325, 110, 1, 2, 62, 295, 304], in order to provide a systematic framework to study the bias in the time step. More precisely, we show in Theorem 4.16 that there exist an integer $p \geq 1$ and a function ψ solution to a Poisson equation (both depending on the numerical scheme at hand and the quadrature rule for the integral), such that, for all $\varphi \in \mathcal{C}$,

$$\int_{\mathcal{X}} \varphi d\mu_{f, \Delta t} = \int_{\mathcal{X}} \varphi d\mu_f + \Delta t^p \int_{\mathcal{X}} \varphi \psi d\mu_f + O(\Delta t^{p+1}). \quad (4.5)$$

This result is very similar to those of weak backward error analysis on invariant probability measures of ergodic processes, see for example [110, 62, 295] and the presentation in Section 1.4.1. Moreover, as the computation of the principal eigenvalue λ (or cumulant function in the large deviations context) of the operator $\mathcal{L} + f$ is one of the main concerns in Feynman–Kac techniques, we provide in Theorem 4.21 the following error estimate:

$$\lambda_{\Delta t} = \frac{1}{\Delta t} \log \left[\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} d\mu_{f, \Delta t} \right] = \lambda + C \Delta t^p + O(\Delta t^{p+1}), \quad (4.6)$$

where $Q_{\Delta t}^f$ is the evolution operator of the discretized dynamics with weight function f . This result, related to Theorem 2.8 in Chapter 2, is interesting since it allows to justify the use of population dynamics methods for discretizations of diffusion processes, see [202, 397, 337] for rare events simulations, and [190] for DMC.

Let us mention that, while the proof of (4.5) relies on previous works concerning error estimates on the invariant measure [400, 398, 62, 110, 295], the novelty of this work lies in taking into account the non-probability conserving feature of the dynamics. With this point of view and, odd as it may seem, formula (4.5) appears as a *consequence* of (4.6), and not conversely. An interpretation of this

fact is that, in order to prove an error estimate on the invariant probability measure of this non probability-conserving dynamics, we must first show that the discretized process creates or destroys probability at a rate correct up to terms small in Δt . Phrased differently, we prove error estimates on the principal eigenvalue before considering the associated eigenvector.

The chapter is organized as follows. Section 4.2 is devoted to general properties of Feynman–Kac semigroups and their discretizations. We then present in Section 4.3 our main results concerning the numerical analysis of the error on the invariant probability measure, depending on the choice of the discretization scheme, before providing numerical applications in Section 4.4. Finally, Section 4.5 proposes possible extensions to this work. The proofs of the most technical results are gathered in Section 4.6.

4.2 Convergence properties of Feynman–Kac semigroups

We present in this section the setting of our study. In particular, we remind convergence results and some useful properties of continuous Feynman–Kac semigroups in Section 4.2.1, as well as convergence results for their discretizations in Section 4.2.2. Although these results are known (and can often be obtained as applications of the results of Chapter 2), we believe that it is useful to gather them here to allow for a self-contained presentation of the numerical analysis framework developed in Section 4.3.

4.2.1 Continuous dynamics

We denote by $(P_t)_{t \geq 0}$ the evolution semigroup associated with the process $(X_t)_{t \geq 0}$ in (4.1): for all $\nu \in \mathcal{P}(\mathcal{X})$ and $\varphi \in \mathcal{C}$,

$$P_t(\nu)(\varphi) = \mathbb{E}_\nu [\varphi(X_t)].$$

Its weighted counterpart is

$$P_t^f(\nu)(\varphi) = \mathbb{E}_\nu \left[\varphi(X_t) e^{\int_0^t f(X_s) ds} \right].$$

The infinitesimal generators of P_t and P_t^f are respectively \mathcal{L} and $\mathcal{L} + f$, where we denote with some abuse of notation by f the multiplication operator by the function f . Whether a statement corresponds to the function f or the associated multiplication operator should be clear from the context. We assume in the sequel that the function f is smooth, so that the associated multiplication operator stabilizes the core \mathcal{C} .

The existence of a spectral gap for the generator $\mathcal{L} + f$ and its adjoint is a key ingredient for our study. Here and in the sequel, and otherwise explicitly mentioned, all operators are considered on the Hilbert space

$$L^2(\mu) = \left\{ \varphi \text{ measurable} \mid \int_{\mathcal{X}} |\varphi|^2 d\mu < +\infty \right\}.$$

Let us insist on the fact that the $L^2(\mu)$ -framework is useful for the numerical analysis of Section 4.3, compared to the $B^\infty(\mathcal{X})$ spaces used in Chapter 2, which motivates the analysis of this section. For a given closed operator T on $L^2(\mu)$, we denote by T^* the adjoint of T in $L^2(\mu)$. In particular,

$$\forall (\phi, \varphi) \in \mathcal{C}, \quad \int_{\mathcal{X}} (T\phi)\varphi d\mu = \int_{\mathcal{X}} \phi(T^*\varphi) d\mu.$$

In this functional framework, the reversibility of the dynamics is equivalent to the self-adjointness of \mathcal{L} on $L^2(\mu)$. As mentioned in Section 1.1, we want our results to be valid for irreversible dynamics. Therefore, we do not assume that \mathcal{L} is symmetric, and this is why we need to distinguish between eigenelements of \mathcal{L} and \mathcal{L}^* . We can then state the following.

Proposition 4.1. *The operator $\mathcal{L} + f$, considered on $L^2(\mu)$, has a real isolated principal eigenvalue λ with associated eigenfunction $h_f \in \mathcal{C}$ normalized as*

$$(\mathcal{L} + f)h_f = \lambda h_f, \quad \int_{\mathcal{X}} h_f d\mu = 1. \tag{4.7}$$

The operator $\mathcal{L}^ + f$ then also admits λ as a real isolated principal eigenvalue, with associated eigenfunction $\ell_f \in \mathcal{C}$ normalized as*

$$(\mathcal{L}^* + f)\ell_f = \lambda \ell_f, \quad \int_{\mathcal{X}} \ell_f d\mu = 1. \tag{4.8}$$

Moreover, the functions h_f and ℓ_f are positive.

The fact that h_f and ℓ_f belong to \mathcal{C} is a consequence of elliptic regularity. Let us emphasize that, as a consequence of (4.8), the measure

$$\mu_f = \ell_f \mu$$

is the only invariant probability measure for the evolution encoded by $P_t^{f-\lambda}$. Moreover, when the underlying diffusion is reversible, *i.e.* $b = -\nabla V$ and $\mu(dx) = Z^{-1} e^{-2V(x)/\sigma^2} dx$, the operator \mathcal{L} is self-adjoint ($\mathcal{L}^* = \mathcal{L}$) so that $h_f = \ell_f$. When $f = 0$, it simply holds $\lambda = 0$ and $h_f = \ell_f = \mathbb{1}$ whatever b .

Proof. It is shown in [199] that the operator $\mathcal{L} + f$ has a real isolated principal eigenvalue when considered as an operator on $C^0(\mathcal{X})$, the space of continuous functions over \mathcal{X} . This can be proved using the Krein–Rutman theorem [141] like in Part II. On the other hand, standard results of spectral theory of elliptic operators on bounded domains show that $\mathcal{L} + f$ on $L^2(\mu)$ has a discrete spectrum, which is bounded above [358]. The first eigenvalue cannot be degenerate since the associated eigenvectors are smooth by elliptic regularity and are therefore also eigenvectors of $\mathcal{L} + f$ considered as an operator on $C^0(\mathcal{X})$. Finally, the positivity of h_f and ℓ_f follows from the fact that the evolution semigroup P_t^f and its adjoint are operators with smooth and positive transition kernels (since the noise is non-degenerate), together with the equalities $P_t^f h_f = e^{\lambda t} h_f$ and $(P_t^f)^* \ell_f = e^{\lambda t} \ell_f$. We refer to Chapter 2 for more details. \square

In what follows, we use the subspaces $L_f^2(\mu)$ and \mathcal{C}_f of functions of average 0 with respect to μ_f :

$$L_f^2(\mu) = \left\{ \varphi \in L^2(\mu) \mid \int_{\mathcal{X}} \varphi d\mu_f = 0 \right\}, \quad \mathcal{C}_f = \left\{ \varphi \in \mathcal{C} \mid \int_{\mathcal{X}} \varphi d\mu_f = 0 \right\}.$$

We also introduce the measure $\hat{\mu}_f = h_f \mu$, the space

$$\hat{\mathcal{C}}_f = \left\{ \varphi \in \mathcal{C} \mid \int_{\mathcal{X}} \varphi d\hat{\mu}_f = 0 \right\},$$

and we denote by

$$\delta_f = \inf \left\{ \lambda - \operatorname{Re}(z), z \in \sigma(\mathcal{L} + f) \setminus \{\lambda\} \right\} > 0 \quad (4.9)$$

the spectral gap of $\mathcal{L} + f$ in $L^2(\mu)$. The fact that the largest eigenvalue λ is a priori non-zero corresponds to a possible creation ($\lambda > 0$) or destruction ($\lambda < 0$) of probability induced by the source term f , which plays the role of an importance sampling function. The statement about the spectral gap in Proposition 4.1 implies the convergence of the Feynman–Kac semigroup (4.3), as stated in the following result.

Proposition 4.2. *There exists $C > 0$ such that, for all $\nu \in \mathcal{P}(\mathcal{X})$ and $\varphi \in L^2(\mu)$,*

$$\forall t \geq 1, \quad \left| \Phi_t^f(\nu)(\varphi) - \int_{\mathcal{X}} \varphi d\mu_f \right| \leq C \|\varphi\|_{L^2(\mu)} e^{-\delta_f t}, \quad (4.10)$$

where δ_f is defined in (4.9).

As made clear in the proof of this result (see Section 4.6.1), it is possible to consider any observable $\varphi \in L^2(\mu)$ even if μ is singular. This is due to the regularizing properties of the underlying diffusion for positive times, and explains why the convergence result is stated only for times $t \geq 1$. The next proposition will be frequently used in this work.

Proposition 4.3. *It holds*

$$\int_{\mathcal{X}} f d\mu_f = \lambda. \quad (4.11)$$

Proof. Integrating both sides of (4.8) on \mathcal{X} ,

$$\int_{\mathcal{X}} f d\mu_f = \int_{\mathcal{X}} \lambda \ell_f d\mu - \int_{\mathcal{X}} \mathcal{L}^* \ell_f d\mu = \lambda \int_{\mathcal{X}} \ell_f d\mu - \int_{\mathcal{X}} \mathcal{L} \mathbb{1} d\mu_f = \lambda,$$

since $\mathcal{L} \mathbb{1} = 0$. \square

A natural corollary of Propositions 4.2 and 4.3 is that the largest eigenvalue of $\mathcal{L} + f$ can be obtained by a long time average of f using the Feynman–Kac semigroup (4.3).

Corollary 4.4. *There exists $C > 0$ such that, for any initial distribution $\nu \in \mathcal{P}(\mathcal{X})$,*

$$\left| \Phi_t^f(\nu)(f) - \lambda \right| \leq C e^{-\delta_f t}.$$

Another important consequence of Proposition 4.1 is the invertibility of the generator and its adjoint over suitable functional spaces.

Proposition 4.5. *The operator $\mathcal{L} + f - \lambda$ is invertible on \mathcal{C}_f , in the sense that, for any $g \in \mathcal{C}_f$, the Poisson equation*

$$(\mathcal{L} + f - \lambda)u = g$$

admits a unique solution $u \in \mathcal{C}_f$, which is denoted by $(\mathcal{L} + f - \lambda)^{-1}g$. Similarly, $\mathcal{L}^ + f - \lambda$ is invertible on $\widehat{\mathcal{C}}_f$.*

The proof of this result can be read in Section 4.6.1. Let us emphasize that the smoothness of f is crucial for this proposition to be true. Note also that the stability of the core of the operator $\mathcal{L} + f$ would be harder to prove for non-compact state spaces, as this is already a non-trivial statement for the Poisson equation with $f = 0$, see [276, 277].

4.2.2 Discretization

We now turn to the discretization of the Feynman–Kac semigroup (4.3). We first define discretization schemes, and show that they are ergodic for some limiting measure under mild assumptions. We also recall the stationarity equation satisfied by this invariant probability measure, which proves crucial for the numerical analysis developed in Section 4.3.

The properties of discretized Feynman–Kac semigroups are related to the properties of the underlying discrete dynamics. The approximation of the continuous dynamics (4.1) is given, for a time Δt , by a Markov chain $(x_n)_{n \in \mathbb{N}}$ such that $x_n \simeq X_{n\Delta t}$. This Markov chain is characterized by the evolution operator $Q_{\Delta t}$ defined as

$$(Q_{\Delta t}\varphi)(x) = \mathbb{E}[\varphi(x_{n+1}) \mid x_n = x]. \quad (4.12)$$

A typical example is the Euler–Maruyama scheme defined by:

$$x_{n+1} = x_n + b(x_n)\Delta t + \sigma\sqrt{\Delta t}G^n, \quad (4.13)$$

where $(G^n)_{n \geq 0}$ is a family of independent and identically distributed standard d -dimensional Gaussian random variables. In order to perform our analysis in Section 4.3, it is convenient to rephrase discretizations of (4.3) such as (4.4) in terms of an evolution operator. For instance, we see that, defining

$$(Q_{\Delta t}^f\varphi)(x) = e^{\Delta t f(x)}(Q_{\Delta t}\varphi)(x), \quad (4.14)$$

the discretization (4.4) reads, for an initial measure ν and a test function φ ,

$$\Phi_{\Delta t, n}^f(\nu)(\varphi) = \frac{\nu((Q_{\Delta t}^f)^n\varphi)}{\nu((Q_{\Delta t}^f)^n\mathbf{1})}. \quad (4.15)$$

We use the definition (4.15) for more general discretizations of (4.3) characterized by an evolution operator $Q_{\Delta t}^f$. Consistency requirements on $Q_{\Delta t}^f$ are made precise in Assumption 4.12 below. This allows us to take into account various integration rules, both for the underlying dynamics and the exponential weights. For instance, the choice

$$(Q_{\Delta t}^f\varphi)(x) = e^{\frac{\Delta t}{2}f(x)} \left[Q_{\Delta t} \left(e^{\frac{\Delta t}{2}f} \varphi \right) \right] (x), \quad (4.16)$$

well-known in the diffusion Monte Carlo community [381, 328, 320, 412], defines the following semi-group:

$$\Phi_{\Delta t, n}^f(\nu)(\varphi) = \frac{\mathbb{E}_\nu \left[\varphi(x_n) e^{\Delta t \sum_{i=0}^{n-1} \frac{f(x_i) + f(x_{i+1})}{2}} \right]}{\mathbb{E}_\nu \left[e^{\Delta t \sum_{i=0}^{n-1} \frac{f(x_i) + f(x_{i+1})}{2}} \right]}.$$

Remark 4.6. *The weighted evolution of $(X_t)_{t \geq 0}$ can be equivalently formulated as the unweighted evolution for the augmented system $(X_t, Z_t)_{t \geq 0}$, where $Z_t \geq 0$ is solution to*

$$dZ_t = Z_t f(X_t) dt, \quad Z_0 = 1.$$

However, Z_t is unbounded and may diverge to $+\infty$. The augmented dynamics $(X_t, Z_t)_{t \geq 0}$ therefore does not have an invariant measure in general, which complicates the analysis of the long time limit. Moreover, a naive discretization like the Euler–Maruyama scheme applied to $(X_t, Z_t)_{t \geq 0}$ reads

$$\begin{cases} x_{n+1} = x_n + b(x_n)\Delta t + \sigma\sqrt{\Delta t}G^n, \\ z_{n+1} = z_n + z_n f(x_n)\Delta t. \end{cases}$$

Observe that the positivity of Z_t may not be preserved during the dynamics if Δt is too large, which is crucial for the numerical scheme to be well-defined. This issue persists in general for other schemes. On the other hand, if x_n is fixed, the process Z_t solving

$$dZ_t = Z_t f(x_n) dt, \quad Z_0 = z_n$$

over a time step Δt admits the exact solution

$$z_{n+1} = z_n e^{f(x_n)\Delta t}.$$

Therefore, a first order splitting between X_t and Z_t leads to the first order integrator (4.14). If we perform a second order splitting between X_t and Z_t , we are back to the second order integration rule prescribed by (4.16). As a result, although considering an extended system $(X_t, Z_t)_{t \geq 0}$ of course makes sense, we see that, in order for the positivity of Z_t to be unconditionally preserved, we are naturally led to the same schemes as for the usual Feynman–Kac dynamics. There is finally a technical restriction with the reformulation of the Feynman–Kac dynamics using the augmented process $(X_t, Z_t)_{t \geq 0}$. The generator \mathcal{L}_{aug} of $(X_t, Z_t)_{t \geq 0}$ is defined, for a test function φ , through $\mathcal{L}_{\text{aug}}\varphi(x, z) = \mathcal{L}\varphi(x, z) + zf(x)\partial_z\varphi(x, z)$. However, the numerical analysis presented in Section 4.3 uses stability properties of the inverse of the generator of the dynamics (see Assumption 4.14 below). While \mathcal{L} is invertible as an operator acting on functions of x , it is much more difficult to define the inverse of \mathcal{L}_{aug} in a general way (think of the case $f = 0$).

In what follows, given that the discrete semigroup defines a measure-valued dynamics, we write for simplicity $\nu_n = \Phi_{\Delta t, n}^f(\nu)$, and we denote again by $B^\infty(\mathcal{X}) = \{\varphi \text{ measurable} \mid \sup_{x \in \mathcal{X}} |\varphi(x)| < +\infty\}$ the space of bounded measurable functions. For a given bounded operator Q on $B^\infty(\mathcal{X})$ and a probability measure $\nu \in \mathcal{P}(\mathcal{X})$, we also denote by νQ the probability measure defined as

$$\forall \varphi \in \mathcal{C}, \quad (\nu Q)(\varphi) = \nu(Q\varphi). \quad (4.17)$$

We start by recalling a one-step formulation of the non-linear dynamics $(\nu_n)_{n \geq 0}$, as suggested *e.g.* in [117] and used in Chapter 2. This formulation is the basis for a stationarity property fundamental in our numerical analysis.

Lemma 4.7. *The sequence of probability measures $\nu_n = \Phi_{\Delta t, n}^f(\nu)$ satisfies the following dynamics:*

$$\nu_{n+1} = \Phi_{\Delta t}^f(\nu_n),$$

where

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \forall \varphi \in \mathcal{C}, \quad \Phi_{\Delta t}^f(\nu)(\varphi) = \frac{\nu(Q_{\Delta t}^f \varphi)}{\nu(Q_{\Delta t}^f \mathbf{1})}. \quad (4.18)$$

Proof. The proof relies on a simple rewriting: for all $\varphi \in \mathcal{C}$,

$$\nu_{n+1}(\varphi) = \frac{\nu\left((Q_{\Delta t}^f)^{n+1}\varphi\right)}{\nu\left((Q_{\Delta t}^f)^{n+1}\mathbf{1}\right)} = \frac{\nu\left((Q_{\Delta t}^f)^n(Q_{\Delta t}^f\varphi)\right)}{\nu\left((Q_{\Delta t}^f)^n\mathbf{1}\right)} \times \frac{\nu\left((Q_{\Delta t}^f)\mathbf{1}\right)}{\nu\left((Q_{\Delta t}^f)(Q_{\Delta t}^f\mathbf{1})\right)} = \frac{\nu_n\left(Q_{\Delta t}^f\varphi\right)}{\nu_n\left(Q_{\Delta t}^f\mathbf{1}\right)},$$

which gives the result. \square

Let us now prove that the measure-valued dynamical process (4.18) admits a limit measure $\mu_{f,\Delta t}$ independent of the initial distribution ν , and that the long time average (4.15) converges to the average with respect to this measure. Although we could use the results of Chapter 2, we follow the strategy of Del Moral and collaborators [114, 117, 115, 113], which relies on the Dobrushin ergodic coefficient of a relevant operator (see Section 4.7). We believe this is instructive, and in the context of a bounded state space it requires less regularity assumptions. For following this approach, we use the following assumption, which is typically satisfied for discretizations associated with the continuous dynamics (4.3) on the torus.

Assumption 4.8. *The operator $Q_{\Delta t}^f$ satisfies a minorization and boundedness condition: there exist $\varepsilon \in (0, 1)$ and $\eta \in \mathcal{P}(\mathcal{X})$ such that, for all non-negative bounded measurable function φ ,*

$$\forall x \in \mathcal{X}, \quad \varepsilon \eta(\varphi) \leq \left(Q_{\Delta t}^f \varphi \right) (x) \leq \varepsilon^{-1} \eta(\varphi). \quad (4.19)$$

Since f is smooth and $\mathcal{X} = \mathbb{T}^d$, the condition (4.19) is satisfied for the evolution operator (4.14) as soon as a condition similar to (4.19) is satisfied for $Q_{\Delta t}$. The latter condition is, in turn, easily seen to be true for the numerical scheme (4.13), with $\eta(dx) = |\mathcal{X}|^{-1} dx$ the normalized Lebesgue measure on \mathcal{X} , see [304, Section 3.3.2]. Similar considerations allow to prove that (4.19) holds for more complicated discretization strategies [219, 304].

We can now recall an important result which ensures the existence and uniqueness of the limiting measure for the discretized Feynman–Kac dynamics. Its proof, taken from [114], is recalled in Section 4.6.2. To state the result, we introduce the total variation distance between two measures $\mu, \nu \in \mathcal{P}(\mathcal{X})$:

$$\|\mu - \nu\|_{\text{TV}} = \sup_{A \subset \mathcal{X}} |\mu(A) - \nu(A)|,$$

where the supremum runs over measurable subsets of \mathcal{X} . Recall that $\mathcal{P}(\mathcal{X})$ is complete for this distance.

Theorem 4.9. *Suppose that Assumption 4.8 holds true. Then the non-linear dynamics (4.18) admits a unique stationary probability measure $\mu_{f,\Delta t}$ which is independent of the initial measure and which is a fixed point of $\Phi_{\Delta t}^f$:*

$$\mu_{f,\Delta t} = \Phi_{\Delta t}^f(\mu_{f,\Delta t}). \quad (4.20)$$

Moreover, for any initial distribution $\nu \in \mathcal{P}(\mathcal{X})$,

$$\|\nu_n - \mu_{f,\Delta t}\|_{\text{TV}} \leq 2(1 - \varepsilon^2)^n. \quad (4.21)$$

Remark 4.10. *Let us emphasize that the prefactor ε in (4.19) typically scales as $\Delta t^{-\frac{d}{2}} \exp(-C_L/\Delta t)$ for some constant $C_L > 0$. Indeed, consider for instance the first order discretization (4.13). Its transition kernel between x and x' reads*

$$Q_{\Delta t}(x, dx') = \frac{1}{(2\pi\sigma^2\Delta t)^{\frac{d}{2}}} \exp\left(-\frac{(x' - x - b(x)\Delta t)^2}{2\sigma^2\Delta t}\right) dx'.$$

We then see that ε scales at dominant order in Δt as $\Delta t^{-\frac{d}{2}} \exp(-C_L/\Delta t)$ for some constant $C_L > 0$ depending on σ and \mathcal{X} , independently on the drift b . Thus, the choice of integrator should not affect significantly the value of ε . We also recall that, if $Q_{\Delta t}^f$ satisfies a uniform version of (4.19) with additional strong Feller and consistency conditions, it is possible to derive uniform in Δt convergence estimates, see Theorem 2.21 in Chapter 2.

As a consequence of Theorem 4.9, if we define a discretization of the Feynman–Kac semigroup (4.3) satisfying Assumption 4.8, the discrete dynamics (4.15) admits an invariant probability measure solution to the fixed point equation (4.20). In view of (4.20) and (4.18), this measure satisfies the following stationarity equation:

$$\forall \varphi \in \mathcal{C}, \quad \int_{\mathcal{X}} Q_{\Delta t}^f \varphi d\mu_{f,\Delta t} = \left(\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} d\mu_{f,\Delta t} \right) \left(\int_{\mathcal{X}} \varphi d\mu_{f,\Delta t} \right). \quad (4.22)$$

In particular, if we define the approximate eigenvalue $\lambda_{\Delta t}$ by

$$e^{\Delta t \lambda_{\Delta t}} = \int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} d\mu_{f,\Delta t}, \quad (4.23)$$

then (4.22) can be rewritten as:

$$\forall \varphi \in \mathcal{C}, \quad \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \varphi \right] d\mu_{f, \Delta t} = 0. \quad (4.24)$$

This is the stationarity equation of the discretized process upon which the analysis in Section 4.3 is built. Let us emphasize that it involves the approximate eigenvalue $\lambda_{\Delta t}$ accounting for the rate of creation of probability of the discretized process, which differs in general from the largest eigenvalue λ of the operator $\mathcal{L}+f$ (which accounts for the rate of creation of probability for the continuous process¹). The numerical analysis of the approximation $\lambda_{\Delta t}$ of λ plays an important role in Section 4.3.

Remark 4.11. *In the case $f \equiv 0$, the measure $\mu_{f, \Delta t} = \mu_{\Delta t}$ is the invariant probability measure of the discretized process without reweighting, and the evolution operator $Q_{\Delta t}$ conserves probability. This also implies that $\lambda_{\Delta t} = 0$. Therefore (4.24) simplifies as*

$$\forall \varphi \in \mathcal{C}, \quad \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t} - 1}{\Delta t} \right) \varphi \right] d\mu_{\Delta t} = 0,$$

which is the standard stationarity equation of the invariant probability measure for discretizations of SDEs [295, 304]. This is because the largest eigenvalue of the discretized evolution operator $Q_{\Delta t}$ is 1, as for the continuous semigroup P_t . As already mentioned in Section 1.4 in the Introduction, this is an important difference with standard numerical analysis for SDEs.

4.3 Numerical analysis of the discretization

We now turn to the main section of the chapter, where we quantify how close $\mu_{f, \Delta t}$, the ergodic measure for the discrete Feynman–Kac dynamics, is from μ_f , the ergodic measure for its continuous counterpart. We also make precise the difference at leading order in Δt . Following a general strategy to study the error on the invariant probability measure of discretizations of stochastic processes dating back to [400] (see Section 1.4 in the Introduction and [110, 295, 304] for more details), we compare the evolution operator $Q_{\Delta t}^f$ with the Feynman–Kac semigroup $e^{\Delta t(\mathcal{L}+f)}$. Although the non probability-conserving feature of the dynamics is an additional difficulty, we obtain in Section 4.3.1 results similar to those of [400, 110, 2, 304] concerning the error on the invariant probability measure. Moreover, we provide in Section 4.3.2 error bounds for estimators of the eigenvalue λ . Finally, we show how to relate the invariant probability measures of different schemes in Section 4.3.3 and discuss in Section 4.3.4 how the Feynman–Kac discretization essentially inherits the properties of the discretization of the underlying unweighted dynamics.

4.3.1 Error estimates on the invariant probability measure

4.3.1.1 Expansions of the discrete evolution operators $Q_{\Delta t}^f$

For unweighted dynamics ($f = 0$), consistency assumptions on the evolution operator $Q_{\Delta t}$ characterizing the discretization rely on an expansion of $Q_{\Delta t}$ in powers of Δt (see the presentation in Section 1.4). More precisely, it is assumed that there exist an integer $p \geq 1$ and differential operators $(\mathcal{A}_k)_{k=1, \dots, p+1}$ such that the evolution operator $Q_{\Delta t}$ of the discrete dynamics admits the following expansion: for all $\varphi \in \mathcal{C}$,

$$Q_{\Delta t} \varphi = \varphi + \Delta t \mathcal{A}_1 \varphi + \Delta t^2 \mathcal{A}_2 \varphi + \dots + \Delta t^{p+1} \mathcal{A}_{p+1} \varphi + \Delta t^{p+2} \mathcal{R}_{\Delta t} \varphi. \quad (4.25)$$

The differential operators \mathcal{A}_k have finite order and smooth coefficients: for any $k \in \{1, \dots, p+1\}$, there exist $m_k \in \mathbb{N}$ and a family of smooth functions $(a_\alpha)_{|\alpha| \leq m_k}$ (with $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}^d$ and $|\alpha| = \alpha_1 + \dots + \alpha_d$) such that

$$\mathcal{A}_k = \sum_{|\alpha| \leq m_k} a_\alpha \partial^\alpha, \quad (4.26)$$

where $\partial^\alpha = \partial_{x_1}^{\alpha_1} \dots \partial_{x_d}^{\alpha_d}$. Moreover, $\mathcal{R}_{\Delta t}$ is an operator uniformly bounded in Δt in the following sense: there exist $\Delta t^* > 0$, $c > 0$ and $m \in \mathbb{N}$ such that

$$\forall \Delta t \in (0, \Delta t^*], \quad \forall \varphi \in \mathcal{C}, \quad \|\mathcal{R}_{\Delta t} \varphi\|_{C^0} \leq c \|\varphi\|_{C^m}, \quad (4.27)$$

¹This is made precise by equation (2.26) in Chapter 2.

where

$$\|\varphi\|_{C^m} = \sup_{|\alpha| \leq m} \sup_{x \in \mathcal{X}} |\partial^\alpha \varphi(x)|. \quad (4.28)$$

The assumptions (4.25) and (4.27) are standard for the numerical analysis of ergodic measures of SDEs [400, 276, 110, 1, 295, 304], and are satisfied for a wide range of explicit and implicit schemes defined on compact domains. A scheme is of weak order p when (4.25) holds with

$$\forall k \in \{1, \dots, p\}, \quad \mathcal{A}_k = \frac{\mathcal{L}^k}{k!},$$

see for instance [332]. Typically, $\mathcal{A}_1 = \mathcal{L}$ for any reasonable discretization scheme.

Besides weak and strong errors, another notion of consistency is the error arising on the invariant probability measure, in situations when the Markov chain associated with $Q_{\Delta t}$ admits an invariant probability measure $\mu_{\Delta t}$. The error between averages with respect to μ and $\mu_{\Delta t}$ are of order at least Δt^p when the scheme is weakly consistent of order p . It can however be of higher order $\Delta t^{p'}$ (with $p' \geq p + 1$) when

$$\forall k \in \{1, \dots, p'\}, \quad \forall \varphi \in \mathcal{C}, \quad \int_{\mathcal{X}} \mathcal{A}_k \varphi d\mu = 0. \quad (4.29)$$

This condition is satisfied by operators which are proportional to powers of \mathcal{L} . See however [3, 295] for examples of situations where \mathcal{A}_k is not a power of \mathcal{L} but the above condition is met.

In the context of Feynman–Kac averages (4.3) where we consider approximations $Q_{\Delta t}^f$ of $e^{\Delta t(\mathcal{L}+f)}$, we generalize the conditions (4.25) and (4.27) as follows.

Assumption 4.12. *There exist an integer $p \geq 1$ and differential operators $(\mathcal{A}_k^f)_{k=1, \dots, p+1}$ of the form (4.26) such that the evolution operator $Q_{\Delta t}^f$ of the Feynman–Kac dynamics admits the following expansion: for all $\varphi \in \mathcal{C}$,*

$$Q_{\Delta t}^f \varphi = \varphi + \Delta t \mathcal{A}_1^f \varphi + \Delta t^2 \mathcal{A}_2^f \varphi + \dots + \Delta t^{p+1} \mathcal{A}_{p+1}^f \varphi + \Delta t^{p+2} \mathcal{R}_{f, \Delta t} \varphi, \quad (4.30)$$

where $\mathcal{R}_{f, \Delta t}$ is a uniformly bounded remainder in the sense of (4.27). We also assume that \mathcal{A}_1^f is such that

$$\mathcal{A}_1^f = \mathcal{A}_1 + f, \quad \mathcal{A}_1 \mathbb{1} = 0, \quad (4.31)$$

where \mathcal{A}_1 is a differential operator. In particular, $\mathcal{A}_1^f \mathbb{1} = f$.

Let us provide an example of such an expansion when $Q_{\Delta t}^f$ is defined by (4.14).

Lemma 4.13. *Assume that (4.25) and (4.27) hold with $\mathcal{A}_1 \mathbb{1} = 0$, and define $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$. Then Assumption 4.12 holds with, for all $k \in \{1, \dots, p + 1\}$,*

$$\mathcal{A}_k^f \varphi = \sum_{m=0}^k \frac{f^m}{m!} \mathcal{A}_{k-m} \varphi. \quad (4.32)$$

Proof. The equality follows by expanding $e^{\Delta t f}$ in Δt and taking the product with the semigroup expansion: there exist Δt^* and $K > 0$ such that

$$Q_{\Delta t}^f \varphi = \left(1 + \Delta t f + \frac{\Delta t^2}{2} f^2 + \dots + \frac{\Delta t^{p+1}}{(p+1)!} f^{p+1} + \Delta t^{p+2} r_{\Delta t, f} \right) \\ \times \left(\varphi + \Delta t \mathcal{A}_1 \varphi + \Delta t^2 \mathcal{A}_2 \varphi + \dots + \Delta t^{p+1} \mathcal{A}_{p+1} \varphi + \Delta t^{p+2} \mathcal{R}_{\Delta t} \varphi \right),$$

with $\|r_{f, \Delta t}\|_{C^0} \leq K$ for $0 < \Delta t \leq \Delta t^*$. Gathering the terms of order Δt^k leads to (4.32) plus a uniformly bounded remainder, which proves the result. \square

Note that, in (4.32), we obtain $\mathcal{A}_1^f = \mathcal{A}_1 + f$ where \mathcal{A}_1 is defined in (4.25). However, there are other ways to construct Feynman–Kac schemes $Q_{\Delta t}^f$, using for instance a splitting strategy. Let us give an example. Assume for instance that the operator \mathcal{L} can be split in two parts: $\mathcal{L} = \mathcal{L}_1 + \mathcal{L}_2$. We can then define a splitting scheme as $Q_{\Delta t} = e^{\Delta t \mathcal{L}_2} e^{\Delta t \mathcal{L}_1}$, and, by discretizing the time integral of f in three parts (using Simpson’s rule) intertwined with $e^{\Delta t \mathcal{L}_2}$ and $e^{\Delta t \mathcal{L}_1}$,

$$Q_{\Delta t}^f = e^{\frac{\Delta t}{6} f} e^{\Delta t \mathcal{L}_2} e^{\frac{2\Delta t}{3} f} e^{\Delta t \mathcal{L}_1} e^{\frac{\Delta t}{6} f}.$$

In this case, we see that the expansion of $Q_{\Delta t}^f$ cannot be derived from the one for $Q_{\Delta t}$ (by a statement similar to (4.32)). The evolution operator $Q_{\Delta t}^f$ nonetheless satisfies Assumption 4.12.

4.3.1.2 Statement of the main result

Before stating our main theorem, we need to introduce the following technical assumptions.

Assumption 4.14 (Stability). *The operators $\mathcal{A}_1 + f - \lambda$ and $\mathcal{A}_1^* + f - \lambda$ are invertible on \mathcal{C}_f and $\widehat{\mathcal{C}}_f$ respectively (in the sense made precise in Proposition 4.5).*

In our setting, a crucial step of the proof consists in building an approximation of the eigenvector h_f to solve an approximate eigenvalue problem for the operator $Q_{\Delta t}^f$. This is an important difference compared to the case $f \equiv 0$, and requires the following assumption.

Assumption 4.15 (Spectral consistency). *The operator $\mathcal{A}_1 + f$, considered on $L^2(\mu)$, admits λ as its largest eigenvalue, with associated eigenvector h_f :*

$$(\mathcal{A}_1 + f)h_f = \lambda h_f.$$

Note that Assumptions 4.14 and 4.15 are immediately met when the schemes are weakly consistent, i.e. $\mathcal{A}_1 = \mathcal{L}$, since Assumption 4.14 is equivalent to Proposition 4.5 while Assumption 4.15 follows from Proposition 4.1. However, it is possible in principle to construct numerical schemes for which $\mathcal{A}_1 \neq \mathcal{L}$, in which case Assumptions 4.14 and 4.15 should be checked directly.

We are now in position to state our main result on the numerical discretization of Feynman–Kac semigroups, which makes precise error estimates à la Talay–Tubaro in the ergodic setting.

Theorem 4.16. *Suppose that Assumptions 4.12, 4.14 and 4.15 hold. Assume also that the operators \mathcal{A}_k^f in (4.30) are such that, for all $k \in \{1, \dots, p\}$, there exists $a_k \in \mathbb{R}$ with*

$$\forall \varphi \in \mathcal{C}, \quad \int_{\mathcal{X}} (\mathcal{A}_k^f \varphi) d\mu_f = a_k \int_{\mathcal{X}} \varphi d\mu_f. \quad (4.33)$$

Define also $\psi \in \mathcal{C}_f$ as

$$\psi = \psi_0 - \int_{\mathcal{X}} \psi_0 d\mu_f, \quad \left\{ \begin{array}{l} (\mathcal{A}_1^* + f - \lambda)(\ell_f \psi_0) = \tilde{g}, \\ \tilde{g} = -(\mathcal{A}_{p+1}^f)^* \ell_f + \ell_f \frac{\int_{\mathcal{X}} \mathcal{A}_{p+1}^f h_f d\mu_f}{\int_{\mathcal{X}} h_f d\mu_f} \in \widehat{\mathcal{C}}_f. \end{array} \right. \quad (4.34)$$

Then, there exist a time step $\Delta t^* > 0$ and a real-valued operator $R_{f,\Delta t}$ (uniformly bounded in Δt in the sense of (4.27)) such that, for any $0 < \Delta t \leq \Delta t^*$ and any $\varphi \in \mathcal{C}$,

$$\int_{\mathcal{X}} \varphi d\mu_{f,\Delta t} = \int_{\mathcal{X}} \varphi d\mu_f + \Delta t^p \int_{\mathcal{X}} \varphi \psi d\mu_f + \Delta t^{p+1} R_{f,\Delta t} \varphi. \quad (4.35)$$

Note that the denominator in the second term on the right-hand side of the definition of \tilde{g} is positive thanks to Proposition 4.1. In general, in (4.30), we expect \mathcal{A}_k^f to be $(\mathcal{L} + f)^k/k!$ (which corresponds to a scheme of weak order k), in which case (4.33) holds for $a_k = \lambda^k/k! \neq 0$ (see (4.48) below for a proof of the latter equality). This factor comes from the fact that $Q_{\Delta t}^f$ does not conserve probability. Indeed, for the evolution operator $Q_{\Delta t}$ of a Markovian dynamics, one always has

$$\int_{\mathcal{X}} Q_{\Delta t} \mathbb{1} d\mu_f = 1.$$

On the other hand, considering (4.30) and applying (4.33) to $\varphi = \mathbb{1}$ leads to

$$\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} d\mu_f = 1 + \Delta t a_1 + \dots + \Delta t^p a_p + \Delta t^{p+1} r_{f,\Delta t},$$

where $r_{f,\Delta t}$ is a remainder term which is uniformly bounded for Δt sufficiently small. This is the manifestation at the discrete level of the fact that, over a time step Δt , the dynamics increases or decreases approximately the probability mass by a factor $e^{\Delta t \lambda}$. The relation (4.33) should thus be compared to the invariance relation (4.29) for $f = 0$.

4.3.1.3 Proof of Theorem 4.16

The proof of Theorem 4.16 relies on four lemmas which allow to easily conclude the proof. We follow the same strategy as for the error analysis of the invariant probability measure proposed in [295, 304] but additional technical difficulties arise due to the non-linearity of the stationarity equation (4.22). The first step (Lemma 4.17) is to construct the leading correction term ψ . We next use a projector in Lemma 4.18 to relate the exact stationary measure μ_f and its approximation $\mu_{f,\Delta t}$. An a priori estimate on the approximate eigenvalue defined in (4.23) is then provided in Lemma 4.19. Finally, an approximate inverse operator is constructed in Lemma 4.20. In the proofs and also in the statements below, the remainders may change from line to line in the calculation, but we do not change the notation for convenience. There are two types of remainders: terms of the form $R_{f,\Delta t}\varphi$ where $R_{f,\Delta t}$ is a differential operator satisfying (4.27), and functions $r_{f,\Delta t}$ such that, for any $k \geq 1$, there is $K > 0$ and Δt^* for which $\|r_{f,\Delta t}\|_{C^k} \leq K$ when $0 < \Delta t \leq \Delta t^*$.

To begin with, we give the expression of the leading correction term ψ . It relies on an approximate reformulation of (4.22) which leads to an expression similar to (4.24) up to a remainder of order Δt^{p+1} .

Lemma 4.17. *Under the assumptions of Theorem 4.16, for any $\varphi \in \mathcal{C}$,*

$$\int_{\mathcal{X}} \left(Q_{\Delta t}^f \varphi \right) (1 + \Delta t^p \psi) d\mu_f = \left(\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} (1 + \Delta t^p \psi) d\mu_f \right) \left(\int_{\mathcal{X}} \varphi (1 + \Delta t^p \psi) d\mu_f \right) + \Delta t^{p+2} R_{f,\Delta t} \varphi, \quad (4.36)$$

where $R_{f,\Delta t}$ is a real-valued uniformly bounded remainder (in the sense of (4.27)) and ψ is defined in (4.34).

The proof of this lemma is presented in Section 4.6.3.1. Defining the approximate eigenvalue $\tilde{\lambda}_{\Delta t}$ by

$$e^{\Delta t \tilde{\lambda}_{\Delta t}} = \int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} (1 + \Delta t^p \psi) d\mu_f, \quad (4.37)$$

equation (4.36) can be rewritten as

$$\forall \varphi \in \mathcal{C}, \quad \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \varphi \right] (1 + \Delta t^p \psi) d\mu_f = \Delta t^{p+1} R_{f,\Delta t} \varphi. \quad (4.38)$$

This expression allows to identify the leading order correction term $\Delta t^p \psi$ in $\mu_{f,\Delta t} - \mu_f$ and can be thought of as the approximate counterpart of (4.24). The second step is to use a projector that on the one hand stabilizes in \mathcal{C}_f the operator appearing in (4.38), and on the other hand relates the exact stationary measure μ_f and its approximation $\mu_{f,\Delta t}$. For this we introduce the following projectors: for all $\phi \in \mathcal{C}$,

$$\Pi \phi = \phi - \int_{\mathcal{X}} \phi d\nu, \quad \Pi_f \phi = \phi - \int_{\mathcal{X}} \phi d\mu_f. \quad (4.39)$$

The operator Π is the $L^2(\mu)$ orthogonal projector on $L_0^2(\mu)$, while Π_f is a projector on $L_f^2(\mu)$ which is not orthogonal for the canonical scalar product on $L^2(\mu)$. However, it is orthogonal on $L^2(\mu_f)$, so that, for all $\phi, \varphi \in \mathcal{C}$,

$$\int_{\mathcal{X}} (\Pi_f \phi) \varphi d\mu_f = \int_{\mathcal{X}} \phi (\Pi_f \varphi) d\mu_f. \quad (4.40)$$

We can then show the following result, whose proof can be found in Section 4.6.3.2.

Lemma 4.18. *Under the assumptions of Theorem 4.16, it holds, for any $\phi \in \mathcal{C}$,*

$$\int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_{f,\Delta t} = \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f + \Delta t^{p+1} R_{f,\Delta t} \phi, \quad (4.41)$$

where $R_{f,\Delta t}$ is a uniformly bounded remainder in the sense of (4.27).

Here, we see that two different operators appear inside the integrals because the factors $e^{\Delta t \tilde{\lambda}_{\Delta t}}$ and $e^{\Delta t \lambda_{\Delta t}}$ are different. The next lemma shows that these quantities are the same up to terms of order Δt^{p+2} . As mentioned earlier, this is an important difference with the analysis in the case $f = 0$. Some a priori estimate on the eigenvalue is required to conclude the proof, whereas, for the

unweighted case, the largest eigenvalue of the evolution operator is 1 with eigenvector $\mathbb{1}$ both for the continuous process and its discretization. The proof, provided in Section 4.6.3.3, relies on building an approximate eigenfunction for the operator $Q_{\Delta t}^f$. Similar estimates were obtained in the diffusion Monte Carlo context in analytically solvable cases in [328].

Lemma 4.19. *Under the assumptions of Theorem 4.16, there exist $\Delta t^* > 0$, $c > 0$ and functions $u_1, \dots, u_p \in \mathcal{C}_f$ such that the function $h_{f,\Delta t} = h_f + \Delta t u_1 + \dots + \Delta t^p u_p$ satisfies*

$$Q_{\Delta t}^f h_{f,\Delta t} = e^{\Delta t \bar{\lambda}_{\Delta t}} h_{f,\Delta t} + \Delta t^{p+2} r_{f,\Delta t}, \quad \int_{\mathcal{X}} h_{f,\Delta t} d\mu = 1, \quad (4.42)$$

where $\|r_{f,\Delta t}\|_{C^0} \leq c$ for all $0 < \Delta t \leq \Delta t^*$. As a consequence, there exist $\Delta t'$ and $C > 0$ such that

$$e^{\Delta t \lambda_{\Delta t}} = e^{\Delta t \bar{\lambda}_{\Delta t}} + \Delta t^{p+2} \tilde{r}_{f,\Delta t}, \quad (4.43)$$

with $|\tilde{r}_{f,\Delta t}| \leq C$ for all $0 < \Delta t \leq \Delta t'$.

Once we have reached this point, it is possible to replace the eigenvalue $e^{\Delta t \lambda_{\Delta t}}$ by $e^{\Delta t \bar{\lambda}_{\Delta t}}$ in Lemma 4.18. The last step is to build an approximate inverse of the operator

$$\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f,$$

as provided in the next lemma (see Section 4.6.3.4 for the proof).

Lemma 4.20. *Under the assumptions of Theorem 4.16, for any $0 < \Delta t \leq \Delta t^*$, there is an operator $S_{\Delta t}^f : \mathcal{C} \rightarrow \mathcal{C}$ for which*

$$\forall \varphi \in \mathcal{C}_f, \quad \Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f S_{\Delta t}^f \varphi = \varphi + \Delta t^{p+1} \mathcal{R}_{f,\Delta t} \varphi, \quad (4.44)$$

where $\mathcal{R}_{f,\Delta t} : \mathcal{C} \rightarrow \mathcal{C}$ is a uniformly bounded remainder in the sense of (4.27), and $S_{\Delta t}^f$ admits the following uniform bounds: for any $k \geq 0$, there exist $C > 0$ and $m \in \mathbb{N}$ (depending on k) such that

$$\forall \Delta t \in (0, \Delta t^*], \quad \left\| S_{\Delta t}^f \varphi \right\|_{C^k} \leq C \|\varphi\|_{C^m}.$$

We now have all the tools to prove Theorem 4.16. First, plugging the estimate (4.43) obtained in Lemma 4.19 in the error expansion (4.41) obtained in Lemma 4.18 leads to, for any $\phi \in \mathcal{C}$,

$$\int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_{f,\Delta t} = \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f + \Delta t^{p+1} R_{f,\Delta t} \phi, \quad (4.45)$$

where $R_{f,\Delta t}$ satisfies (4.27). We next consider the approximate inverse operator $S_{\Delta t}^f$ built in Lemma 4.20, and set $\phi = S_{\Delta t}^f \Pi_f \varphi$ in (4.45). Therefore, for any $\varphi \in \mathcal{C}$,

$$\int_{\mathcal{X}} (\Pi_f \varphi) d\mu_{f,\Delta t} = \int_{\mathcal{X}} (\Pi_f \varphi) (1 + \Delta t^p \psi) d\mu_f + \Delta t^{p+1} \tilde{R}_{f,\Delta t} \varphi = \Delta t^p \int_{\mathcal{X}} (\Pi_f \varphi) \psi d\mu_f + \Delta t^{p+1} \tilde{R}_{f,\Delta t} \varphi,$$

where $\tilde{R}_{f,\Delta t}$ satisfies (4.27). Since ψ has average 0 with respect to μ_f , this gives

$$\int_{\mathcal{X}} \varphi d\mu_{f,\Delta t} = \int_{\mathcal{X}} \varphi d\mu_f + \Delta t^p \int_{\mathcal{X}} \varphi \psi d\mu_f + \Delta t^{p+1} \tilde{R}_{f,\Delta t} \varphi,$$

which concludes the proof of Theorem 4.16.

4.3.2 Alternative error estimate for the principal eigenvalue

We present in this section a useful application of Theorem 4.16, which provides an error estimate for the approximation of the principal eigenvalue λ of the operator $\mathcal{L} + f$. Such an estimate is particularly useful for computing large deviations functions, as motivated in Chapter 3. We also illustrate this analysis in Chapter 5 below. The choice $\varphi \equiv f$ allows to compute the principal eigenvalue by ergodic averages, as shown in Proposition 4.3 and Corollary 4.4. As a result, this eigenvalue can be approximated using Theorem 4.16, whose application to $\varphi \equiv f$ gives

$$\int_{\mathcal{X}} f d\mu_{f,\Delta t} = \lambda + \Delta t^p \int_{\mathcal{X}} f\psi d\mu_f + \Delta t^{p+1} r_{f,\Delta t},$$

where $r_{f,\Delta t}$ is uniformly bounded for Δt small enough. Although this formula can be used in simulations to estimate λ , we present an error estimate for an alternative approximation more commonly used in practice. We will also see in Section 4.3.4 that this alternative formula can be more accurate than the estimate based on averaging f .

Theorem 4.21. *Suppose that Assumption 4.12 holds, with a numerical scheme consistent at order p (that is, $\mathcal{A}_k^f = (\mathcal{L} + f)^k/k!$ for $1 \leq k \leq p$). Then there exist $\Delta t^* > 0$ and $C > 0$ such that*

$$\lambda_{\Delta t} = \frac{1}{\Delta t} \log \left[\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} d\mu_{f,\Delta t} \right] = \lambda + \Delta t^p \left(\lambda_{p+1} - \frac{\lambda^{p+1}}{(p+1)!} \right) + \Delta t^{p+1} r_{f,\Delta t}, \quad (4.46)$$

with $|r_{f,\Delta t}| \leq C$ for any $0 < \Delta t \leq \Delta t^*$, and

$$\lambda_{p+1} = \int_{\mathcal{X}} \mathcal{A}_{p+1}^f \mathbb{1} d\mu_f + \int_{\mathcal{X}} f\psi d\mu_f. \quad (4.47)$$

This result is important since it implies that we can approximate the eigenvalue λ by computing $\lambda_{\Delta t}$, which is proportional to the logarithm of the average creation of probability over a time step Δt (given by $Q_{\Delta t}^f \mathbb{1}$) at stationarity. This is the reason why we need the coefficients a_k to be correct up to order p (i.e. $a_k = \lambda^k/k!$) since they represent the creation of probability of the discretized process. The estimate (4.46) justifies the use of population based dynamics [202, 397, 337] when the underlying continuous diffusions are discretized in time. We illustrate the error estimate (4.46) in the numerical simulations reported in Section 4.4.

Proof. We use Lemma 4.19 to prove the theorem, which highlights the importance of this result in our context. In all this proof, $r_{f,\Delta t}$ denotes a smooth function which may change from line to line, but whose C^0 norm is always uniformly bounded for sufficiently small time steps Δt . From the definition (4.23) and the estimate (4.43),

$$\lambda_{\Delta t} = \frac{1}{\Delta t} \log (e^{\Delta t \lambda_{\Delta t}}) = \frac{1}{\Delta t} \log (e^{\Delta t \bar{\lambda}_{\Delta t}} + \Delta t^{p+2} r_{f,\Delta t}).$$

Expanding $e^{\Delta t \bar{\lambda}_{\Delta t}}$ defined in (4.37) in powers of Δt and recalling that $\mathcal{A}_1^f \mathbb{1} = f$,

$$\begin{aligned} \lambda_{\Delta t} &= \frac{1}{\Delta t} \log \left[\int_{\mathcal{X}} \left(1 + \Delta t f + \Delta t^2 \mathcal{A}_2^f \mathbb{1} + \dots + \Delta t^{p+1} \mathcal{A}_{p+1}^f \mathbb{1} \right) (1 + \Delta t^p \psi) d\mu_f + \Delta t^{p+2} r_{f,\Delta t} \right] \\ &= \frac{1}{\Delta t} \log \left[1 + \Delta t \lambda + \dots + \Delta t^p \frac{\lambda^p}{p!} + \Delta t^{p+1} \int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \mathbb{1} + f\psi \right) d\mu_f + \Delta t^{p+2} r_{f,\Delta t} \right], \end{aligned}$$

where we used that $\int_{\mathcal{X}} \psi d\mu_f = 0$ and, in view of (4.8),

$$\forall k \in \{1, \dots, p\}, \quad \int_{\mathcal{X}} \mathcal{A}_k^f \mathbb{1} d\mu_f = \int_{\mathcal{X}} \left[\frac{(\mathcal{L} + f)^k}{k!} \mathbb{1} \right] \ell_f d\mu = \int_{\mathcal{X}} \left[\frac{(\mathcal{L}^* + f)^k}{k!} \ell_f \right] d\mu = \frac{\lambda^k}{k!}. \quad (4.48)$$

Therefore,

$$\begin{aligned} \lambda_{\Delta t} &= \frac{1}{\Delta t} \log \left[e^{\Delta t \lambda} - \sum_{k=p+2}^{+\infty} \Delta t^k \frac{\lambda^k}{k!} + \Delta t^{p+1} \int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \mathbb{1} - \frac{\lambda^{p+1}}{(p+1)!} + f\psi \right) d\mu_f + \Delta t^{p+2} r_{f,\Delta t} \right] \\ &= \frac{1}{\Delta t} \log \left[e^{\Delta t \lambda} + \Delta t^{p+1} \int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \mathbb{1} - \frac{\lambda^{p+1}}{(p+1)!} + f\psi \right) d\mu_f + \Delta t^{p+2} r_{f,\Delta t} \right]. \end{aligned}$$

Given that $e^{\Delta t \lambda}$ is uniformly bounded for $0 < \Delta t \leq \Delta t^*$ and equal to 1 at leading order in Δt , we obtain, by expanding the logarithm,

$$\lambda_{\Delta t} = \lambda + \Delta t^p e^{-\Delta t \lambda} \left[\int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \mathbb{1} + f\psi \right) d\mu_f - \frac{\lambda^{p+1}}{(p+1)!} \right] + \Delta t^{p+1} r_{f,\Delta t}.$$

The result then follows from $e^{-\Delta t \lambda} = 1 + \Delta t r_{\lambda,\Delta t}$ and the definition (4.47) of λ_{p+1} . \square

4.3.3 TU Lemma

In the context of splitting schemes, it may be useful to relate the invariant probability measures of two numerical schemes differing by the ordering of the applied operators. This is the purpose of a result called ‘‘TU lemma’’ in [295], which we adapt to our context in Lemma 4.22. We then state a similar version of this lemma for the eigenvalues of two such schemes in Proposition 4.23. We will see in Section 4.3.4 that this last result can be combined with Theorem 4.21 to show that the schemes (4.14) and (4.16) both provide second order estimates of the principal eigenvalue λ using (4.46), when the discretization of the process $Q_{\Delta t}$ is weakly consistent of order 2.

Lemma 4.22. *Consider two numerical schemes for the Feynman–Kac dynamics with associated evolution operators $Q_{\Delta t}^f$ and $\tilde{Q}_{\Delta t}^f$ satisfying Assumption 4.8, and denote by $\mu_{f,\Delta t}$ and $\tilde{\mu}_{f,\Delta t}$ respectively the associated ergodic measures in the sense of Theorem 4.9. Assume that the evolution operators are related by two operators $T_{\Delta t}^f$ and $U_{\Delta t}^f$, bounded on $B^\infty(\mathcal{X})$, as:*

$$\forall n \geq 1, \quad \left(\tilde{Q}_{\Delta t}^f \right)^n = T_{\Delta t}^f \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f. \quad (4.49)$$

Then, for any $\varphi \in \mathcal{C}$,

$$\int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t} = \frac{\int_{\mathcal{X}} U_{\Delta t}^f \varphi d\mu_{f,\Delta t}}{\int_{\mathcal{X}} U_{\Delta t}^f \mathbb{1} d\mu_{f,\Delta t}}. \quad (4.50)$$

For the TU lemma stated in [295], the typical case of application corresponds to $Q_{\Delta t} = U_{\Delta t} T_{\Delta t}$ and $\tilde{Q}_{\Delta t} = T_{\Delta t} U_{\Delta t}$, with two Markov operators $T_{\Delta t}$ and $U_{\Delta t}$. In this case, the relation (4.49) holds with a power $n - 1$ on the right-hand side. For Feynman–Kac semigroups, $T_{\Delta t}^f$ and $U_{\Delta t}^f$ are a priori such that $T_{\Delta t}^f \mathbb{1} \neq \mathbb{1}$ and $U_{\Delta t}^f \mathbb{1} \neq \mathbb{1}$. A typical case of interest is $Q_{\Delta t}^f = e^{\Delta t f} Q_{\Delta t}$ and $\tilde{Q}_{\Delta t}^f = Q_{\Delta t} (e^{\Delta t f} \cdot)$, in which case (4.49) is satisfied with $T_{\Delta t}^f = e^{-\Delta t f}$ and $U_{\Delta t}^f = e^{\Delta t f}$.

Proof. For any $\nu \in \mathcal{P}(\mathcal{X})$ and any $\varphi \in \mathcal{C}$,

$$\begin{aligned} \tilde{\Phi}_{\Delta t,n}^f(\nu)(\varphi) &= \frac{\nu \left(\left(\tilde{Q}_{\Delta t}^f \right)^n \varphi \right)}{\nu \left(\left(\tilde{Q}_{\Delta t}^f \right)^n \mathbb{1} \right)} = \frac{\nu \left(T_{\Delta t}^f \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \varphi \right)}{\nu \left(T_{\Delta t}^f \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \mathbb{1} \right)} \\ &= \frac{\left(\nu T_{\Delta t}^f \right) \left(\left(Q_{\Delta t}^f \right)^n \mathbb{1} \right)}{\left(\nu T_{\Delta t}^f \right) \left(\left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \mathbb{1} \right)} \times \frac{\left(\nu T_{\Delta t}^f \right) \left(\left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \varphi \right)}{\left(\nu T_{\Delta t}^f \right) \left(\left(Q_{\Delta t}^f \right)^n \mathbb{1} \right)} = \frac{\Phi_{\Delta t,n}^f(\nu_1)(U_{\Delta t}^f \varphi)}{\Phi_{\Delta t,n}^f(\nu_1)(U_{\Delta t}^f \mathbb{1})}, \end{aligned}$$

where $\nu_1 \in \mathcal{P}(\mathcal{X})$ is defined by

$$\forall \phi \in \mathcal{C}, \quad \nu_1(\phi) = \frac{\nu \left(T_{\Delta t}^f \phi \right)}{\nu \left(T_{\Delta t}^f \mathbb{1} \right)}.$$

The result then follows from the ergodic limits

$$\lim_{n \rightarrow +\infty} \Phi_{\Delta t,n}^f(\nu_1)(\varphi) = \int_{\mathcal{X}} \varphi d\mu_{f,\Delta t}, \quad \lim_{n \rightarrow +\infty} \tilde{\Phi}_{\Delta t,n}^f(\nu)(\varphi) = \int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t},$$

as provided by Theorem 4.9. \square

In our framework, the approximate principal eigenvalue $\lambda_{\Delta t}$ is another important feature of a discretization scheme. In fact, under an additional assumption on the operators $T_{\Delta t}^f$ and $U_{\Delta t}^f$, schemes related by (4.49) share the same approximate eigenvalues in the sense of (4.23). This is made precise in the following proposition (see Section 4.6.4 for the proof).

Proposition 4.23. *Fix a time step $\Delta t > 0$ and consider a numerical scheme for the Feynman–Kac dynamics corresponding to an evolution operator $Q_{\Delta t}^f$ satisfying Assumption 4.8, with associated invariant measure $\mu_{f,\Delta t}$ given by Theorem 4.9, and eigenvalue $\lambda_{\Delta t}$ defined by (4.23). Consider next a second scheme corresponding to an operator $\tilde{Q}_{\Delta t}^f$ related to $Q_{\Delta t}^f$ by (4.49), with operators $U_{\Delta t}^f$ and $T_{\Delta t}^f$ bounded on $B^\infty(\mathcal{X})$ and for which there exists $\alpha > 0$ such that, for any $\varphi \in \mathcal{C}$ with $\varphi \geq 0$,*

$$\alpha\varphi \leq U_{\Delta t}^f\varphi \leq \alpha^{-1}\varphi, \quad \alpha\varphi \leq T_{\Delta t}^f\varphi \leq \alpha^{-1}\varphi. \quad (4.51)$$

Then, $\tilde{Q}_{\Delta t}^f$ satisfies Assumption 4.8, and its invariant probability measure is denoted by $\tilde{\mu}_{f,\Delta t}$. Moreover, its associated eigenvalue $\tilde{\lambda}_{\Delta t}$ defined by

$$\tilde{\lambda}_{\Delta t} = \frac{1}{\Delta t} \log \left[\int_{\mathcal{X}} \tilde{Q}_{\Delta t}^f \mathbf{1} d\tilde{\mu}_{f,\Delta t} \right], \quad (4.52)$$

is such that

$$\tilde{\lambda}_{\Delta t} = \lambda_{\Delta t}.$$

The eigenvalue $\tilde{\lambda}_{\Delta t}$ should not be mistaken in this context with the definition (4.37), which serves as an intermediate in the proof of Theorem 4.16. A careful inspection of the proof shows that it would be possible to consider a slightly different assumption (4.51).

Remark 4.24. *Although Proposition 4.23 may look odd at first sight, it has a natural interpretation in terms of matrices. Indeed, if $A \in \mathbb{R}^{d \times d}$ and $B \in \mathbb{R}^{d \times d}$ are two square matrices with nonnegative entries, the products AB and BA share the same real principal eigenvalue. One can show this by the following argument. For any matrix $M \in \mathbb{R}^{d \times d}$ with nonnegative entries, the spectral radius*

$$\rho(M) = \lim_{n \rightarrow +\infty} \|M^n\|^{\frac{1}{n}}$$

is an eigenvalue of M (see [383]). This eigenvalue is the equivalent of the principal eigenvalue for the operator $\mathcal{L} + f$ since it is the eigenvalue of the matrix M with the largest real part. It is easy to see that $\rho(AB) = \rho(BA)$ by noting that

$$\rho(AB) = \lim_{n \rightarrow +\infty} \|(AB)^n\|^{\frac{1}{n}} = \lim_{n \rightarrow +\infty} \|A(BA)^{n-1}B\|^{\frac{1}{n}} \leq \lim_{n \rightarrow +\infty} \|A\|^{\frac{1}{n}} \|(BA)^{n-1}\|^{\frac{1}{n}} \|B\|^{\frac{1}{n}} = \rho(BA).$$

This leads to $\rho(AB) \leq \rho(BA)$, and, by symmetry, $\rho(BA) \leq \rho(AB)$; hence $\rho(AB) = \rho(BA)$. In the same way, evolution operators related by (4.49) share the same principal eigenvalue even though, a priori, they do not share the same invariant probability measures. The proof of Proposition 4.23, presented in Section 4.6.4, follows a path similar to the one used here for matrices.

4.3.4 Second order schemes

We now turn to second order schemes for Feynman–Kac dynamics. They are the most interesting ones in practice, since they can provide an important improvement in the accuracy for a relatively cheap computational overhead. Moreover, in our case, they can be straightforwardly built from second order schemes for the dynamics (4.1), as a consequence of Theorem 4.16.

Lemma 4.25. *Suppose that (4.25) and (4.27) hold with the following expansion for $Q_{\Delta t}$:*

$$\forall \varphi \in \mathcal{C}, \quad Q_{\Delta t}\varphi = \varphi + \Delta t \mathcal{L}\varphi + \Delta t^2 \frac{\mathcal{L}^2\varphi}{2} + \Delta t^3 \mathcal{A}_3\varphi + \Delta t^4 \mathcal{R}_{\Delta t}\varphi, \quad (4.53)$$

where \mathcal{A}_3 is a differential operator with smooth coefficients and $\mathcal{R}_{\Delta t}$ satisfies (4.27). Then the operator $Q_{\Delta t}^f$ defined by

$$\forall \varphi \in \mathcal{C}, \quad Q_{\Delta t}^f\varphi = e^{\frac{\Delta t}{2}f} Q_{\Delta t} \left(e^{\frac{\Delta t}{2}f} \varphi \right),$$

satisfies Assumption 4.12 with $p = 2$:

$$\forall \varphi \in \mathcal{C}, \quad Q_{\Delta t}^f \varphi = \varphi + \Delta t(\mathcal{L} + f)\varphi + \Delta t^2 \frac{(\mathcal{L} + f)^2 \varphi}{2} + \Delta t^3 \mathcal{A}_3^f \varphi + \Delta t^4 \mathcal{R}_{f, \Delta t} \varphi, \quad (4.54)$$

where

$$\mathcal{A}_3^f \varphi = \mathcal{A}_3 \varphi + \frac{f^3 \varphi}{6} + \frac{\mathcal{L}(f^2 \varphi)}{8} + \frac{\mathcal{L}^2(f \varphi)}{4} + \frac{f \mathcal{L}^2 \varphi}{4} + \frac{f \mathcal{L}(f \varphi)}{4} + \frac{f^2 \mathcal{L} \varphi}{8},$$

and $\mathcal{R}_{f, \Delta t}$ satisfies (4.27).

The interpretation of this result is the following: when we have a scheme consistent at order 2 for the dynamics with $f = 0$, we immediately obtain a second order scheme for the Feynman–Kac dynamics by using the corresponding Markov chain and a trapezoidal rule for the time integral in the exponential. Thanks to the consistency at order one ($\mathcal{A}_1^f = \mathcal{L} + f$) and Propositions 4.1 and 4.5, the assumptions of Theorems 4.16 and 4.21 are immediately satisfied with $p = 2$.

Proof. The expression of \mathcal{A}_3^f can be obtained by a direct computation or with the Baker–Campbell–Hausdorff formula (see [215]), which is a convenient way to perform the algebra allowing to make precise the various terms in expansions in powers of Δt . Let us sketch how this is done, and refer to [295] for strategies of proof in order to make the expansions below rigorous. First,

$$Q_{\Delta t} = e^{\Delta t \mathcal{L}} + \Delta t^3 \left(\mathcal{A}_3 - \frac{\mathcal{L}^3}{6} \right) + \dots$$

and, by the Baker–Campbell–Hausdorff formula,

$$e^{\Delta t f/2} e^{\Delta t \mathcal{L}} e^{\Delta t f/2} = e^{S_{\Delta t}}, \quad S_{\Delta t} = \Delta t(\mathcal{L} + f) + \frac{\Delta t^3}{12} \left(-\frac{1}{2} [f, [f, \mathcal{L}]] + [\mathcal{L}, [\mathcal{L}, f]] \right),$$

where $[A, B] = AB - BA$ denotes the commutator of two operators A and B . Therefore,

$$e^{\Delta t f/2} e^{\Delta t \mathcal{L}} e^{\Delta t f/2} = I_d + \Delta t(\mathcal{L} + f) + \frac{\Delta t^2}{2} (\mathcal{L} + f)^2 + \frac{\Delta t^3}{6} (\mathcal{L} + f)^3 + \frac{\Delta t^3}{12} \left(-\frac{1}{2} [f, [f, \mathcal{L}]] + [\mathcal{L}, [\mathcal{L}, f]] \right) + \dots$$

The conclusion then follows from

$$e^{\Delta t f/2} Q_{\Delta t} e^{\Delta t f/2} = e^{\Delta t f/2} e^{\Delta t \mathcal{L}} e^{\Delta t f/2} + \Delta t^3 \left(\mathcal{A}_3 - \frac{\mathcal{L}^3}{6} \right) + \dots$$

upon developping the commutators. \square

When we are interested in the computation of the principal eigenvalue with Theorem 4.21, we can in fact show that the left-point integration (4.14) is sufficient for $\lambda_{\Delta t}$ to be correct at order 2 if $Q_{\Delta t}$ is consistent at order 2 (*i.e.* (4.53) holds). In particular, the discretization scheme for the Feynman–Kac dynamics need not be consistent at order 2 for the eigenvalue to be correct at order 2 (in the same way that the invariant probability measure for discretizations of ergodic SDEs can be correct at order 2 even if the discretization itself is only weakly consistent at order 1, see [3, 295, 304]). This consequence of Proposition 4.23 is made precise in the following proposition.

Proposition 4.26. *Consider an evolution operator $Q_{\Delta t}$ with the following family of discretizations for the Feynman–Kac dynamics:*

$$Q_{\Delta t}^{f, \delta} = e^{(1-\delta)f\Delta t} Q_{\Delta t} e^{\delta f\Delta t}, \quad \delta \in [0, 1].$$

Suppose that Assumption 4.8 holds for at least one of these schemes, and denote by $\lambda_{\Delta t}^\delta$ their associated eigenvalues as in (4.23). Then, $\lambda_{\Delta t}^\delta$ is independent of δ . Moreover, when $Q_{\Delta t}$ satisfies (4.53), the eigenvalue $\lambda_{\Delta t}^\delta$ satisfies (4.46) with $p = 2$ for any $\delta \in [0, 1]$.

Proof. The proof is a simple application of Proposition 4.23. Consider the scheme $Q_{\Delta t}^{f, \delta}$ for a fixed $\delta \in [0, 1]$ and the scheme $\tilde{Q}_{\Delta t}^f = e^{\Delta t \frac{f}{2}} Q_{\Delta t} e^{\Delta t \frac{f}{2}}$, which corresponds to a trapezoidal approximation of the integral. We can assume without loss of generality that $\tilde{Q}_{\Delta t}^f$ satisfies Assumption 4.8. Then, $Q_{\Delta t}^{f, \delta}$ is related to $\tilde{Q}_{\Delta t}^f$ through (4.49) for the corresponding operators:

$$U_{\Delta t}^f = e^{(\delta - \frac{1}{2})\Delta t f}, \quad T_{\Delta t}^f = e^{(\frac{1}{2} - \delta)\Delta t f}.$$

The operators $U_{\Delta t}^f$ and $T_{\Delta t}^f$ are bounded on B^∞ and satisfy (4.51) with $\alpha = e^{-\Delta t \|f\|_{B^\infty}/2}$. Therefore, by Proposition 4.23, the eigenvalue $\lambda_{\Delta t}^\delta$ associated to $Q_{\Delta t}^{f,\delta}$ is equal to $\tilde{\lambda}_{\Delta t}$, the eigenvalue associated to $\tilde{Q}_{\Delta t}^f$, and thus does not depend on δ . Moreover, by Lemma 4.25, if $Q_{\Delta t}$ satisfies (4.53), $\tilde{Q}_{\Delta t}^f$ satisfies the assumptions of Theorem 4.21 with $p = 2$. This shows that the eigenvalue $\lambda_{\Delta t}^\delta$ satisfies (4.46) with $p = 2$ whatever the integration rule (*i.e.* for any $\delta \in [0, 1]$). \square

Remark 4.27. Proposition 4.26 shows that the eigenvalue $\lambda_{\Delta t}$ can be correct at order two even though the scheme only has weak order one. One may also wonder whether it is also possible to have second order convergence on the invariant measure when $Q_{\Delta t}$ corresponds to a scheme of weak order one. As mentioned in Section 4.3.1.1 this is the case when $f = 0$, see the examples in [295]. Perturbative arguments for small f however show that this extra cancellation on the invariant measure cannot happen for a non-constant f , which we detail below.

To be more specific, we consider a numerical scheme $Q_{\Delta t}$ for an SDE with generator \mathcal{L} such that

$$Q_{\Delta t} = I_d + \Delta t \mathcal{L} + \frac{\Delta t^2}{2} \mathcal{A}_2 + O(\Delta t^3),$$

with $\mathcal{A}_2 \neq \mathcal{L}^2$ but

$$\int_{\mathcal{X}} \mathcal{A}_2 \varphi d\mu = 0.$$

We know such schemes exists from [295], which is due to particular algebraic cancellations. For exploring whether this situation is still possible when $f \neq 0$, we use a perturbative argument with a potential εf for $\varepsilon \ll 1$. Denoting by $\lambda_\varepsilon, h_\varepsilon$ the spectral elements (instead of $\lambda_{\varepsilon f}, h_{\varepsilon f}$), it holds

$$(\mathcal{L}^* + \varepsilon f)h_\varepsilon = \lambda_\varepsilon h_\varepsilon. \quad (4.55)$$

Expanding in powers of ε , we find $h_\varepsilon = 1 + \varepsilon u_1 + \varepsilon^2 u_2 + O(\varepsilon^3)$, and $\lambda_\varepsilon = \varepsilon \lambda_1 + \varepsilon^2 \lambda_2 + O(\varepsilon^3)$. Plugging these expressions in (4.55), we obtain that $\lambda_1 = \lambda$ and u_1 is solution to

$$-\mathcal{L}^* u_1 = \Pi_0 f,$$

where Π_0 is a projection operator on functions of average 0 with respect to μ . In particular, we have $u_1 = -(\mathcal{L}^*)^{-1} \Pi_0 f$. Now, a natural way to obtain a second order discretization for the Feynman–Kac dynamics is to set

$$Q_{\Delta t}^f = e^{\varepsilon \frac{f \Delta t}{2}} Q_{\Delta t} e^{\varepsilon \frac{f \Delta t}{2}} = I_d + \Delta t (\mathcal{L} + f) + \frac{\Delta t^2}{2} \mathcal{A}_2^\varepsilon + O(\Delta t^3),$$

with

$$\mathcal{A}_2^\varepsilon = \mathcal{A}_2 + \varepsilon f \mathcal{L} + \varepsilon \mathcal{L} f + \varepsilon^2 f^2. \quad (4.56)$$

We used the notation $\mathcal{A}_2^\varepsilon$ instead of $\mathcal{A}_2^{\varepsilon f}$ compared to the notation in Section 4.3 for simplicity. If $\mathcal{A}_2 = \mathcal{L}^2$, we are in the setting described in Section 4.3.4. Otherwise we should check that condition (4.33) holds at second order, namely

$$\int_{\mathcal{X}} (\mathcal{A}_2^\varepsilon \varphi) h_\varepsilon d\mu = a_2^\varepsilon \int_{\mathcal{X}} \varphi h_\varepsilon d\mu,$$

with $a_2^\varepsilon = a_{2,0} + \varepsilon a_{2,1} + O(\varepsilon^2)$. Expanding the various terms and identifying each order in ε , we find at dominant order in ε the condition

$$\int_{\mathcal{X}} \mathcal{A}_2 \varphi d\mu = a_{2,0} \int_{\mathcal{X}} \varphi d\mu,$$

which imposes $a_{2,0} = 0$ by assumption on the numerical scheme. At first order in ε , we find the equation

$$\int_{\mathcal{X}} (\mathcal{A}_2 \varphi) u_1 d\mu + \int_{\mathcal{X}} [(f \mathcal{L} + \mathcal{L} f) \varphi] d\mu = a_{2,0} \int_{\mathcal{X}} \varphi u_1 d\mu + a_{2,1} \int_{\mathcal{X}} \varphi d\mu.$$

Taking the adjoint in $L^2(\mu)$, since $a_{2,0} = 0$, we are led to

$$\mathcal{A}_2^* u_1 + (f \mathcal{L} + \mathcal{L} f)^* \mathbf{1} = a_{2,1},$$

for some $a_{2,1} \in \mathbb{R}$. The above equation rewrites, since $\mathcal{L}^* \mathbf{1} = 0$ and $u_1 = -(\mathcal{L}^*)^{-1} \Pi_0 f$,

$$[-\mathcal{A}_2^* (\mathcal{L}^*)^{-1} \Pi_0 + \mathcal{L}^*] f = a_{2,1}. \quad (4.57)$$

In the case $f \equiv \mathbb{1}$, the equation (4.57) boils down to

$$\mathcal{L}^*\mathbb{1} = a_{2,1},$$

which is indeed satisfied for $a_{2,1} = 0$. Similarly, (4.57) holds true for $a_{2,1} = 0$ when $\mathcal{A}_2 = \mathcal{L}^2$. On the other hand, when $\mathcal{A}_2 \neq \mathcal{L}^2$, it is impossible for (4.57) to be satisfied for any f . As a result, it seems difficult to construct numerical schemes with weak order one and an extra cancellation on the invariant measure, apart from very particular choices of f .

4.4 Numerical application

The goal of this section is to illustrate the error estimates presented in Section 4.3 on a toy example. For this, we consider (4.1) over the one dimensional torus $\mathcal{X} = \mathbb{T}$ with possibly non-gradient drifts:

$$dX_t = (-V'(X_t) + \gamma) dt + \sigma dB_t, \quad (4.58)$$

where V is a smooth potential and $\gamma \in \mathbb{R}$. Let us emphasize that a constant force is not the gradient of a smooth periodic function. We first make precise in Section 4.4.1 the Monte Carlo algorithm used to compute the Feynman–Kac averages. We next describe in Section 4.4.2 a Galerkin method to compute reference values for the properties of interest. Note that such a discretization method can be used only for low-dimensional systems; but, when it can be used, it typically provides more accurate results than stochastic methods. Finally, we present our numerical results in Section 4.4.3.

4.4.1 Monte Carlo discretization

Discretization of the underlying SDE. The Euler–Maruyama discretization of the dynamics (4.58) is given by:

$$x_{n+1} = x_n + (-V'(x_n) + \gamma)\Delta t + \sigma\sqrt{\Delta t}G^n, \quad (4.59)$$

where G^n are independent and identically distributed one-dimensional standard Gaussian variables. It is well known that this scheme is weakly consistent of order one (see for instance [332, 110]). In order to test our results on a second order scheme, we use a discretization proposed *e.g.* in [1, 436, 177, 409]:

$$x_{n+1} = x_n - V' \left(x_n + (-V'(x_n) + \gamma) \frac{\Delta t}{2} + \frac{1}{2} \sigma \sqrt{\Delta t} G^n \right) \Delta t + \gamma \Delta t - \frac{\sigma^2}{8} V'''(x_n) \Delta t^2 + \sigma \sqrt{\Delta t} G^n. \quad (4.60)$$

It can be proved that this scheme is of weak order 2.

Weighted dynamics. Once the underlying SDE has been discretized, a Monte Carlo scheme for approximating the associated Feynman–Kac semigroup (4.15) has to be devised. Several methods have been successfully applied in order to compute Feynman–Kac averages, generally referred to as *Sequential Monte Carlo* or *Population Monte Carlo* methods [139, 113, 301], see Section 1.4.2. For simplicity and numerical efficiency, we present here a population method with multinomial resampling. More precisions on this family of algorithms are available in [139], see also [301, Chapter 6] in the context of free energy computation and [221] in the context of diffusion Monte Carlo.

The algorithm relies on a dynamics run over a set of replicas of the system. At each step, the replicas are updated according to the dynamics prescribed by the evolution operator $Q_{\Delta t}$, and are assigned an importance weight depending on the choice of discretization rule for the integral. The replicas are then resampled following a multinomial distribution with their respective weights, before computing the desired averages. This technique prevents the variance of the estimator to increase exponentially in time, a common problem when computing directly quantities such as (4.15). We now make precise the algorithm.

Consider a population of M replicas $(x^m)_{m=1,\dots,M}$ initially distributed according to some probability measure $\nu^{\otimes M}$ over \mathcal{X}^M and evolving through a Markov kernel $Q_{\Delta t}$ with time step $\Delta t > 0$. We denote by $\chi_{\Delta t} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ a weight function to be chosen later on. The algorithm consists in repeating for each time $0 \leq n < N_{\text{iter}}$ the following steps:

- (1) For $m \in \{1, \dots, M\}$, evolve the m^{th} replica as $\tilde{x}_{n+1}^m \sim Q_{\Delta t}(x_n^m, \cdot)$;
- (2) Compute the weight of each replica as $w_n^m = e^{\chi_{\Delta t}(x_n^m, \tilde{x}_{n+1}^m)}$;

(3) Compute the total creation of mass as

$$P_n = \sum_{m=1}^M w_n^m,$$

and the normalized probability vector $p_n \in \mathbb{R}^m$ with components $p_n^m = w_n^m/P_n$, for $m \in \{1, \dots, M\}$;

(4) Resample the replicas $(\tilde{x}_{n+1}^m)_{m=1}^M$ according to the multinomial distribution associated with p_n , which defines a new set of replicas $(x_{n+1}^m)_{m=1}^M$;

(5) Compute the estimator

$$\hat{\varphi}_n = \frac{1}{M} \sum_{m=1}^M \varphi(x_{n+1}^m).$$

Until now, we did not specify the choice of function $\chi_{\Delta t}$, which depends on the discretization rule for the integral in (4.3). In practice, given a discretization of the SDE characterized by an operator $Q_{\Delta t}$, we use the schemes defined by the left point integration $e^{\Delta t f} Q_{\Delta t}$, and by the trapezoidal integration $e^{\Delta t \frac{f}{2}} Q_{\Delta t} e^{\Delta t \frac{f}{2}}$. They correspond respectively to the choices:

$$\chi_{\Delta t}(x, x') = \Delta t f(x) \quad \text{and} \quad \chi_{\Delta t}(x, x') = \Delta t \left(\frac{f(x) + f(x')}{2} \right). \quad (4.61)$$

The principal eigenvalue of the operator $\mathcal{L} + f$ is then estimated with (4.46) through

$$\lambda_{\Delta t} = \frac{1}{\Delta t} \log \left[\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} \, d\mu_{f, \Delta t} \right] \approx \frac{1}{\Delta t} \log \left[\frac{1}{MN_{\text{iter}}} \sum_{n=0}^{N_{\text{iter}}-1} P_n \right], \quad (4.62)$$

while the average of φ is estimated by

$$\int_{\mathcal{X}} \varphi \, d\mu_{f, \Delta t} \approx \frac{1}{N_{\text{iter}}} \sum_{n=0}^{N_{\text{iter}}-1} \hat{\varphi}_n, \quad (4.63)$$

where the \approx sign indicates the approximation arising from the finiteness of the number M of replicas and of the number N_{iter} of steps. We do not take these errors into account and ensure numerically that they are sufficiently small in our simulations to observe the bias due to the time step (this bias being quite small in practice, this also motivates to study a one-dimensional model, see the numerical results below). The reader interested in the convergence rates of this type of algorithm when $M \rightarrow +\infty$ and $N_{\text{iter}} \rightarrow +\infty$ is referred *e.g.* to [117, 139, 377, 164].

4.4.2 Galerkin discretization

We now make precise the Galerkin method that can be used to estimate $\lambda_{\Delta t}$ and $\int_{\mathcal{X}} \varphi \, d\mu_{f, \Delta t}$. This discretization provides reference values for the Monte Carlo method described in Section 4.4.1. In particular, when $V = 0$ and $\gamma = 0$, the two methods should give the same result since the Euler scheme (4.59) is exact in law in this specific case.

Choice of the Galerkin basis. Since we work with periodic functions, we consider the Galerkin subspace $\text{Span}\{e_{-N}, \dots, e_N\}$ with

$$e_k(x) = e^{2i\pi kx}.$$

The generator of the SDE (4.58) reads

$$\mathcal{L} = (-V' + \gamma)\partial_x + \frac{\sigma^2}{2}\partial_x^2.$$

The operators \mathcal{L}^\dagger and f are represented in this Galerkin subspace by the matrices $L^N, B^N \in \mathbb{C}^{(2N+1) \times (2N+1)}$ defined as

$$\forall j, k \in \{-N, \dots, 0, \dots, N\}, \quad L_{j,k}^N = \int_{\mathcal{X}} e_j(\mathcal{L}^\dagger e_k), \quad B_{j,k}^N = \int_{\mathcal{X}} f e_j e_k.$$

The value of N is chosen sufficiently large for all results to be converged with respect to this parameter. The only source of error in the quantities we compute then arises from the finiteness of the time step $\Delta t > 0$, and possibly numerical quadratures to evaluate certain integrals. Our experience shows that $N = 30$ is already sufficient for the applications described in Section 4.4.3.

References quantities for $\Delta t = 0$. The invariant probability measure μ_f satisfies the eigenvalue problem $(\mathcal{L}^\dagger + f)\mu_f = \lambda\mu_f$. We compute a reference approximation λ_0^N to λ by computing the eigenvalue of $L^N + B^N$ with the largest real part:

$$(L^N + B^N)\mathcal{V}_{f,0}^N = \lambda_0^N \mathcal{V}_{f,0}^N.$$

The associated eigenvector allows to construct the following approximation of μ_f :

$$\mu_{f,0}^N = \sum_{k=-N}^N [\mathcal{V}_{f,0}^N]_k e_k.$$

The normalization condition $[\mathcal{V}_{f,0}^N]_0 = 1$ ensures that $\mu_{f,0}^N$ has a total mass 1. Averages of observables φ are then estimated by computing the following integral

$$\int_{\mathcal{X}} \varphi(x) \mu_{f,0}^N(x) dx$$

using a one-dimensional quadrature rule.

Reference quantities for $\Delta t > 0$. We next approximate the evolution operators of the first order scheme $e^{\Delta t f} Q_{\Delta t}$ and of the second order one $e^{\Delta t \frac{f}{2}} Q_{\Delta t} e^{\Delta t \frac{f}{2}}$, respectively as

$$Q_{\Delta t,1}^{f,N} = e^{\Delta t B^N} e^{\Delta t L^N}, \quad Q_{\Delta t,2}^{f,N} = e^{\Delta t \frac{B^N}{2}} e^{\Delta t L^N} e^{\Delta t \frac{B^N}{2}}. \quad (4.64)$$

For each value of Δt , we construct the above matrices, and compute their respective principal eigenvalues $\Lambda_{\Delta t,1}^N, \Lambda_{\Delta t,2}^N$ and eigenvectors $\mathcal{V}_{f,\Delta t}^{N,1}, \mathcal{V}_{f,\Delta t}^{N,2} \in \mathbb{C}^{2N+1}$ by diagonalization (still with the normalization condition $[\mathcal{V}_{f,\Delta t}^{N,j}]_0 = 1$ for $j = 1, 2$). We then consider the following approximations of the principal eigenvalue λ of the Feynman–Kac operator $\mathcal{L} + f$, based on (4.46):

$$\lambda_{\Delta t,1}^N = \frac{1}{\Delta t} \log \Lambda_{\Delta t,1}^N, \quad \lambda_{\Delta t,2}^N = \frac{1}{\Delta t} \log \Lambda_{\Delta t,2}^N. \quad (4.65)$$

Averages of φ with respect to the invariant probability measure are approximated by the following quantity, using the eigenvectors $\mathcal{V}_{f,\Delta t}^{N,1}$ and $\mathcal{V}_{f,\Delta t}^{N,2}$: for $j = 1, 2$,

$$\int_{\mathcal{X}} \varphi(x) \mu_{f,\Delta t}^{N,j}(x) dx, \quad \mu_{f,\Delta t}^{N,j} = \sum_{k=-N}^N [\mathcal{V}_{f,\Delta t}^{N,j}]_k e_k. \quad (4.66)$$

In view of Theorem 4.16, we expect the average of φ to converge linearly in Δt for the first order scheme when $\Delta t \rightarrow 0$, and quadratically for the second order scheme. We also use the TU-lemma to show that, by appropriately correcting the first order scheme, we recover the same results as for the second order scheme. More precisely, we apply (4.50) with $U_{\Delta t}^f = e^{\Delta t \frac{f}{2}}$, which leads to the following approximation of the average (estimated in practice using a numerical quadrature):

$$\frac{\int_{\mathcal{X}} e^{\Delta t \frac{f(x)}{2}} \varphi(x) \mu_{f,\Delta t}^{N,1}(x) dx}{\int_{\mathcal{X}} e^{\Delta t \frac{f(x)}{2}} \mu_{f,\Delta t}^{N,1}(x) dx}. \quad (4.67)$$

On the other hand, from Proposition 4.26, the eigenvalues $\lambda_{\Delta t,1}^N$ and $\lambda_{\Delta t,2}^N$ should be equal, and therefore $\lambda_{\Delta t,1}^N$ need not be corrected.

4.4.3 Numerical results

Zero-potential case. We first choose $V = 0$, $\sigma = \sqrt{2}$, $f(x) = (\cos(2\pi x))^2$ and $\varphi(x) = \exp(\cos(2\pi x))$. As mentioned earlier, in this case, the Euler scheme (4.59) is exact in law, so that the only source of error arises from the integration of the exponential weight. We consider the dynamics represented by the operator $e^{\Delta t f} Q_{\Delta t}$ and $e^{\Delta t \frac{f}{2}} Q_{\Delta t} e^{\Delta t \frac{f}{2}}$ with $Q_{\Delta t} = e^{\Delta t \mathcal{L}}$, and first compare the results of the Galerkin discretization discussed in Section 4.4.2. The results reported in Figure 4.1 confirm our predictions: the averages of φ converge at first and second order for the first order and second order




Figure 4.1 – Estimated error on the principal eigenvalue (left) and on the average of φ with respect to the invariant probability measure (right) as a function of the time step, by Galerkin approximation. The eigenvalues are computed with (4.65). The first and second order averages of φ correspond to (4.66) with $j = 1$ and $j = 2$ respectively. The first order-TU scheme is computed with (4.67). The dashed lines show reference first and second order convergences.

Galerkin schemes respectively; while the eigenvalues are the same, as expected from Proposition 4.26, and so both converge at second order. In this case, the numerical method based on (4.46) is therefore more accurate than the one based on averaging f with (4.35) to compute the principal eigenvalue λ , which would lead to errors of order 1 in the time step (numerical results not shown here).

We next consider the Monte Carlo scheme presented in Section 4.4.1, taking $M = 5 \times 10^4$ and an integration time $T = 5 \times 10^2$, with $N_{\text{iter}} = \lfloor \frac{T}{\Delta t} \rfloor$ for each time step Δt . We use half of the time for burn-in, and average in time over the second half of the simulation. Moreover, for each value of Δt , we run 30 realizations in order to reduce the variance of the estimator and to estimate error bars on the Monte Carlo estimates (not displayed on the pictures). The choice of the function $\chi_{\Delta t}$ depends on the scheme through (4.61). We compare in Figure 4.2 the results of the Monte Carlo algorithm with the Galerkin approximation, which serves as a reference. The agreement is very good, up to small errors arising from the finiteness of the population and of the simulation time. This result was expected since, given that the integration by the Euler scheme is exact in law in this case, the Monte Carlo method must match exactly the Galerkin approximation provided N , N_{iter} and M are all sufficiently large.

Situation with a strong potential. We next show an application with a non-zero drift by setting $V(x) = \cos(2\pi x)$ and $\gamma = 1$. Let us recall that this dynamics is non-reversible since a constant function is not the gradient of a smooth periodic potential. The other parameters are left unchanged. Concerning the Galerkin approximation, we consider the two schemes described in Section 4.4.2, and characterized by the matrices defined in (4.64). For these schemes, the eigenvalues are the same and converge at second order (so we only consider one scheme), while the averages of φ converge at first and second order respectively.

For the Monte Carlo algorithm described in Section 4.4.1, we consider the three following schemes:

- $Q_{\Delta t}$ is discretized with the Euler scheme (4.59), and $\chi_{\Delta t}(x, x') = \Delta t f(x)$ is chosen as the left point integration; in this case, the eigenvalue and the average of φ converge at order one, so the scheme is referred to as *first order*.
- $Q_{\Delta t}$ is discretized with the second order scheme (4.60), and we set

$$\chi_{\Delta t}(x, x') = \Delta t f(x).$$

In this case, the eigenvalue converges at second order whereas the average of φ converges at first order only, so the scheme is referred to as *hybrid scheme*.

- $Q_{\Delta t}$ is discretized with the second order scheme (4.60), and we set

$$\chi_{\Delta t}(x, x') = \Delta t \left(\frac{f(x) + f(x')}{2} \right),$$

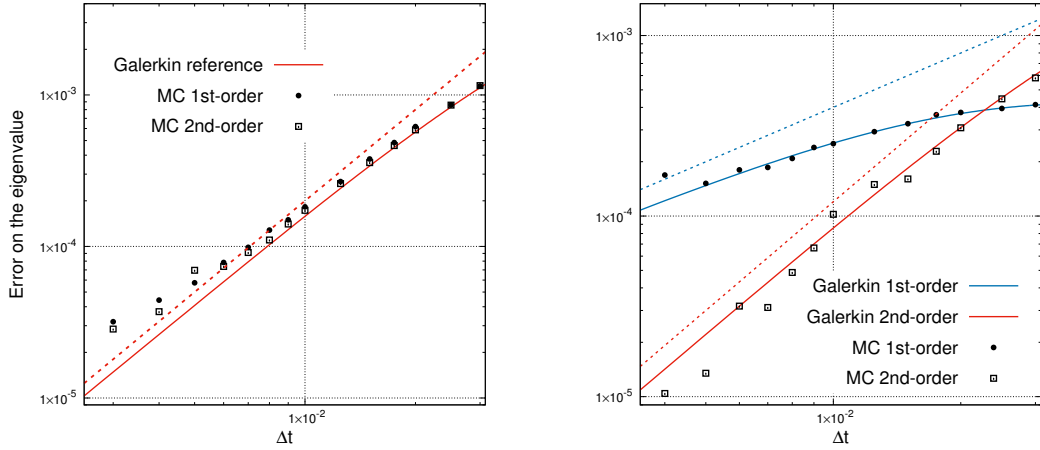


Figure 4.2 – Estimation of the error for the principal eigenvalue (left) and the average of φ with respect to the invariant probability measure (right) as a function of the time step, by Monte Carlo simulation (with comparison to Galerkin, see Figure 4.1). The eigenvalues and the averages of φ are computed with (4.62) and (4.63) respectively. The first and second order schemes are relative to the choice of the weight $\chi_{\Delta t}$ in (4.61). The dashed lines show reference first and second convergences. For very small values of the error on the eigenvalue, we observe the bias due to the finite size of the population.

which corresponds to a trapezoidal rule for the time integral; in this case, both the eigenvalue and the average of φ converge at order two, so we refer to this scheme as *second order*.

We present the numerical results obtained with the various schemes we consider in Figures 4.3 (eigenvalues) and 4.4 (averages of φ):

- Concerning the eigenvalues computed with the Monte Carlo method, we indeed observe first order convergence for the first order scheme, and second order convergence for the hybrid and second order schemes. In particular, the results of the hybrid and the second order scheme are exactly the same. The Galerkin method also converges at second order, but with a much smaller prefactor. This is due to the fact that in this case most of the error is due to the discretization of the dynamics rather than the discretization of the time integral.
- Concerning the average of φ , the first order scheme converges at order one, while the hybrid and second order scheme converge at order two. We would have expected the hybrid scheme to converge at first order but, once again, this is due to the fact that most of the error is due to the discretization of the dynamics, and not to the time integral – as shown by the results of the Galerkin method, which amounts to observing the error due to the discretization of the time integral only. We indeed observe first and second order convergence for the Galerkin approximation, but we see that the error is orders of magnitude smaller than the one of the Monte Carlo approximation. This explains why the Monte Carlo hybrid and second order schemes seem to provide the same results.

Situation with a weak potential. In order to obtain a better trade-off between the error due to the discretization of the dynamics and of the time integral, we run simulations with the same parameters as in the previous situation but with a smaller potential energy $V(x) = 0.02 \cos(2\pi x)$. The results are the following:

- All the eigenvalues now seem to converge at second order (see Figure 4.5 (left)). This is due to the fact that the error due to the discretization of the dynamics is very small, and that the discretization of the time integral, which gives the dominant error term, always leads to an effective second order convergence.
- The behaviour of the average of φ is more interesting (see Figure 4.5 (right)). The Galerkin first and second order schemes provide first and second order convergence respectively. The hybrid scheme exhibits a first order convergence, that matches the Galerkin first order scheme for small time steps. This result can be expected since the two schemes match at order one.

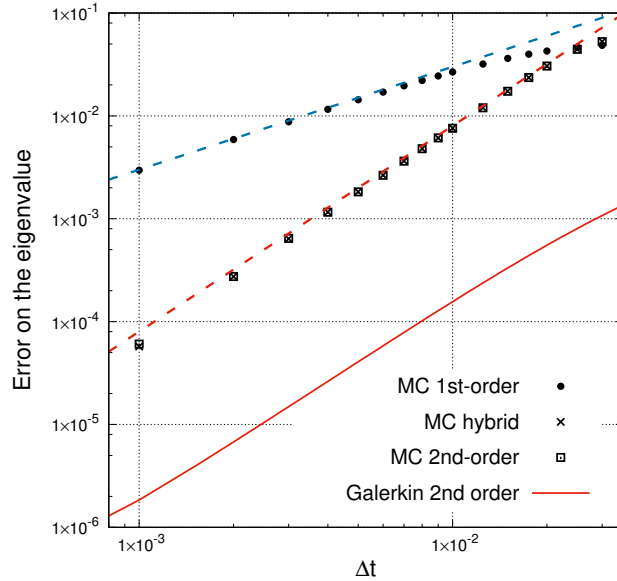


Figure 4.3 – Estimation of the error on the principal eigenvalue as a function of the time step, by Monte Carlo simulation and Galerkin approximation, for $V(x) = \cos(2\pi x)$. The Monte Carlo estimates of the eigenvalues are computed with (4.62), while the Galerkin approximations of the eigenvalues are obtained with (4.65) for $j = 1$ and $j = 2$. The dashed lines show reference linear and quadratic convergences to zero.

The first order scheme also converges at first order but with a larger prefactor, which is due to the discretization of the dynamics. On the other hand, the second order Monte Carlo scheme converges at second order, like the Galerkin second order scheme.

Conclusion. The numerical applications we presented show the validity of our analysis on a simple test case. However, we observe numerically that the prefactor of the leading error term depends on the choice of parameters. This has the consequence that some schemes may effectively seem to exhibit an improved order of convergence than expected, while they actually have a small prefactor at leading order, depending on the discretization at hand. This observation also motivates the study of a one-dimensional model: not only can the Galerkin discretization be made sufficiently accurate by considering a very large number of basis functions, but we can also run sufficiently long Monte Carlo simulations in order for the statistical error to be negligible compared to the bias arising from the time step discretization. Although the order of convergence would be harder to observe for higher dimensional systems, the framework is still applicable and we refer to [313] and references therein for examples in high dimension.

4.5 Discussion

The analysis and simulations we performed in this work were done for SDEs with a non-degenerate noise on a torus. We however believe that most of our results could be extended to more general settings. In particular, we wonder whether the results of Section 4.3 can be applied for unbounded dynamics satisfying Lyapunov type conditions, as in Part II. However, in the functional framework of Chapter 3, the stability property (Assumption 4.14), which is crucial for our analysis to hold, should be rephrased as the invariance of a functional space (containing unbounded functions) under the action of $\ell_f^{-1}(\mathcal{L}^* + f - \lambda)^{-1}(\ell_f \cdot)$. In the case $f = 0$, which does not involve the eigenvector ℓ_f (since $h_0 = \mathbb{1}$), this is already a quite technical result to obtain (see [276, 277]). Here, the presence of the eigenvector ℓ_f adds a significant difficulty, which leaves the situation open.

Finally, in the context of large deviations, we would like to obtain error estimates on the rate function, which is the Legendre transform of the cumulant function λ , see for instance Corollary 3.12 in Chapter 3. However, transferring our error analysis to the rate function is difficult because of the supremum involved, and we believe it is also an interesting open problem.




Figure 4.4 – Estimation of the error on the average of φ with respect to the invariant probability measure as a function of the time step, by Monte Carlo simulation (left) and Galerkin approximation (right), for $V(x) = \cos(2\pi x)$. The Monte Carlo estimates of the averages of φ are computed with (4.63), while the Galerkin approximation is obtained with (4.66) for $j = 1$ and $j = 2$. The dashed lines show reference linear and quadratic convergences to zero.

4.6 Proofs

4.6.1 Proof of the results of Section 4.2.1

Let us first give a result which shows that it suffices to prove Proposition 4.2 for probability measures which admit a positive and bounded density with respect to the Lebesgue measure. This result relies on the regularizing properties of the underlying diffusion.

Lemma 4.28. *For any $\alpha > 0$, denote by $\mathcal{P}_\alpha(\mathcal{X})$ the subspace of probability measures which admit a smooth density with respect to the Lebesgue measure, and whose density is bounded below by $\alpha > 0$ and bounded above by $1/\alpha$. Then there exists $\alpha_* > 0$ such that $\Phi_1(\nu) \in \mathcal{P}_{\alpha_*}(\mathcal{X})$ for any $\nu \in \mathcal{P}(\mathcal{X})$.*

Proof. Note that, for any $\varphi \in \mathcal{C}$,

$$\left(P_t^f \varphi\right)(x) = \int_{\mathcal{X}} p_t^f(x, x') \varphi(x') dx',$$

where p_t^f is the integral kernel of the semigroup $e^{t(\mathcal{L}+f)}$. By parabolic regularity (see for instance [174]), the integral kernel is smooth for any $t > 0$. It is also positive when $f = 0$ by the irreducibility properties of the underlying non-degenerate diffusion: there exists $\varepsilon > 0$ such that (setting $t = 1$)

$$\forall (x, x') \in \mathcal{X}^2, \quad \varepsilon \leq p_1^0(x, x') \leq \frac{1}{\varepsilon}.$$

Given that f is bounded, a similar property holds for p_t^f : there exists $\alpha > 0$ such that

$$\forall (x, x') \in \mathcal{X}^2, \quad \sqrt{\alpha} \leq p_1^f(x, x') \leq \frac{1}{\sqrt{\alpha}}.$$

Since, for any bounded measurable function φ ,

$$\Phi_1(\nu)(\varphi) = \frac{\nu(P_1^f \varphi)}{\nu(P_1^f \mathbf{1})} = \frac{1}{\nu(P_1^f \mathbf{1})} \int_{\mathcal{X}} \int_{\mathcal{X}} \varphi(x') p_1^f(x, x') \nu(dx) dx',$$

it follows that $\Phi_1(\nu)$ has a smooth density with respect to the Lebesgue measure, denoted by $F_{1,\nu}$:

$$F_{1,\nu}(x) = \frac{1}{\nu(P_1^f \mathbf{1})} \int_{\mathcal{X}} p_1^f(x', x) \nu(dx').$$

Moreover, since $\nu(P_1^f \mathbf{1}) \geq \sqrt{\alpha}$, it holds

$$\alpha \leq F_{1,\nu} \leq \frac{1}{\alpha},$$

which gives the claimed result. \square




Figure 4.5 – Estimation of the error on the principal eigenvalue (left) and the average of φ with respect to the invariant probability measure (right) as a function of the time step, for $V(x) = 0.02 \cos(2\pi x)$. The Monte Carlo estimates of the averages of φ are computed with (4.63), while the Galerkin approximations are obtained with (4.66) for $j = 1$ and $j = 2$. The dashed lines show reference first and second order convergences. For very small values of the error on the eigenvalue, we observe the bias due to the finite size of the population. We also observe that the error on the average of φ becomes noisier below 10^{-4} .

We can now provide the proof of Proposition 4.2.

Proof of Proposition 4.2. In view of the semigroup property $\Phi_t^f(\nu) = \Phi_{t-1}^f(\Phi_1^f(\nu))$ when $t \geq 1$, it is sufficient by Lemma 4.28 to prove the result for measures $\nu \in \mathcal{P}_\alpha(\mathcal{X})$, where $\alpha > 0$. The proof is conducted in two steps: we first prove a convergence result for the linear semigroup $P_t^{f-\lambda}$ in $L^2(\mu)$ and any times $t \geq 0$, and then rely on the fact that any probability measure in $\mathcal{P}_\alpha(\mathcal{X})$ is equivalent to μ to obtain (4.10).

Introduce the projector (different from the one defined in (4.39))

$$\widehat{\Pi}_f \varphi = \varphi - h_f \frac{\langle \varphi, \ell_f \rangle_{L^2(\mu)}}{\langle h_f, \ell_f \rangle_{L^2(\mu)}} = \varphi - h_f \frac{\int_{\mathcal{X}} \varphi d\mu_f}{\int_{\mathcal{X}} h_f d\mu_f}.$$

A simple computation shows that $\widehat{\Pi}_f$ commutes with $\mathcal{L} + f$ and $P_t^{f-\lambda}$. It is easily seen that the spectrum of the operator $\widehat{\Pi}_f(\mathcal{L} + f - \lambda)\widehat{\Pi}_f$ is

$$\sigma(\mathcal{L} + f - \lambda) \setminus \{0\} \subset \{z \in \mathbb{C}, \operatorname{Re}(z) \leq -\delta_f\},$$

and that the associated semigroup satisfies $e^{t\widehat{\Pi}_f(\mathcal{L}+f-\lambda)\widehat{\Pi}_f} = P_t^{f-\lambda}\widehat{\Pi}_f$. By the Hille–Yosida theorem (see for instance [351, 107]), there exists a constant $C > 0$ such that, for any $\varphi \in L^2(\mu)$,

$$\forall t \geq 0, \quad \left\| P_t^{f-\lambda} \widehat{\Pi}_f \varphi \right\|_{L^2(\mu)} \leq C e^{-\delta_f t} \|\varphi\|_{L^2(\mu)}. \quad (4.68)$$

We now show that (4.68) implies the convergence result (4.10) for the class of probability measures $\mathcal{P}_\alpha(\mathcal{X})$. For a given $\nu \in \mathcal{P}_\alpha(\mathcal{X})$ and $\varphi \in L^2(\mu)$,

$$\begin{aligned} \nu \left[\left| P_t^{f-\lambda} \widehat{\Pi}_f \varphi \right| \right] &= \int_{\mathcal{X}} \left| P_t^{f-\lambda} \widehat{\Pi}_f \varphi \right| d\nu \leq \frac{1}{\alpha} \int_{\mathcal{X}} \left| P_t^{f-\lambda} \widehat{\Pi}_f \varphi \right| \frac{d\mu}{\inf_{\mathcal{X}} \mu} \\ &\leq \frac{1}{\alpha \inf_{\mathcal{X}} \mu} \left\| P_t^{f-\lambda} \widehat{\Pi}_f \varphi \right\|_{L^2(\mu)} \leq \frac{C}{\alpha \inf_{\mathcal{X}} \mu} e^{-\delta_f t} \|\varphi\|_{L^2(\mu)}, \end{aligned}$$

where we used a Cauchy-Schwarz inequality on $L^2(\mu)$ to go from the first to the second line. The latter computation shows that, for any $\nu \in \mathcal{P}_\alpha(\mathcal{X})$ and $\varphi \in L^2(\mu)$, there are functions a_t, b_t for which

$$\nu \left(P_t^{f-\lambda} \varphi \right) = \frac{\int_{\mathcal{X}} h_f d\nu}{\int_{\mathcal{X}} h_f d\mu_f} \int_{\mathcal{X}} \varphi d\mu_f + a_t, \quad \nu \left(P_t^{f-\lambda} \mathbb{1} \right) = \frac{\int_{\mathcal{X}} h_f d\nu}{\int_{\mathcal{X}} h_f d\mu_f} + b_t,$$

with $|a_t| \leq K\|\varphi\|_{L^2(\mu)}e^{-\delta_f t}$ and $|b_t| \leq Ke^{-\delta_f t}$ for some constant $K > 0$ independent of ν and φ . Moreover, there exists $\varepsilon > 0$ such that $\varepsilon \leq h_f \leq 1/\varepsilon$. Note also that $|b_t| \leq \varepsilon^2$ for $t \geq \ln(K/\varepsilon^2)/\delta_f$ and that

$$\frac{\int_{\mathcal{X}} h_f d\nu}{\int_{\mathcal{X}} h_f d\mu_f} \geq \varepsilon^2.$$

Since

$$\Phi_t^f(\nu)(\varphi) = \frac{\nu(P_t^f \varphi)}{\nu(P_t^f \mathbf{1})} = \frac{\nu(P_t^{f-\lambda} \varphi)}{\nu(P_t^{f-\lambda} \mathbf{1})},$$

it follows that, for $t \geq \ln(2K/\varepsilon)/\delta_f$,

$$\begin{aligned} \left| \Phi_t^f(\nu)(\varphi) - \int_{\mathcal{X}} \varphi d\mu_f \right| &= \left| \frac{\int_{\mathcal{X}} h_f d\nu \int_{\mathcal{X}} \varphi d\mu_f + a_t \int_{\mathcal{X}} h_f d\mu_f}{\int_{\mathcal{X}} h_f d\nu + b_t \int_{\mathcal{X}} h_f d\mu_f} - \int_{\mathcal{X}} \varphi d\mu_f \right| = \left| \frac{(a_t - b_t \int_{\mathcal{X}} \varphi d\mu_f) \int_{\mathcal{X}} h_f d\mu_f}{\int_{\mathcal{X}} h_f d\nu + b_t \int_{\mathcal{X}} h_f d\mu_f} \right| \\ &\leq \frac{1}{\frac{\int_{\mathcal{X}} h_f d\nu}{\int_{\mathcal{X}} h_f d\mu_f} - |b_t|} \left(|a_t| + |b_t| \left| \int_{\mathcal{X}} \varphi d\mu_f \right| \right) \\ &\leq \frac{K}{2\varepsilon^2} \left(\|\varphi\|_{L^2(\mu)} + \left| \int_{\mathcal{X}} \varphi d\mu_f \right| \right) e^{-\delta_f t}. \end{aligned}$$

The inequality

$$\left| \int_{\mathcal{X}} \varphi d\mu_f \right| = \left| \int_{\mathcal{X}} \varphi \ell_f d\mu \right| \leq \|\ell_f\|_{L^2(\mu)} \|\varphi\|_{L^2(\mu)}$$

allows to obtain the desired conclusion. \square

Let us conclude this section with the proof of Proposition 4.5.

Proof of Proposition 4.5. The exponential convergence result (4.68) implies that the operator $\widehat{\Pi}_f(\mathcal{L} + f - \lambda)\widehat{\Pi}_f$ is invertible on $\text{Ran}(\widehat{\Pi}_f) = L_f^2(\mu)$ with inverse given by

$$\left(\widehat{\Pi}_f(\mathcal{L} + f - \lambda)\widehat{\Pi}_f \right)^{-1} = - \int_0^{+\infty} P_t^{f-\lambda} \widehat{\Pi}_f dt.$$

The solution to $(\mathcal{L} + f - \lambda)u = g$ with $g \in L_f^2(\mu)$ then admits a unique solution in $L_f^2(\mu)$. By elliptic regularity, $u \in \mathcal{C}_f$ when $g \in \mathcal{C}_f$. The result for $\mathcal{L}^* + f - \lambda$ can be obtained by a similar reasoning. \square

Note that, alternatively, it would have been possible to resort to the Fredholm alternative to prove Proposition 4.5.

4.6.2 Proof of Theorem 4.9

Theorem 4.9 is a rewriting of [114, Corollary 2.5], which is stated in the context of a finite state space. In order for the paper to be self-contained, we prove Theorem 4.9 in our setting of continuous but compact state space, and in the simplified case of a time-homogeneous Markov chain, adapting the arguments of [114]. The idea is to prove some contraction property using the Dobrushin coefficient defined in Section 4.7 and the reformulation (4.70) below of the semigroup. We work on the space of probability measures $\mathcal{P}(\mathcal{X})$ endowed with the total variation distance.

Define the weights

$$g_n = (Q_{\Delta t}^f)^n \mathbf{1},$$

and the Markov operator S_n as

$$(S_n \varphi)(x) = \frac{Q_{\Delta t}^f(g_n \varphi)(x)}{(Q_{\Delta t}^f g_n)(x)}. \quad (4.69)$$

The dynamics (4.15) can then be rephrased as

$$\nu_n(\varphi) = \frac{\nu(g_n(K_n \varphi))}{\nu(g_n)}, \quad K_{n+1} = S_n K_n, \quad K_0 = I_d. \quad (4.70)$$

This equality can be proved by induction. The result is clear for $n = 0$. For $n = 1$, we have (with $\nu_0 = \nu$)

$$\frac{\nu(g_1(K_1\varphi))}{\nu(g_1)} = \frac{\nu(Q_{\Delta t}^f \mathbb{1} S_0\varphi)}{\nu(Q_{\Delta t}^f \mathbb{1})} = \frac{\nu(Q_{\Delta t}^f \varphi)}{\nu(Q_{\Delta t}^f \mathbb{1})} = \nu_1(\varphi).$$

Assuming that ν_n satisfies (4.70) at rank n , using (4.15) and recalling the definition (4.17), it holds

$$\nu_{n+1}(\varphi) = \frac{\nu((Q_{\Delta t}^f)^{n+1}(\varphi))}{\nu((Q_{\Delta t}^f)^{n+1}\mathbb{1})} = \frac{(\nu Q_{\Delta t}^f)((Q_{\Delta t}^f)^n(\varphi))}{(\nu Q_{\Delta t}^f)((Q_{\Delta t}^f)^n\mathbb{1})} = \Phi_{\Delta t, n}^f(\nu Q_{\Delta t}^f)(\varphi),$$

so that, using the recursion hypothesis and $Q_{\Delta t}^f g_n = g_{n+1}$, it follows

$$\begin{aligned} \nu_{n+1}(\varphi) &= \frac{(\nu Q_{\Delta t}^f)(g_n(K_n\varphi))}{(\nu Q_{\Delta t}^f)(g_n)} = \frac{\nu(Q_{\Delta t}^f(g_n(K_n\varphi)))}{\nu(Q_{\Delta t}^f g_n)} \\ &= \frac{\nu(Q_{\Delta t}^f(g_n)S_n(K_n\varphi))}{\nu(Q_{\Delta t}^f g_n)} = \frac{\nu(g_{n+1}(K_{n+1}\varphi))}{\nu(g_{n+1})}, \end{aligned}$$

which concludes the recurrence.

We next introduce the family of operators $T_n : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$ defined by:

$$\forall \nu \in \mathcal{P}(\mathcal{X}), \quad \forall \varphi \in \mathcal{C}, \quad (\nu T_n)(\varphi) = \frac{\nu(g_n\varphi)}{\nu(g_n)},$$

so that from (4.70) we have $\nu_n = \Phi_{\Delta t, n}^f(\nu) = \nu T_n K_n$. Using the definitions of Section 4.7, we obtain, for two initial measures $\mu, \nu \in \mathcal{P}(\mathcal{X})$,

$$\|\mu_n - \nu_n\|_{\text{TV}} = \|\mu T_n K_n - \nu T_n K_n\|_{\text{TV}} \leq \|K_n\| \|\mu T_n - \nu T_n\|_{\text{TV}}.$$

Given that $T_n : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$, we can bound $\|\mu T_n - \nu T_n\|_{\text{TV}}$ by 2. The next step consists in studying the contraction induced by the operator $K_n = S_{n-1}S_{n-2}\dots S_1$, where S_k is defined in (4.70). We have

$$\|K_n\| \leq \prod_{k=0}^{n-1} \|S_k\|,$$

so that, using the relationship (4.92),

$$\|\mu_n - \nu_n\|_{\text{TV}} \leq 2 \prod_{k=0}^n (1 - \alpha(S_k)).$$

The last step consists in using Assumption 4.8 in order to obtain a lower bound on $\alpha(S_k)$ independent of k . First, for all $x \in \mathcal{X}$ and $A \subset \mathcal{X}$,

$$S_k(x, A) = \frac{Q_{\Delta t}^f(g_k \mathbb{1}_A)(x)}{Q_{\Delta t}^f(g_k)(x)} \geq \varepsilon^2 \frac{\eta(g_k \mathbb{1}_A)}{\eta(g_k)}.$$

Then, it follows from definition (4.91) that

$$\alpha(S_k) = \inf_{\substack{x, x' \in \mathcal{X} \\ \{A_i\}_{1 \leq i \leq m} \subset \mathcal{X}}} \left\{ \sum_{i=1}^m \min(S_k(x, A_i), S_k(x', A_i)) \right\} \geq \varepsilon^2 \inf_{\{A_i\}_{1 \leq i \leq m} \subset \mathcal{X}} \left\{ \frac{\eta(g_k \sum_{i=1}^m \mathbb{1}_{A_i})}{\eta(g_k)} \right\} = \varepsilon^2,$$

since the infimum is taken over partitions $(A_i)_{i=1}^m$ of \mathcal{X} . As a result, we obtain that, for all measures $\nu, \nu' \in \mathcal{P}(\mathcal{X})$,

$$\left\| \Phi_{\Delta t, n}^f(\nu) - \Phi_{\Delta t, n}^f(\nu') \right\|_{\text{TV}} \leq 2(1 - \varepsilon^2)^n. \quad (4.71)$$

Setting $\nu' = \nu_m$ for $m \in \mathbb{N}$ and using the semigroup property, we get

$$\left\| \Phi_{\Delta t, n}^f(\nu) - \Phi_{\Delta t, n+m}^f(\nu) \right\|_{\text{TV}} \leq 2(1 - \varepsilon^2)^n, \quad (4.72)$$

so that $(\nu_n)_{n \geq 1}$ is a Cauchy sequence in $\mathcal{P}(\mathcal{X})$. By completeness of $\mathcal{P}(\mathcal{X})$ for the total variation norm, we can conclude that, for any initial measure ν , there exists $\mu_{f,\Delta t}$ such that $\nu_n \rightarrow \mu_{f,\Delta t}$ in total variation norm. Then using the one step formulation of the dynamics (4.18) and the semigroup property, we obtain with the choice $\nu' = \Phi_{\Delta t}^f(\nu)$,

$$\left\| \Phi_{\Delta t,n}^f(\nu) - \Phi_{\Delta t}^f(\Phi_{\Delta t,n}^f(\nu)) \right\|_{\text{TV}} \leq 2(1 - \varepsilon^2)^n,$$

so that, taking $n \rightarrow \infty$ and using the continuity of $\Phi_{\Delta t}^f$ on $\mathcal{P}(\mathcal{X})$ endowed with the total variation norm, it follows that $\mu_{f,\Delta t} = \Phi_{\Delta t}^f(\mu_{f,\Delta t})$. Passing to the limit $m \rightarrow +\infty$ in (4.72),

$$\|\nu_n - \mu_{f,\Delta t}\|_{\text{TV}} \leq 2(1 - \varepsilon^2)^n.$$

Finally, it follows from (4.71) that the limit $\mu_{f,\Delta t}$ does not depend on the initial measure ν .

4.6.3 Proofs related to Theorem 4.16

4.6.3.1 Proof of Lemma 4.17

The idea is to approximate at leading order the stationary measure $\mu_{f,\Delta t}$ as $(1 + \Delta t^p \psi)\mu_f$, since we expect the invariant probability measure to be correct at order p . We start from the stationarity equation (4.36) and search for a function $\psi \in \mathcal{C}$ and a remainder $R_{f,\Delta t} : \mathcal{C} \rightarrow \mathbb{R}$ satisfying (4.27) such that, for all $\phi \in \mathcal{C}$,

$$\int_{\mathcal{X}} (Q_{\Delta t}^f \phi)(1 + \Delta t^p \psi) d\mu_f - \left(\int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1}(1 + \Delta t^p \psi) d\mu_f \right) \left(\int_{\mathcal{X}} \phi(1 + \Delta t^p \psi) d\mu_f \right) = \Delta t^{p+2} R_{f,\Delta t} \phi. \quad (4.73)$$

In view of the expansion (4.30) of $Q_{\Delta t}^f$ and of the invariance relation (4.33), the first term of the left hand side is

$$\begin{aligned} & \int_{\mathcal{X}} \left(\phi + \Delta t \mathcal{A}_1^f \phi + \dots + \Delta t^p \mathcal{A}_p^f \phi + \Delta t^{p+1} \mathcal{A}_{p+1}^f \phi + \Delta t^{p+2} \mathcal{R}_{f,\Delta t} \phi \right) (1 + \Delta t^p \psi) d\mu_f \\ &= (1 + a_1 \Delta t + \dots + a_p \Delta t^p) \int_{\mathcal{X}} \phi d\mu_f + \Delta t^p \int_{\mathcal{X}} \phi \psi d\mu_f + \Delta t^{p+1} \int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \phi + \psi(\mathcal{A}_1^f \phi) \right) d\mu_f + \Delta t^{p+2} R_{f,\Delta t} \phi, \end{aligned}$$

where $R_{f,\Delta t}$ gathers the terms of order at least $p+2$, and is uniformly bounded in Δt for $0 < \Delta t \leq \Delta t^*$ in the sense of (4.27) when $\psi \in \mathcal{C}$. On the other hand, the second term on the left hand side of (4.73) can be written as, using again (4.33),

$$\begin{aligned} & \left(1 + \Delta t a_1 + \dots + \Delta t^p a_p + \Delta t^p \int_{\mathcal{X}} \psi d\mu_f + \Delta t^{p+1} \int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \mathbb{1} + \psi(\mathcal{A}_1^f \mathbb{1}) \right) d\mu_f \right) \int_{\mathcal{X}} \phi(1 + \Delta t^p \psi) d\mu_f \\ & \quad + \Delta t^{p+2} R_{f,\Delta t} \phi \\ &= (1 + \Delta t a_1 + \dots + \Delta t^p a_p) \int_{\mathcal{X}} \phi d\mu_f + \Delta t^p \left(\int_{\mathcal{X}} \phi d\mu_f \int_{\mathcal{X}} \psi d\mu_f + \int_{\mathcal{X}} \phi \psi d\mu_f \right) \\ & \quad + \Delta t^{p+1} \left(\int_{\mathcal{X}} \mathcal{A}_{p+1}^f \mathbb{1} d\mu_f \int_{\mathcal{X}} \phi d\mu_f + a_1 \int_{\mathcal{X}} \phi \psi d\mu_f + \int_{\mathcal{X}} \psi(\mathcal{A}_1^f \mathbb{1}) d\mu_f \int_{\mathcal{X}} \phi d\mu_f \right) + \Delta t^{p+2} R_{f,\Delta t} \phi, \end{aligned}$$

where $R_{f,\Delta t}$ is uniformly bounded in Δt in the sense of (4.27) when $\psi \in \mathcal{C}$. We can now equate the different orders in powers of Δt on both sides of (4.73) and choose ψ such that only a remainder of order $p+2$ remains. The terms $a_k \Delta t^k \int_{\mathcal{X}} \phi d\mu_f$ cancel, so the first non-trivial condition to be satisfied to eliminate terms of order Δt^p reads

$$\int_{\mathcal{X}} \phi \psi d\mu_f = \left(\int_{\mathcal{X}} \phi d\mu_f \right) \left(\int_{\mathcal{X}} \psi d\mu_f \right) + \int_{\mathcal{X}} \phi \psi d\mu_f.$$

This equality is satisfied for all $\phi \in \mathcal{C}$ if and only if (take *e.g.* $\phi = \psi$)

$$\int_{\mathcal{X}} \psi d\mu_f = 0. \quad (4.74)$$

The condition arising from the equality of terms of order Δt^{p+1} is

$$\int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \phi + \psi(\mathcal{A}_1^f \phi) \right) \ell_f d\mu = a_1 \int_{\mathcal{X}} \phi \psi d\mu_f + \left(\int_{\mathcal{X}} \left((\mathcal{A}_{p+1}^f \mathbb{1}) + \psi(\mathcal{A}_1^f \mathbb{1}) \right) \ell_f d\mu \right) \int_{\mathcal{X}} \phi d\mu_f.$$

Using that $\mathcal{A}_1^f \mathbb{1} = f$ along with condition (4.33), we have $a_1 = \lambda$. In addition, taking adjoints in $L^2(\mu)$ and recalling $\mathcal{A}_1^f = \mathcal{A}_1 + f$,

$$\int_{\mathcal{X}} \phi \left((\mathcal{A}_{p+1}^f)^* \ell_f + (\mathcal{A}_1^* + f - \lambda)(\ell_f \psi) \right) d\mu = \left(\int_{\mathcal{X}} \left((\mathcal{A}_{p+1}^f)^* \ell_f + (\mathcal{A}_1^* + f)(\ell_f \psi) \right) d\mu \right) \int_{\mathcal{X}} \phi d\mu_f.$$

Moreover, in view of (4.74), one can subtract $(\lambda \int_{\mathcal{X}} \psi \ell_f d\mu) (\int_{\mathcal{X}} \phi d\mu_f)$ from the right hand side of last equation. Finally, we obtain the following equation (with unknown ψ): for all $\phi \in \mathcal{C}$,

$$\int_{\mathcal{X}} \phi \left((\mathcal{A}_{p+1}^f)^* \ell_f + (\mathcal{A}_1^* + f - \lambda)(\ell_f \psi) \right) d\mu = \left(\int_{\mathcal{X}} (\mathcal{A}_{p+1}^f)^* \ell_f + (\mathcal{A}_1^* + f - \lambda)(\ell_f \psi) d\mu \right) \int_{\mathcal{X}} \phi d\mu_f. \quad (4.75)$$

By Assumption 4.14, the operator $\mathcal{A}_1^* + f - \lambda$ is invertible on $\widehat{\mathcal{C}}_f$ and leaves this space invariant. We can therefore define a solution ψ_0 to the following equation:

$$\begin{cases} (\mathcal{A}_1^* + f - \lambda)(\ell_f \psi_0) = \tilde{g}, \\ \tilde{g} = -(\mathcal{A}_{p+1}^f)^* \ell_f + \ell_f \frac{\int_{\mathcal{X}} \left((\mathcal{A}_{p+1}^f)^* \ell_f \right) h_f d\mu}{\int_{\mathcal{X}} h_f \ell_f d\mu} \in \widehat{\mathcal{C}}_f. \end{cases} \quad (4.76)$$

The function $\ell_f \psi_0$ is uniquely defined in $\widehat{\mathcal{C}}_f$ by Assumption 4.14 since \tilde{g} has average 0 with respect to $\hat{\mu}_f$, and one can check that it is indeed solution of (4.75). Since the eigenvector ℓ_f is regular with $\ell_f > 0$, the function ψ_0 belongs to \mathcal{C} . However, ψ_0 is not a priori of average 0 with respect to μ_f , so that condition (4.74) is not satisfied. We can however consider the function $\psi_\alpha = \psi_0 + \alpha$, which is still such that (4.75) holds. The choice $\alpha = -\int_{\mathcal{X}} \psi_0 d\mu_f$ ensures that (4.74) is satisfied. This provides the solution (4.34) and concludes the proof.

4.6.3.2 Proof of Lemma 4.18

We start by considering (4.24) and (4.38) for $\varphi = \Pi_f \phi$ with $\phi \in \mathcal{C}$:

$$\int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_{f, \Delta t} = 0, \quad (4.77)$$

and

$$\int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f = \Delta t^{p+1} R_{f, \Delta t} \phi. \quad (4.78)$$

We next stabilize the operator in \mathcal{C}_f by another application of the projector Π_f . First,

$$\begin{aligned} & \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_{f, \Delta t} \\ &= \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_{f, \Delta t} - \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f \\ &= - \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f, \end{aligned} \quad (4.79)$$

thanks to (4.77). Second, since ψ has average 0 with respect to μ_f ,

$$\begin{aligned} & \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f \\ &= \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f - \left(\int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f \right) \int_{\mathcal{X}} (1 + \Delta t^p \psi) d\mu_f \\ &= \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f - \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f. \end{aligned}$$

In view of (4.78), the first term of the right hand side of the above equation is a remainder of order Δt^{p+1} . Therefore,

$$\begin{aligned} & \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f = - \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f + \Delta t^{p+1} R_{f, \Delta t} \phi \\ & = - \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f + \Delta t^{p+1} R_{f, \Delta t} \phi + \left(\frac{e^{\Delta t \bar{\lambda}_{\Delta t}} - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \int_{\mathcal{X}} \Pi_f \phi d\mu_f \\ & = - \int_{\mathcal{X}} \left[\left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_f + \Delta t^{p+1} R_{f, \Delta t} \phi, \end{aligned} \quad (4.80)$$

since $\Pi_f \phi$ has average 0 with respect to μ_f . Combining (4.80) with (4.79),

$$\begin{aligned} \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \lambda_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] d\mu_{f, \Delta t} & = \int_{\mathcal{X}} \left[\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \bar{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f \phi \right] (1 + \Delta t^p \psi) d\mu_f \\ & \quad + \Delta t^{p+1} R_{f, \Delta t} \phi, \end{aligned}$$

where $R_{f, \Delta t}$ satisfies (4.27). This concludes the proof of the lemma.

4.6.3.3 Proof of Lemma 4.19

The first part of the proof of Lemma 4.19 consists in constructing an approximate eigenvector $h_{f, \Delta t}$ of h_f for the evolution operator $Q_{\Delta t}^f$. We use to this end Assumption 4.15 and (4.33), as well as the definition of the leading order correction ψ in (4.34). More precisely, we consider $h_{f, \Delta t} = u_0 + \Delta t u_1 + \dots + \Delta t^p u_p \in \mathcal{C}$ and look for functions $u_1, \dots, u_p \in \mathcal{C}_f$ and $u_0 \in \mathcal{C}$ with $\int_{\mathcal{X}} u_0 d\mu = 1$ such that

$$Q_{\Delta t}^f h_{f, \Delta t} = e^{\Delta t \bar{\lambda}_{\Delta t}} h_{f, \Delta t} + \Delta t^{p+2} r_{f, \Delta t}, \quad (4.81)$$

with $\|r_{f, \Delta t}\|_{B^\infty} \leq C$ for $0 < \Delta t \leq \Delta t^*$. Recall that, by (4.37),

$$e^{\Delta t \bar{\lambda}_{\Delta t}} = \int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} (1 + \Delta t^p \psi) d\mu_f.$$

Expanding the left hand side of (4.81) using (4.30) leads to

$$Q_{\Delta t}^f h_{f, \Delta t} = \sum_{k=0}^{p+1} \Delta t^k \mathcal{A}_k^f h_{f, \Delta t} + \Delta t^{p+2} \mathcal{R}_{f, \Delta t} h_{f, \Delta t} = \sum_{k=0}^{p+1} \Delta t^k \sum_{m=0}^k \mathcal{A}_m^f u_{k-m} + \Delta t^{p+2} \mathcal{R}_{f, \Delta t} h_{f, \Delta t}, \quad (4.82)$$

with the convention $\mathcal{A}_0^f = I_d$ and $u_{p+1} = 0$. The right hand side of (4.81) can be expanded as

$$\begin{aligned} & e^{\Delta t \bar{\lambda}_{\Delta t}} h_{f, \Delta t} \\ & = \left[\int_{\mathcal{X}} \left(1 + \Delta t \mathcal{A}_1^f \mathbb{1} + \dots + \Delta t^{p+1} \mathcal{A}_{p+1}^f \mathbb{1} + \Delta t^{p+2} \mathcal{R}_{f, \Delta t} \mathbb{1} \right) (1 + \Delta t^p \psi) d\mu_f \right] (u_0 + \Delta t u_1 + \dots + \Delta t^p u_p) \\ & = \left[1 + \Delta t \int_{\mathcal{X}} \mathcal{A}_1^f \mathbb{1} d\mu_f + \dots + \Delta t^p \int_{\mathcal{X}} \mathcal{A}_p^f \mathbb{1} d\mu_f + \Delta t^{p+1} \int_{\mathcal{X}} \left(\mathcal{A}_{p+1}^f \mathbb{1} + \psi \mathcal{A}_1^f \mathbb{1} \right) d\mu_f + \Delta t^{p+2} r_{f, \Delta t} \right] \\ & \quad \times (u_0 + \Delta t u_1 + \dots + \Delta t^p u_p) \\ & = \sum_{k=0}^{p+1} \Delta t^k \sum_{m=0}^k \lambda_m u_{k-m} + \Delta t^{p+2} r_{f, \Delta t}, \end{aligned} \quad (4.83)$$

where we introduced $\lambda_0 = 1$,

$$\forall m \in \{1, \dots, p\}, \quad \lambda_m = \int_{\mathcal{X}} \mathcal{A}_m^f \mathbb{1} d\mu_f, \quad (4.84)$$

and λ_{p+1} is defined in (4.47):

$$\lambda_{p+1} = \int_{\mathcal{X}} \mathcal{A}_{p+1}^f \mathbb{1} d\mu_f + \int_{\mathcal{X}} f \psi d\mu_f.$$

We see from (4.33) that $\lambda_m = a_m$ for $m \in \{1, \dots, p\}$, with in particular $\lambda_1 = \int_{\mathcal{X}} f d\mu_f = \lambda$.

We now build the functions u_m by induction. Let us show the first steps of the recurrence, before proceeding to the general argument. Plugging (4.82) and (4.83) in (4.81), the equality of terms of order 1 leads to the trivial equality $u_0 = u_0$. Equating terms of order Δt gives

$$\mathcal{A}_1^f u_0 + \mathcal{A}_0^f u_1 = \lambda_1 u_0 + \lambda_0 u_1,$$

so that, using $\mathcal{A}_0^f = I_d$, $\lambda_0 = 1$, $\mathcal{A}_1^f = \mathcal{A}_1 + f$ and $\lambda_1 = \lambda$,

$$(\mathcal{A}_1 + f)u_0 = \lambda u_0.$$

In view of Assumption 4.15, we can conclude that $u_0 = h_f$. The identification of terms of order Δt^2 in (4.82)-(4.83) leads to

$$\mathcal{A}_2^f u_0 + \mathcal{A}_1^f u_1 + \mathcal{A}_0^f u_2 = \lambda_2 u_0 + \lambda_1 u_1 + \lambda_0 u_2,$$

which can be rewritten as

$$(\mathcal{A}_1 + f - \lambda)u_1 = g_{1,0}, \quad g_{1,0} = -\mathcal{A}_2^f h_f + \lambda_2 h_f, \quad (4.85)$$

where the expression of λ_2 is given by (4.84) when $p \geq 2$ and by (4.47) when $p = 1$. In order to prove that (4.85) is well-posed, it is sufficient to show that $g_{1,0}$ belongs to \mathcal{C}_f . We show in fact in the sequel that each function u_k is solution to a Poisson equation similar to (4.85) with a right-hand side that always belongs to \mathcal{C}_f .

Let us now present the inductive construction to any order, until the terminal case $k = p$, showing in particular the well-posedness of the equations defining each mode u_k . This construction is reminiscent of techniques used to build the expansion of the invariant probability measure in Δt in related works, in particular [110]. Suppose that we have built functions $u_0, \dots, u_k \in \mathcal{C}_f$ for some $k \geq 1$. Inserting again (4.82) and (4.83) into (4.81) and equating terms of order Δt^{k+1} then leads to

$$\sum_{m=0}^{k+1} \mathcal{A}_{k+1-m}^f u_m = \sum_{m=0}^{k+1} \lambda_{k+1-m} u_m. \quad (4.86)$$

For $m = k + 1$, we have $\mathcal{A}_0^f u_{k+1} = u_{k+1}$ on the left hand side and $\lambda_0 u_{k+1} = u_{k+1}$ on the right hand side, so that the terms of order $k + 1$ compensate. Taking aside the terms of order $m = k$ leads to the equation:

$$(\mathcal{A}_1 + f - \lambda)u_k = \sum_{m=0}^{k-1} g_{k,m}, \quad g_{k,m} = -\mathcal{A}_{k+1-m}^f u_m + \lambda_{k+1-m} u_m. \quad (4.87)$$

A sufficient condition for the solution u_k to exist in \mathcal{C}_f is that $g_{k,m} \in \mathcal{C}_f$ for $m \in \{0, \dots, k-1\}$. For $m \in \{1, \dots, k-1\}$, a sufficient condition for that is that $\mathcal{A}_{k+1-m}^f u_m$ has average 0 with respect to μ_f , which is clear from (4.33) and the fact that $u_m \in \mathcal{C}_f$. It therefore only remains to show that $g_{k,0} = -\mathcal{A}_{k+1}^f h_f + \lambda_{k+1} h_f$ belongs to \mathcal{C}_f . Two cases have to be distinguished here:

- (a) if $k < p$, then $k + 1 \leq p$ and we can still use the invariance relation (4.33) applied to $\phi \equiv h_f$, along with the fact that $\lambda_{k+1} = a_{k+1}$:

$$\int_{\mathcal{X}} g_{k,0} d\mu_f = - \int_{\mathcal{X}} (\mathcal{A}_{k+1}^f h_f) d\mu_f + \lambda_{k+1} \int_{\mathcal{X}} h_f d\mu_f = -a_{k+1} \int_{\mathcal{X}} h_f d\mu_f + a_{k+1} \int_{\mathcal{X}} h_f d\mu_f = 0.$$

- (b) in the terminal case $k = p$, we cannot use (4.33) and λ_{p+1} has a different expression (recall (4.47)). Let us compute this expression explicitly. In view of (4.34),

$$\int_{\mathcal{X}} f \psi d\mu_f = \int_{\mathcal{X}} f \psi_0 \ell_f d\mu - \left(\int_{\mathcal{X}} \psi_0 d\mu_f \right) \left(\int_{\mathcal{X}} f d\mu_f \right),$$

and, given that $f \psi_0 \ell_f = \tilde{g} + \lambda \ell_f \psi_0 - \mathcal{A}_1^*(\ell_f \psi_0)$ and $\int_{\mathcal{X}} f d\mu_f = \lambda$,

$$\int_{\mathcal{X}} f \psi d\mu_f = \int_{\mathcal{X}} \tilde{g} d\mu + \lambda \int_{\mathcal{X}} \psi_0 \ell_f d\mu - \int_{\mathcal{X}} \mathcal{A}_1^*(\ell_f \psi_0) d\mu - \lambda \int_{\mathcal{X}} \psi_0 \ell_f d\mu.$$

Since $\mathcal{A}_1 \mathbb{1} = 0$,

$$\int_{\mathcal{X}} \mathcal{A}_1^*(\ell_f \psi_0) d\mu = \int_{\mathcal{X}} (\mathcal{A}_1 \mathbb{1}) \ell_f \psi_0 d\mu = 0.$$

Finally, using the expression of \tilde{g} in (4.34) and $\int_{\mathcal{X}} \ell_f d\mu = 1$,

$$\begin{aligned} \int_{\mathcal{X}} f \psi d\mu_f &= \int_{\mathcal{X}} \tilde{g} d\mu = - \int_{\mathcal{X}} (\mathcal{A}_{p+1}^f)^* \ell_f d\mu + \int_{\mathcal{X}} \ell_f d\mu \frac{\int_{\mathcal{X}} (\mathcal{A}_{p+1}^f h_f) d\mu_f}{\int_{\mathcal{X}} h_f d\mu_f} \\ &= - \int_{\mathcal{X}} \mathcal{A}_{p+1}^f \mathbb{1} d\mu_f + \frac{\int_{\mathcal{X}} (\mathcal{A}_{p+1}^f h_f) d\mu_f}{\int_{\mathcal{X}} h_f d\mu_f}. \end{aligned}$$

From this calculation, we obtain, with (4.47),

$$\lambda_{p+1} = \int_{\mathcal{X}} \mathcal{A}_{p+1}^f \mathbb{1} d\mu_f + \int_{\mathcal{X}} f \psi d\mu_f = \frac{\int_{\mathcal{X}} \mathcal{A}_{p+1}^f h_f d\mu_f}{\int_{\mathcal{X}} h_f d\mu_f},$$

so that

$$\int_{\mathcal{X}} g_{p,0} d\mu_f = - \int_{\mathcal{X}} (\mathcal{A}_{p+1}^f h_f) d\mu_f + \lambda_{p+1} \int_{\mathcal{X}} h_f d\mu_f = 0.$$

Therefore, for any $k \in \{0, \dots, p\}$ and any $m \in \{0, \dots, k-1\}$, it holds $g_{k,m} \in \mathcal{C}_f$. This allows to conclude that the equations (4.87) are well-posed in \mathcal{C}_f and (4.81) is satisfied.

We are now in position to conclude the proof. Inserting (4.81) in the stationarity equation (4.22),

$$\int_{\mathcal{X}} Q_{\Delta t}^f h_{f,\Delta t} d\mu_{f,\Delta t} = \int_{\mathcal{X}} \left(e^{\Delta t \tilde{\lambda}_{\Delta t}} h_{f,\Delta t} + \Delta t^{p+2} r_{f,\Delta t} \right) d\mu_{f,\Delta t} = e^{\Delta t \lambda_{\Delta t}} \int_{\mathcal{X}} h_{f,\Delta t} d\mu_{f,\Delta t},$$

so that

$$e^{\Delta t \lambda_{\Delta t}} = e^{\Delta t \tilde{\lambda}_{\Delta t}} + \Delta t^{p+2} \frac{\int_{\mathcal{X}} r_{f,\Delta t} d\mu_{f,\Delta t}}{\int_{\mathcal{X}} h_{f,\Delta t} d\mu_{f,\Delta t}}.$$

At this stage, it suffices to prove that the remainder term is uniformly of order Δt^{p+2} for Δt sufficiently small. We note to this end that $h_{f,\Delta t} = h_f + \Delta t u_1 + \dots + \Delta t^p u_p$, where the functions u_1, \dots, u_p are regular and $h_f > 0$. Given that the state space \mathcal{X} is compact, there exists $\varepsilon > 0$ such that $h_f \geq \varepsilon > 0$. This implies in particular that there exists $\Delta t' > 0$ such that, for any $0 < \Delta t \leq \Delta t'$, it holds $h_{f,\Delta t} \geq \varepsilon/2 > 0$. We also know that there exists $\Delta t^* > 0$ and $C > 0$ such that, for any $0 < \Delta t \leq \Delta t^*$, it holds $\|r_{f,\Delta t}\|_{B^\infty} \leq C$. As a result, for $0 < \Delta t \leq \min(\Delta t', \Delta t^*)$,

$$\left| \frac{\int_{\mathcal{X}} r_{f,\Delta t} d\mu_{f,\Delta t}}{\int_{\mathcal{X}} h_{f,\Delta t} d\mu_{f,\Delta t}} \right| \leq \frac{\int_{\mathcal{X}} |r_{f,\Delta t}| d\mu_{f,\Delta t}}{\int_{\mathcal{X}} h_{f,\Delta t} d\mu_{f,\Delta t}} \leq \frac{2C}{\varepsilon},$$

which gives the claimed result.

4.6.3.4 Proof of Lemma 4.20

We follow the strategy outlined in [295, 304], which uses a truncated inverse series expansion. The first step is to use the expansion of the eigenvalue $e^{\Delta t \tilde{\lambda}_{\Delta t}}$ as in the proof of Lemma 4.19:

$$\begin{aligned} e^{\Delta t \tilde{\lambda}_{\Delta t}} &= \int_{\mathcal{X}} Q_{\Delta t}^f \mathbb{1} (1 + \Delta t^p \psi) d\mu_f \\ &= \int_{\mathcal{X}} \left(1 + \Delta t \mathcal{A}_1^f \mathbb{1} + \Delta t^2 \mathcal{A}_2^f \mathbb{1} + \dots + \Delta t^{p+1} \mathcal{A}_{p+1}^f \mathbb{1} + \Delta t^{p+2} \mathcal{R}_{f,\Delta t} \mathbb{1} \right) (1 + \Delta t^p \psi) d\mu_f \\ &= 1 + \Delta t \lambda + \Delta t^2 \lambda_2 + \dots + \Delta t^p \lambda_p + \Delta t^{p+1} \lambda_{p+1} + \Delta t^{p+2} r_{f,\Delta t}, \end{aligned}$$

where the coefficients λ_m are defined in (4.84)-(4.47), and there exists $C > 0$ such that $|r_{f,\Delta t}| \leq C$ for $0 < \Delta t \leq \Delta t^*$. This expression, combined with the expansion (4.30) of $Q_{\Delta t}^f$ leads to:

$$\Pi_f \left(\frac{Q_{\Delta t}^f - e^{\Delta t \tilde{\lambda}_{\Delta t}}}{\Delta t} \right) \Pi_f = A + \Delta t B_{\Delta t} + \Delta t^{p+1} R_{f,\Delta t},$$

with

$$A = \Pi_f(\mathcal{A}_1 + f - \lambda)\Pi_f, \quad B_{\Delta t} = \Pi_f(\mathcal{A}_2^f - \lambda_2)\Pi_f + \dots + \Delta t^{p-1}\Pi_f(\mathcal{A}_{p+1}^f - \lambda_{p+1})\Pi_f.$$

The operator A is invertible on \mathcal{C}_f by Assumption 4.14. Now we are back to the setting of [295, 304] and it suffices to write the formal series expansion of the inverse of $A + \Delta t B_{\Delta t} = (I_d + \Delta t B_{\Delta t} A^{-1})A$ up to order p by setting

$$\tilde{S}_{\Delta t}^f = A^{-1} \sum_{n=0}^p (-1)^n (B_{\Delta t} A^{-1})^n,$$

and then only retaining the terms of order at most Δt^{p+1} in this expression. More precisely, denoting $C_k = \Pi_f(\mathcal{A}_k^f - \lambda_k)\Pi_f$, we find

$$S_{\Delta t}^f = A^{-1} - \Delta t A^{-1} C_2 A^{-1} + \Delta t^2 (A^{-1} C_2 A^{-1} C_2 A^{-1} - A^{-1} C_3 A^{-1}) + \Delta t^3 C_3 + \dots + \Delta t^p C_p,$$

where the operators C_k are defined using the operators C_k and A^{-1} . The operator $S_{\Delta t}^f$ is well defined and leaves \mathcal{C}_f invariant since each C_k consists in a finite number of applications of operators of the form $C_k A^{-1}$ and a final application of $A^{-1} = \Pi_f(\mathcal{A}_1 + f - \lambda)^{-1}\Pi_f$. It is then easy to check that, by construction, the equality (4.44) is satisfied.

4.6.4 Proof of Proposition 4.23

We first show that, if $Q_{\Delta t}^f$ satisfies Assumption 4.8 with a reference probability measure η , then $\tilde{Q}_{\Delta t}^f$ satisfies Assumption 4.8 with the same measure η . By Assumption 4.8, there exist $\varepsilon > 0$ and a measure $\eta \in \mathcal{P}(\mathcal{X})$ such that, for any bounded measurable nonnegative function φ ,

$$\varepsilon \eta(\varphi) \leq Q_{\Delta t}^f \varphi \leq \varepsilon^{-1} \eta(\varphi), \tag{4.88}$$

so that, applying $U_{\Delta t}^f$ on the right of $Q_{\Delta t}^f$ and $T_{\Delta t}^f$ on the left,

$$\varepsilon \eta(U_{\Delta t}^f \varphi) T_{\Delta t}^f \mathbf{1} \leq T_{\Delta t}^f Q_{\Delta t}^f U_{\Delta t}^f \varphi \leq \varepsilon^{-1} \eta(U_{\Delta t}^f \varphi) T_{\Delta t}^f \mathbf{1}.$$

Using (4.51) leads to

$$\varepsilon \alpha^2 \eta(\varphi) \leq \tilde{Q}_{\Delta t}^f \varphi \leq \alpha^{-2} \varepsilon^{-1} \eta(\varphi),$$

so that $\tilde{Q}_{\Delta t}^f$ satisfies Assumption 4.8. In view of Theorem 4.9, the scheme $\tilde{Q}_{\Delta t}^f$ admits a unique invariant probability measure $\tilde{\mu}_{f,\Delta t}$ and an eigenvalue $\tilde{\lambda}_{\Delta t}$ defined by (4.52). Now, integrating (4.88) with respect to $\mu_{f,\Delta t}$ and using (4.22) gives

$$\varepsilon \eta(\varphi) \leq e^{\Delta t \lambda_{\Delta t}} \mu_{f,\Delta t}(\varphi) \leq \varepsilon^{-1} \eta(\varphi).$$

The same reasoning holds for $\tilde{\mu}_{f,\Delta t}$. There exists therefore $\varepsilon' > 0$ for which the following inequalities hold in the sense of positive measures:

$$\varepsilon' \eta \leq \mu_{f,\Delta t} \leq \frac{1}{\varepsilon'} \eta, \quad \varepsilon' \eta \leq \tilde{\mu}_{f,\Delta t} \leq \frac{1}{\varepsilon'} \eta. \tag{4.89}$$

We are now in position to prove the equality of the eigenvalues $\lambda_{\Delta t}$ and $\tilde{\lambda}_{\Delta t}$ defined respectively by (4.23) and (4.52). From (4.22), it holds, for any $\varphi \in \mathcal{C}$,

$$\int_{\mathcal{X}} (Q_{\Delta t}^f)^n \varphi d\mu_{f,\Delta t} = \left(\int_{\mathcal{X}} Q_{\Delta t}^f \mathbf{1} d\mu_{f,\Delta t} \right) \left(\int_{\mathcal{X}} (Q_{\Delta t}^f)^{n-1} \varphi d\mu_{f,\Delta t} \right) = \left(\int_{\mathcal{X}} Q_{\Delta t}^f \mathbf{1} d\mu_{f,\Delta t} \right)^n \left(\int_{\mathcal{X}} \varphi d\mu_{f,\Delta t} \right).$$

Applying this last relation to $U_{\Delta t}^f \varphi$ for $\varphi \in \mathcal{C}$ and using the definition of $\lambda_{\Delta t}$,

$$\int_{\mathcal{X}} (Q_{\Delta t}^f)^n U_{\Delta t}^f \varphi d\mu_{f,\Delta t} = \left(\int_{\mathcal{X}} Q_{\Delta t}^f \mathbf{1} d\mu_{f,\Delta t} \right)^n \left(\int_{\mathcal{X}} U_{\Delta t}^f \varphi d\mu_{f,\Delta t} \right) = e^{n \Delta t \lambda_{\Delta t}} \int_{\mathcal{X}} U_{\Delta t}^f \varphi d\mu_{f,\Delta t}.$$

Similarly,

$$\int_{\mathcal{X}} \left(\tilde{Q}_{\Delta t}^f \right)^n \varphi d\tilde{\mu}_{f,\Delta t} = \left(\int_{\mathcal{X}} \tilde{Q}_{\Delta t}^f \mathbb{1} d\tilde{\mu}_{f,\Delta t} \right)^n \left(\int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t} \right) = e^{n\Delta t \tilde{\lambda}_{\Delta t}} \int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t}.$$

It then follows that, for any positive $\varphi \in \mathcal{C}$,

$$e^{n\Delta t(\lambda_{\Delta t} - \tilde{\lambda}_{\Delta t})} = \frac{\int_{\mathcal{X}} \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \varphi d\mu_{f,\Delta t}}{\int_{\mathcal{X}} \left(\tilde{Q}_{\Delta t}^f \right)^n \varphi d\tilde{\mu}_{f,\Delta t}} \times \frac{\int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t}}{\int_{\mathcal{X}} U_{\Delta t}^f \varphi d\mu_{f,\Delta t}} = \frac{\int_{\mathcal{X}} \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \varphi d\mu_{f,\Delta t}}{\int_{\mathcal{X}} T_{\Delta t}^f \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \varphi d\tilde{\mu}_{f,\Delta t}} \times \frac{\int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t}}{\int_{\mathcal{X}} U_{\Delta t}^f \varphi d\mu_{f,\Delta t}}. \quad (4.90)$$

It remains to note that the right hand side of (4.90) is uniformly bounded in n . Indeed, denoting by $\phi_n = \left(Q_{\Delta t}^f \right)^n U_{\Delta t}^f \varphi$ for a positive $\varphi \in \mathcal{C}$, we obtain using (4.51) and (4.89):

$$0 \leq \frac{\int_{\mathcal{X}} \phi_n d\mu_{f,\Delta t}}{\int_{\mathcal{X}} T_{\Delta t}^f \phi_n d\tilde{\mu}_{f,\Delta t}} \leq \frac{\int_{\mathcal{X}} \phi_n d\mu_{f,\Delta t}}{\int_{\mathcal{X}} \alpha \phi_n d\tilde{\mu}_{f,\Delta t}} \leq \frac{\int_{\mathcal{X}} \phi_n (\varepsilon')^{-1} d\eta}{\alpha \int_{\mathcal{X}} \phi_n \varepsilon' d\eta} \leq \frac{1}{\alpha (\varepsilon')^2},$$

this bound being independant of n . Similarly,

$$0 \leq \frac{\int_{\mathcal{X}} \varphi d\tilde{\mu}_{f,\Delta t}}{\int_{\mathcal{X}} U_{\Delta t}^f \varphi d\mu_{f,\Delta t}} \leq \frac{1}{\alpha (\varepsilon')^2}.$$

Therefore, the right-hand side of (4.90) is uniformly bounded for all $n \geq 0$, which proves that $\lambda_{\Delta t} \leq \tilde{\lambda}_{\Delta t}$ by taking the limit $n \rightarrow +\infty$. A similar reasoning leads to $\tilde{\lambda}_{\Delta t} \leq \lambda_{\Delta t}$, hence $\lambda_{\Delta t} = \tilde{\lambda}_{\Delta t}$.

Acknowledgments

We are grateful to Mathias Rousset for his help in understanding Feynman–Kac models. We also thank Frédéric Cérou, Jonathan Mattingly, Julien Roussel, Hugo Touchette and Jonathan Weare for fruitful discussions, and the anonymous referees for their useful comments on the published version [182].

4.7 Markov contractions and Dobrushin coefficients

Denoting by $\mathcal{M}(\mathcal{X})$ is the set of measures over \mathcal{X} , we define $\mathcal{M}_0(\mathcal{X}) = \{\eta \in \mathcal{M}(\mathcal{X}) \mid \eta(\mathcal{X}) = 0\}$ the set of (unsigned) measures with zero mass. The contraction norm of a Markov operator $Q : \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{X})$ is

$$\|Q\| = \sup_{\eta \in \mathcal{M}_0(\mathcal{X})} \frac{\|\eta Q\|_{\text{TV}}}{\|\eta\|_{\text{TV}}} = \sup_{\mu, \nu \in \mathcal{P}(\mathcal{X})} \frac{\|\mu Q - \nu Q\|_{\text{TV}}}{\|\mu - \nu\|_{\text{TV}}},$$

the second equality coming from the fact that all elements in $\mathcal{M}_0(\mathcal{X})$ are proportional to the difference of two probability measures. In particular,

$$\|\mu Q - \nu Q\|_{\text{TV}} \leq \|Q\| \|\mu - \nu\|_{\text{TV}}.$$

A fundamental tool [115, 113, 114] for the study of Feynman–Kac type semigroups (4.15) and introduced by Dobrushin [127, 128] is the so-called Dobrushin ergodic coefficient, which can be defined for a Markov operator Q as:

$$\alpha(Q) = \inf_{\substack{x, x' \in \mathcal{X} \\ \{A_i\}_{1 \leq i \leq m} \subset \mathcal{X}}} \left\{ \sum_{i=1}^m \min(Q(x, A_i), Q(x', A_i)) \right\}, \quad (4.91)$$

where the infimum in the last equality runs over points $x, x' \in \mathcal{X}$ and all partitions $(A_i)_{i=1}^m$ of \mathcal{X} . If we interpret $Q(x, A_i)$ as the probability of going from x into the set A_i , we see that this coefficient provides information on the mixing properties of the operator Q . The link between this coefficient and the contraction properties of Q is made precise by the following relationship [127, 128]:

$$\|Q\| = 1 - \alpha(Q). \quad (4.92)$$

As a result, a minorization condition on Q translates into a contraction of the operator through its ergodic coefficient $\alpha(Q)$. Relation (4.92) is essentially obtained by a Hahn decomposition of measures of zero mass, as made precise in [127, 128].

CHAPTER 5

ADAPTIVE SAMPLING OF LARGE DEVIATIONS FUNCTIONS

The material for this chapter has been published in the Journal of Statistical Physics [184].

Contents

5.1	Motivation	179
5.2	Framework	180
5.2.1	Model and notation	180
5.2.2	Large deviations functions and driven process	182
5.3	Adaptive algorithm	184
5.3.1	Time discretization	184
5.3.2	Stochastic approximation and annealing	185
5.3.3	Spatial projection	186
5.3.4	Algorithm	186
5.3.5	Remarks	187
5.4	Applications	188
5.4.1	Ornstein–Uhlenbeck process	188
5.4.2	Periodic diffusion	191
5.5	Conclusion	192

Abstract. We introduce and test an algorithm that adaptively estimates the large deviation functions introduced in Chapter 3. These functions play an important role for predicting the probability and pathways of rare events in stochastic processes, as well as for understanding the physics of nonequilibrium systems driven in steady states by external forces and reservoirs. The algorithm uses methods from risk-sensitive and feedback control to estimate from a single trajectory a new process, called the driven process, known to be efficient for importance sampling, as presented in Section 1.4.2. Its advantages compared to other simulation techniques, such as splitting or cloning, are discussed and illustrated with simple equilibrium and nonequilibrium diffusion models.

5.1 Motivation

We consider in this paper the problem of estimating large deviations functions introduced in Section 1.3.1 and studied precisely in Chapter 3. These functions have a wide range of applications in engineering and physical sciences, where they are used to predict the probability of rare events [388, 119, 121] and to understand how these events arise via transition paths or modified processes [250, 96, 97]. Large deviations theory also underlies now much of the research on nonequilibrium systems driven in steady states by non-conservative forces or boundary reservoirs [406, 122, 228]. In this context, large deviations functions play the role of nonequilibrium potentials, similar to the entropy or free energy,

that characterize the steady state and fluctuations of physical quantities, such as energy or particle currents exchanged with reservoirs, as well as the presence of symmetries [227] and phase transitions in fluctuations [197, 233, 170, 8].

Recently, many efforts have been devoted to developing numerical algorithms for estimating large deviations functions that go beyond the direct sampling of probabilities, which require prohibitively large samples, as explained in Section 1.4.2. The most popular algorithms are variations of two basic approaches used in rare event simulations, namely:

1. Splitting [109, 79, 13] or cloning [212, 202, 289] algorithms, which use population dynamics to estimate probabilities or generating functions that have a multiplicative structure in time;
2. Importance sampling [75, 256, 15] (including transition path sampling [50]), which modifies the process to be simulated, so as to transform rare events into typical events that can be simulated efficiently.

These aspects are detailed in Section 1.4, and we have seen an application of cloning algorithm in Chapter 4. Deterministic methods not based a priori on sampling can also be used, including spatial discretizations of various spectral and optimal control representations of large deviations functions, which work well for low-dimensional systems, in addition to action minimization methods, which can be applied in the low-noise or low-temperature limit [241, 414, 208] as illustrated in Chapter 6.

In this chapter, we propose an algorithm that combines spectral methods with importance sampling to efficiently estimate large deviations functions in an adaptive way. The core of the algorithm comes from recent works on learning algorithms for risk-sensitive control of Markov chains [58, 5, 28, 57], which we adapt to continuous-time diffusion processes and to the problem of estimating large deviations functions. The algorithm works by estimating or learning “on the fly” a modified process, called the auxiliary or *driven process*, which corresponds to the process that is asymptotically equivalent to the original process conditioned on the rare event of interest [97] or, alternatively, to the exponential tilting of the original process, known to be efficient for importance sampling [98] (see the discussion in Section 1.4.2). This modified process is given by a principal eigenvalue problem related to the Feynman–Kac equation or, equivalently, by a stochastic optimal control problem [98] that we solve iteratively using stochastic approximation and feedback control methods.

The main advantage of this algorithm, compared to splitting or cloning, is that it does not require the simulation of many copies of the considered process – it runs on one long trajectory of that process, modified with a feedback-reinforcement rule, to adaptively learn the driven process, thereby reducing significantly the complexity of estimating large deviations functions. The calculation of error bars for the estimated quantities is also simplified compared to other techniques, as the algorithm is based on simple time averages and stochastic approximations [94, 372, 30]. Finally, the errors incurred by discretizing continuous-time processes and functionals can be analysed in a precise way, in principle, via Feynman–Kac semigroups as made precise in Chapter 4.

We discuss these advantages and test the algorithm in Section 5.4 with simple equilibrium and nonequilibrium diffusions, after introducing the general model and notations in Section 5.2 and the algorithm in Section 5.3. The results at this point are preliminary and are presented as a proof of concept of the algorithm. More detailed results about the time-discretization and sampling errors will be addressed in future works, together with more complex applications involving interacting particle systems and higher-dimensional diffusions.

5.2 Framework

5.2.1 Model and notation

We still consider in this chapter an ergodic diffusion $(X_t)_{t \geq 0}$ evolving in a state space $\mathcal{X} \subset \mathbb{R}^d$ with appropriate boundary conditions according to the following stochastic differential equation (SDE):

$$dX_t = b(X_t) dt + \sigma dB_t, \quad (5.1)$$

where $b : \mathcal{X} \rightarrow \mathbb{R}^d$ is a smooth drift function, $(B_t)_{t \geq 0}$ is an m -dimensional Brownian motion, and σ is a $d \times m$ matrix, assumed to be constant for simplicity (see [97] for a treatment of diffusions with multiplicative noise). The generator of this diffusion reads

$$\mathcal{L} = b \cdot \nabla + \frac{1}{2} \nabla \cdot D \nabla, \quad (5.2)$$

where \cdot denotes the scalar product and

$$D = \sigma \sigma^T,$$

with T as the transpose, is the diffusion matrix, assumed to be positive definite (we thus use the notation D instead of S in Chapter 3). This is the generator of the evolution semigroup $(P_t)_{t \geq 0}$, defined by

$$P_t \varphi(x) = \mathbb{E}[\varphi(X_t) | X_0 = x], \quad (5.3)$$

for all time $t \geq 0$ and any smooth test function φ . The dual \mathcal{L}^\dagger of \mathcal{L} in the space $L^2(dx)$ of square-integrable functions with respect to the Lebesgue measure is the generator of the Fokker–Planck equation

$$\partial_t \rho_t = \mathcal{L}^\dagger \rho_t, \quad (5.4)$$

which gives the evolution of the probability density ρ_t of X_t starting from some initial density ρ_0 for X_0 .

Our goal here is to study the fluctuations of time-integrated functionals of X_t , called *observables*, having the general form

$$A_t = \frac{1}{t} \int_0^t f(X_s) ds + \frac{1}{t} \int_0^t g(X_s) \circ dX_s, \quad (5.5)$$

where $f : \mathcal{X} \rightarrow \mathbb{R}$ and $g : \mathcal{X} \rightarrow \mathbb{R}^d$ are reasonably smooth functions (*e.g.* continuous) and \circ denotes the Stratonovich product [350]¹. Such a functional defined over the time horizon $[0, t]$ can represent, for example, a control cost associated with the state X_t and its increments [95] or a physical quantity integrated in time, such as the work performed on a particle by external forces or the heat exchanged by a particle with its environment [385].

Assuming that the process is ergodic with respect to an invariant measure $\mu(dx) = \rho^*(x) dx$ with smooth density ρ^* , we have almost surely

$$A_t \xrightarrow[t \rightarrow \infty]{} \int_{\mathcal{X}} f(x) \rho^*(x) dx + \int_{\mathcal{X}} g(x) \cdot J^*(x) dx = a^*, \quad (5.6)$$

where, for any $x \in \mathcal{X}$,

$$J^*(x) = b(x) \rho^*(x) - \frac{D}{2} \nabla \rho^*(x) \quad (5.7)$$

is the stationary current field associated with ρ^* [98]. The theory of large deviations [119] refines this ergodic theorem, generalized here with the additional g term, by providing estimates for the rate at which the probability distribution of A_t concentrates on its ergodic value a^* . Such estimates can be derived under general conditions² and take, in the simplest case, the form

$$\lim_{t \rightarrow \infty} -\frac{1}{t} \log \mathbb{P}(A_t \in B) = \min_{a \in B} I(a), \quad (5.8)$$

for any Borel subset B of \mathbb{R} , where $I : \mathbb{R} \rightarrow [0, +\infty]$ is a positive function such that $I(a^*) = 0$. When this limit exists, A_t is said to satisfy the *large deviations principle* (LDP) with rate function I , and we refer to Section 1.2 for an introductory presentation and precise definitions. Formally, this means that

$$\mathbb{P}(A_t \in da) = e^{-tI(a) + o(t)} da, \quad (5.9)$$

where $o(t)$ denotes corrections in the exponential that grow slower than linearly in t . Thus, we see that the rate function provides useful information about the fluctuations of A_t : the likelihood that $A_t = a$ decays exponentially with time for all $a \neq a^*$, since $I(a) > 0$ in this case, and converges otherwise to 1 as $t \rightarrow \infty$, since $I(a^*) = 0$. Moreover, the rate function is in general not a parabola, meaning that it describes fluctuations that are generally not Gaussian. In this sense, large deviations theory is often seen as an extension of both the ergodic theorem, which describes the concentration of A_t towards its mean, and the central limit theorem, which describes the local Gaussian fluctuations of A_t around its mean [119]. This is explained more in detailed in Section 1.1.2.

¹In order to respect the notation of Chapter 3, we should maybe write $L_t(f, g)$ for the time-integrated functional, but we believe A_t is more digest.

²We refer to Chapter 3 for proofs when $g = 0$, but the situation for $g \neq 0$ still seems to be open from a mathematical perspective.

5.2.2 Large deviations functions and driven process

In practice, the rate function $I(a)$ can be calculated in many different ways other than by direct sampling, which requires exponentially large samples with t [75]. The most common method proceeds from the *scaled cumulant generating function* (SCGF), defined for $k \in \mathbb{R}$ by

$$\lambda(k) = \lim_{t \rightarrow \infty} \frac{1}{t} \log \mathbb{E} [e^{ktA_t}]. \quad (5.10)$$

By the Gärtner-Ellis Theorem [119], the Legendre–Fenchel transform of this function yields the rate function:

$$I(a) = \sup_{k \in \mathbb{R}} \{ka - \lambda(k)\}, \quad (5.11)$$

provided, essentially, that $\lambda(k)$ is differentiable; see Chapter 3.

For the SDE (5.1) and the additive functional (5.5), the SCGF is known to be given by the principal eigenvalue of the operator

$$\mathcal{L}_k = b \cdot (\nabla + kg) + \frac{1}{2}(\nabla + kg) \cdot D(\nabla + kg) + kf, \quad (5.12)$$

which is the generator of the *Feynman–Kac semigroup* P_t^k , defined by

$$P_t^k \varphi(x) = \mathbb{E} [\varphi(X_t) e^{tkA_t} \mid X_0 = x], \quad (5.13)$$

for any smooth function φ ; see Appendix A.2 of [97]³. In the end, the rate function can thus be computed by solving the spectral problem

$$\mathcal{L}_k h_k = \lambda(k) h_k, \quad (5.14)$$

where $\lambda(k)$ is the principal eigenvalue of \mathcal{L}_k and h_k its corresponding eigenvector. This holds provided that this operator has reasonable spectral properties, made precise in the following assumption.

Assumption 5.1. *The operator \mathcal{L}_k defined in (5.12) acting on $L^2(\mu)$ has an isolated largest eigenvalue $\lambda(k)$. Its multiplicity is one and it is associated with a regular eigenvector $h_k \in L^2(\mu)$ such that for all $x \in \mathcal{X}$, $h_k(x) > 0$.*

This assumption holds for many systems, in particular when \mathcal{X} is bounded [40, 182] or when b and g are gradient fields with appropriate growth conditions [304, Section 2.5]; we refer to Chapter 2 for more details. In practice, the spectral problem (5.14) can be solved numerically using standard projection or discretization (Galerkin) methods, which work well for low-dimensional systems [93], or more involved real space renormalization methods when dealing with higher-dimensional systems [204]. Note that, for $g = 0$, \mathcal{L}_k is the usual Feynman–Kac generator with source term kf . Moreover, for $k = 0$, $\mathcal{L}_0 = \mathcal{L}$ is the generator of the SDE, so that $\lambda(0) = 0$ and $h_0 = 1$, the constant unit function.

The numerical method that we propose in the next section attempts to estimate the spectral elements $\lambda(k)$ and h_k in a different way using the fact that they are solutions of the family of eigenproblems

$$P_t^k h_k = e^{t\lambda(k)} h_k, \quad \forall t > 0, \quad (5.15)$$

which can be approximated stochastically. The method also exploits a connection between large deviations and control theory showing that $\lambda(k)$ is the ergodic limit of an optimal control cost satisfying a stochastic Hamilton–Jacobi–Bellman equation [187], and that h_k determines the controlled diffusion achieving the optimal cost, which is the driven process mentioned in the introduction. These results are explained in detail in [98] (see also the references therein and the computations in Section 1.4.2); for the purpose of this chapter, we only state them without proofs.

The controlled diffusion, denoted by $(\tilde{X}_t)_{t \geq 0}$, satisfies the SDE

$$d\tilde{X}_t = b_k(\tilde{X}_t) dt + \sigma dB_t, \quad (5.16)$$

where

$$b_k(x) = b(x) + D[kg(x) + \nabla \log h_k(x)] \quad (5.17)$$

³We should again write $P_t^{kf,kg}$, $\lambda_{f,g}(k)$ and $I_{f,g}$ to match the notation of Part II but prefer the shorthand notation P_t^k , $\lambda(k)$ and I since f and g are fixed.

is the optimal control drift defined for all $x \in \mathcal{X}$, which is motivated in Section 1.4.2 in the case $g = 0$. Under Assumption 5.1, this modified diffusion is ergodic with respect to a new invariant measure μ_k , whose density is

$$\rho_k(x) = h_k(x)l_k(x), \quad (5.18)$$

where l_k is the dual of h_k in $L^2(dx)$ satisfying $\mathcal{L}_k^\dagger l_k = \lambda(k)l_k$ [97, Section 3.3]⁴. These two functions are normalized such that

$$\int_{\mathcal{X}} h_k(x)l_k(x) dx = 1, \quad \int_{\mathcal{X}} l_k(x) dx = 1. \quad (5.19)$$

Because of the change of process, the observable A_t must concentrate on a new value, which can be shown to be given by $\lambda'(k)$, that is,

$$A_t \xrightarrow[t \rightarrow \infty]{} \lambda'(k) = a_k, \quad (5.20)$$

almost surely with respect to $(\tilde{X}_t)_{t \geq 0}$. Similarly, the control cost

$$C_t = kA_t - R_t, \quad (5.21)$$

where

$$R_t = \frac{1}{2t} \int_0^t [b(\tilde{X}_s) - b_k(\tilde{X}_s)] \cdot D^{-1}[b(\tilde{X}_s) - b_k(\tilde{X}_s)] ds, \quad (5.22)$$

reaches in the ergodic limit the value $\lambda(k)$, so that

$$C_t \xrightarrow[t \rightarrow \infty]{} \lambda(k) \quad (5.23)$$

almost surely under $(\tilde{X}_t)_{t \geq 0}$ [98, Section 4]⁵. Finally, it can be shown by Legendre duality that the rate function at the value $a_k = \lambda'(k)$ is given by the ergodic limit of R_t above, leading to

$$R_t \xrightarrow[t \rightarrow \infty]{} I(a_k), \quad (5.24)$$

almost surely with respect to $(\tilde{X}_t)_{t \geq 0}$ [98, Section 4]. Note that the diffusion is not modified for $k = 0$, so that $a_0 = a^*$, $\lambda'(0) = a^*$, and $I(a^*) = 0$.

These ergodic limits provide direct estimators of the SCGF and the rate function, based on a single trajectory of the driven process, which can be simulated for different values of the parameter $k \in \mathbb{R}$. For the SCGF, there are in fact three possible estimators:

1. From (5.20): the value of A_t , integrated numerically with the condition $\lambda(0) = 0$;
2. From (5.23): the value of C_t ;
3. The eigenvalue returned by the algorithm proposed in Section 5.3.

In practice, we find that the last estimator is more stable, although the first and second are more adapted to obtain error bars.

For the rate function, we have two possible estimators:

1. The Legendre transform (5.11) of the SCGF, given in parametric form by

$$I(a_k) = ka_k - \lambda(k), \quad (5.25)$$

where a_k is either estimated from A_t or by taking the numerical derivative of $\lambda(k)$;

2. From (5.24): the value of R_t obtained for the value of A_t , giving the couple (A_t, R_t) .

⁴The notation is slightly different than that of Chapter 4, see for instance Proposition 4.1.

⁵Note the similarity between the cost function R_t and the action functional (1.38).

In practice, we find that the first estimator based on the Legendre transform is more reliable. By comparison, the computation of $I(a)$ based on (5.22) involves the optimal drift b_k and, therefore, the logarithmic derivative of h_k , which is more difficult to estimate in a stable way.

In all cases, error bars can be constructed from the same trajectory by estimating, in principle, the variance of A_t , C_t and R_t using covariance techniques for Markov processes [94, 372, 30]. This is an advantage over splitting and cloning algorithms, for which the calculation of errors bars is difficult, as they involve correlated copies or “clones” of the simulated process [377, 337, 338].

In closing, it is interesting to note that the driven process can also be interpreted as the change of process in the importance sampling of the probability $\mathbb{P}(A_t \in B)$ that is optimal in the sense of logarithmic or asymptotic efficiency [75]. Therefore, it can be used not only to estimate the SCGF and rate function, but also to estimate the actual probability $\mathbb{P}(A_t \in B)$ in an efficient way [58]. The optimal change of process in this case is known to correspond to the exponential tilting of the original process [75], which is a time-dependent process in general; see Appendix D of [97]. In the ergodic limit, this process converges to a homogeneous process corresponding exactly to the driven process (5.16). We refer to [98] for more details about these results.

5.3 Adaptive algorithm

We are now ready to present the algorithm for estimating the SCGF and the rate function of A_t for continuous-time diffusions. The algorithm is based, as mentioned, on prior algorithms proposed in [58, 5, 28, 57] for Markov chains and exploits the fact that

$$e^{-t\lambda(k)}(P_t^k \varphi) \xrightarrow[t \rightarrow \infty]{} h_k \int_{\mathcal{X}} \varphi(x) l_k(x) dx, \quad (5.26)$$

for any smooth test function φ . This result follows under Assumption 5.1; see [97, Section III.B] or the results of Chapter 2 for a proof. The algorithm that we propose works from this limit by approximating the action of the Feynman–Kac semigroup $(P_t^k)_{t \geq 0}$ in a stochastic way as a time average computed over a long trajectory of the diffusion. Moreover, it continuously modifies the diffusion as we estimate h_k to construct the driven process $(\tilde{X}_t)_{t \geq 0}$, which underlies the estimators of the large deviations functions.

The algorithm is presented next. Our contribution compared to [58, 5, 28, 57] is to consider general time-continuous processes, time-additive functionals of these processes that depend on both their state and increments and, more importantly, to construct the driven process explicitly so as to estimate the SCGF and the rate function adaptively using the estimators introduced in the previous section.

5.3.1 Time discretization

The first step required in the algorithm is to discretize in time the SDE (5.1) and its associated Feynman–Kac semigroup (5.13) by transforming the Markov diffusion $(X_t)_{t \geq 0}$ into a Markov chain $(x_n)_{n \in \mathbb{N}}$ with a small time step. Many discretization schemes can be used for the SDE; see, for instance, [270]. Here we use the standard Euler–Maruyama scheme with constant time step Δt , given by

$$x_{n+1} = x_n + b(x_n)\Delta t + \sigma\sqrt{\Delta t}\xi_n, \quad (5.27)$$

where $(\xi_n)_{n \geq 0}$ is a sequence of independent standard d -dimensional Gaussian random variables. The corresponding discretization of the evolution semigroup P_t over a time step Δt is denoted by $Q_{\Delta t}$, so that

$$Q_{\Delta t}\varphi(x) = \mathbb{E}[\varphi(x_{n+1})|x_n = x], \quad (5.28)$$

for any test function φ and $x \in \mathcal{X}$. We refer again to [270, 304] and Section 1.4.1 for more information about the discretization of SDEs and their weak error analysis.

We have seen in Chapter 4 that many discretizations also exist for the Feynman–Kac semigroup P_t^k . Here, we use the natural scheme where the diffusion is discretized as above and the integral of A_t is discretized as a Riemann sum with the left-point rule for the integral involving f and the mid-point rule for the Stratonovich integral involving g . The action of P_t^k is thus replaced by

$$Q_{\Delta t}^k\varphi(x) = \mathbb{E} \left[e^{k \left[f(x_n)\Delta t + g\left(\frac{x_{n+1}+x_n}{2}\right) \cdot (x_{n+1}-x_n) \right]} \varphi(x_{n+1}) \mid x_n = x \right]. \quad (5.29)$$

For our purposes, h_k will be approximated by recursive applications of $Q_{\Delta t}^k$, based on the following assumption.

Assumption 5.2. *There exist a time step $\Delta t^* > 0$ and $p > 0$ such that, for $0 < \Delta t \leq \Delta t^*$,*

$$Q_{\Delta t}^k h_{k,\Delta t} = e^{\Delta t \lambda_{\Delta t}(k)} h_{k,\Delta t}, \quad (5.30)$$

where

$$h_{k,\Delta t} = h_k + O(\Delta t^p), \quad \lambda_{\Delta t}(k) = \lambda(k) + O(\Delta t^p). \quad (5.31)$$

This assumption means that the time-discretized operator $Q_{\Delta t}^k$ admits h_k as an approximate eigenvector with approximate eigenvalue $\lambda(k)$. This applies, for example, when \mathcal{X} is compact. In this case, precise estimates for the errors in (5.31) are obtained for $g = 0$ in Chapter 4. We believe these error estimates can be extended to $g \neq 0$ under appropriate modification of the assumptions, but the situation when \mathcal{X} is unbounded is difficult.

In the following, we will drop the subscript Δt on $h_{k,\Delta t}$ and $\lambda_{\Delta t}(k)$ to simplify the notations, and will present the algorithm essentially as if h_k were an exact eigenvector of $Q_{\Delta t}^k$ with exact eigenvalue $\lambda(k)$. However, we should keep in mind that this is only approximately true due to the errors in Δt . We will comment on this in Section 5.4 with specific numerical examples.

5.3.2 Stochastic approximation and annealing

The main ingredient of the algorithm is the limit (5.26) of the Feynman–Kac semigroup, which shows that h_k and $\lambda(k)$ can be computed by successively applying $Q_{\Delta t}^k$ to an initial guess φ , so as to obtain

$$(Q_{\Delta t}^k)^n \varphi \sim e^{n\Delta t \lambda(k)} h_k \int_{\mathcal{X}} \varphi(x) l_k(x) dx \quad (5.32)$$

as $n \rightarrow \infty$. To perform this iteration, which is a functional version of the well-known *power method* for matrices [120], we apply a stochastic approximation [353, 30, 28] whereby the expectation appearing in the action of $Q_{\Delta t}^k$ is replaced by the iterates of the Markov chain:

$$\begin{aligned} Q_{\Delta t}^k \varphi(x_n) &= \mathbb{E} \left[e^{k \left[f(x_n) \Delta t + g \left(\frac{x_{n+1} + x_n}{2} \right) \cdot (x_{n+1} - x_n) \right]} \varphi(x_{n+1}) \mid x_n \right] \\ &\approx e^{k \left[f(x_n) \Delta t + g \left(\frac{x_{n+1} + x_n}{2} \right) \cdot (x_{n+1} - x_n) \right]} \varphi(x_{n+1}), \end{aligned} \quad (5.33)$$

where x_{n+1} is a random variable distributed according to $Q_{\Delta t}(x_n, \cdot)$. This approximation is known to reproduce the expectation as a statistical average in the ergodic limit $n \rightarrow \infty$ [5, 28, 30].

In our case, we simulate not the Markov chain $(x_n)_{n \in \mathbb{N}}$ but a modified chain, corresponding to the discretization of the driven process (5.16), which we express as

$$\tilde{x}_{n+1} = \tilde{x}_n + [b(\tilde{x}_n) + F_n(\tilde{x}_n)] \Delta t + \sigma \sqrt{\Delta t} \xi_n, \quad (5.34)$$

where

$$F_n = D(kg + \nabla \log h_k^n) \quad (5.35)$$

is the extra biasing force derived, according to (5.17), from the estimate h_k^n of h_k at time n . In this case, the evolution (5.33) is modified by the Girsanov formula [350] to

$$\begin{aligned} Q_{\Delta t}^k \varphi(\tilde{x}_n) &= \mathbb{E} \left[e^{k \left[f(\tilde{x}_n) \Delta t + g \left(\frac{\tilde{x}_{n+1} + \tilde{x}_n}{2} \right) \cdot (\tilde{x}_{n+1} - \tilde{x}_n) \right]} \varphi(\tilde{x}_{n+1}) R_n(\tilde{x}_n, \tilde{x}_{n+1}) \mid \tilde{x}_n = x \right] \\ &\approx e^{k \left[f(\tilde{x}_n) \Delta t + g \left(\frac{\tilde{x}_{n+1} + \tilde{x}_n}{2} \right) \cdot (\tilde{x}_{n+1} - \tilde{x}_n) \right]} \varphi(\tilde{x}_{n+1}) R_n(\tilde{x}_n, \tilde{x}_{n+1}), \end{aligned} \quad (5.36)$$

where

$$R_n(\tilde{x}_n, \tilde{x}_{n+1}) = \exp \left(-\frac{1}{2\sigma^2} F_n^2(\tilde{x}_n) \Delta t - \sqrt{\frac{\Delta t}{\sigma^2}} F_n(\tilde{x}_n) \cdot \xi_n \right) \quad (5.37)$$

is the Radon–Nikodym derivative of the transition kernel of x_n with respect to that of \tilde{x}_n .

In the end, we also apply an annealing scheme, commonly used in stochastic approximations, which consists in replacing the update rule (5.32), defined by $\varphi^0 = \varphi$ and $\varphi^{n+1} = Q_{\Delta t}^k \varphi^n$, by the scheme

$$\varphi^{n+1} = \varphi^n + a_n (Q_{\Delta t}^k \varphi^n - \varphi^n), \quad (5.38)$$

where $(a_n)_{n \in \mathbb{N}}$ is a decreasing sequence, often called the adaption or learning sequence, which acts as a smoothing parameter, filtering here the noisy update of the eigenfunction. This sequence is usually chosen in such a way that

$$\sum_{n \geq 0} a_n = \infty, \quad \sum_{n \geq 0} a_n^2 < \infty, \quad (5.39)$$

with the understanding that a_n should not be decreased too slowly, so as to limit noise, nor too fast, so as to reach the “correct” fixed point. These conditions can be relaxed under stability assumptions or by an averaging procedure [353]; see [30] for more details.

5.3.3 Spatial projection

The iteration just described for approximating h_k can be performed numerically by discretizing the state space \mathcal{X} into small cells (grid discretization). For high-dimensional systems, however, it is more convenient to use a Galerkin-type approximation of the eigenvalue and eigenfunction [28], obtained by projecting the problem onto a set of basis functions $\{\phi_j\}_{j=1}^M$ with $\phi_j : \mathcal{X} \rightarrow \mathbb{R}$ [93]. Let us denote by $\mathcal{H}_M = \text{Span}\{\phi_j\}_{j=1}^M$ the space spanned by these functions. Then the M -dimensional eigenproblem that we need to solve is

$$Q_{\Delta t}^k \tilde{h} = e^{\Delta t \lambda} \tilde{h}, \quad \tilde{h} \in \mathcal{H}_M, \quad (5.40)$$

using, for notational convenience, the same symbols \tilde{h} and λ for the exact and the projected spectral elements. We also drop from now on the parameter k , which will be implicit.

The eigenfunction h is expressed in that basis as

$$\tilde{h}(x) = \sum_{j=1}^M \alpha_j \phi_j(x) = \alpha^T \phi(x), \quad (5.41)$$

where $\alpha = [\alpha_1, \dots, \alpha_M]^T$ and $\phi(x) = [\phi_1(x), \dots, \phi_M(x)]^T$. Multiplying (5.40) by ϕ_i for $i \in \{1, \dots, M\}$ and integrating over any measure η on \mathcal{X} yields

$$\sum_{j=1}^M \alpha_j \int_{\mathcal{X}} \phi_i(Q_{\Delta t}^k \phi_j) d\eta = e^{\Delta t \lambda} \sum_{j=1}^M \alpha_j \int_{\mathcal{X}} \phi_i \phi_j d\eta, \quad i \in \{1, \dots, M\}. \quad (5.42)$$

As a result, we see that the vector of coefficients $\alpha \in \mathbb{R}^M$ is the principal solution of the eigenproblem

$$A\alpha = \Lambda B\alpha, \quad \lambda = \frac{1}{\Delta t} \log \Lambda, \quad (5.43)$$

where

$$A = \int_{\mathcal{X}} \phi(Q_{\Delta t}^k \phi^T) d\eta, \quad B = \int_{\mathcal{X}} \phi \phi^T d\eta. \quad (5.44)$$

Note that the matrix B is invertible as soon as the $\{\phi_j\}_{j=1}^M$ form a linearly independent family.

5.3.4 Algorithm

We are now ready to describe all the steps of the algorithm that estimates the principal eigenvalue of \mathcal{L}_k and its corresponding eigenfunction. For a fixed $k \in \mathbb{R}$, we initiate the process at a position $x_0 \in \mathcal{X}$ and define a first approximation \tilde{h}^0 of h as

$$\tilde{h}^0 = (\alpha^0)^T \phi, \quad (5.45)$$

where ϕ is the vector of basis functions and α^0 is the initial vector of coefficients, chosen such that h^0 is constant (although prior information could be incorporated at this stage). At each iteration, we then perform the following steps:

1. Draw a new position \tilde{x}_{n+1} according to the Markov chain (5.34);
2. Compute the extra bias F_n according to (5.35), which in the function basis takes the form

$$F_n = D \left(kg + \frac{\sum_{j=1}^M \alpha_j^n \nabla \phi_j}{\sum_{j=1}^M \alpha_j^n \phi_j} \right); \quad (5.46)$$

3. Compute the Girsanov weight R_n according to (5.37);
4. Compute the matrices A_{n+1} and B_{n+1} using the formulae

$$\begin{aligned} A_{n+1} &= \frac{1}{n+1} \sum_{m=0}^n e^{k[f(\tilde{x}_m)\Delta t + g\left(\frac{\tilde{x}_{m+1} + \tilde{x}_m}{2}\right) \cdot (\tilde{x}_{m+1} - \tilde{x}_m)]} \phi(\tilde{x}_m) \phi(\tilde{x}_{m+1})^T R_m(\tilde{x}_m, \tilde{x}_{m+1}), \\ B_{n+1} &= \frac{1}{n+1} \sum_{m=0}^n \phi(\tilde{x}_m) \phi(\tilde{x}_m)^T, \end{aligned} \quad (5.47)$$

which follow by projecting (5.36) and (5.44), respectively;

5. Update the coefficient vector α^n , giving the decomposition of the iterate h^n , as

$$\alpha^{n+1} = \alpha^n + a_n \left(\frac{B_n^{-1} A_n}{\tilde{h}^n(x_0)} - I_d \right) \alpha^n, \quad (5.48)$$

where I_d is the identity matrix of size M [28];

6. Estimate the eigenvalue as

$$\lambda^{n+1} = \frac{n}{n+1} \lambda^n + \frac{1}{(n+1)\Delta t} \log(\tilde{h}^n(\tilde{x}_0)), \quad (5.49)$$

which follows from (5.43).

Repeating these steps, it can be proved that the iterates h^n and λ^n converge to the solution of the spectral problem (5.40), following the analysis found in [58, 5, 28]; see also [30]. Moreover, because $h^n \rightarrow h$, the Markov chain $(\tilde{x}_n)_{n \in \mathbb{N}}$ samples in the long run the Euler–Maruyama discretization of the driven process (5.16).

5.3.5 Remarks

The following are technical remarks worth noting about the algorithm:

1. The matrix A_n can be updated at each step using

$$A_{n+1} = \frac{n}{n+1} A_n + \frac{1}{n+1} e^{k[f(\tilde{x}_n)\Delta t + g\left(\frac{\tilde{x}_{n+1} + \tilde{x}_n}{2}\right) \cdot (\tilde{x}_{n+1} - \tilde{x}_n)]} \phi(\tilde{x}_n) \phi(\tilde{x}_{n+1})^T R_n, \quad (5.50)$$

instead of the sum shown in (5.47). Similarly, the matrix B_n can be updated, following [28], using the Sherman–Morrison–Woodbury formula, which leads here to

$$B_{n+1}^{-1} = \frac{n+1}{n} B_n^{-1} - \frac{n+1}{n} \frac{B_n^{-1} \phi(\tilde{x}_n) \phi(\tilde{x}_n)^T B_n^{-1}}{n + \phi(\tilde{x}_n)^T B_n^{-1} \phi(\tilde{x}_n)}. \quad (5.51)$$

The advantage of this formula is that the computation required for updating the coefficient α^n scales with the number M of basis vectors as M^2 , as is common in the power method [120], whereas the typical cost of inverting the matrix B scales as M^3 .

2. We normalize the eigenfunction at every iteration by setting $\tilde{h}(\tilde{x}_0) = 1$ at an arbitrary location, taken here to be the initial state \tilde{x}_0 [58]. This prevents the norm of the eigenfunction from diverging or decaying to zero, as is common in the power method, and provides an estimate of $\lambda(k)$ through (5.49). This normalization step can be based on other norms, at the expense of computing integrals.
3. The algorithm can be run with the original (unbiased) process $(X_t)_{t \geq 0}$, but the estimation of the matrices A and B in this case typically suffers from large statistical errors due to the high variance of the underlying estimator. This is a known problem related to the estimation of Feynman–Kac functionals and exponential integrals in general [231, 230, 433, 374]. Biasing the dynamics with the driven process $(\tilde{X}_t)_{t \geq 0}$ reduces this variance in an optimal way (in the sense of asymptotic or logarithmic efficiency) by forcing the exploration of the process in important regions of the state space where the integrand of the generating function $\mathbb{E}[e^{ktA_t}]$ is largest [98]. This issue is explained more precisely by the computations of Section 1.4.2.

4. As the dynamics is biased towards the driven process $(\tilde{X}_t)_{t \geq 0}$, the vector α^n of basis coefficients representing h_k converges towards the solution of the eigenproblem (5.43), with the matrices A and B computed as ergodic averages under $(\tilde{X}_t)_{t \geq 0}$, so that $\eta = \mu_k$ in (5.44). Even for a small number of basis functions, it would be a priori impossible to compute these matrices by numerical quadrature. This shows that the algorithm can be used to obtain good approximations of h_k even for high-dimensional systems, provided that enough basis functions are used to represent the support of μ_k , which is typically concentrated on a subset of \mathcal{X} .
5. The algorithm is stable despite the fact that it includes the Girsanov reweighting factor, which is exponential in time. The reason for this stability, already noted in [5], is that the Girsanov weight is computed and accumulated incrementally over single time steps in (5.47).
6. The learning sequence $(a_n)_{n \in \mathbb{N}}$ is chosen here in the following way:
 - (a) For $0 \leq n \leq N_1$, we take $a_n = 0$, so there is no adaption at the level of h and λ , although the matrices A and B are evolved – this “burn-in” period allows for a better initial guess of the various functions estimated through their ergodic averages;
 - (b) For $N_1 < n \leq N_1 + N_{\text{iter}} - N_2$, we take $a_n = 1$, that is, the full information of the process is taken into account;
 - (c) For $N_1 + N_{\text{iter}} - N_2 < n \leq N_1 + N_{\text{iter}}$, we take $a_n = C/(n - (N_1 + N_{\text{iter}} - N_2))$ with constant $C > 0$, so the process learns less with time, smoothing the noise in the long run. In the following, we choose $C = 1$.

The times N_1 and N_2 can be fixed or can be chosen dynamically according to some stopping rule.

7. In practice, we can perform independent simulations with different values of k to obtain an interpolation of the SCGF over some range, say, $[k_{\min}, k_{\max}]$, which can then be used to obtain the rate function by Legendre transform. Alternatively, we can do a simulation in which k is slowly increased from $k = 0$ in a “quasi-static” way, so as to adaptively update the biasing force over a range of values for k [341] or to reach some prescribed value of the SCGF for an unknown k [58].

5.4 Applications

We apply in this section the algorithm to two simple test cases involving one-dimensional diffusions. The first is the Ornstein–Uhlenbeck process, for which the large deviations functions of the area per unit time are known exactly [97], while the second is a driven diffusion on the circle, often used in physics as a model of nonequilibrium systems, including Josephson junctions perturbed by thermal noise and Brownian particles controlled by external forces [366, 360, 100]. We discuss for both the convergence and efficiency of the algorithm.

5.4.1 Ornstein–Uhlenbeck process

The first example that we consider is the mean area or mean position

$$A_t = \frac{1}{t} \int_0^t X_s ds \quad (5.52)$$

of the Ornstein–Uhlenbeck process on \mathbb{R} satisfying the SDE

$$dX_t = -2\theta X_t dt + \sqrt{2} dB_t, \quad (5.53)$$

where $\theta > 0$. In the notations of Section 5.2, we thus have $b(x) = -2\theta x$, $\sigma = \sqrt{2}$, $f(x) = x$ and $g(x) = 0$, so that

$$\mathcal{L} = -2\theta x \frac{d}{dx} + \frac{d^2}{dx^2}, \quad (5.54)$$

and $\mathcal{L}_k = \mathcal{L} + kx$.




Figure 5.1 – Large deviations of the mean area for the Ornstein–Uhlenbeck process. Left: Evolution of the estimated SCGF in time. Right: Evolution of the biasing potential $V_k(x)$ in time for the linear basis (curves from top to bottom), compared with the final estimate for the Hermite basis and the exact result.

For this process and observable, it can be checked [97] that the SCGF, corresponding to the dominant eigenvalue of \mathcal{L}_k , is

$$\lambda(k) = \frac{k^2}{4\theta^2}, \quad (5.55)$$

so that $I(a) = \theta^2 a^2$ from (5.25). This is expected, since the integral of a Gaussian process is also Gaussian. Moreover, the associated dominant eigenfunction is

$$h_k(x) = e^{\frac{k}{2\theta}x}, \quad (5.56)$$

leading with (5.17) to the optimal drift

$$b_k(x) = -2\theta x + 2(\log h_k(x))' = -2\theta x + \frac{k}{\theta}. \quad (5.57)$$

This shows that a fluctuation of A_t is created in an optimal way by adding a constant to the drift, which “moves” the Gaussian stationary density of the Ornstein–Uhlenbeck process

$$\rho^*(x) = \sqrt{\frac{\theta}{\pi}} e^{-\theta x^2} \quad (5.58)$$

to

$$\rho_k(x) = \sqrt{\frac{\theta}{\pi}} e^{-\theta(x-m_k)^2}, \quad m_k = \frac{k}{2\theta^2}, \quad (5.59)$$

leading to

$$A_t \xrightarrow[t \rightarrow \infty]{} m_k = a_k \quad (5.60)$$

almost surely with respect to \tilde{X}_t , in agreement with (5.20).

We compare the algorithm against these exact results using a simple mesh discretization (first-order finite elements), defined by the basis functions

$$\phi_j(x) = \begin{cases} \frac{x-x_j}{\delta} + 1, & x \in [x_j - \delta, x_j], \\ -\frac{x-x_j}{\delta} + 1, & x \in [x_j, x_j + \delta], \\ 0, & \text{otherwise,} \end{cases} \quad (5.61)$$

where the points x_j define the centers of each “cell” of width $\delta > 0$. In the following, we refer to this basis simply as the *linear basis*. To illustrate the flexibility of the algorithm, we also perform simulations using a Hermite polynomial basis, which forms a complete orthonormal basis in $L^2(\rho^*)$. We run the algorithm with $k = 1$, $\theta = 1/2$, $x_0 = 0$, final time $T = 2 \times 10^4$ and intermediate times $T_1 = N_1\Delta t = T_2 = N_2\Delta t = 2 \times 10^3$ with $\Delta t = 5 \times 10^{-3}$, which a standard time step used in simulations relative to the basic timescale of the dynamics, corresponding here to $1/(2\theta) = 1$. For the linear basis, we use $M = 61$ equally spaced cells with $\delta = 0.25$ around $x = 0$, whereas for the Hermite basis we use only $M = 10$ basis functions.

Figure 5.1 illustrates the results of a typical simulation, starting on the left with the evolution of the estimate of the SCGF, given by (5.49), as time increases. We observe a very good agreement in the long run with the exact value, which for the parameters used is equal to 1, with a faster convergence observed for the Hermite basis compared to the linear basis. As mentioned before, we can also recover the SCGF by recording the stationary value of A_t , which corresponds as above to $m_k = a_k = \lambda'(k)$, and numerically integrate the result in k from $\lambda(0) = 0$. The results obtained are similar to those obtained from the eigenvalue estimate (5.49), and are not shown for this reason.

To understand the convergence of the eigenvalue at the process level, we show in the right plot of Figure 5.1 the evolution of the effective potential V_k associated with the modified drift b_k according to

$$b_k(x) = -V_k'(x). \quad (5.62)$$

We show the results for both the linear basis and the Hermite basis, with the zero of the potential arbitrarily set at $V_k(0) = 0$. We see from these that the tail of the potential takes longer to be estimated correctly, as the process starts to explore values away from $m_0 = 0$. After the convergence time, there is a good agreement with the parabola

$$V_k(x) = \theta(x - m_k)^2, \quad (5.63)$$

which is the exact solution predicted by (5.59). The small errors are due to the finite time step Δt used and the finite basis function set. The small “wiggles” seen in the potential obtained with the linear basis come from the fact that this basis is piecewise linear, so that V_k is piecewise logarithmic. The potential obtained with the Hermite basis is smoother, as expected.

Repeating the simulations for other values of k , we can recover $\lambda(k)$, as shown in the left plot of Figure 5.2. The result is in good agreement with the exact solution (5.55). Estimating $\lambda'(k)$ with (5.20) leads to an estimated rate function, by the Legendre transform (5.25), which also agrees well with the exact rate function, shown in the right plot of Figure 5.2. By comparison, the rate function obtained from the time average (5.24) of the extra biasing force (with the Hermite basis) is not as good: it lies above the exact rate function and shows a larger offset or error as k increases, which does not decrease by reducing Δt or increasing the number of basis functions.

To understand this error, we show in the left plot of Figure 5.3 the evolution of the extra biasing force F_n for $k = 1$, estimated at the current location of the system, which should approach the constant 2, following (5.57). The evolution is noisy, as can be seen, which is expected, since F_n is estimated by the logarithmic derivative

$$F_n(x) = 2 \frac{(h^n(x))'}{h^n(x)}, \quad (5.64)$$

computed in the function basis from (5.46). The derivative amplifies the Monte Carlo errors inherent in the estimate h^n . In addition, the denominator often takes small values away from the mean position m_k , which makes the estimation of the optimal force still more difficult. The right plot of Figure 5.3 shows that the noise on F_n is considerably filtered out by the time average underlying the estimator R_t of the rate function, although a bias remains even after the convergence time of the SCGF, which leads to the offset seen in Figure 5.2. The results in both cases are more noisy for the linear basis because the derivative is not continuous across the different cells.

The offset on R_t remains more or less constant by running independent simulations (error bars on 30 simulations are too small to show), and so appears to be a systematic error or bias. Many factors can account for this bias. First, the rate function estimated from the limit shown in (5.24) is known to be an upper bound on the true rate function [98], which is tight if and only if the modified drift estimated in the simulation is the optimal control drift (5.17). Here, F_n is not constant, as predicted from (5.57), so we expect the estimate R_t to lie above its expected value. Second, we have noticed in simulations that A_t underestimates a_k for large k , which has the effect of further “pushing” the estimate




Figure 5.2 – Large deviations of the mean area for the Ornstein–Uhlenbeck process. Left: Estimated SCGF compared with the exact result. Right: Estimated rate function based on the Legendre transform of the SCGF and on the averaging of the biasing force (using the Hermite basis), compared with the exact result.

of the rate function above $I(a)$. This is most likely due again to F_n being non-constant. Finally, the time average of F_n involves the ratio of two functions, according to (5.64), which fluctuate in time via the updating of α^n in (5.46). As a result, we expect this additional randomness to artificially increase the second moment of F_n , leading to a further bias.

It is difficult to isolate these factors, and all, in fact, seem to play a role. In future works, it would be interesting to study the bias observed in R_t by computing, for example, the time average of F_n^2 using the final estimate of h_k rather than the time-evolved estimate h^n . Different annealing sequences or averaging techniques could also be used to “filter” the time average of F_n^2 and mitigate the bias on R_t . For now, the most efficient and reliable way to compute the rate function is to use the estimator based on the SCGF and its Legendre transform.

5.4.2 Periodic diffusion

We have motivated in the Introduction our interest for nonequilibrium dynamics. We thus consider for the second test a diffusion on the unit circle satisfying the SDE

$$dX_t = (-V'(X_t) + \gamma)dt + \sqrt{2}dB_t, \quad (5.65)$$

with the periodic potential $V(x) = \cos(2\pi x)$ and $\gamma \in \mathbb{R}$ a constant drive. For $\gamma \neq 0$, the total drift cannot be expressed as the gradient of a smooth periodic function, so $(X_t)_{t \geq 0}$ is a nonequilibrium process violating detailed balance [366]. The observable studied for this process is the winding number

$$A_t = \frac{1}{t} \int_0^t dX_s, \quad (5.66)$$

calculated with the real rather than periodic state, which can be interpreted physically as the mean velocity or current of a Brownian particle moving around the circle [410]. In the notation of Section 5.2, we have $b(x) = -V'(x) + \gamma$, $\sigma = \sqrt{2}$, $f(x) = 0$, and $g(x) = 1$.

The SCGF and the rate function of this observable are not known exactly, but Galerkin approximations can easily be found by Fourier series, using the basis functions $\phi_j(x) = e^{j2\pi i x}$ with




Figure 5.3 – Large deviations of the mean area of the Ornstein–Uhlenbeck process. Left: Biasing force F_n at the current state \tilde{x}_n of the process at time n . Right: Evolution of the time average R_t , which gives the rate function according to (5.24).

$j \in \mathbb{Z}$ [326, 340, 410]. This basis is used to compute reference values for the SCGF and the rate function, devoid of Δt errors, by projecting the spectral problem (5.14) in Fourier space and by ensuring that enough basis functions are used. We find in practice that $M = 41$ Fourier modes are sufficient. We use the same Fourier basis for the algorithm, also with $M = 41$ modes, in addition to $T = 2 \times 10^4$ and $T_1 = T_2 = 2 \times 10^3$ for the integration times, as in the first test.

We show in Figure 5.4 the results of a typical simulation for $\gamma = 1$ and $k = 1$. The left plot in this figure shows the evolution in time of the estimated SCGF from the eigenvalue iteration (5.49), while the right plot shows the final estimated eigenfunction h_k for different time steps Δt , compared with the reference Fourier solution. We see that the results for the eigenvalue and the eigenvector significantly depart from the reference solutions for $\Delta t = 0.05$, but converge to them as Δt is decreased, in accordance with Assumption 5.2 and the error estimates obtained in Chapter 4. For the same time step used for the Ornstein–Uhlenbeck process, namely, $\Delta t = 5 \times 10^{-3}$, we see no notable difference between the estimated eigenfunction and the reference values, leading to a precise estimation of the SCGF. The convergence of the eigenvalue here is much faster than for the Ornstein–Uhlenbeck process because the space explored is compact, being limited to $[0, 1]$ with periodic boundary conditions.

As for the Ornstein–Uhlenbeck process, we can repeat these simulations over a range of values for k to obtain the SCGF and the rate function. The left plot of Figure 5.5 shows that the SCGF is in good agreement with the Fourier solution for $\Delta t = 5 \times 10^{-3}$, and so is the rate function estimated by Legendre transform. However, as seen before, the rate function estimated with the time average R_t of the biasing force shows an offset, although smaller this time, which comes from the noisy estimation of F_n . As before, the estimator of the rate function that should be used is the one based on the Legendre transform of the SCGF, with a_k estimated by A_t .

5.5 Conclusion

We have presented a new algorithm for estimating the large deviations functions of time-integrated observables of Markov processes, which characterize the likelihood of their fluctuations in the long-time limit. The algorithm draws on earlier results on stochastic control [58, 5, 28, 57], and works by adaptively estimating the principal eigenvalue and eigenfunction of a spectral problem related to the large deviations problem. The adaptive part consists in modifying the process considered, using feedback and reinforcement learning, so as to reach the so-called driven process, which is known to




Figure 5.4 – Current large deviations of the driven periodic diffusion. Left: Evolution of the estimated SCGF with time for different time steps Δt . Right: Estimated eigenfunction h_k for different time steps Δt compared with the reference Fourier solution.

be optimal for the purpose of estimating large deviations functions using importance sampling [98]. In this sense, the algorithm relates to many adaptive importance sampling methods that have been proposed recently for rare event simulations; see, for instance, [150, 339, 46, 261]. It is also closely related to diffusion Monte Carlo methods [190, 377, 313], which attempt to estimate the ground-state wavefunction of a many-body quantum system by simulating a related stochastic process.

The proposed algorithm can be applied to diffusions, as illustrated here, but also to Markov chains in discrete or continuous time, and works for both equilibrium and nonequilibrium systems, that is, reversible and non-reversible processes, respectively. Moreover, although the test cases presented here are simple, they clearly show that the algorithm has the potential to improve upon other simulation methods, especially cloning methods, since it runs on a single simulation of the process and provides information about how fluctuations are created in time by constructing the driven process in a non-parametric form and with no prior information. A modification of the original cloning algorithm was proposed recently [337] to construct the driven process, but it is based on a different feedback rule that compares two time-dependent histograms, whose estimation is noisy and requires a large number of clones. The results presented here show that a single clone, evolving over a long-enough time, is sufficient. This obviously cuts the computational complexity of estimating large deviations, but also simplifies, as mentioned, the error and convergence analyses of the algorithm.

As with any new proposal, more work is needed to understand the benefits and limitations of the algorithm, to test its applicability to realistic systems, and to benchmark it against other numerical methods. Of particular importance is to derive precise error estimates associated with the space and time discretizations of the spectral elements, large deviations functions, and the driven process. Estimates for the discretization errors in Δt for the SCGF can be found in Chapter 4, but these discretization errors are also present in cloning algorithms when applied to diffusions, so they are not specific to the algorithm presented here. The use of Galerkin discretizations requires further investigations, particularly in the low-noise limit and for processes involving many interacting particles, and we mention that recent work on optimal control problems show that alternative representations can be used, such as neural networks [223, 154, 224], which is an interesting path to follow.




Figure 5.5 – Current large deviations for the driven periodic diffusion. Left: Estimated SCGF compared with the reference Fourier solution. Right: Estimated rate function based on the Legendre transform and on the averaging of the biasing force, compared with the reference Fourier solution.

Acknowledgements

We are grateful to Florian Angeletti for useful discussions in the initial phase of this work and to Hadrien Vroylandt for kindly pointing out to us an error in Eq. (5.51), corrected in this version. The research for this chapter was supported in part by the International Centre for Theoretical Sciences (ICTS) during a visit for participating in the program “large deviations theory in statistical physics: Recent advances and future challenges” (Code: ICTS/Prog-ldt/2017/8).

CHAPTER 6

CONSTRUCTING APPROXIMATE CONTROLS FOR SMALL TEMPERATURE SYSTEMS

The material for this chapter is a work in preparation with E. Vanden-Eijnden and T. Gräfkke [185].

Contents

6.1	Motivation	195
6.2	Optimal control and low temperature limit	196
6.2.1	A direct approach	196
6.2.2	Optimal tilting on path space	197
6.3	Instanton and optimal bias	198
6.3.1	First order control	198
6.3.2	Second order control	200
6.3.3	Consistency with optimal control	201
6.4	Numerical applications	203
6.5	Conclusion	205

Abstract. The computation of free energies is a common issue in statistical physics, in particular for systems at low temperature. A natural technique to compute such high dimensional integrals is to resort to Monte Carlo simulations. However, these techniques generally suffer from a very high variance in the low temperature regime because the expectation is dominated by high values corresponding to rare trajectories, which is the situation of interest in practice. We propose here a general strategy to build approximate controls in order to reduce the variance of such estimators. Our construction builds upon the instanton theory, presented in Section 1.1.3, and can be realized offline. We show that the resulting estimator behaves well in the small temperature regime, and that the approximate control is actually an expansion of the optimal control around the instanton.

6.1 Motivation

This work is concerned with variance reduction for estimating free energy-like functions arising in statistical physics. We will consider low temperature finite time problems [155, 118, 157, 414], which have been introduced in Section 1.1.3 in the Introduction. In such a situation, high variance arises because the observable takes large values along very rare trajectories [156].

We have motivated in Section 1.4 two ways for reducing the variance of naive estimators for this kind of problem. One is to introduce a bias in the dynamics, so that rare trajectories become more likely under the new dynamics [155, 231] – a strategy sometimes referred to as tilting. However, the expectations at hand have to be corrected with a Girsanov weight, and it is actually not obvious

that the additional drift decreases the variance, unless the optimal control is exactly known. For high-dimensional systems, it is hopeless to compute this optimal control to a high degree of accuracy, and poor approximations may deteriorate the quality of the estimator.

Another strategy is to resort to population dynamics [212, 116, 113, 109, 289, 71], which is another instance of importance sampling. The idea here is to run a series of systems in parallel, and to select the ones that realize the rare trajectories. We have used this method to compute cumulant functions in Chapter 4 above, and there are actually various ways to design a selection mechanism, some provably behaving better than others [377, 12, 140]. However, it is a well-known fact that, in high-dimension, the number of replicas needed for performing accurate computations becomes very large.

It is also possible to combine the two approaches, see [337] for an application to the computation of long time large deviations functions. However, for a general system, one cannot compute the exact optimal control, nor run enough replicas in parallel to obtain accurate free energy estimates. We focus here on the construction of approximate optimal controls. Since several equivalent expressions are available for the optimal control, many approximation techniques have been developed, including cross-entropy methods [433], milestoning [231], martingale based techniques [312], model reduction [232] and forward-backward SDE [265], or more recently machine learning based algorithms [223, 154, 224]. However, once an approximation of the optimal control has been estimated, one needs to correct the expectation with the appropriate Girsanov weight. Here as well, it is not guaranteed in general that these techniques lead to a variance reduction.

The goal of this chapter is to provide a simple way to construct approximate controls, and to show that they have a finite variance in the small temperature regime. We rely for this on the theory of instantons [155, 156, 209, 211], which we recall in Section 1.1.3 in Part I. Our technique allows to compute the approximate control offline, contrarily to the previous work [414], which required on the fly update.

Our presentation is organized as follows. Section 6.2 recalls well-known facts about zero variance estimators with an optimal control. Since this optimal control is not accessible in practice, we turn in Section 6.3 to the construction of approximate optimal control that behave well in the small temperature regime. Section 6.4 provides some illustrative applications.

6.2 Optimal control and low temperature limit

6.2.1 A direct approach

In this study, we focus on computing expectations of exponential functionals for which naive estimators typically have a large variance. We first recall the framework of Section 1.1.3 and consider, for a fixed time $T > 0$,

$$A_\varepsilon = \mathbb{E}_x \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right], \quad (6.1)$$

where $(X_t^\varepsilon)_{t \geq 0}$ is solution to the following stochastic differential equation

$$dX_t^\varepsilon = b(X_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma(X_t^\varepsilon) dB_t. \quad (6.2)$$

Here, $(B_t)_{t \geq 0}$ is a d -dimensional Brownian motion, and $b : \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ and $f : \mathbb{R}^d \rightarrow \mathbb{R}$ are smooth functions. In (6.1), \mathbb{E}_x denotes the expectation with respect to all trajectories solving (6.2) and starting at $x \in \mathbb{R}^d$ at $t = 0$. Note that we could also consider time-dependent functions b and σ without additional difficulty. The generator of the dynamics (6.2) reads

$$\mathcal{L} = b \cdot \nabla + \varepsilon \frac{\sigma \sigma^T}{2} : \nabla^2, \quad (6.3)$$

where σ^T stands for the transpose of the matrix σ , \cdot is the scalar product in \mathbb{R}^d . In what follows, we again use the possibly position-dependent diffusion matrix

$$D = \sigma \sigma^T, \quad (6.4)$$

which is assumed to be positive definite everywhere in \mathbb{R}^d .

A motivation for studying (6.1) is the computation of the free energy

$$Z_\varepsilon = \varepsilon \log A_\varepsilon \quad (6.5)$$

in the small temperature regime. It is known by large deviations arguments that $Z_\varepsilon \rightarrow Z_0$ as $\varepsilon \rightarrow 0$, see for instance [193, 211] and the discussion in Section 1.1.3. This is why we are interested in numerically estimating (6.1) for $\varepsilon \ll 1$. In this situation, the expectation in (6.1) is dominated by very large values realized by rare trajectories.

The goal of this chapter is to propose a biasing strategy to reduce the variance of the naive estimator for (6.1) obtained by running independent samples according to (6.2). We will also need to discretize in time the SDE (6.2), but we refer the reader to the vast literature on the subject already discussed in Chapter 4. Before presenting our strategy, we first recall some well-known facts about optimal control for variance reduction of (6.1).

6.2.2 Optimal tilting on path space

We present the biasing framework used throughout. For this, we first consider a modification $(\tilde{X}_t^\varepsilon)_{t \geq 0}$ of the diffusion (6.2) given by

$$d\tilde{X}_t^\varepsilon = b(\tilde{X}_t^\varepsilon) dt + D(\tilde{X}_t^\varepsilon) \nabla g(t, \tilde{X}_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma(\tilde{X}_t^\varepsilon) dB_t, \quad (6.6)$$

where $g : \mathbb{R}_+ \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth function, and D is defined in (6.4). Assuming that the control has a gradient form is natural since, as shown in Proposition 6.2 below, the optimal control is gradient. For this, we need the following integration by part formula.

Proposition 6.1. *Consider the dynamics (6.2) and (6.6). Then (6.1) rewrites*

$$A_\varepsilon = e^{g(0,x)} \mathbb{E}_x \left[\exp \left(\frac{1}{\varepsilon} [f(\tilde{X}_T^\varepsilon) - g(T, \tilde{X}_T^\varepsilon)] + \frac{1}{\varepsilon} \int_0^T \alpha(t, \tilde{X}_t^\varepsilon) dt \right) \right], \quad (6.7)$$

where

$$\alpha(t, x) = \partial_t g(t, x) + \mathcal{L}g(t, x) + \frac{1}{2} |\sigma \nabla g|^2(t, x). \quad (6.8)$$

The proof relies on the Girsanov theorem and the gradient structure of the drift, and closely follows the computations in Section 1.4.2 in the context of long time problems.

Proof. First write the Girsanov formula for the path change of measure [262, 362] between the processes $(X_t^\varepsilon)_{t \geq 0}$ and $(\tilde{X}_t^\varepsilon)_{t \geq 0}$:

$$A_\varepsilon = \mathbb{E}_x \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right] = \mathbb{E}_x \left[e^{\frac{1}{\varepsilon} f(\tilde{X}_T^\varepsilon) - \frac{1}{2\varepsilon} \int_0^T |\sigma \nabla g|^2(t, \tilde{X}_t^\varepsilon) dt - \frac{1}{\sqrt{\varepsilon}} \int_0^T \sigma \nabla g(t, \tilde{X}_t^\varepsilon) dB_t} \right]. \quad (6.9)$$

Write now Itô formula over a trajectory of $(\tilde{X}_t^\varepsilon)_{t \geq 0}$ using the generator (6.3):

$$dg(t, \tilde{X}_t^\varepsilon) = (\partial_t g + \mathcal{L}g + \nabla g \cdot D \nabla g)(t, \tilde{X}_t^\varepsilon) dt + \sqrt{\varepsilon} \sigma \nabla g(t, \tilde{X}_t^\varepsilon) dB_t.$$

Integrating in time and dividing by ε , the above equation becomes

$$-\frac{1}{\sqrt{\varepsilon}} \int_0^T \sigma \nabla g(t, \tilde{X}_t^\varepsilon) dB_t = -\frac{g(T, \tilde{X}_T^\varepsilon) - g(0, \tilde{X}_0)}{\varepsilon} + \frac{1}{\varepsilon} \int_0^T (\partial_t g + \mathcal{L}g + |\sigma \nabla g|^2)(t, \tilde{X}_t^\varepsilon) dt.$$

Inserting the above equation into (6.9) leads to the conclusion. \square

The idea of Proposition 6.1 is that the Girsanov theorem can be rewritten without martingale term upon changing the value of the observable f at final time. A direct consequence is the expression of the optimal bias and the corresponding zero variance estimator.

Proposition 6.2. *Define the function $\psi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ as*

$$\psi_\varepsilon(t, x) = \mathbb{E}_{t,x} \left[e^{\frac{1}{\varepsilon} f(X_T^\varepsilon)} \right], \quad (6.10)$$

and

$$g_\varepsilon(t, x) = \varepsilon \log \psi_\varepsilon(t, x). \quad (6.11)$$

Then, by setting g to g_ε in (6.6), the estimator (6.7) has zero variance, namely

$$A_\varepsilon = \psi_\varepsilon(0, x). \quad (6.12)$$

In (6.11), $\mathbb{E}_{t,x}$ means that the dynamics (6.2) is started at time t from position x . The proof of Proposition 6.2 is a simple consequence of the Feynman–Kac formula.

Proof. We first note, using the Feynman–Kac formula [259, Theorem 21.1], that ψ_ε is solution to the following backward PDE (for simplicity we drop the subscript ε for the rest of the proof):

$$\begin{cases} \partial_t \psi + \mathcal{L}\psi = 0 \\ \psi(T, x) = e^{\frac{1}{\varepsilon} f(x)}, \quad \forall x \in \mathbb{R}^d. \end{cases} \quad (6.13)$$

Defining $g = \varepsilon \log \psi$, we see that g satisfies

$$\begin{aligned} \partial_t g &= \varepsilon \frac{\partial_t \psi}{\psi} = \varepsilon \frac{-\mathcal{L}\psi}{\psi} \\ &= -\varepsilon e^{-g/\varepsilon} \mathcal{L} e^{g/\varepsilon} = \varepsilon \left(-\varepsilon^{-1} b \cdot \nabla g - e^{-g/\varepsilon} \frac{\sigma \sigma^T}{2} : (\nabla(e^{g/\varepsilon} \nabla g)) \right) \\ &= \left(-\mathcal{L}g - \frac{1}{2} |\sigma \nabla g|^2 \right). \end{aligned}$$

Using the terminal condition in (6.13) shows that g is solution to

$$\begin{cases} \partial_t g + \mathcal{L}g + \frac{1}{2} |\sigma \nabla g|^2 = 0 \\ g(T, x) = f(x), \quad \forall x \in \mathbb{R}^d. \end{cases} \quad (6.14)$$

As a result, (6.14) ensures that $\alpha(t, x) = 0$ in Proposition 6.1. The terminal condition then leads to the conclusion. \square

In general, Proposition 6.2 cannot be used as such for numerical applications, since estimating the optimal control is still more difficult than solving the initial problem of estimating (6.1). However, this result serves as a guide to design approximate controls that are easier to compute and behave well in the small ε regime. We explain in the next section how to build such approximate controls from transition path theory, which is the main contribution of this work.

6.3 Instanton and optimal bias

Since the optimal control g_ε given by Proposition 6.2 cannot be computed in general, the goal of this section is to present a method for constructing a good approximation of this function from the transition path theory [155, 156, 209, 211]. In Section 6.3.1, we propose a natural position-independent drift built on the *instanton* of the system, that is a low temperature reaction path. We show in particular that in the small noise limit our estimator has a finite variance, whereas a naive estimator has (or a general approximation of g_ε may have) a variance which increases exponentially with ε^{-1} . We then propose in Section 6.3.2 a second order approximation, which adds an inhomogeneous harmonic force to the control. Finally, in Section 6.3.3 we show that this procedure actually boils down to a Taylor expansion in $\sqrt{\varepsilon}$ of the optimal drift g_ε around the instanton, and show that an approximation of any order can be computed with this strategy.

In this section, we assume for convenience that the noise is additive, that is $\sigma \in \mathbb{R}^{d \times d}$ is constant and so the diffusion matrix (6.4) is constant, although this can be generalized to situations where σ is position-dependent. We may also assume that b , σ and f depend explicitly on time upon minor modifications.

6.3.1 First order control

If the control g_ε is difficult to estimate in practice, we can in general have access to an instanton, or reaction path, which stands for the zero temperature most likely path of fluctuation for the dynamics. This has been explained in Section 1.1.3, and we recall here some elements concerning instantons [211].

Considering (6.1) in the limit $\varepsilon \rightarrow 0$, we know that A_ε obeys the following scaling

$$A_\varepsilon \asymp \exp \left(\varepsilon^{-1} f(\phi_T) - \frac{1}{2} \varepsilon^{-1} \int_0^T \theta_s \cdot D\theta_s \, ds \right) \quad (6.15)$$

where \asymp denotes logarithmic equivalence, and $(\phi_t, \theta_t)_{t \in [0, T]}$ solves the forward-backward system

$$\begin{cases} \dot{\phi}_t = b(\phi_t) + D\theta_t, & \phi_0 = x, \\ \dot{\theta}_t = -(\nabla b)(\phi_t) \cdot \theta_t, & \theta_T = \nabla f(\phi_T). \end{cases} \quad (6.16)$$

The reaction path $(\phi_t)_{t \in [0, T]}$ corresponds to a typical path realizing fluctuations in (6.1). In the small noise limit, we know by the Freidlin–Wentzell theory [193, Section 3] that the trajectories of $(X_t^\varepsilon)_{t \in [0, T]}$ dominating the expectation (6.1) concentrate exponential fast on this path for the uniform convergence norm under mild conditions, and we refer to the discussion in Sections 1.1.3 and 1.3.2 for more details.

When such a reaction path is known, it is natural to consider the biased dynamics

$$d\tilde{X}_t^\varepsilon = b(\tilde{X}_t^\varepsilon) dt + D\theta_t dt + \sqrt{\varepsilon} \sigma dB_t, \quad (6.17)$$

which corresponds to the dynamics (6.6) with the choice

$$g(t, x) = \theta_t \cdot (x - \phi_t), \quad (6.18)$$

in place of the optimal control g_ε defined in (6.11). This dynamics is natural because, in the small ε limit, $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ satisfies the following expansion:

$$\tilde{X}_t^\varepsilon = \phi_t + \sqrt{\varepsilon} \zeta_t + O(\varepsilon), \quad (6.19)$$

where

$$d\zeta_t = \nabla b(\phi_t) \zeta_t dt + \sigma dB_t. \quad (6.20)$$

This can be checked by directly expanding $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ around the instanton $(\phi_t)_{t \in [0, T]}$ in (6.17), and studied rigorously with the tools from [193, Section 2.2]. This means that, in the small temperature regime, $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ follows the path $(\phi_t)_{t \in [0, T]}$ perturbed by an inhomogeneous Ornstein–Uhlenbeck process at scale $\sqrt{\varepsilon}$.

Using such a control, (6.1) can be estimated with Proposition 6.1. The weighting coefficient reads

$$\alpha(t, x) = \frac{|\sigma \theta_t|^2}{2} + \theta_t \cdot (b(x) - \dot{\phi}_t) + \dot{\theta}_t \cdot (x - \phi_t),$$

which may be rewritten thanks to (6.16) as

$$\alpha(t, x) = -\frac{|\sigma \theta_t|^2}{2} + \theta_t \cdot (b(x) - b(\phi_t)) + \dot{\theta}_t \cdot (x - \phi_t). \quad (6.21)$$

We can then consider the small ε scaling of the estimator (6.7) for such a choice of control g . Inserting the expansion (6.19) into (6.21) we obtain

$$\begin{aligned} \varepsilon^{-1} \alpha(t, \tilde{X}_t^\varepsilon) &= -\frac{|\sigma \theta_t|^2}{2\varepsilon} + \varepsilon^{-\frac{1}{2}} \nabla b(\phi_t) \theta_t \cdot \zeta_t + \varepsilon^{-\frac{1}{2}} \dot{\theta}_t \cdot \zeta_t + O(1) \\ &= -\frac{1}{2\varepsilon} |\sigma \theta_t|^2 + O(1), \end{aligned} \quad (6.22)$$

where we used the equations (6.16) to obtain the second line. Finally from (6.16) and (6.18) we have (using the boundary condition at $t = T$ in (6.16))

$$\varepsilon^{-1} g(T, \tilde{X}_T^\varepsilon) = \varepsilon^{-1/2} \theta_T \cdot \zeta_T + O(1) = \varepsilon^{-1/2} \nabla f(\phi_T) \cdot \zeta_T + O(1), \quad (6.23)$$

and

$$\varepsilon^{-1} f(\tilde{X}_T^\varepsilon) = \varepsilon^{-1} f(\phi_T) + \varepsilon^{-1/2} \nabla f(\phi_T) \cdot \zeta_T + O(1), \quad (6.24)$$

so that it holds $\varepsilon^{-1} (f(\tilde{X}_T^\varepsilon) - g(T, \tilde{X}_T^\varepsilon)) = \varepsilon^{-1} f(\phi_T) + O(1)$. Therefore (6.7) becomes, in the small ε regime,

$$A_\varepsilon = \exp \left(\varepsilon^{-1} f(\phi_T) - \frac{1}{2} \varepsilon^{-1} \int_0^T \theta_t \cdot D\theta_t dt \right) \mathbb{E} \left[\exp(O(1)) \right]. \quad (6.25)$$

This expression is consistent with (6.15), and this shows that the estimator has a finite variance in the low noise regime. As a conclusion, the pseudo-potential g defined in (6.18) is a good guess to try to reduce the variance of (6.1). Before showing that g is an approximation of the optimal pseudo-potential g_ε , it is natural to turn to the next order approximation.

6.3.2 Second order control

The method described above can actually be extended using a higher order ansatz. For this, we come back to the dynamics (6.6) for an abstract function g . From the expansion in (6.19), we expect the trajectory $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ to concentrate on the path $(\phi_t)_{t \in [0, T]}$ up to a $\sqrt{\varepsilon}$ -Ornstein–Uhlenbeck process. If we expand ∇g correspondingly, we expect the extra drift to take the form

$$\nabla g(t, \tilde{X}_t^\varepsilon) = \nabla g(t, \phi_t + \tilde{X}_t^\varepsilon - \phi_t) = \nabla g(t, \phi_t) + \nabla^2 g(t, \phi_t) \cdot (\tilde{X}_t^\varepsilon - \phi_t) + O(\varepsilon). \quad (6.26)$$

Moreover, from Section 6.2.2 (in particular Proposition 6.2) we would like g to approximate the solution to the Hamilton–Jacobi equation

$$\partial_t g + \mathcal{L}g + \frac{1}{2}|\sigma \nabla g|^2 = 0, \quad g(T, x) = f(x),$$

as $\varepsilon \rightarrow 0$. Since $\mathcal{L} = b \cdot \nabla + \varepsilon D : \nabla^2 / 2$, we may approximate the latter equation by

$$\partial_t g + b \cdot \nabla g + \frac{1}{2}|\sigma \nabla g|^2 = 0, \quad g(T, x) = f(x). \quad (6.27)$$

One can check that the instanton defined in (6.16) corresponds to the choice $g(t, x) = \theta(t) \cdot (x - \phi_t)$ where (ϕ, θ) is such that g solves (6.27), see [211]. However, in this case, it holds $\nabla^2 g = 0$ so there is no improvement in considering the expansion (6.26). This naturally suggests to consider instead a quadratic ansatz for g

$$g(t, x) = \theta_t \cdot (x - \phi_t) + \frac{1}{2}(x - \phi_t) \cdot K_t(x - \phi_t), \quad (6.28)$$

where $K_t \in \mathbb{R}^{d \times d}$ is a time dependent matrix to be determined. For this, let us compute the equation satisfied by $\nabla^2 g$ from (6.27), which rewrites componentwise

$$\partial_t g + b_k \partial_k g + \frac{1}{2} D_{jk} \partial_j g \partial_k g = 0,$$

where we used Einstein’s notation for summation over repeated indices. Taking the derivative with respect to an index $i \in \{1, \dots, d\}$ leads to

$$\partial_t \partial_i g + \partial_i b_k \partial_k g + b_k \partial_{ik}^2 g + D_{jk} \partial_j g \partial_{ik}^2 g = 0,$$

where we used the symmetry of the matrix $D \in \mathbb{R}^{d \times d}$. Taking then the derivative with respect to another index $l \in \{1, \dots, d\}$ shows that

$$\partial_t \partial_{il}^2 g + \partial_{il}^2 b_k \partial_k g + \partial_i b_k \partial_{kl}^2 g + \partial_l b_k \partial_{ik}^2 g + b_k \partial_{ikl}^3 g + D_{jk} \partial_{jl}^2 g \partial_{ik}^2 g + D_{jk} \partial_j g \partial_{ikl}^3 g = 0. \quad (6.29)$$

The above equation becomes, in vectorial form,

$$\partial_t \nabla^2 g + \nabla^2 b \cdot \nabla g + (\nabla b)^T \nabla^2 g + \nabla^2 g \nabla b + (\nabla^2 g)^T D \nabla^2 g + b \nabla^3 g + (\nabla^3 g)^T D \nabla g = 0, \quad (6.30)$$

where the equation is evaluated at (t, x) . Note that $\nabla^2 b$ is a third order tensor whose contraction with the vector ∇g is defined by (6.29). From (6.6) and (6.26), we want to find a time-dependent evolution equation for $\nabla^2 g(t, \phi_t)$. We therefore compute

$$\begin{aligned} \frac{d}{dt} \nabla^2 g(t, \phi(t)) &= \partial_t \nabla^2 g + \dot{\phi}(t) \cdot \nabla^3 g \\ &= -(\nabla^2 b)^T \nabla g - (\nabla b)^T \nabla^2 g - \nabla^2 g \nabla b - (\nabla^2 g)^T D \nabla^2 g \\ &\quad - b \nabla^3 g - (\nabla^3 g)^T D \nabla g + \dot{\phi} \nabla^3 g, \end{aligned} \quad (6.31)$$

where the right hand side is evaluated at (t, ϕ_t) . Considering the ansatz given by (6.28), we have $\nabla g(t, \phi_t) = \theta_t$, $\nabla^2 g(t, \phi_t) = K_t$ and $\nabla^3 g(t, \phi_t) = 0$, so (6.31) provides the following equation for K :

$$\begin{cases} \partial_t K_t + (\nabla b)^T K_t + K_t \nabla b + (\nabla^2 b)^T \theta_t + K_t^T D K_t = 0, \\ K_T = \nabla^2 f(\phi_T), \end{cases} \quad (6.32)$$

where ∇b and $\nabla^2 b$ are evaluated at ϕ_t . This is an instance of an algebraic Riccati equation [283], which is an interesting feature compared to the more standard instanton result of Section 6.3.1.

Indeed, in addition to simply guiding the dynamics along the path $(\phi_t)_{t \in [0, T]}$ through a position-independent drift like in (6.17), an additional harmonic force drives the system back to the instanton. The coefficients of this harmonic force are given by the matrix $(K_t)_{t \in [0, T]}$ solution to (6.32), which can be thought of as the next order expansion of the instanton in the setting presented in Section 1.1.3. In particular, these coefficients are unsigned, see the applications in Section 6.4. Note that Riccati equations recurrently appear in optimal control theory [257], so it is no surprise to encounter such an equation in our approximation procedure.

Finally, the dynamics $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ corresponding to the second order ansatz for g reads

$$d\tilde{X}_t^\varepsilon = b(\tilde{X}_t^\varepsilon) dt + (D\theta_t + DK_t(\tilde{X}_t^\varepsilon - \phi_t)) dt + \sqrt{\varepsilon}\sigma dB_t. \quad (6.33)$$

Although (6.32) may seem long to compute, we recall that this is an offline computation, and once ϕ_t has been computed, (6.32) is an ordinary differential equation in $\mathbb{R}^{d \times d}$. Next, once K has been estimated, there is small additional cost in running (6.33) compared to (6.2), the overhead being that of computing a linear force with known coefficients. However, since (6.32) is nonlinear, singularities may appear in finite time, and we assume throughout that the various parameters of the model allow to define its solution, see [283]. Finally, since $\tilde{X}_t^\varepsilon - \phi_t$ is of order $\sqrt{\varepsilon}$, the second correction term in (6.33) is also of order $\sqrt{\varepsilon}$.

6.3.3 Consistency with optimal control

From the expansions performed in Sections 6.3.1 and 6.3.2, we can expect to build an approximation of order $M \in \mathbb{N}^*$ through

$$g^M(t, x) = \sum_{k=1}^M T_k(t) \odot (x - \phi_t)^{\otimes k}, \quad (6.34)$$

where $\otimes k$ stands for the k^{th} order tensorization of a d -dimensional vector, \odot the k^{th} order contraction and, for each $k \geq 1$, $(T_k(t))_{t \in [0, T]}$ is a time dependent k^{th} order tensor. A natural question is to describe the limiting object

$$g^\infty(t, x) = \sum_{k=1}^{+\infty} T_k(t) \odot (x - \phi_t)^{\otimes k},$$

provided by the large M limit. Since this approximation solves the zero- ε Hamilton–Jacobi equation (6.27), and $\tilde{X}_t^\varepsilon - \phi_t$ is typically of order $\sqrt{\varepsilon}$ along a tilted trajectory, we expect g^∞ to be the expansion of the optimal drift g_ε defined in (6.11) in powers of $\sqrt{\varepsilon}$. In other words, we can interpret the approximations g defined in Sections 6.3.1 and 6.3.2 as the first terms of an expansion of g_ε in powers of $\sqrt{\varepsilon}$ around the instanton. This idea is made more precise by the following proposition, which concerns the second order expansion, namely $k = 2$.

Proposition 6.3. *Consider the drift g_ε defined in (6.11), and the second order approximation*

$$g(t, x) = \theta_t \cdot (x - \phi_t) + \frac{1}{2}(x - \phi_t) \cdot K_t(x - \phi_t), \quad (6.35)$$

where $(\phi_t, \theta_t)_{t \in [0, T]}$ is defined in (6.16) and $(K_t)_{t \in [0, T]}$ satisfies (6.32). Then, in the small ε limit, for any $t \in [0, T]$ and $x \in \mathbb{R}^d$, it holds

$$g_\varepsilon(t, x) = f(\phi_T) - \int_t^T \theta_s \cdot D\theta_s ds + g(t, x) + O(\varepsilon^{3/2}). \quad (6.36)$$

Proposition 6.3 is a consistency result for the approximations built in Sections 6.3.1 and 6.3.2. Indeed, the first part of (6.36) is simply the action functional appearing in (6.15) truncated over $[t, T]$. Since it does not depend on the position, it does not influence the drift ∇g . With the notation of the expansion (6.34), Proposition 6.3 actually reads $T_1 = \theta$ and $T_2 = K$. This result does not seem standard in the literature [193, 211], and it is an interesting question to find a general structure for the higher order terms of the expansion. In particular, as noted in Section 6.3.2, it is somehow natural for K to solve a Riccati equation since it corresponds to a linearization of the control force [257]. We may wonder if there is a particular structure for the higher order tensors T_k arising in the decomposition (6.34).

The proof displayed below clearly shows that such a higher order expansion can be performed formally, leading to an expression like (6.36) with a smaller remainder. However, proving mathematically

this result (by using techniques as those described in Section 1.1.3) may be difficult, the first task being to find a general expression for the tensors T_k defined in (6.34).

Proof. The idea is to rewrite the Feynman–Kac mode ψ_ε defined in (6.10) with Proposition 6.1 in order to exhibit the leading behaviour. Consider the dynamics $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ defined in (6.6) with g given by (6.35). Adapting Proposition 6.1 starting from any $t \geq 0$ we have

$$\begin{aligned} \psi_\varepsilon(t, x) &= \mathbb{E}_{t, x} \left[e^{\frac{1}{\varepsilon}(f(\tilde{X}_T^\varepsilon) - g(T, \tilde{X}_T^\varepsilon) + g(t, \tilde{X}_t^\varepsilon)) + \frac{1}{\varepsilon} \int_t^T \alpha(s, \tilde{X}_s) ds} \right] \\ &= e^{\frac{g(t, x)}{\varepsilon}} \mathbb{E}_{t, x} \left[e^{\frac{1}{\varepsilon}(f(\tilde{X}_T^\varepsilon) - g(T, \tilde{X}_T^\varepsilon)) + \frac{1}{\varepsilon} \int_t^T \alpha(s, \tilde{X}_s) ds} \right], \end{aligned} \quad (6.37)$$

where the function α is defined in (6.8). We now perform an expansion inside (6.37) like that of Section 6.3.1 but with an additional order of accuracy in ε and starting from an arbitrary position. First, when g is the second order potential (6.35), the process $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ admits the following expansion:

$$\tilde{X}_t^\varepsilon = \phi_t + \sqrt{\varepsilon} \zeta_t + \varepsilon \beta_t + O(\varepsilon^{3/2}), \quad (6.38)$$

where

$$\begin{cases} d\zeta_t = (\nabla b(\phi_t) + K_t) dt + \sigma dB_t, \\ d\beta_t = (\nabla b(\phi_t)\beta_t + \frac{1}{2}\zeta_t \cdot \nabla^2 b(\phi_t)\zeta_t + K_t\beta_t) dt. \end{cases}$$

This follows by expanding $(\tilde{X}_t^\varepsilon)_{t \in [0, T]}$ around the path $(\phi_t)_{t \in [0, T]}$ and identifying the terms of different orders in ε . Note that we will actually not need the precise expression for the processes $(\zeta_t)_{t \in [0, T]}$ and $(\beta_t)_{t \in [0, T]}$ in what follows.

We first consider the terminal terms in (6.37). Using (6.38), we obtain

$$\begin{aligned} f(\tilde{X}_T^\varepsilon) &= f(\phi_T) + \sqrt{\varepsilon} \nabla f(\phi_T) \cdot \zeta_T + \varepsilon \left[\nabla f(\phi_T) \cdot \beta_T + \frac{1}{2} \zeta_T \cdot \nabla^2 f(\phi_T) \zeta_T \right] + O(\varepsilon^{3/2}) \\ g(T, \tilde{X}_T^\varepsilon) &= \sqrt{\varepsilon} \theta_T \cdot \zeta_T + \varepsilon \left[\theta_T \cdot \beta_T + \frac{1}{2} \zeta_T \cdot K_T \zeta_T \right] + O(\varepsilon^{3/2}). \end{aligned}$$

The terminal conditions for θ and K lead to

$$f(\tilde{X}_T^\varepsilon) - g(T, \tilde{X}_T^\varepsilon) = f(\phi_T) + O(\varepsilon^{3/2}). \quad (6.39)$$

It remains to study the integral term in (6.37). For this, we follow the computations leading to (6.22) but going one order further in ε , and still neglecting the term of order ε in the generator \mathcal{L} . This leads to

$$\alpha(t, \tilde{X}_t^\varepsilon) = \partial_t g(t, \tilde{X}_t^\varepsilon) + b(\tilde{X}_t^\varepsilon) \cdot \nabla g(t, \tilde{X}_t^\varepsilon) + \frac{1}{2} |\sigma \theta|^2. \quad (6.40)$$

It then holds (we omit the dependency of ϕ , θ and K on time for conciseness)

$$\begin{aligned} \partial_t g(t, x) &= \dot{\theta} \cdot (x - \phi) - \theta \cdot \dot{\phi} + \frac{1}{2} (x - \phi) \cdot \dot{K} (x - \phi) - (x - \phi) \cdot K \dot{\phi} \\ \nabla g(t, x) &= \theta + K(x - \phi). \end{aligned}$$

As a result, (6.40) reads

$$\begin{aligned} \alpha(t, \tilde{X}_t^\varepsilon) &= \dot{\theta} \cdot (\tilde{X}_t^\varepsilon - \phi) - \theta \dot{\phi} + \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot \dot{K} (\tilde{X}_t^\varepsilon - \phi) - \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot K \dot{\phi} - \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot K^T \dot{\phi} \\ &\quad + b(\tilde{X}_t^\varepsilon) \cdot (\theta + K(\tilde{X}_t^\varepsilon - \phi)) + \frac{1}{2} |\sigma(\theta + K(\tilde{X}_t^\varepsilon - \phi))|^2, \end{aligned}$$

which may be reorganized as follows (using (6.16) for the time derivatives of ϕ and θ):

$$\begin{aligned} \alpha(t, \tilde{X}_t^\varepsilon) &= -|\sigma \theta|^2 + \frac{1}{2} |\sigma \theta|^2 + \frac{1}{2} \theta \cdot DK(\tilde{X}_t^\varepsilon - \phi) + \frac{1}{2} \theta \cdot DK^T(\tilde{X}_t^\varepsilon - \phi) + \frac{1}{2} |\sigma K(\tilde{X}_t^\varepsilon - \phi)|^2 + \theta(b(\tilde{X}_t^\varepsilon) - b(\phi)) \\ &\quad - \theta \cdot \nabla b(\phi)(\tilde{X}_t^\varepsilon - \phi) + \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot \dot{K}(\tilde{X}_t^\varepsilon - \phi) + \frac{1}{2} (b(\tilde{X}_t^\varepsilon) - b(\phi)) \cdot K(\tilde{X}_t^\varepsilon - \phi) \\ &\quad + \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot K^T(b(\tilde{X}_t^\varepsilon) - b(\phi)) - \frac{1}{2} \theta \cdot DK(\tilde{X}_t^\varepsilon - \phi) - \frac{1}{2} \theta \cdot DK^T(\tilde{X}_t^\varepsilon - \phi) \\ &= -\frac{1}{2} |\sigma \theta|^2 + \frac{1}{2} |\sigma K(\tilde{X}_t^\varepsilon - \phi)|^2 + \theta(b(\tilde{X}_t^\varepsilon) - b(\phi)) - \theta \cdot \nabla b(\phi)(\tilde{X}_t^\varepsilon - \phi) \\ &\quad + \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot \dot{K}(\tilde{X}_t^\varepsilon - \phi) + \frac{1}{2} (b(\tilde{X}_t^\varepsilon) - b(\phi)) \cdot K(\tilde{X}_t^\varepsilon - \phi) + \frac{1}{2} (\tilde{X}_t^\varepsilon - \phi) \cdot K^T(b(\tilde{X}_t^\varepsilon) - b(\phi)). \end{aligned}$$

Inserting the expansion (6.38) in the above equation leads to

$$\begin{aligned} \alpha(t, \tilde{X}_t^\varepsilon) = & -\frac{1}{2}|\sigma\theta|^2 + \frac{\varepsilon}{2}|\sigma\zeta_t K|^2 + \sqrt{\varepsilon}\theta \cdot \nabla b(\phi)\zeta_t + \varepsilon\theta \cdot \nabla b(\phi)\beta_t + \frac{\varepsilon}{2}\zeta_t \cdot \theta \nabla^2 b(\phi)\zeta_t \\ & - \theta \cdot \nabla b(\phi)(\sqrt{\varepsilon}\zeta_t + \varepsilon\beta_t) + \frac{\varepsilon}{2}\zeta_t \cdot \dot{K}\zeta_t + \frac{1}{2}\varepsilon\zeta_t \cdot \nabla b(\phi)K\zeta_t + \frac{1}{2}\varepsilon\zeta_t \cdot K^T \nabla b(\phi)\zeta_t. \end{aligned}$$

Now, we may identify the terms of various orders in $\sqrt{\varepsilon}$ in the above equation. At leading order it remains $-|\sigma\theta|^2/2$, which was expected. At order $\sqrt{\varepsilon}$, the terms are

$$\theta \cdot \nabla b(\phi)\zeta_t - \theta \cdot \nabla b(\phi)\zeta_t = 0.$$

We now turn to the terms of order ε , which read

$$\frac{1}{2}\zeta_t \cdot KDK\zeta_t + \theta \cdot \nabla b(\phi)\beta_t - \theta \cdot \nabla b(\phi)\beta_t + \frac{1}{2}\zeta_t \cdot \dot{K}\zeta_t + \frac{1}{2}\zeta_t \cdot \theta \nabla^2 b(\phi)\zeta_t + \frac{1}{2}\zeta_t \cdot K^T \nabla b(\phi)\zeta_t + \frac{1}{2}\zeta_t \cdot \nabla b(\phi)^T K\zeta_t.$$

We see that the terms proportional to β_t cancel, while the quadratic product in ζ_t factors out, so it remains:

$$\dot{K} + KDK + K^T \nabla b(\phi) + \nabla b^T(\phi)K + \theta \nabla^2 b(\phi),$$

which is equal to 0 since K is solution to (6.32). Gathering the results show that (6.40) reads in the small ε limit

$$\alpha(t, \tilde{X}_t^\varepsilon) = -\frac{1}{2}|\sigma\theta_t|^2 + O(\varepsilon^{3/2}).$$

Plugging the above estimate together with (6.39) into (6.37), we obtain

$$\psi_\varepsilon(t, x) = e^{\frac{g(t, x)}{\varepsilon}} \mathbb{E}_{t, x} \left[e^{\frac{1}{\varepsilon} \left(f(\phi_T) - \frac{1}{2} \int_t^T |\sigma\theta_s|^2 ds + O(\varepsilon^{3/2}) \right)} \right]$$

As a result, the optimal control (6.11) admits the following expansion

$$g_\varepsilon(t, x) = g(t, x) + f(\phi_T) - \frac{1}{2} \int_t^T |\sigma\theta_s|^2 ds + O(\varepsilon^{3/2}),$$

in the small ε regime, where g is defined in (6.35). This concludes the proof of Proposition 6.3. \square

6.4 Numerical applications

We start with the simple situation in dimension one where the drift is given by $b = -\nabla V$ with $V(x) = |x|^2/2$ and $\sigma(x) = 1$. In particular, when $f(x) = x$ we know that the optimal drift is actually equal to θ_s . This has two consequences. First, this means that the first order approximation described in Section 6.3.1 actually provides the zero variance estimator of Proposition 6.2. We should therefore obtain an estimator with variance close to zero. Then, the second order term should not be of any help, so the matrix K defined in Section 6.3.2 should be zero. For the numerical simulations, we consider $T = 2$, $\Delta t = 0.005$, $N = 5000$ realizations of the dynamics, and $x = -1$ for the initial condition. The instanton ϕ , its momentum θ and the matrix K are computed offline. In Figure 6.1, we plot ϕ , θ , K together with realizations of the unbiased process, and the processes with first and second order corrections. We note in particular that the matrix K indeed vanishes, showing that the first order control is the optimal one. In Table 6.1, we present the value estimated by the different estimators with the corresponding relative error (empirical standard deviation over empirical average). We see that the naive estimator becomes very ineffective in the low temperature regime, while the other two produces near zero variance estimators. The remaining variance is actually due to the time discretization of the Girsanov weights.

We consider next the double well potential defined by

$$V(x) = \frac{(x^2 - 1)^2}{4}, \tag{6.41}$$

and leave the other parameters unchanged. Since the process starts from the left well, a typical fluctuation leading to high values of (6.1) corresponds to a trajectory crossing to the right well. We plot the various quantities of interesting in Figure 6.2. We observe now that the matrix K is non-trivial,




Figure 6.1 – Instanton and control matrix with examples of realizations for the free dynamics and those with first and second order controls, for $\varepsilon = 0.05$.

Estimator-relative error / ε	Naive estimator	First order	Second order
1	1.127 – 0.83	$1.117 \cdot 10^{-15}$	$1.117 \cdot 10^{-15}$
0.5	1.26 – 1.2	$1.248 \cdot 10^{-15}$	$1.1248 \cdot 10^{-15}$
0.1	2.98 – 9.2	$3.025 \cdot 10^{-15}$	$3.025 \cdot 10^{-15}$
0.005	7.25 – 12.1	$9.15 \cdot 10^{-14}$	$9.15 \cdot 10^{-14}$

Table 6.1 – Estimated values of the free energy (6.1) at various temperatures for the different estimators, together with relative error, in the case of a quadratic potential.

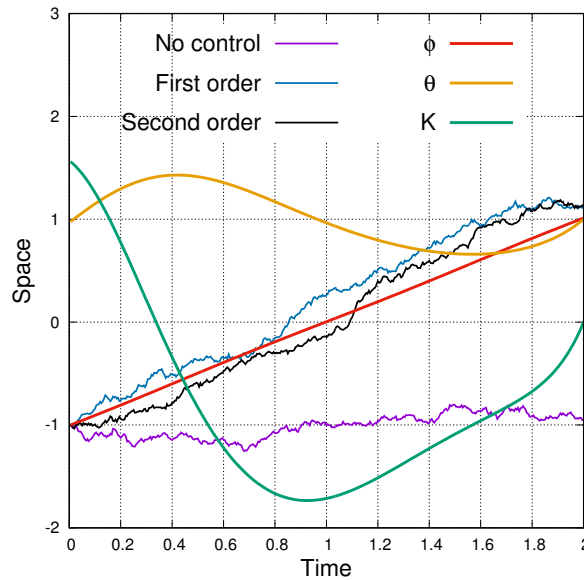


Figure 6.2 – Instanton and control matrix for $\varepsilon = 0.05$.

and even becomes negative. Thus the harmonic force corresponding to the second order correction becomes *repulsive* from the instanton at some moments. Table 6.2 shows that the naive estimator rapidly becomes irrelevant, while the controlled dynamics behave very well at small temperatures. In particular, the second order correction substantially reduces the relative error as $\varepsilon \rightarrow 0$, but deteriorates the quality of the estimator when ε is of order one.

Estimator-relative error / ε	Naive estimator	First order	Second order
1	0.925 – 1.05	0.950 – 1.254	0.745 – 2.62
0.5	0.747 – 2.6	0.72 – 1.52	0.691 – 5.85
0.1	0.0055 – 38	0.231 – 1.25	0.237 – 1.34
0.05	1.2×10^{-7} – 18	0.103 – 1.5	0.102 – 0.51
0.01	10^{-39} – 32	1.96×10^{-4} – 1.26	1.97×10^{-4} – 0.17

Table 6.2 – Estimated values of the free energy (6.1) at various temperatures for the different estimators, together with relative error, in the case of the quartic potential (6.41).

As motivated in Section 1.1, we are also interested in nonequilibrium dynamics. We thus turn to a two-dimensional example, which means in particular that $K \in \mathbb{R}^{2 \times 2}$, and consider for this the potential

$$V(x, y) = \frac{1}{4}(x^2 - 1)^2 + \frac{y^2}{2},$$

and the nonequilibrium drift

$$b(x, y) = -\nabla V(x, y) + C \begin{pmatrix} x \\ y \end{pmatrix} \quad \text{with} \quad C = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}. \quad (6.42)$$

The observable is set to $f(x, y) = x + y$ and the process is started from $(x_0, y_0) = (-1, -1)$. In Figure 6.3, we present the three types of trajectories and the instanton ϕ . Interestingly, the unbiased trajectory remains trapped close to the bottom of one well of V , which is due to the nonequilibrium forcing. The biased trajectory escapes from this well by performing a kind of loop in the direction of the non-gradient drift. Figure 6.4 shows the evolution of the various components of ϕ , θ and K with time. We see that K has non-zero outer diagonal terms, and that some coefficients also become negative at some point. Table 6.3 presents the results at various temperatures, leading to the same conclusion as the one dimensional systems. This proves that our method is suitable for variance reduction of nonequilibrium systems.




Figure 6.3 – Trajectories with the instanton ϕ for $\varepsilon = 10^{-2}$.

6.5 Conclusion

Free energies are difficult to compute in practice due to the high variance of naive estimators in the small temperature regime. In particular, the expectation (6.1) is typically dominated by very large values taken over rare trajectories. One way to improve the standard Monte Carlo estimator based




Figure 6.4 – Study of the instanton’s coefficients for the two-dimensional dynamics defined by the drift (6.42). Left: values of the instanton and momentum. Right: values of the coefficients of the matrix K (by symmetry it holds $K_{21} = K_{12}$).

Estimator-relative error / ε	Naive estimator	First order	Second order
0.1	307 – 20	1077 – 0.54	1058 – 0.61
0.05	9.18×10^4 – 50	1.186×10^6 – 0.66	1.187×10^6 – 0.50
0.025	3.15×10^6 – 48	1.475×10^{12} – 0.61	1.474×10^{12} – 0.43
0.01	4.71×10^{11} – 50	2.81×10^{30} – 0.47	2.82×10^{30} – 0.20
0.005	2.50×10^{16} – 70	8.62×10^{60} – 0.44	8.61×10^{60} – 0.15

Table 6.3 – Estimated values of the free energy (6.1) at various temperatures for the different estimators, together with relative error, for the two dimensional drift (6.42).

on drawing independent trajectories is to estimate the optimal bias which leads to a zero variance estimator. However, one is not ensured that any such estimator of the optimal control indeed reduces the variance in practice.

We have shown in this chapter that it is possible, thanks to the transition path theory, to construct approximate controls that behave well in the small temperature regime. Since this is the regime we are interested in, this is a nice feature of the approximation we build. Applications to (possibly nonequilibrium) low dimensional systems prove the capacity of the method to reduce the variance in this regime, even for nonequilibrium systems. The second order approximation provides better results when the temperature becomes smaller and smaller. Interestingly, we show formally that the polynomial terms of our approximate control correspond to an expansion of the optimal control around the instanton. Providing a mathematical proof of these computations is an interesting challenge.

Part IV

Coulomb gases and random matrices

CHAPTER 7

SAMPLING COULOMB GASES WITH HAMILTONIAN MONTE CARLO

The material for this chapter has been published in the Journal of Statistical Physics [82].

Contents

7.1	Presentation	209
7.1.1	Boltzmann–Gibbs measures	210
7.1.2	Static energy and equilibrium measures	211
7.1.3	Two remarkable gases from random matrix theory	211
7.2	Simulating log-gases and Coulomb gases	213
7.2.1	Standard sampling methods	214
7.2.2	Hybrid Monte Carlo algorithm	215
7.3	Numerical experiments on remarkable models	218
7.3.1	Case study: 1D	218
7.3.2	Case study: 2D	218
7.3.3	Case study: 3D	218

Abstract. Coulomb and log-gases are exchangeable singular Boltzmann–Gibbs measures appearing in mathematical physics at many places. They are also related to random matrix theory, as motivated in Chapter 1.1.4. We explore experimentally an efficient numerical method for simulating such gases. It is an instance of the Hybrid or Hamiltonian Monte Carlo algorithm, in other words a Metropolis–Hastings algorithm with proposals produced by a kinetic or underdamped Langevin dynamics. This algorithm has excellent numerical behavior despite the singular interaction, in particular when the number of particles gets large. It is more efficient than the well known overdamped version previously used for such problems, and allows new numerical explorations. It suggests for instance to conjecture a universality of the Gumbel fluctuation at the edge of beta Ginibre ensembles for all beta.

7.1 Presentation

We explore the numerical simulation of Coulomb gases and log-gases by mean of Hybrid or Hamiltonian Monte Carlo algorithms (HMC) [142, 245]. Such algorithms consist basically in using a discretization of the (underdamped) Langevin dynamics, presented in Section 1.1, to produce proposals for Metropolis–Hastings algorithms. This can be viewed as a way to add momentum to a Monte Carlo interacting particle system. The basic outcome of this exploratory work is that HMC algorithms have remarkably good numerical behavior for such gases despite the singularity of the interactions. Such algorithms scale well with the dimension of the system, see [37, 63]. They are therefore more efficient than the

tamed overdamped version already explored in the literature for instance in [311]. In this paper, we benchmark the capability of the algorithm to reproduce known results efficiently, and we make it ready to explore new conjectures.

Another advantage of this approach is that it can be adapted to take into account a sub-manifold constraint [302]. We use this technique to simulate constrained gases in Chapter 8. We will illustrate this technique to study numerically the behaviour of the spectrum of random matrices with prescribed trace or determinant, which is difficult to achieve by direct sampling of matrices.

For the sake of completeness, we should mention that there are remarkable alternative simulation algorithms which are not based on a diffusion process, such as the ones based on piecewise deterministic Markov processes (PDMP), see for instance [260] and [415].

7.1.1 Boltzmann–Gibbs measures

We are interested in interacting particle systems subject to an external field and experiencing singular pair interactions. In order to encompass Coulomb gases as well as log-gases from random matrix theory (see the motivation in Section 1.1.4), we introduce a vector subspace \mathcal{X} of dimension d of \mathbb{R}^n , with $n \geq 2$ and $n \geq d \geq 1$. The particles belong to \mathcal{X} , and \mathbb{R}^n is understood as a physical ambient space. We equip \mathcal{X} with the trace of the Lebesgue measure of \mathbb{R}^n , denoted by dx . The external field and the pair interaction are respectively denoted by $V : \mathcal{X} \mapsto \mathbb{R}$ and $K : \mathcal{X} \mapsto (-\infty, +\infty]$, and belong to $C^2(\mathcal{X})$ functions, with $K(x) < \infty$ for all $x \neq 0$. For any $N \geq 2$, we consider the probability measure P_N on $\mathcal{X}^N = \mathcal{X} \times \cdots \times \mathcal{X}$ defined by

$$P_N(dx) = \frac{e^{-\beta_N H_N(x_1, \dots, x_N)}}{Z_N} dx_1 \cdots dx_N, \quad (7.1)$$

where $\beta_N > 0$ is a parameter,

$$Z_N = \int_{\mathcal{X}^N} e^{-\beta_N H_N(x_1, \dots, x_N)} dx_1 \cdots dx_N$$

is the normalizing factor, and

$$H_N(x_1, \dots, x_N) = \frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{2N^2} \sum_{i \neq j} K(x_i - x_j)$$

is usually called energy or Hamiltonian of the system¹. We assume that β_N , V , and K are chosen in such a way that $Z_N < \infty$ for any N . The law P_N is invariant by permutation of the coordinates x_1, \dots, x_N (exchangeable), and H_N depends only on the empirical measure

$$\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}.$$

Therefore P_N is also the law of a random empirical measure encoding a cloud of indistinguishable particles x_1, \dots, x_N . We emphasize that the particles live on the space $\mathcal{X}^N = \mathcal{X} \times \cdots \times \mathcal{X}$ of dimension dN . The parameter n serves as the physical dimension of the ambient space, for the Coulomb gas setting described next.

For any $m \geq 1$ and $x \in \mathbb{R}^m$, we denote by $|x| = \sqrt{x_1^2 + \cdots + x_m^2}$ the Euclidean norm of x . This matches the absolute value when $m = 1$ and the modulus when $m = 2$, $\mathbb{R}^2 \equiv \mathbb{C}$.

7.1.1.1 Coulomb gases

The notion of Coulomb gas is based on elementary electrostatics. Here the vector subspace \mathcal{X} is interpreted as a conductor. It corresponds to taking $K = g$ where g is the Coulomb kernel or Green function in the physical space \mathbb{R}^n . More precisely, recall that the Green function g in \mathbb{R}^n , $n \geq 2$, is defined for all $x \in \mathbb{R}^n$, $x \neq 0$, by

$$g(x) = \begin{cases} \log \frac{1}{|x|} & \text{if } n = 2, \\ \frac{1}{|x|^{n-2}} & \text{if } n \geq 3. \end{cases}$$

¹Note however that this is *not* the Hamiltonian of the Langevin dynamics described in Section 1.1, since there is no momentum here.

This function is the fundamental solution of the Poisson equation, namely, denoting by Δ the Laplace operator in \mathbb{R}^n and by δ_0 the Dirac mass at 0, we have, in the sense of distributions,

$$-\Delta g = c\delta_0, \quad \text{with} \quad c = \begin{cases} 2\pi & \text{if } n = 2, \\ (n-2)|\mathbb{S}^{n-1}| = \frac{n(n-2)\pi^{n/2}}{\Gamma(1+n/2)} & \text{if } n \geq 3. \end{cases}$$

The physical interpretation in terms of electrostatics is as follows: $H_N(x_1, \dots, x_N)$ is the electrostatic energy of a configuration of N electrons in \mathbb{R}^n lying on \mathcal{X} at positions x_1, \dots, x_N , in an external field given by the potential V . The Green function or Coulomb kernel g expresses the Coulomb repulsion which is a two body singular interaction. The probability measure P_N can be seen as a Boltzmann–Gibbs measure, β_N playing the role of an inverse temperature. The probability measure P_N is known as a Coulomb gas or as a one-component plasma, see for instance [387] and references therein.

7.1.1.2 Log-gases

A log-gas corresponds to choosing $d = n$ and a logarithmic interaction K whatever the value of n is, namely

$$K(x) = \log \frac{1}{|x|} = -\frac{1}{2} \log(x_1^2 + \dots + x_d^2), \quad x \in S.$$

Coulomb gases and log-gases coincide when $d = n = 2$. In dimension $d = n \geq 3$, log-gases are natural and classical objects of approximation theory and can be seen as limiting Riesz potentials, namely $\lim_{\alpha \rightarrow 0} \frac{1}{\alpha} (|x|^{-\alpha} - 1)$, see for instance [391, 390, 387].

7.1.2 Static energy and equilibrium measures

Under natural assumptions over V and K , typically when $\beta_N \gg N$ and V beats K at infinity, it is well known, see for instance [84, 386] and references therein, that P_N almost surely, the empirical measure

$$\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$$

tends as $N \rightarrow \infty$ to a non random probability measure, the equilibrium measure

$$\mu_\star = \arg \inf \mathcal{E}.$$

The measure μ_\star is the unique minimizer of the strictly convex lower semi-continuous “energy” \mathcal{E} defined by

$$\mu \mapsto \mathcal{E}(\mu) = \int_{\mathcal{X}} V d\mu + \iint_{\mathcal{X} \times \mathcal{X}} K(x-y) \mu(dx) \mu(dy).$$

This concentration phenomenon is described in more details in Section 1.3.3. When $K = g$ is the Coulomb kernel, the quantity $\mathcal{E}(\mu)$ is the electrostatic energy of the distribution of charges μ , formed by the sum of the electrostatic potential coming from the external electric field V with the Coulomb self repulsion by mean of the Coulomb kernel g . Note that $\mathcal{E}(\mu) = \infty$ if μ has a Dirac mass due to the singularity of g . An Euler–Lagrange variational analysis reveals that when $\mathcal{X} = \mathbb{R}^d$ and V is smooth, convex, and grows faster than g at infinity then the equilibrium probability measure μ_\star is compactly supported and has density proportional to ΔV , see [84] and references therein. Table 7.1 gives examples of equilibrium measures in this Coulomb setting. We refer to [284, 237, 380, 386, 387] for old and new potential theory from this analytic point of view. Moreover, quite a few equilibrium measures are known for log-gases beyond Coulomb gases, see for instance [87].

Actually it can be shown that essentially if $\beta_N \gg N$ and V beats g at infinity then under $(P_N)_N$ the sequence of random empirical measures $(\mu_N)_N$ satisfies a large deviation principle with speed β_N and good rate function \mathcal{E} , see [84, 196, 32]. This will be used in Chapter 8. Concentration of measure inequalities are also available, see [85] and references therein.

7.1.3 Two remarkable gases from random matrix theory

Let us give a couple of famous gases from random matrix theory that will serve as benchmark for our algorithm. They correspond to $n = 2$ because the Lebesgue measure on a matrix translates via the

d	\mathcal{X}	n	V	μ_*	Nickname
1	\mathbb{R}	2	$\infty \mathbf{1}_{\text{interval}^c}$	arcsine	
1	\mathbb{R}	2	x^2	semicircle	GUE
2	\mathbb{R}^2	2	$ x ^2$	uniform on a disc	Ginibre
$d \geq 3$	\mathbb{R}^d	d	$ x ^2$	uniform on a ball	
$d \geq 3$	\mathbb{R}^d	d	radial	radial in a ring	

Table 7.1 – Examples of equilibrium measures for Coulomb gases, see [380, 84].

Jacobian of the change of variable to a Vandermonde determinant on the eigenvalues, giving rise to the two-dimensional Coulomb kernel inside the exponential via the identity

$$\prod_{i < j} |x_i - x_j| = \exp \left(\sum_{i < j} \log |x_i - x_j| \right).$$

Hence the name “log-gases”. A good reference on this subject is [188] and we refer to [162, 145, 188, 189, 254] for more examples of Coulomb gases related to random matrix models. Coulomb gases remain interesting in any dimension n beyond random matrices, see [386, 387].

7.1.3.1 Beta-Hermite model

This model corresponds to

$$d = 1, \quad n = 2, \quad \mathcal{X} = \mathbb{R}, \quad V(x) = \frac{x^2}{2\beta}, \quad K(x) = -\log |\cdot|, \quad \beta_N = N^2\beta, \quad \beta \in (0, \infty).$$

This means that the particles evolve on the line \mathbb{R} with Coulomb interactions given by the Coulomb kernel in \mathbb{R}^2 . For $\beta = 2$, it becomes the famous Gaussian Unitary Ensemble (GUE), which is the distribution of the eigenvalues of random $N \times N$ Hermitian matrices distributed according to the Gaussian probability measure with density proportional to $H \mapsto e^{-N\text{Tr}(H^2)}$. Beyond the case $\beta = 2$, the cases $\beta = 1$ and $\beta = 4$ correspond respectively to Gaussian random matrices with real and quaternionic entries. Following [145], for all $\beta \in (0, \infty)$, the measure P_N is also the distribution of the eigenvalues of special random $N \times N$ Hermitian tridiagonal matrices with independent but non identically distributed entries. Back to the case $\beta = 2$, the law P_N writes

$$(x_1, \dots, x_N) \in \mathbb{R}^N \mapsto e^{-\frac{N}{2} \sum_{i=1}^N x_i^2} \prod_{i < j} (x_i - x_j)^2. \quad (7.2)$$

In this case, the Coulomb gas P_N has a determinantal structure, making it integrable or exactly solvable for any $N \geq 2$, see [327, 188]. This provides in particular a formula for the density of the mean empirical spectral distribution $\mathbb{E}\mu_N$ under P_N , namely

$$x \in \mathbb{R} \mapsto \frac{e^{-\frac{N}{2}x^2}}{\sqrt{2\pi N}} \sum_{\ell=0}^{N-1} H_\ell^2(\sqrt{N}x), \quad (7.3)$$

where $(H_\ell)_{\ell \geq 0}$ are the Hermite polynomials which are the orthonormal polynomials for the standard Gaussian distribution $\mathcal{N}(0, 1)$. The equilibrium measure μ_* in this case is the Wigner semicircle distribution with the following density with respect to the Lebesgue measure:

$$x \in \mathbb{R} \mapsto \frac{\sqrt{4 - x^2}}{2\pi} \mathbf{1}_{x \in [-2, 2]}. \quad (7.4)$$

A plot of μ_* and $\mathbb{E}\mu_N$ is provided in Figure 7.1, together with our simulations. We refer to [290] for a direct proof of convergence of (7.3) to (7.4) as $N \rightarrow \infty$. Beyond the case $\beta = 2$, the equilibrium measure μ_* is still a Wigner semicircle distribution, scaled by β , supported by the interval $[-\beta, \beta]$, but up to our knowledge we do not have a formula for the mean empirical spectral distribution $\mathbb{E}\mu_N$, except when β is an even integer, see [145].

7.1.3.2 Beta-Ginibre model

This model corresponds to

$$d = 2, \quad n = 2, \quad \mathcal{X} = \mathbb{R}^2, \quad V(x) = \frac{|x|^2}{\beta}, \quad K(x) = -\log|x|, \quad \beta_N = N^2\beta, \quad \beta \in (0, \infty).$$

In this case, the particles move in \mathbb{R}^2 with a Coulomb repulsion of dimension 2 – it is therefore a Coulomb gas. As for the GUE, the law P_N can be written as

$$(x_1, \dots, x_N) \in (\mathbb{R}^2)^N \mapsto e^{-N \sum_{i=1}^N |x_i|^2} \prod_{i < j} |x_i - x_j|^\beta. \quad (7.5)$$

When $\beta = m$ for an even integer $m \in \{2, 4, \dots\}$, the law of this gas matches the Laughlin wavefunction modeling the fractional quantum Hall effect (FQHE), see for instance [175].

For $\beta = 2$, this gas, known as the complex Ginibre Ensemble, matches the distribution of the eigenvalues of random $N \times N$ complex matrices distributed according to the Gaussian probability measure with density proportional to $M \mapsto e^{-N \operatorname{Tr}(MM^*)}$ where $M^* = \overline{M}^\top$. In this case P_N has a determinantal structure, see [327, 188]. This provides a formula for the density of the mean empirical spectral distribution $\mathbb{E}\mu_N$ under P_N , namely

$$x \in \mathbb{R}^2 \mapsto \frac{e^{-N|x|^2}}{\pi} \sum_{\ell=0}^{N-1} \frac{|\sqrt{N}x|^{2\ell}}{\ell!}, \quad (7.6)$$

which is the analogue of (7.3) for the Gaussian Unitary Ensemble. Moreover, if Y_1, \dots, Y_N are independent and identically distributed Poisson random variables of mean $|x|^2$ for some $x \in \mathbb{R}^2$, then (7.6) writes

$$x \in \mathbb{R}^2 \mapsto \frac{1}{\pi} \mathbb{P} \left(\frac{Y_1 + \dots + Y_N}{N} < 1 \right).$$

As $N \rightarrow \infty$, by the law of large numbers, it converges to $1/\pi$ if $|x| < 1$ and to 0 if $|x| > 1$, while by the central limit theorem it converges to $1/(2\pi)$ if $|x| = 1$. It follows that $\mathbb{E}\mu_N$ converges weakly as $N \rightarrow \infty$ to the uniform distribution on the disk, with density

$$x \in \mathbb{R}^2 \mapsto \frac{\mathbf{1}_{|x| < 1}}{\pi}, \quad (7.7)$$

which is the equilibrium measure μ_* . When N is finite, the numerical evaluation of (7.6) is better done by mean of the Gamma law. Namely, by induction and integration by parts, (7.6) writes

$$x \in \mathbb{R}^2 \mapsto \frac{1}{\pi(N-1)!} \int_{N|x|^2}^{\infty} u^{N-1} e^{-u} du = \frac{\Gamma(N, N|x|^2)}{\pi},$$

where Γ is the normalized incomplete Gamma function and where we used the identity

$$e^{-r} \sum_{\ell=0}^{N-1} \frac{r^\ell}{\ell!} = \frac{1}{(N-1)!} \int_r^{\infty} u^{N-1} e^{-u} du.$$

Note that $t \mapsto 1 - \Gamma(N, t)$ is the cumulative distribution function of the Gamma distribution with shape parameter N and scale parameter 1. Figure 7.4 illustrates the difference between the limiting distribution (7.7) and the mean empirical spectral distribution (7.6) for a finite N . Beyond the case $\beta = 2$, we no longer have a formula for the density of $\mathbb{E}\mu_N$, but a simple scaling argument reveals that the equilibrium measure μ_* is in this case the uniform distribution on the centered disk of radius $\sqrt{\beta/2}$.

7.2 Simulating log-gases and Coulomb gases

Regarding simulation of log-gases or Coulomb gases such as (7.1), it is natural to use the random matrix models when they are available. There exist also methods specific to determinantal processes which cover the log-gases of random matrix theory with $\beta = 2$, see [246, 382, 347, 111, 27, 285, 226]. Beyond these specially structured cases, a great variety of methods are available for simulating Boltzmann–Gibbs measures, such as overdamped Langevin diffusion algorithm, Metropolis–Hastings

algorithm, Metropolis adjusted Langevin algorithm (MALA), and kinetic versions called Hybrid or Hamiltonian Monte Carlo (HMC) which are based on a kinetic (or underdamped) Langevin diffusion, see for instance [73, 304]. Other possibilities exist, such as Nosé–Hoover dynamics [255] or piecewise deterministic Markov processes [64].

Two difficulties arise when sampling measures as (7.1). First, the Hamiltonian H_N involves all couples, so the computation of forces and energy scales quadratically with the number of particles. A natural way to circumvent this numerical problem is to use clusterization procedures such as the “fast multipole methods”, see for instance [243]. A second difficult feature of such a Hamiltonian is the singularity of the interacting function K , which typically results in numerical instability. A standard stabilization procedure is to “tame” the dynamics [247, 74], which is the strategy adopted in [311]. However, this smoothing of the force induces a supplementary bias in the invariant measure, as shown in [74] for regular Hamiltonians. This requires using small time steps, hence long computations. In the present note, we explore for the first time the usage of HMC for general Coulomb gases in the context of random matrices, in the spirit of [393], the difficulty being the singularity of the interaction. This method has the advantage of sampling the exact invariant measure (7.1), while allowing to choose large time steps, which reduces the overall computational cost [176].

In Section 7.2.1, we review standard methods for sampling measures of the form $e^{-\beta_N H_N}$, before presenting in detail the HMC algorithm in Section 7.2.2.

7.2.1 Standard sampling methods

To simplify and from now on, we suppose the support set \mathcal{X} in (7.1) to be \mathbb{R}^d . We introduce the methods based on the overdamped Langevin dynamics (already described in Section 1.1). To sample approximately (7.1), the idea is to exploit the fact that P_N in (7.1) is the reversible invariant probability measure of the Markov diffusion process $(X_t)_{t \geq 0}$ solution to the stochastic differential equation:

$$dX_t = -\alpha_N \nabla H_N(X_t) dt + \sqrt{2 \frac{\alpha_N}{\beta_N}} dB_t, \quad (7.8)$$

or in other words

$$X_t = X_0 - \alpha_N \int_0^t \nabla H_N(X_s) ds + \sqrt{2 \frac{\alpha_N}{\beta_N}} B_t,$$

where $(B_t)_{t \geq 0}$ is a standard Brownian motion on \mathcal{X}^N and $\alpha_N > 0$ is an arbitrary time scaling parameter (for instance $\alpha_N = 1$ or $\alpha_N = \beta_N$). The infinitesimal generator associated with (7.8) is

$$Lf = \frac{\alpha_N}{\beta_N} \Delta f - \alpha_N \nabla H_N \cdot \nabla f.$$

The difficulty in solving (7.8) lies in the fact that the energy H_N involves a singular interaction K , which may lead the process to explode. Actually, under certain conditions on β_N and V , the equation (7.8) is well posed, the process $(X_t)_{t \geq 0}$ is well defined, and

$$X_t \xrightarrow[t \rightarrow \infty]{\text{Law}} P_N,$$

for all non-degenerate initial condition X_0 . See for instance [9, 169, 88] for the case of Beta-Hermite case known as the Dyson Ornstein–Uhlenbeck process, and [51] for the Beta-Ginibre case. We do not discuss these delicate aspects in this note. A convergence in Cesàro mean is provided by the ergodic theorem for additive functionals,

$$\frac{1}{t} \int_0^t \delta_{X_s} ds \xrightarrow[t \rightarrow \infty]{\text{weak}} P_N$$

almost surely or, for any test function $f \in L^1(P_N)$,

$$\frac{1}{t} \int_0^t f(X_s) ds \xrightarrow[t \rightarrow \infty]{} \int_{\mathcal{X}} f dP_N,$$

almost surely. It is also possible to accelerate the convergence by adding a divergence free term in the dynamics (7.8), see for instance [146, 300] and references therein. This modification keeps the same invariant distribution but produces a non-reversible dynamics.

This method of simulation is referred to as an “unadjusted Langevin algorithm”, a terminology which will be clarified later on. In practice, one cannot simulate the continuous stochastic process $(X_t)_{t \geq 0}$ solution to (7.8), and resorts to a numerical integration with a finite time step Δt . A typical choice is the Euler–Maruyama scheme [271, 332], which reads

$$x_{k+1} = x_k - \nabla H_N(x_k) \alpha_N \Delta t + \sqrt{2 \frac{\alpha_N}{\beta_N} \Delta t} G_k, \quad (7.9)$$

where (G_k) is a family of independent and identically distributed standard Gaussian variables, and x_k is an approximation of $X_{k\Delta t}$. Note that α_N and Δt play the same role here. However, because of the singularity of H_N , this sampling scheme leads to important biases in practice, and (7.9) may even lack an invariant measure [324, Section 6]. One way to stabilize the dynamics is to use a tamed version of (7.9), which typically takes the following form:

$$x_{k+1} = x_k - \frac{\nabla H_N(x_k) \alpha_N \Delta t}{1 + |\nabla H_N(x_k)| \alpha_N \Delta t} + \sqrt{2 \frac{\alpha_N}{\beta_N} \Delta t} G_k. \quad (7.10)$$

This strategy is used in [311] but, as noted by the authors, the number of time steps needed to run a trajectory of fixed time T as N increases scales as $\Delta t \sim N^{-2}$, which makes the study of large systems difficult.

Another strategy is to add a selection step at each iteration. This is the idea of the Metropolis Adjusted (overdamped) Langevin Algorithm (MALA) [373], which prevents irrelevant moves with a Metropolis step. One can also view the MALA algorithm as a Metropolis algorithm in which the proposal is produced by using a one step discretization of the Langevin dynamics (7.8). Let us make this precise; more details can be found *e.g.* in [373, 370].

Algorithm 7.1 (Metropolis Adjusted (overdamped) Langevin Algorithm – MALA). *Let $Q_{\Delta t}$ be the Gaussian transition kernel associated to the Markov chain of the Euler discretization (7.9) of the dynamics (7.8). For each step k ,*

1. Draw a proposal \tilde{x}_{k+1} according to the kernel $Q_{\Delta t}(x_k, \cdot)$;
2. Compute the probability

$$p_k = 1 \wedge \frac{Q_{\Delta t}(\tilde{x}_{k+1}, x_k) e^{-\beta_N H_N(\tilde{x}_{k+1})}}{Q_{\Delta t}(x_k, \tilde{x}_{k+1}) e^{-\beta_N H_N(x_k)}}; \quad (7.11)$$

3. Set

$$x_{k+1} = \begin{cases} \tilde{x}_{k+1} & \text{with probability } p_k, \\ x_k & \text{with probability } 1 - p_k. \end{cases}$$

Note that the “reversed” kernel $Q_{\Delta t}(\cdot, x)$ is Gaussian only if H_N is a quadratic form. Note also that if the proposal kernel K is symmetric in the sense that $Q_{\Delta t}(x, y) = Q_{\Delta t}(y, x)$ for all x, y then it disappears in (7.11), and it turns out that this is the case for the Hybrid Monte Carlo algorithm described next (up to momentum reversal)!

A natural issue with these algorithms is the choice of Δt : if it is too large, an important fraction of the proposed moves will be rejected, hence poor convergence properties; conversely, if Δt is too small, many steps will be accepted but the physical elapsed time will be small, hence a large variance for a fixed number of iterations. Algorithm 7.1 actually has a nice scaling of the optimal time step Δt with the dimension of the system. Indeed, it can be shown that it scales as $\Delta t \sim N^{-\frac{1}{3}}$, at least for product measures (see [371] and references therein). Although this algorithm is already efficient, we propose to use a kinetic version with further advantages.

7.2.2 Hybrid Monte Carlo algorithm

Hybrid Monte Carlo is built on Algorithm 7.1, but using a kinetic version of (7.8). For this, a momentum variable is introduced so as to improve the exploration of the space, as discussed in the Introduction. Namely, let $U_N : \mathbb{R}^{dN} \rightarrow \mathbb{R}$ be smooth and such that $e^{-\beta_N U_N}$ is Lebesgue integrable. Let $(X_t, Y_t)_{t \geq 0}$ be the diffusion process on $\mathbb{R}^{dN} \times \mathbb{R}^{dN}$ solution to the stochastic differential equation

$$\begin{cases} dX_t = \alpha_N \nabla U_N(Y_t) dt, \\ dY_t = -\alpha_N \nabla H_N(X_t) dt - \gamma_N \alpha_N \nabla U_N(Y_t) dt + \sqrt{2 \frac{\gamma_N \alpha_N}{\beta_N}} dB_t, \end{cases} \quad (7.12)$$

where $(B_t)_{t \geq 0}$ is a dN -dimensional Brownian motion, and $\gamma_N > 0$ is an arbitrary parameter which plays the role of a friction, and which may depend a priori on N and $(X_t)_{t \geq 0}$, even if we do not use this possibility here². In addition, H_N and β_N are as in (7.1), while U_N plays the role of a generalized kinetic energy [393]. This dynamics admits the following generator:

$$Lf = \underbrace{-\alpha_N \nabla H_N(x) \cdot \nabla_y f + \alpha_N \nabla U_N(y) \cdot \nabla_x f}_{\mathcal{L}_{\text{ham}}} + \underbrace{\frac{\gamma_N \alpha_N}{\beta_N} \Delta_y f - \gamma_N \alpha_N \nabla U_N(y) \cdot \nabla_y f}_{\mathcal{L}_{\text{FD}}}, \quad (7.13)$$

where \mathcal{L}_{ham} is known as the Hamiltonian part while \mathcal{L}_{FD} is called the fluctuation-dissipation part, see the discussion in Section 1.1. The dynamics leaves invariant the product Boltzmann–Gibbs measure

$$\rho_N = P_N \otimes \omega_N \quad \text{where} \quad \omega_N(dy) = \frac{e^{-\beta_N U_N(y)}}{Z'_N} dy,$$

see for instance [393] and the discussion in Section 1.1. In other words

$$\rho_N(dx, dy) = \frac{e^{-\beta_N \tilde{H}_N(x, y)}}{Z_N Z'_N} dx dy \quad \text{with} \quad \tilde{H}_N(x, y) = H_N(x) + U_N(y). \quad (7.14)$$

As for the overdamped dynamics, the ergodic theorem for additive functionals gives

$$\frac{1}{t} \int_0^t \delta_{(X_s, Y_s)} ds \xrightarrow[t \rightarrow \infty]{\text{weak}} \rho, \quad \text{almost surely.}$$

When $U_N(y) = \frac{1}{2}|y|^2$ then $Y_t = dX_t/dt$, and in this case X_t and Y_t can be interpreted respectively as the *position* and the *velocity* of a system of N points in \mathcal{X} at time t . In this case we say that U_N is the *kinetic energy*. For simplicity, we specialize in what follows to this “physical” or “kinetic” case and refer to [393] for more possibilities.

Remark 7.2 (Stability of singular Langevin dynamics). *Studying theoretically the long time behaviour of (7.12) is made difficult by the singularity of the interaction kernel K . We have seen in Chapter 3 that this can be addressed through Lyapunov function techniques, and a natural choice for the Lyapunov function is*

$$W(x, y) = \exp\left(aH_N(x) + a\frac{|y|^2}{2} + \varepsilon x \cdot y\right),$$

for some $a, \varepsilon > 0$ (see Lemma 3.23 in Chapter 3). The cross term $x \cdot y$ ensures the transmission of the dissipation from momenta to positions when particles are far from the origin. However, because of the singular interaction kernel, the system may blow up if two particles come close together. In the case of Coulomb gases, it has been shown recently [316] that a good choice of Lyapunov function is actually of the form

$$W(x, y) = \exp\left(aH_N(x) + a\frac{|y|^2}{2} + \varepsilon x \cdot y + \varepsilon' \frac{x \cdot y}{|x|}\right),$$

for some $a, \varepsilon, \varepsilon' > 0$. The extra factor ensures the transmission of dissipation when two particles get close. Using the results of [316], we can actually apply the results of Chapter 3 to obtain large deviations results in time for the dynamics (7.12) when N is fixed.

As before, to simulate $(X_t, Y_t)_{t \geq 0}$, one can discretize (7.12) and sample from a trajectory. This will provide a proposal for the HMC scheme as the Euler discretization (7.9) did for Algorithm 7.1. A good way of doing this is a splitting procedure. First, one integrates the Hamiltonian part *i.e.* the operator \mathcal{L}_{ham} in (7.13), which amounts to a standard Hamiltonian dynamics, before integrating the fluctuation-dissipation part *i.e.* the operator \mathcal{L}_{FD} in (7.13). For discretizing the Hamiltonian dynamics over a time step, a standard approach is the Verlet integrator [215, 301], which we describe now. For a time step $\Delta t > 0$, this scheme reads, starting from a state (x_k, y_k) at time k :

$$\begin{cases} y_{k+\frac{1}{2}} = y_k - \nabla H_N(x_k) \alpha_N \frac{\Delta t}{2}, \\ x_{k+1} = x_k + y_{k+\frac{1}{2}} \alpha_N \Delta t, \\ \tilde{y}_{k+1} = y_{k+\frac{1}{2}} - \nabla H_N(x_{k+1}) \alpha_N \frac{\Delta t}{2}. \end{cases}$$

²In the previous chapters, we have used the notation $(q_t, p_t)_{t \geq 0}$ to denote the Langevin dynamics. Here we write $(X_t, Y_t)_{t \geq 0}$ to avoid notational conflicts with the letters p and q , which are used for p -Wasserstein topologies in Chapter 8.

This corresponds to updating the velocity over half a time step, then the positions over a time step, and again the velocity over half a time-step. Given that this scheme only corresponds to the Hamiltonian part, it remains to integrate the fluctuation-dissipation part, corresponding to \mathcal{L}_{FD} in (7.13). For a quadratic kinetic energy, it is a simple Ornstein–Uhlenbeck process whose variance can be computed explicitly. Therefore, we add to the previous scheme the following velocity update which comes from the Mehler formula³:

$$y_{k+1} = \eta \tilde{y}_{k+1} + \sqrt{\frac{1 - \eta^2}{\beta_N}} G_k, \quad \eta = e^{-\gamma_N \alpha_N \Delta t},$$

where G_k is a standard Gaussian random variable. Like the numerical scheme (7.9), because of the singularity of the interactions, this integrator may not have an invariant measure [324], or its invariant measure may be a poor approximation of ρ_N depending on the time step [295]. Note that, here again, α_N and Δt play the same role.

Hybrid or Hamiltonian Monte Carlo (HMC) methods, built on the later integration, appeared in theoretical physics in lattice quantum chromodynamics with [142], see also [375], and are still actively studied in applied mathematics, see for instance [37, 393, 301, 242, 63, 153, 105] and references therein. The HMC algorithm can be thought of in a sense as a special Metropolis Adjusted (underdamped) Langevin Algorithm. Indeed, inspired by the MALA Algorithm 7.1, a way to avoid the stability problem of the discretization of the kinetic Langevin dynamics mentioned above is to add an acceptance-rejection step. A surprising advantage of this approach is that the Verlet integration scheme is time reversible up to momenta reversal [301, Sec. 2.1.3 and eq. (2.11)], hence when computing the acceptance probability as in (7.11), the transition kernel does not appear. Note that the Verlet algorithm has been widely used for years by statistical physicists, and goes back to the historical works of Verlet [418] and Levesque and Verlet [310, 309]. Let us now describe the algorithm.

Algorithm 7.3 (HMC). *Start from a configuration (x_0, y_0) and perform the following steps for each time $k \geq 0$:*

1. *Update the velocities with*

$$\tilde{y}_k = \eta y_k + \sqrt{\frac{1 - \eta^2}{\beta_N}} G_k, \quad \eta = e^{-\gamma_N \alpha_N \Delta t}.$$

2. *Run one step of the Verlet scheme:*

$$\begin{cases} \tilde{y}_{k+\frac{1}{2}} = \tilde{y}_k - \nabla H_N(x_k) \alpha_N \frac{\Delta t}{2}, \\ \tilde{x}_{k+1} = x_k + \tilde{y}_{k+\frac{1}{2}} \alpha_N \Delta t, \\ \tilde{y}_{k+1} = \tilde{y}_{k+\frac{1}{2}} - \nabla H_N(x_{k+1}) \alpha_N \frac{\Delta t}{2}. \end{cases} \quad (7.15)$$

3. *Compute the probability ratio*

$$p_k = 1 \wedge \exp \left[-\beta_N \left(H_N(\tilde{x}_{k+1}) + \frac{|\tilde{y}_{k+1}|^2}{2} - H_N(x_k) - \frac{|y_k|^2}{2} \right) \right].$$

4. *Set*

$$(x_{k+1}, y_{k+1}) = \begin{cases} (\tilde{x}_{k+1}, \tilde{y}_{k+1}) & \text{with probability } p_k, \\ (x_k, -y_k) & \text{with probability } 1 - p_k. \end{cases}$$

As noted in the various references above, the Metropolis step acts as a corrector on the energy conservation of the Hamiltonian step. In this, it helps avoiding irrelevant moves, while enhancing the exploration capacities of the dynamics through the speed variable. A more precise argument in favor of this algorithm is the scaling of the time step Δt with respect to the system size N . Indeed, as shown in [37] for product measures, the optimal scaling is as $\Delta t \sim N^{-\frac{1}{4}}$, which makes the algorithm appealing for large systems. Since the Hamiltonian computational cost scales as N^2 , we see that the

³The Mehler formula states that the Ornstein–Uhlenbeck process $(Z_t)_{t \geq 0}$ in \mathbb{R}^n solution of the stochastic differential equation $dZ_t = \sqrt{2\sigma^2} dB_t - r Z_t dt$ satisfies $\text{Law}(Z_{t+s} | Z_s = z) = \mathcal{N}(ze^{-rt}, \frac{1 - e^{-2rt}}{r} \sigma^2 I_n)$.

cost of the algorithm for a fixed time T and $N = \lceil T/\Delta t \rceil$ is in $\mathcal{O}(N^{\frac{9}{4}})$, which has to be compared to the $\mathcal{O}(N^4)$ cost reached in [311]. Finally, the parameter γ_N can also be tuned in order to optimize the speed of convergence – we leave this point here and stick to $\gamma_N = 1$.

The control of the error or rate of convergence for the HMC algorithm is the subject of active research, see for instance [293] and [153, 60] for some results under structural assumptions.

From a practical point of view, the algorithm can be tested in the following way. First, when only the Hamiltonian part of the dynamics is integrated with the Verlet scheme (7.15), it can be checked that the energy variation over one time step scales as Δt^3 as $\Delta t \rightarrow 0$. Then, if the selection step is added, the rejection rate should also scale as Δt^3 . When the momentum resampling is added, this rejection rate scaling should not change. For completeness, we illustrate some of these facts in Section 8.4.

7.3 Numerical experiments on remarkable models

In this section, we start testing Algorithm 7.3 for the two cases described in Section 7.1.3. Since the equilibrium measures are known for any $N \geq 2$, we will be able to compare accurately our results with the expected one. We will also consider models for which the empirical spectral distribution and the equilibrium distribution are not known. We remind that when $\mathcal{X} = \mathbb{R}^d$ with $d \geq 1$ we have the following formulas that hold in any dimension:

$$\nabla |x|^2 = 2x, \quad \nabla \log \frac{1}{|x|} = -\frac{x}{|x|^2}, \quad \nabla \frac{1}{|x|} = -\frac{x}{|x|^3}.$$

7.3.1 Case study: 1D

We test the numerical method by looking at the mean empirical distribution in the case of the Gaussian Unitary Ensemble (7.2) with $\beta = 2$, $N = 8$, for which the exact expression of $\mathbb{E}\mu_N$ under P_N is provided by (7.3). The results in Figure 7.1 show a very good agreement between the exact result and the algorithm. For completeness, we study the rejection rate of the algorithm as Δt goes to zero, as mentioned at the end of Section 7.2.2. More precisely, we compute over a trajectory the rate of rejected moves in the Step (4) of Algorithm 7.3. The logarithmic plot in Figure 7.2 shows a linear fit with a slope of about 3.1, which confirms the expected scaling in Δt^3 .

We also study the quartic confinement potential $V(x) = x^4/4$, as in [311]. In this case, the empirical spectral distribution is not known, but the equilibrium distribution has density with respect to the Lebesgue measure given by

$$x \in \mathbb{R} \mapsto (2a^2 + x^2) \frac{\sqrt{4a^2 - x^2}}{2\pi} \mathbb{1}_{x \in [-2a, 2a]}, \quad a = 3^{-\frac{1}{4}}.$$

The results of the numerical simulations, see Figure 7.3, show a good agreement with the equilibrium measure when N is large. Note that a tridiagonal random matrix model is known but it does not have independent entries, see [278, Proposition 2.1].

7.3.2 Case study: 2D

We next consider in Figure 7.4 the mean empirical distribution in the case of the Complex Ginibre Ensemble (7.5) with $\beta = 2$, $N = 8$. In this case, we also know a theoretical formula for $\mathbb{E}\mu_N$ under P_N , given by (7.6). For completeness, we investigate the scaling of the relative energy difference in the Step (3) of Algorithm 7.3 (by turning off the selection procedure of Step (4)). The logarithmic plot in Figure 7.5 shows a slope of about 2.9, which confirms the expected scaling in Δt^3 that corresponds to the error of energy conservation, over one time step, of the Verlet integrator (7.15).

We explore next in Figure 7.6 the Gumbel fluctuation at the edge, which is proved for $\beta = 2$ and conjectured for $\beta \neq 2$, see [365, 86, 143] (note that in this case we have a formula for μ_\star but not for $\mathbb{E}\mu_N$ under P_N). One could also explore the crystallization phenomenon, see [45] and references therein.

7.3.3 Case study: 3D

In Figure 7.7, we finally turn to the Coulomb gas which corresponds to $\mathcal{X} = \mathbb{R}^3$, $d = n = 3$, $V = |\cdot|^2/\beta$, $K = 1/|\cdot|$ and to the log-gas for which $K = -\log|\cdot|$. In the first case the equilibrium measure μ_\star is




Figure 7.1 – Study of the Gaussian Unitary Ensemble with $N = 8$ (top) and $N = 50$ (bottom). The solid line is the plot of the limiting spectral distribution (7.4) while the dashed line is the plot of the mean empirical distribution (7.3). The bars form the histogram of simulations obtained using our HMC algorithm. This algorithm was run once with final-time $T = 10^6$ and time-step $\Delta t = 0.5$. The histogram was produced by looking at the last half of the trajectory and retaining the positions each 1000 time-steps, producing n values, g namely $\approx 8 \times 10^3$ and $\approx 5 \times 10^4$ respectively.

uniform on the centered ball of \mathbb{R}^d of radius $(\beta(d-2)/2)^{1/d}$, see for instance [84, Corollary 1.3], while in the second case the equilibrium measure is not known yet, see however [87]. In both cases we do not have a formula for $\mathbb{E}\mu_N$ under P_N . One could study the fluctuation at the edge, which is conjectured to be Gumbel, just like for the complex Ginibre ensemble in 2D.

Acknowledgments

We are also grateful to Thomas Leblé and Laure Dumaz for their comments on the first version of this work.




Figure 7.2 – Evolution of the rejection rate in Algorithm 7.3 as Δt goes to zero, for the Gaussian Unitary Ensemble with $N = 50$, $\beta = 2$ and $T = 10^5$ (in log-log coordinate).




Figure 7.3 – Study of the quartic confinement with $N = 8$ (top) and $N = 50$ (bottom). The solid line is the plot of the limiting spectral distribution (7.4). The bars form the histogram of simulations obtained using our HMC algorithm. This algorithm was run once with final-time $T = 10^6$ and time-step $\Delta t = 0.5$. The histogram was produced by looking at the last half of the trajectory and retaining the positions each 1000 time-steps, producing n values namely $\approx 8 \times 10^3$ and $\approx 5 \times 10^4$ respectively. We do not have a formula for the mean empirical distribution for this model. This gas describes the law of the eigenvalues of a random symmetric tridiagonal matrix model but its entries are not independent, see [278, Proposition 2].




Figure 7.4 – Study of the complex Ginibre ensemble with $N = 8$ (top) and $N = 50$ (bottom). The solid line is the plot of the limiting spectral distribution (7.7) while the dashed line is the plot of the mean empirical distribution (7.6), both as functions of the radius $|z|$ and scaled by 2π (in order to obtain a radial density). The bars form the histogram of simulations obtained using our HMC algorithm. This algorithm was run 40 times with final-time $T = 10^5$ and time-step $\Delta t = 0.1$. The histogram was produced by looking at the last halves of the 40 trajectories and retaining the positions each 10000 time-steps, producing n values namely $\approx 16 \times 10^3$ and $\approx 10^5$ respectively.




Figure 7.5 – Evolution of the energy difference in Algorithm 7.3 as Δt goes to zero, for the Complex Ginibre Ensemble with $N = 50$, $\beta = 2$ and $T = 10^3$ (in log-log coordinate).

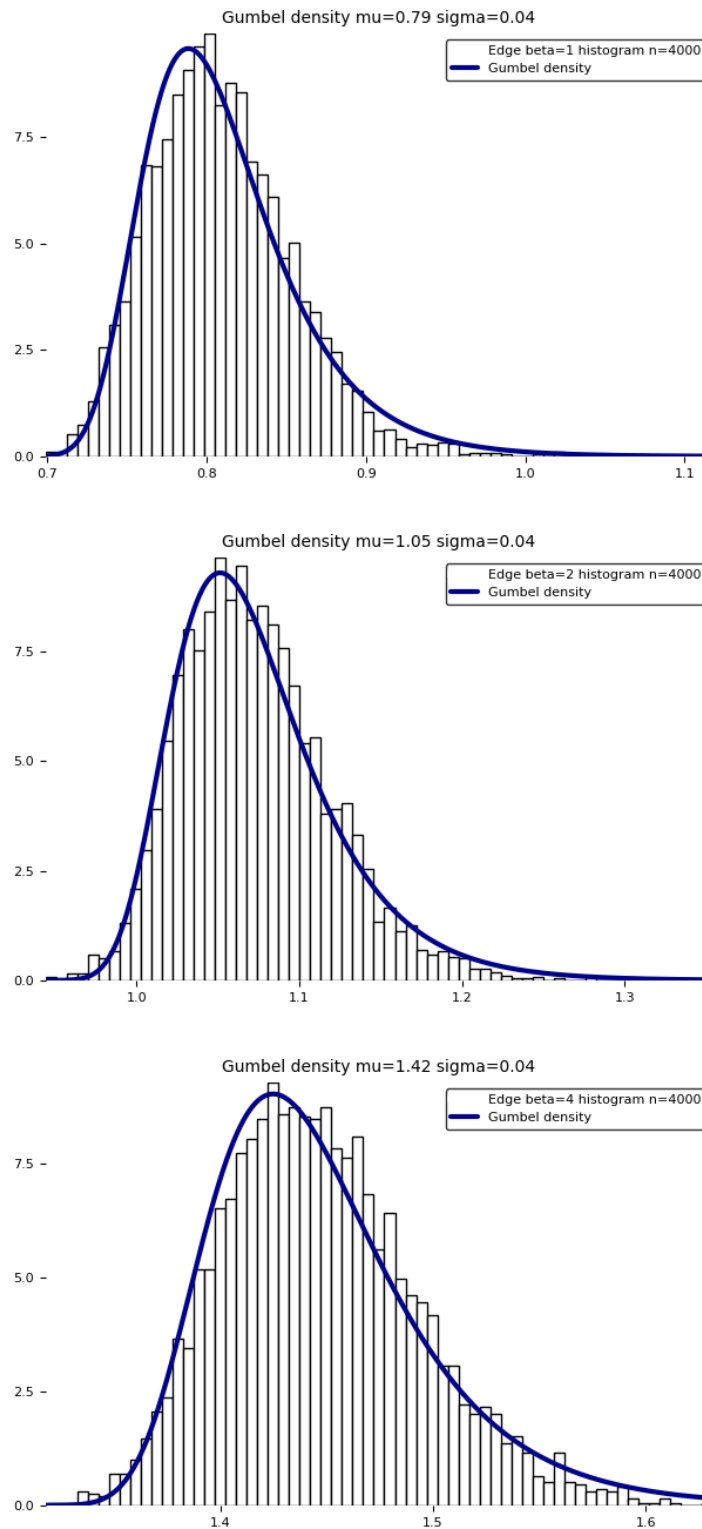


Figure 7.6 – Study of the fluctuation of the largest particle in modulus for the β complex Ginibre ensemble with $N = 50$, in the cases $\beta \in \{1, 2, 4\}$. The solid line is the plot of the fit with a translation-scale Gumbel distribution. The Gumbel fluctuation is proved only in the case $\beta = 2$, see [365, 86]. These simulations suggest to conjecture that the Gumbel fluctuation is valid for any $\beta > 0$. The simulation matches pretty well the edge support at $\sqrt{\beta/2}$ and suggests that the variance is not very sensitive to β .

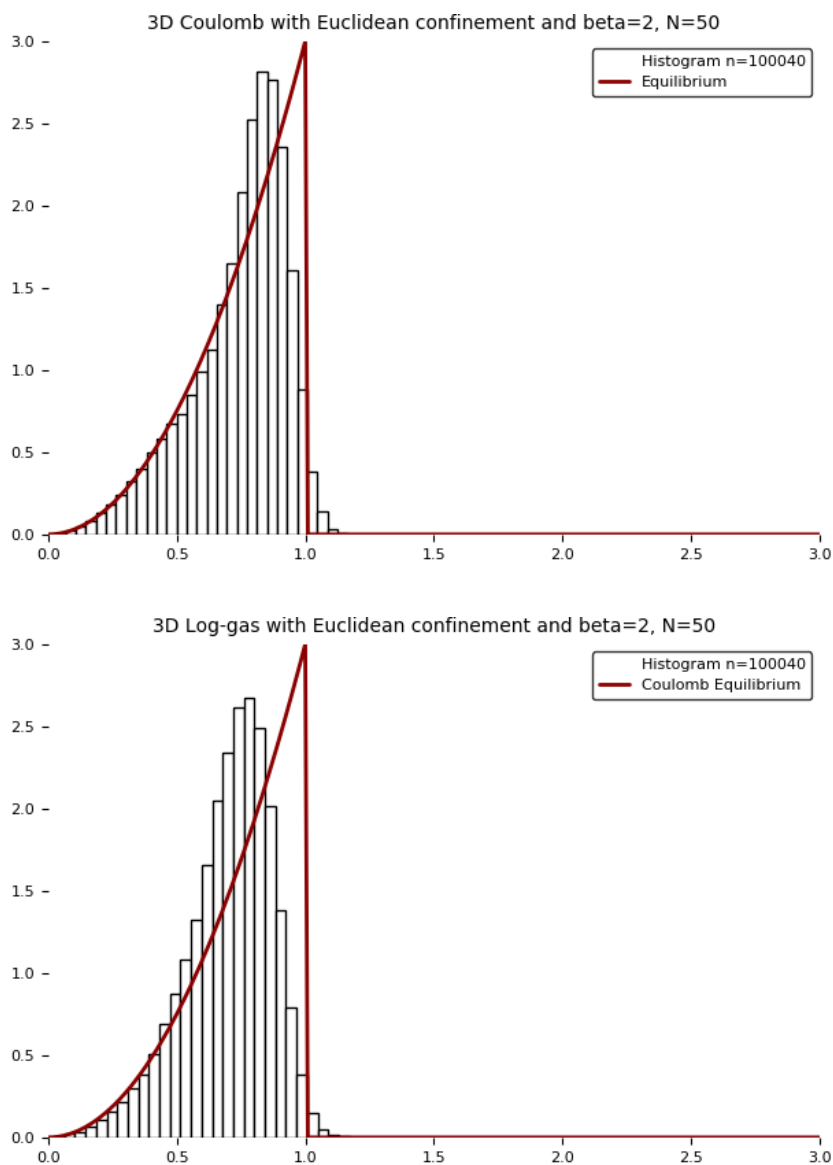


Figure 7.7 – Study of the 3D Coulomb case (top) and 3D Log-gas (bottom) with Euclidean confinement and $\beta = 2$ and $N = 50$. Equilibrium measure in solid line and histogram obtained with our HMC algorithm with $N = 50$ and same simulation parameters as for Figure 7.4. In contrast with the GUE case and the Ginibre case, we do not have a formula for the mean empirical distribution at fixed N for both cases, and for the Log-gas (bottom) the equilibrium measure is not known.

CHAPTER 8

COULOMB GASES UNDER CONSTRAINT: SOME THEORETICAL AND NUMERICAL RESULTS

The material for this chapter has been released on arXiv [83].

Contents

8.1	Motivation	228
8.2	From conditioning to shifting: quadratic confinement with linear constraint	231
8.3	A general conditioning framework	233
8.3.1	Gibbs conditioning	233
8.3.2	Linear statistics	235
8.3.3	Quadratic statistics	237
8.4	Numerical illustration	239
8.4.1	Description of the algorithm	239
8.4.2	Numerical results	243
8.5	Proofs	248
8.5.1	Proof of Theorem 8.1	248
8.5.2	Proofs of Section 8.3.1	250
8.5.3	Proof of Theorem 8.9	252
8.5.4	Proof of Theorem 8.16	256

Abstract. We consider Coulomb gas models for which the empirical measure typically concentrates, when the number of particles gets large, on an equilibrium measure minimizing an electrostatic energy. We study the behavior when the gas is conditioned on a rare event. We first show that the special case of quadratic confinement and linear constraint is exactly solvable due to a remarkable factorization, and the conditioning has then the simple effect of shifting the cloud of particles without deformation. To address more general cases, we perform a theoretical asymptotic analysis relying on the Gibbs conditioning principle presented in Section 1.2.2. The technical part amounts to establishing that the conditioning ensemble is an I -continuity set of the energy. This leads to characterize the conditioned equilibrium as the solution of a modified variational problem. For simplicity, we focus on linear statistics and on quadratic statistics constraints. Finally, we illustrate numerically our predictions and explore cases in which no explicit solution is known. For this, we use a Generalized Hybrid Monte Carlo algorithm (which generalizes the algorithm used in Chapter 7 by taking into a submanifold constraint) for sampling from the conditioned distribution for a finite but large system.

8.1 Motivation

This section contains the main elements of the considered model, some motivations and the plan of the chapter. We consider here the so called Coulomb gas model already encountered in Section 1.1.4 and Chapter 7, which shows an interesting behaviour in the limit of a large number of particles, see for instance [84, 387, 287]¹. As we have seen at several places above, the model consists in a set of random particles $X_{N,1}, \dots, X_{N,N}$ for $N \geq 2$, where each $X_{N,i}$ belongs to $\mathcal{X} = \mathbb{R}^d$ for some physical dimension $d \geq 2$. The particles interact through the Coulomb kernel $g : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by

$$g(x) = \begin{cases} \log \frac{1}{|x|}, & \text{if } d = 2, \\ \frac{1}{(d-2)|x|^{d-2}}, & \text{if } d \geq 3. \end{cases}$$

This denomination comes from the equation satisfied by the interaction g . Indeed, denoting by δ_0 the Dirac mass at 0, g solves in the sense of distributions the following Poisson problem:

$$-\Delta g = c_d \delta_0, \quad \text{with } c_d = \text{surface}(\{x \in \mathbb{R}^d : |x| = 1\}) = 2 \frac{\pi^{d/2}}{\Gamma(d/2)}. \quad (8.1)$$

Note that $\lim_{|x| \rightarrow +\infty} g(x) = 0$ if $d \geq 3$, while $\lim_{|x| \rightarrow +\infty} g(x) = +\infty$ if $d = 2$. In (8.1) $\Delta = \sum_{i=1}^d \partial_i^2$ denotes the Laplacian operator in \mathbb{R}^d . In addition to this pair interaction, the particles are subject to a confining potential $V : \mathbb{R}^d \rightarrow \mathbb{R}$ assumed to be lower semi-continuous and such that

$$\liminf_{|x| \rightarrow +\infty} (V(x) - 2\mathbf{1}_{d=2} \log |x|) > -\infty. \quad (8.2)$$

Following [85] or [386], under this assumption, we can define the electrostatic energy on $\mathcal{P}(\mathbb{R}^d)$ by

$$\mathcal{E}(\mu) = \iint_{\mathbb{R}^d \times \mathbb{R}^d} \left(g(x-y) + \frac{V(x) + V(y)}{2} \right) \mu(dx) \mu(dy). \quad (8.3)$$

This makes sense in $\mathbb{R} \cup \{+\infty\}$ since the integrand is bounded from below thanks to the assumption (8.2) on V . Moreover for all $\mu \in \mathcal{P}(\mathbb{R}^d)$ such that $\int \log(1 + |x|) \mathbf{1}_{d=2} \mu(dx) < \infty$, we have

$$\mathcal{E}(\mu) = \iint_{\mathbb{R}^d \times \mathbb{R}^d} g(x-y) \mu(dx) \mu(dy) + \int_{\mathbb{R}^d} V(x) \mu(dx). \quad (8.4)$$

The functional \mathcal{E} has a unique minimizer on $\mathcal{P}(\mathbb{R}^d)$ called the *equilibrium measure* [85, 386]

$$\mu_\star = \operatorname{argmin}_{\mathcal{P}(\mathbb{R}^d)} \mathcal{E}. \quad (8.5)$$

It has compact support, and if moreover V has a Lipschitz continuous derivative then it has density

$$\frac{\Delta V}{2c_d} \quad (8.6)$$

on the interior of its support. In particular if V is proportional to $|\cdot|^2$ then μ_\star is uniform on a ball. The compactness of the support of μ_\star comes from the *strong confinement* assumption (8.2). Note that it is possible to consider *weakly confining potentials* for which the equilibrium measure still exists but is no longer compactly supported, see for instance the spherical ensemble in [225, 86].

Let $X_N = (X_{N,1}, \dots, X_{N,N})$ be a random vector of $(\mathbb{R}^d)^N$ with law

$$P_N(dx) = \frac{e^{-\beta_N H_N(x_1, \dots, x_N)}}{Z_N} dx_1 \cdots dx_N, \quad (8.7)$$

where $\beta_N > 0$ satisfies

$$\lim_{N \rightarrow \infty} \frac{\beta_N}{N} = +\infty,$$

¹Compared to the previous chapter, we focus on this model (compared to more general gases) because it is more amenable to a theoretical analysis, but we believe most of our results can be generalized at the price of some additional technical difficulties.

and

$$H_N(x_1, \dots, x_N) = \frac{1}{N} \sum_{i=1}^N V(x_i) + \frac{1}{N^2} \sum_{i \neq j} g(x_i - x_j). \quad (8.8)$$

This makes sense only if

$$Z_N = \int_{(\mathbb{R}^d)^N} e^{-\beta_N H_N(x_1, \dots, x_N)} dx_1 \cdots dx_N < \infty, \quad (8.9)$$

and this is the case when V satisfies

$$\int_{\mathbb{R}^d} e^{-\frac{\beta_N}{N}(V(x) - 2\mathbb{1}_{d=2} \log(1+|x|))} dx < \infty. \quad (8.10)$$

As we have seen in Chapter 7 and before in Section 1.1.4, this model is standard in mathematical physics: P_N is a Boltzmann–Gibbs measure modelling a gas of particles, called here a *Coulomb gas*, at inverse temperature β_N and with Hamiltonian H_N . The law P_N is exchangeable in the sense that H_N is symmetric in x_1, \dots, x_N . Indeed, it depends on x_1, \dots, x_N only via the empirical measure, namely, P_N almost surely,

$$H_N = \int_{\mathbb{R}^d} V(x) \mu_N(dx) + \iint_{\neq} g(x-y) \mu_N(dx) \mu_N(dy) \quad \text{with} \quad \mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}, \quad (8.11)$$

where the double integration runs over $\{(x, y) \in \mathbb{R}^d \times \mathbb{R}^d \mid x \neq y\}$. A heuristic reasoning suggests that, if $\beta_N \rightarrow +\infty$ fast enough, under P_N the empirical measure μ_N should concentrate in the limit $N \rightarrow +\infty$ on the equilibrium measure μ_* that minimizes the energy \mathcal{E} in (8.4)–(8.5). This is intuited from the Laplace principle given the expression (8.7) for P_N , where H_N is defined by (8.11). This intuition can be made rigorous through a large deviations principle (LDP), which can be established in this case and many others, see for instance [84, 148, 196] and the references therein. In particular, the case $d = 2$ with quadratic confinement V corresponds to the well-known Ginibre ensemble for random matrices [188, 169], see Section 1.1.4. We mention again that we could consider more general interactions, such as Riesz kernels [84, 287], discontinuous [148] or weak [225] confinement (see Section 1.3.3), but we stick to this setting for ease of presentation. The technical requirements needed for extending our proofs will be pointed out throughout the chapter, and cases not covered by the theoretical analysis will be investigated numerically in Section 8.4.

As large deviations are concerned with probabilities of rare fluctuations, it is possible to consider the empirical measure of the random gas conditioned on such a fluctuation. There has been a number of works on the behaviour of such gases conditioned on having an unusual proportion of the particles lying in some region of the space. As an example, for $d = 2$ and V quadratic, [14] reformulates the conditioned equilibrium measure through an obstacle problem. On the other hand [201, 200] consider the rare situation in which there is a “hole” in the distribution, in other words no particle around zero. Finally [318, 319] consider the one dimensional Wigner situation in which an abnormal proportion of particles lie on one side of the real line. Explicit expressions can be obtained in the latter case. The study of such conditionings is motivated by questions arising in theoretical physics, see for instance the discussion in [319].

While the above mentioned works bring substantial contributions to the understanding of conditioned random gas distributions, they also motivate further questions. Indeed, one may consider more general constraints, like conditioning on the barycenter of the cloud being far away from the origin. This may be of interest for both theoretical [14] and practical purposes (if one wants to filter out noise conditioned on some rare event [66]). Moreover, the numerical methods proposed in [201, 200, 319] do not seem adapted to sampling the empirical distribution conditioned on some event – since this event is typically rare, naive sampling is generally not efficient. The goal of this work is therefore to investigate some theoretical results on such conditioned Coulomb gases, as well as providing an efficient algorithm to sample conditioned distributions. The algorithm we use is a generalization of the Hamiltonian Monte Carlo algorithm used in Chapter 7 which incorporates a submanifold constraint.

Mathematically, our aim is to consider the particles $Y_N = (Y_{N,1}, \dots, Y_{N,N})$ in $(\mathbb{R}^d)^N$ such that

$$Y_N \sim \text{Law}(X_N \mid \xi_N(X_N) \leq 0), \quad (8.12)$$

where $\xi_N : (\mathbb{R}^d)^N \rightarrow \mathbb{R}$, and to consider the limiting behaviour of the empirical measure

$$\frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}},$$

as $N \rightarrow +\infty$, depending on the confinement potential V and the constraint ξ_N . Instead of an inequality constraint like (8.12), we may instead consider an equality constraint

$$Y_N \sim \text{Law}(X_N \mid \xi_N(X_N) = 0).$$

We will generally consider inequality constraints since they naturally lead to a Gibbs conditioning principle. Equality constraints could be considered as well by an additional limiting procedure, see [119, Section 7.3] and the discussion in Section 8.3. We could also set ξ_N to be \mathbb{R}^m -valued for some $m \geq 2$ but we restrict to one dimensional constraints for ease of exposition. The cases studied in [14, 201, 200, 318, 319] correspond to the choice

$$\xi_N(x_1, \dots, x_N) = \mu_N(\mathbb{1}_U) - c,$$

for some measurable set $U \subset \mathbb{R}^d$ and constant $c \in \mathbb{R}$. We will study in this chapter more general *linear statistics* of the form

$$\xi_N(x_1, \dots, x_N) = \mu_N(\varphi), \tag{8.13}$$

for some constraint function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying growth conditions, see Section 8.3.2. A particular case of interest is when the constraint function φ is itself linear, namely:

$$\varphi(x) = x \cdot v - c, \tag{8.14}$$

for $v \in \mathbb{R}^d$ and $c \in \mathbb{R}$. Indeed, when φ is chosen according to (8.14) and V is *quadratic*, the equilibrium measure under conditioning is the unconditioned one translated in the direction of v . We provide a simple proof of this result in Section 8.2. We next turn to more general constraints in Section 8.3, proving first an abstract Gibbs conditioning principle in Section 8.3.1. When considering linear statistics, we prove in Section 8.3.2 that conditioning P_N boils down to modifying the confinement potential V . In Section 8.3.3 we consider the case of quadratic statistics, which modifies the interaction kernel g .

In order to validate our theoretical results and explore cases in which explicit solutions are not available, we also propose an original method for sampling the law of Y_N for a fixed N . Based on the Hamiltonian Monte Carlo (HMC) method used in [82] for sampling Gibbs measures associated to Coulomb and Log-gases, we describe and implement the generalized Hamiltonian Monte Carlo algorithm proposed in [303] for sampling probability measures on submanifolds. The method, detailed in Section 8.4.1, shows remarkable performance, and the results presented in Section 8.4.2 are in agreement with the theory.

Notation

We introduce and recall some notation used throughout the chapter. For all $d \geq 1$ we denote by $|x| = (x_1^2 + \dots + x_d^2)^{1/2}$ the Euclidean norm and by $x \cdot y = x_1 y_1 + \dots + x_d y_d$ the scalar product on \mathbb{R}^d . We denote by $\mathcal{P}(\mathbb{R}^d)$ the set of probability measures on \mathbb{R}^d and, for all $p \geq 1$, by $\mathcal{P}_p(\mathbb{R}^d)$ those probability measures having finite p -moments in the sense that $|\cdot|^p$ is integrable. For any measure $\mu \in \mathcal{P}(\mathbb{R}^d)$, the support of μ is defined as $\text{supp}(\mu) = \mathbb{R}^d \setminus A$, where A is the largest open set such that $\mu(A) = 0$ (which may be empty). For all measurable $f : \mathbb{R}^d \rightarrow \mathbb{R}$, we define

$$\|f\|_\infty = \sup_{x \in \mathbb{R}^d} |f(x)| \quad \text{and} \quad \|f\|_{\text{Lip}} = \sup_{x \neq y} \frac{|f(x) - f(y)|}{|x - y|}.$$

We define the bounded-Lipschitz distance on $\mathcal{P}(\mathbb{R}^d)$ by

$$d_{\text{BL}}(\mu, \nu) = \sup_{\substack{\|f\|_\infty \leq 1 \\ \|f\|_{\text{Lip}} \leq 1}} \int_{\mathbb{R}^d} f d(\mu - \nu).$$

For all $p \geq 1$, we define the p -Wasserstein² distance on $\mathcal{P}_p(\mathbb{R}^d)$ by

$$d_{\text{W}_p}(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu, \nu)} \iint_{\mathbb{R}^d \times \mathbb{R}^d} |x - y|^p \pi(dx, dy) \right)^{\frac{1}{p}},$$

²Or Monge, or Kantorovich, or transportation distance.

where $\Pi(\mu, \nu)$ is the set of probability measures on the product space $\mathbb{R}^d \times \mathbb{R}^d$ with marginal distributions μ and ν . Following [419], the Kantorovich–Rubinstein duality theorems state that

$$d_{W_1}(\mu, \nu) = \sup_{\|f\|_{\text{Lip}} \leq 1} \int_{\mathbb{R}^d} f d(\mu - \nu) \quad \text{and} \quad d_{W_p}(\mu, \nu)^p = \sup_{\substack{f \in L^1(\mu), g \in L^1(\nu) \\ f(x) \leq g(y) + |x-y|^p}} \left(\int_{\mathbb{R}^d} f d\mu - \int_{\mathbb{R}^d} g d\nu \right). \quad (8.15)$$

For any $p \geq 1$, we say that a function f is dominated by $|x|^p$ when³

$$\|f\|_{\infty, p} = \sup_{x \in \mathbb{R}^d} \frac{|f(x)|}{1 + |x|^p} < \infty.$$

The p -Wasserstein topology is the one induced on $\mathcal{P}_p(\mathbb{R}^d)$ by d_{W_p} . If $(\nu_N)_N$ is a sequence in $\mathcal{P}(\mathbb{R}^d)$ then $\lim_{N \rightarrow \infty} d_{\text{BL}}(\nu_N, \nu) = 0$ if and only if $\lim_{N \rightarrow \infty} \int f d\nu_N = \int f d\nu$ for all bounded continuous $f : \mathbb{R}^d \rightarrow \mathbb{R}$. For all $p \geq 1$ and all sequence $(\nu_N)_N$ in $\mathcal{P}_p(\mathbb{R}^d)$, we have $\lim_{N \rightarrow \infty} d_{W_p}(\nu_N, \nu) = 0$ if and only if $\lim_{N \rightarrow \infty} d_{\text{BL}}(\nu_N, \nu) = 0$ and $\lim_{N \rightarrow \infty} \int |x|^p d\nu_N = \int |x|^p d\nu$. In other words d_{BL} metrizes weak convergence, while d_{W_p} metrizes weak convergence plus convergence of the p -moment, see [419].

We denote $X \sim P$ to say that the random variable X has law P , and $X \stackrel{d}{=} Y$ to say that the random variables X and Y have same law. A sequence of random variables (X_1, \dots, X_N) is exchangeable if for any permutation σ of $\{1, \dots, N\}$ it holds $(X_1, \dots, X_N) \stackrel{d}{=} (X_{\sigma(1)}, \dots, X_{\sigma(N)})$.

We now recall the large deviations definitions from Section 1.2.1 in Part I. We say that a sequence Y_1, \dots, Y_N of random variables taking value in a metric space (\mathcal{Y}, d) satisfies a large deviations principle at speed β_N if, for any measurable set $A \subset \mathcal{Y}$, it holds

$$-\inf_{\mu \in \bar{A}} I(\mu) \leq \liminf_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P}(Y_N \in A) \leq \overline{\lim}_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P}(Y_N \in A) \leq -\inf_{\mu \in \bar{A}} I(\mu), \quad (8.16)$$

where the interior and closure are taken with respect to the topology induced by d , while $I : \mathcal{Y} \rightarrow [0, +\infty]$ is lower semicontinuous, and called the rate function. If I has compact level sets for the topology induced by d , we say that I is a good rate function (see Definitions 1.3 and 1.4 in Part I).

We finally recall some elements of potential theory. The interaction energy

$$J(\mu, \nu) = \iint_{\mathbb{R}^d \times \mathbb{R}^d} g(x-y) \mu(dx) \nu(dy) \quad (8.17)$$

is well defined for any signed measures μ, ν with compact supports and takes values in $\mathbb{R} \cup \{+\infty\}$, see [284, Chapter I]. We use the abuse of notation $J(\mu) = J(\mu, \mu)$ for the associated quadratic form. Moreover, for any compact set $K \subset \mathbb{R}^d$, J attains its infimum over probability measures supported on K . This value is called the *capacity* of the set K [284, Chapter II]. A measurable set A has *positive capacity* if it contains a compact set K and a measure μ with $\text{supp}(\mu) \subset K$ and such that $J(\mu) < +\infty$. Otherwise, the set is said to have *null capacity*. A property is said to hold *quasi-everywhere* if it is satisfied on a set whose complementary has null capacity. Although inner and outer capacities should be considered, we know these notions coincide for Borel sets on \mathbb{R}^d , see [284, Theorem 2.8]. Denoting by $\mathcal{P}_c(\mathbb{R}^d)$ the set of compactly supported probability measures, in accordance with [284, Theorems 1.15 and 1.16], for any $\mu, \nu \in \mathcal{P}_c(\mathbb{R}^d)$ with $J(\mu) < +\infty$ and $J(\nu) < +\infty$, it holds $J(\mu - \nu) = 0$ if and only if $\mu = \nu$.

8.2 From conditioning to shifting: quadratic confinement with linear constraint

This section is devoted to the particular case where $V(x) = |x|^2$ and the constraint is chosen according to (8.13)-(8.14). This corresponds to the Ginibre ensemble of random matrices if $d = 2$, as presented in Section 1.1.4. The following theorem states that this special case is exactly solvable: the conditioning has the effect of a shift without deformation, due to a remarkable factorization. The proof, presented in Section 8.5.1, is quite elegant. It is inspired from the seemingly unrelated work [88]. The result by itself appears as a special case of the general variational approach presented in Section 8.3 below.

³We use a simpler notation for functions of bounded growth compared to Part II because the only cost (Lyapunov) function we use here is $W(x) = 1 + |x|^p$.

Theorem 8.1 (From conditioning to shifting). *Let $d, N \geq 2$ and $V = |\cdot|^2$, so that (8.9) holds. Let $X_N = (X_{N,1}, \dots, X_{N,N})$ and P_N be as in (8.7). Then the equilibrium measure μ_\star is the uniform law on the centered ball of \mathbb{R}^d of radius 1, and moreover, almost surely, for all $p \geq 1$,*

$$\lim_{N \rightarrow \infty} d_{W_p} \left(\frac{1}{N} \sum_{i=1}^N \delta_{X_{N,i}}, \mu_\star \right) = 0, \tag{8.18}$$

regardless of the way we define the random variables X_N on the same probability space. Now let $v \in \mathbb{R}^d$ with $|v| = 1$, $c \in \mathbb{R}$, choose $\varphi(x) = x \cdot v - c$ and consider $Y_N = (Y_{N,1}, \dots, Y_{N,N})$ with

$$Y_N \sim \text{Law} \left(X_N \left| \frac{\varphi(X_{N,1}) + \dots + \varphi(X_{N,N})}{N} = 0 \right. \right).$$

Then

$$Y_N \stackrel{d}{=} X_N + \left(c - \frac{X_{N,1} + \dots + X_{N,N}}{N} \cdot v \right) (v, \dots, v).$$

Moreover, denoting $\mu^\varphi = \delta_{cv} * \mu_\star$, we have that almost surely and for all $p \geq 1$, it holds

$$\lim_{N \rightarrow \infty} d_{W_p} \left(\frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}}, \mu^\varphi \right) = 0, \tag{8.19}$$

regardless of the way we define the random variables Y_N on the same probability space.

The proof of Theorem 8.1 relies crucially on the quadratic nature of the confinement potential, but remains valid whatever the pair interaction, beyond Coulomb gases, as far as it is translation invariant⁴. More general linear projections can be used. Indeed, if we choose $\varphi(x) = p(x) - c$ where p is a linear projection over a subspace $E \subset \mathbb{R}^d$ of dimension m and $c \in \mathbb{R}^m$, the result still holds.

Remark 8.2 (Ginibre random matrices and Hermite ensembles). *Let M be an $N \times N$ random matrix from the complex Ginibre ensemble, which is a Hermite ensemble (quadratic confinement). Its eigenvalues have law P_N with $d = 2$, $\beta_N = N^2$, and, $V(x) = |x|^2$ and $g(x) = -\log|x|$, see for instance [188, Ch. 15]. Let $\mathbb{R}^2 = \mathbb{C}$, $v \in \mathbb{R}^2$ with $|v| = 1$ and $c \in \mathbb{R}$. The assumptions of Theorem 8.1 are satisfied and the constraint in terms of matrices reads $\text{Tr}(M) \cdot v = nc$, where we identify again \mathbb{C} with \mathbb{R}^2 . More precisely (8.19) holds and the conditioned equilibrium measure reads*

$$\mu^\varphi(dz) = \frac{\mathbb{1}_{|z-cv| \leq 1}}{\pi} dz.$$

The same reasoning can be applied to other matrix ensembles, and more generally to β -ensembles (corresponding to $\beta_N = (\beta/2)N^2$) under quadratic confinement and pairwise repulsion. In the special case of the complex Ginibre ensemble, the entries of M are independent, in particular the diagonal and off-diagonal are independent, and we could deduce the result from the universality theorem on the circular law [402] together with the Gibbs conditioning principle for independent Gaussian variables. Some numerical experiments are provided in Section 8.4.2.

In practice, we would like to consider a non-quadratic confinement and a non-linear constraint function φ . The numerical applications presented in Section 8.4.2 show a much wider range of behaviour than shifting the equilibrium measure. It turns out that the conditioning mechanism is an instance of the *Gibbs conditioning principle* from large deviations theory presented in Section 1.2.2. The purpose of the next section is to provide proofs in this direction, which allow to derive the conditioned equilibrium measure in more general contexts, of which Theorem 8.1 appears as a particular case.

⁴Theorem 8.1 thus shows a spectacular property of *Gaussian ensembles* (in the sense of quadratic confinement). We have already observed two times in this thesis the particular structure of Gaussian variables. First, in Section 1.2.2, we saw that independent Gaussian variables have both small and large Gaussian deviations. Next, in Chapter 3, we noticed that the Orstein–Uhlenbeck process has the particularity of fluctuating like independent Gaussian variables (see Remark 3.11). We can only observe the beautiful regularity of Gaussian objects without more precise statements.

8.3 A general conditioning framework

As is known from the seminal work of Ben Arous and Guionnet [29], large deviations theory provides a natural framework to study the concentration of empirical measures of the spectrum of random matrices and, beyond, of singularly interacting particles systems. We refer in particular to [84, 148, 33] and references therein for recent accounts. Since large deviations theory is concerned with estimating probabilities of rare events, conditioning on such a rare event is a natural direction to follow. This procedure is generally referred to as *Gibbs conditioning principle* or *maximum entropy principle*. This principle is explained for instance in [357, Section 6.3] and [119, Section 7.3].

When no conditioning is considered, we know that under mild assumptions the empirical measure associated to the Gibbs measure (8.7) satisfies a LDP with rate function \mathcal{E} . When the random gas is considered under conditioning on an appropriate rare event, the Gibbs conditioning principle states that the resulting conditioned empirical measures concentrate on a minimizer of \mathcal{E} under constraint. Proofs of this fact in our context are presented in Section 8.3.1. Next, Section 8.3.2 studies the corresponding constrained minimization problem for linear statistics, while Section 8.3.3 is concerned with quadratic statistics.

8.3.1 Gibbs conditioning

The goal of this section is to present an abstract Gibbs conditioning principle and apply it to the Coulomb gas model. Most works considered hitherto Gibbs principles associated to Sanov's theorem [119, 357, 306], in other words in absence of interaction, showing that by conditioning the empirical measure, the resulting equilibrium measure minimizes the rate function under constraint. The same strategy can actually be applied to any exchangeable system satisfying a large deviations principle provided the conditioning set is an I -continuity set, following for instance [104], [119, Section 1.2] and [357, Section 5.3]. This is the purpose of the next proposition, which can be of independent interest and is already presented in Part I (Proposition 1.13). The proof is postponed to Section 8.5.2.

Proposition 8.3 (A Gibbs conditioning principle). *Suppose that Y_1, \dots, Y_N are random variables taking values in a metric space (\mathcal{Y}, d) satisfying a large deviations principle at speed $(\beta_N)_N$, and with good rate function I . Consider a closed set B which is I -continuous in the sense that*

$$\inf_B I = \inf_B I < +\infty. \quad (8.20)$$

Then, the set of minimizers

$$\mathcal{I}_B = \left\{ y \in \mathcal{Y} : I(y) = \inf_B I \right\} \quad (8.21)$$

is a non-empty closed subset of B . Moreover, for any $\varepsilon > 0$, setting

$$A_\varepsilon = \{y \in \mathcal{Y} : d(y, \mathcal{I}_B) > \varepsilon\},$$

there exists $c_\varepsilon > 0$ such that

$$\overline{\lim}_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P} \left(Y_N \in A_\varepsilon \mid Y_N \in B \right) \leq -c_\varepsilon. \quad (8.22)$$

In particular if we define a random variable $Y'_N \sim \text{Law}(Y_N \mid Y_N \in B)$ for all N then almost surely $\lim_{N \rightarrow \infty} d(Y'_N, \mathcal{I}_B) = 0$ regardless of the way we define the Y'_N 's on the same probability space. Further in particular if $\mathcal{I}_B = \{y_B\}$ is a singleton, then almost surely it holds $\lim_{N \rightarrow \infty} Y'_N = y_B = \min_B I$.

In other words, (8.22) entails that the variables Y_N conditioned on being in B concentrate on a minimizer of I over B (see Theorem 8.7 below). In order to use Proposition 8.3 in the Coulomb gas setting, we start by recalling a LDP associated with the Coulomb gas model (see [84, 148]). In order to consider unbounded constraints in what follows, we make the following assumption.

Assumption 8.4 (Growth condition). *There exist $a > 0$, $R \in \mathbb{R}$ and $q > 1$ such that*

$$\forall x \in \mathbb{R}^d, \quad V(x) \geq a|x|^q - R.$$

The above growth condition not only ensures that V satisfies (8.2), but also allows to consider a finer topology for the LDP, see [148, Theorem 1.8]. It could certainly be relaxed under appropriate modifications. In particular, Assumption 8.4 shows that [148, Assumption C'1] is satisfied for any

function of the form $|x|^p$ for $1 < p < q$, so [148, Lemma 1.1] applies. Assumption 8.4 thus has the following consequence. Consider an exponent $p \in (1, q)$. Then, under P_N defined in (8.7), the empirical measure

$$\mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$$

satisfies a large deviations principle (8.16) in the p -Wasserstein topology at speed $(\beta_N)_N$ and with the following good rate function:

$$\mathcal{E}_\star = \mathcal{E} - \inf_{\mathcal{P}_p(\mathbb{R}^d)} \mathcal{E}, \quad \text{with} \quad \inf_{\mathcal{P}_p(\mathbb{R}^d)} \mathcal{E} > -\infty,$$

where \mathcal{E} is defined in (8.3). The energy \mathcal{E} has additional nice properties, which we recall below for convenience.

Proposition 8.5 (Properties of the electrostatic energy). *Let \mathcal{E} be as in (8.3). Suppose that Assumption 8.4 holds, and take some $p \in (1, q)$. Then, denoting $D_\mathcal{E} = \{\mu \in \mathcal{P}(\mathbb{R}^d) : \mathcal{E}(\mu) < +\infty\}$ the domain of \mathcal{E} , the following properties are satisfied:*

- $D_\mathcal{E}$ is convex and \mathcal{E} is convex on $D_\mathcal{E}$;
- $D_\mathcal{E} \subset \mathcal{P}_p(\mathbb{R}^d)$ and there exists a unique $\mu_\star \in \mathcal{P}_c(\mathbb{R}^d)$ such that

$$\mathcal{E}(\mu_\star) = \inf_{\mathcal{P}(\mathbb{R}^d)} \mathcal{E} = \inf_{\mathcal{P}_p(\mathbb{R}^d)} \mathcal{E} = \inf_{\mathcal{P}_c(\mathbb{R}^d)} \mathcal{E}; \tag{8.23}$$

- the minimizer $\mu_\star \in \mathcal{P}_c(\mathbb{R}^d)$ satisfies the Euler–Lagrange conditions (where $C_\star = \mathcal{E}(\mu_\star)$)

$$\begin{cases} 2g \ast \mu_\star + V = C_\star, & \text{quasi – everywhere in } \text{supp}(\mu_\star), \\ 2g \ast \mu_\star + V \geq C_\star, & \text{quasi – everywhere.} \end{cases} \tag{8.24}$$

The domain $D_\mathcal{E}$ is not empty since it contains for instance measures with smooth density over a compact support. For convenience, we recall a proof of these classical results in Section 8.5.2.

Remark 8.6 (Going beyond Coulomb gases and convexity). *The LDP presented here holds for a much larger range of models than the Coulomb gas setting, see for instance [148]. However, the assumptions in [148] do not ensure the convexity of the rate function \mathcal{E} , which poses problems when it comes to identifying the equilibrium measure – we thus stick to this setting here. In practice, the convexity of \mathcal{E} is derived from a Bochner-type positivity of the interaction potential, see [84].*

We are now in position to apply Proposition 8.3 to the Coulomb gas model.

Theorem 8.7 (Gibbs conditioning for Coulomb gases). *Let \mathcal{E} be as in (8.3). Suppose that Assumption 8.4 holds and take some $p \in (1, q)$. Consider a closed set $B \subset \mathcal{P}(\mathbb{R}^d)$ such that*

$$\inf_{\overset{\circ}{B}} \mathcal{E} = \inf_B \mathcal{E} < +\infty, \tag{8.25}$$

where the interior is taken with respect to the p -Wasserstein topology. Then the set of minimizers

$$\mathcal{E}_B = \left\{ \mu \in \mathcal{P}(\mathbb{R}^d) : \mathcal{E}(\mu) = \inf_B \mathcal{E} \right\} \tag{8.26}$$

is a non-empty closed subset of B . Moreover, if $X_N \sim P_N$ is as in (8.7), and if $Y_N = (Y_{N,1}, \dots, Y_{N,N})$ is such that

$$Y_N \sim \text{Law}(X_N \mid \mu_N \in B), \quad \text{with} \quad \mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{X_{N,i}},$$

then almost surely it holds

$$\lim_{N \rightarrow \infty} d_{W_p} \left(\frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}}, \mathcal{E}_B \right) = 0,$$

regardless of the way we define the random variables Y_N on the same probability space.

The proof of Theorem 8.7, which can be found in Section 8.5.2, is an instance of the Gibbs conditioning principle provided by Proposition 8.3: the conditioned empirical measure concentrates almost surely on a minimizer of \mathcal{E} over B . It is alluring to consider a more general LDP for the conditioned empirical measure as in [281], but the arguments proposed in [281] do not fit the case of our singular rate function, in particular $D_{\mathcal{E}}$ has empty interior. Our strategy is then to restrict to an I -continuity set B satisfying (8.25), which allows to use the lower bound of the LDP, very similarly to [357].

Next, rather than aiming at the greatest generality, we consider the case of linear and quadratic statistics constraints, for which I -continuity can be proved and the resulting equilibrium measure can be identified in terms of a modified version of (8.24).

8.3.2 Linear statistics

As explained in the introduction, the case of *linear statistics* is of particular importance. This motivates focusing first on conditioning sets B of the form

$$B = \{\nu \in \mathcal{P}_p(\mathbb{R}^d) : \nu(\varphi) \leq 0\}, \tag{8.27}$$

for some measurable function $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$. This kind of constraint was studied for example in [14, 201, 200, 318, 319]. In particular, the Ginibre case with $\varphi = \mathbf{1}_U - c$ for some measurable set $U \subset \mathbb{R}^2$ and $c \in \mathbb{R}$ is considered in [14]. The choice $\varphi(x) = c - x \cdot v$ for $v \in \mathbb{R}^d$ has been treated in Section 8.2 for the related equality constraint. We consider here more general potentials V and constraint functions φ . The next assumption on φ ensures that B is suitable for conditioning.

Assumption 8.8.

- Assumption 8.4 holds for some $q > 1$;
- $\|\varphi\|_{\text{Lip}} < +\infty$ and thus $\|\varphi\|_{\infty,p} < +\infty$ for all $p \in (1, q)$;
- there exists $\mu_- \in D_{\mathcal{E}}$ such that $\mu_-(\varphi) < 0$;
- there exists $\mu_+ \in D_{\mathcal{E}}$ such that $\mu_+(\varphi) > 0$.

The existence of μ_- means that the set B has non empty interior, while that of μ_+ implies that $B \neq \mathcal{P}_p(\mathbb{R}^d)$, so that the constraint is not trivial. Since the Gibbs principle relies on B being an I -continuity set, we provide a fine analysis of the minimization of \mathcal{E} over the set B defined in (8.27). We prove in particular that the minimizer is unique with compact support, and we characterize it through an integral equation similar to (8.24) with an additional Lagrange multiplier. The proof of this result is presented in Section 8.5.3.

Theorem 8.9 (Variational characterization). *Let $\mu_* \in \mathcal{P}_c(\mathbb{R}^d)$ be the unconstrained equilibrium measure as in Proposition 8.5, and let B the set defined in (8.27). Suppose that Assumption 8.8 holds, for some $q > 1$ and $p \in (1, q)$. Then B is closed in the p -Wasserstein topology and*

$$\inf_{\tilde{B}} \mathcal{E} = \inf_B \mathcal{E} < +\infty. \tag{8.28}$$

Moreover

$$\mathcal{E}_B = \left\{ \mu \in \mathcal{P}(\mathbb{R}^d) : \mathcal{E}(\mu) = \inf_B \mathcal{E} \right\} = \{\mu^\varphi\},$$

where μ^φ has compact support and is solution to, for some $\alpha \geq 0$,

$$\begin{cases} 2g * \mu^\varphi + V + \alpha\varphi = C_\varphi, & \text{quasi - everywhere in } \text{supp}(\mu^\varphi), \\ 2g * \mu^\varphi + V + \alpha\varphi \geq C_\varphi, & \text{quasi - everywhere,} \end{cases} \tag{8.29}$$

with $C_\varphi = \mathcal{E}(\mu^\varphi)$. Finally, one of the two following conditions holds:

- $\mu_* \in B$ and $\alpha = 0$;
- $\mu_* \notin B$, in which case $\mu^\varphi(\varphi) = 0$ and $\alpha > 0$ in other words the constraint is saturated and the Lagrange multiplier is activated.

We have now the following consequence of Theorems 8.7 and 8.9.

Corollary 8.10 (From conditioning to confinement deformation). *Suppose that Assumption 8.8 holds with $q > 1$ and $p \in (1, q)$. Consider a Coulomb gas $X_N = (X_{N,1}, \dots, X_{N,N}) \sim P_N$ as in (8.7). Introduce $Y_N = (Y_{N,1}, \dots, Y_{N,N})$ with law given by*

$$Y_N \sim \text{Law} \left(X_N \left| \frac{1}{N} \sum_{i=1}^N \varphi(X_{N,i}) \leq 0 \right. \right).$$

Let μ^φ be as in Theorem 8.9. Then almost surely it holds

$$\lim_{N \rightarrow \infty} d_{W_p} \left(\frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}}, \mu^\varphi \right) = 0, \tag{8.30}$$

regardless of the way we define the random variables Y_N on the same probability space.

Theorem 8.9 and Corollary 8.10 show that conditioning on a linear statistics is equivalent to changing the confinement potential V into $V + \alpha\varphi$ where $\alpha \geq 0$ is a constant determined by the constraint. If $\mu_\star \in B$, $\alpha = 0$ and the conditioning produces no effect. Note also that the global Lipschitz condition on φ in Assumption 8.8 could possibly be relaxed. For instance, if $\|\varphi\|_{\infty,p} < +\infty$ for some $p \in (1, q)$ we expect that a minimizing measure has compact support, and that assuming φ locally Lipschitz suffices to prove Theorem 8.9. We leave these refinements to further studies.

Remark 8.11 (Equality constraints). *When considering conditioning principles, one is often interested in equality constraints. It is not obvious at first sight to consider a set B defined by an equality constraint, since its interior may well be empty. A common strategy is to use a limiting procedure by introducing nested sets [119]. This is unnecessary here since we observe in Theorem 8.9 that either the equilibrium measure lies in B , either the constraint is saturated.*

Remark 8.12 (Projection). *The conditioned equilibrium measure μ^φ can be interpreted as an instance of entropic projection. These projections have been studied for a long time in the context of the Sanov theorem, in other words independent particles or equivalently product measures without interaction at all, see [305] and the references therein. Theorem 8.9 is therefore a precise study of such a projection in the context of Coulomb gases under linear statistics constraint, where the entropy is replaced by the electrostatic energy \mathcal{E} . These remarks also apply to Section 8.3.3.*

Remark 8.13 (Formula for constrained equilibrium measure under regularity assumptions). *Suppose that Assumption 8.8 holds and that V and φ have Lipschitz continuous derivatives. Then the conditioned equilibrium measure μ^φ which appears in Theorem 8.9 and in Corollary 8.10 satisfies*

$$\begin{cases} \mu^\varphi = \frac{\Delta V + \alpha \Delta \varphi}{2c_d}, & \text{almost everywhere in } \text{supp}(\mu^\varphi), \\ \mu^\varphi = 0, & \text{almost everywhere outside } \text{supp}(\mu^\varphi). \end{cases} \tag{8.31}$$

Indeed, it suffices to apply the Laplacian to both sides of (8.29) and use (8.1), and we refer for example to [386, Proposition 2.22] for the technical details.

It is now possible to come back to the translation phenomenon described in Section 8.2 through the energetic approach considered in the present section.

Alternative proof of Theorem 8.1. Using (8.31) under the assumptions of Theorem 8.1, we have $\Delta V = 2d$ and $\Delta \varphi = 0$, so that μ^φ is constant and equal to d/c_d on its support. It then remains to show that this support is indeed a ball of correct center and radius. For this, we observe that, since $|v| = 1$,

$$V(x) + \alpha\varphi(x) = |x|^2 + \alpha(c - x \cdot v) = \left| x - \frac{\alpha v}{2} \right|^2 + \frac{\alpha^2}{4} + \alpha c,$$

so that the effective confining potential is quadratic with variance $1/2$ and center $x_0 = \alpha v/2$. By radial symmetry around x_0 , μ^φ must be a uniform distribution on a ball $B(x_0, r)$ centered at x_0 with radius $r > 0$. In order to find the value of α , we write the constraint

$$|B(x_0, r)|^{-1} \int_{B(x_0, r)} x \cdot v \, dx = c.$$

The left hand side of the above equation reads, by symmetry,

$$|B(x_0, r)|^{-1} \int_{B(x_0, r)} (x - x_0) \cdot v \, dx + x_0 \cdot v = |B(x_0, r)|^{-1} \int_{B(0, r)} x \cdot v \, dx + x_0 \cdot v = x_0 \cdot v.$$

Since $x_0 = \alpha v/2$ and $|v|^2 = 1$ we obtain

$$\alpha = 2c,$$

which leads to $x_0 = cv$. Finally, the value of μ^ν over its support is d/c_d , where c_d is the surface of the sphere in dimension d . Since the volume of the sphere of radius r is equal to rc_d/d , we obtain that $r = 1$ and we reach the conclusion of Theorem 8.1. \square

8.3.3 Quadratic statistics

Once the linear statistics case has been studied, it is natural to turn to more general constraints. Considering second order statistics is a first step in this direction, which motivates to consider sets of the form, for $p > 1$,

$$B = \{\nu \in \mathcal{P}_p(\mathbb{R}^d) : Q(\nu) \leq 0\}, \quad (8.32)$$

where Q is the “quadratic form”

$$Q : \mu \in \mathcal{P}_p(\mathbb{R}^d) \mapsto \iint_{\mathbb{R}^d \times \mathbb{R}^d} \psi(x, y) \mu(dx) \mu(dy), \quad (8.33)$$

and $\psi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a prescribed function. For any $\mu \in \mathcal{P}(\mathbb{R}^d)$, we denote by

$$U_\mu^\psi : x \in \mathbb{R}^d \mapsto \int_{\mathbb{R}^d} \psi(x, y) \mu(dy)$$

the “potential” generated by μ for the interaction ψ , whenever this makes sense. We now make some assumptions on the interaction ψ for the functional Q to define an I -continuity set B in (8.32).

Assumption 8.14.

- Assumption 8.4 holds for some $q > 1$;
- There is $C_{\text{Lip}} > 0$ such that, for any $\mu \in \mathcal{P}_p(\mathbb{R}^d)$,

$$\|U_\mu^\psi\|_{\text{Lip}} \leq C_{\text{Lip}}, \quad (8.34)$$

and thus $\|U_\mu^\psi\|_{\infty, p} < +\infty$ for all $p \in (1, q)$;

- ψ is symmetric, i.e. $\psi(x, y) = \psi(y, x)$ for all $x, y \in \mathbb{R}^d$;
- Q is convex;
- there exists $\mu_- \in D_\mathcal{E}$ such that $Q(\mu_-) < 0$;
- there exists $\mu_+ \in D_\mathcal{E}$ such that $Q(\mu_+) > 0$.

Before turning to the minimization under constraint, let us present a sufficient condition of convolution nature for a function ψ to satisfy (8.34).

Proposition 8.15 (Sufficient condition for (8.34)). *Assume that $\psi(x, y) = \phi(x - y)$ for a function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfying $\|\phi\|_{\text{Lip}} < +\infty$ and thus $\|\phi\|_{\infty, p} < +\infty$. Then (8.34) holds with $C_{\text{Lip}} = \|\phi\|_{\text{Lip}}$.*

Proof of Proposition 8.15. For all $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ and $x, x' \in \mathbb{R}^d$, it holds

$$|U_\nu^\psi(x) - U_\nu^\psi(x')| \leq \int_{\mathbb{R}^d} |\phi(x - y) - \phi(x' - y)| \nu(dy) \leq \|\phi\|_{\text{Lip}} \int_{\mathbb{R}^d} |x - y - (x' - y)| \nu(dy) = \|\phi\|_{\text{Lip}} |x - x'|.$$

We thus obtain (8.34) with $C_{\text{Lip}} = \|\phi\|_{\text{Lip}}$. \square

In addition to the regularity ensured by Proposition 8.15, the convexity of Q is an important part of Assumption 8.14. Conditions for this convexity to hold for an interaction of the form $\psi(x, y) = \phi(x - y)$, which is related to Bochner-type positivity, are discussed at length in [31, 272, 84]. In particular, the choice $\psi = g$ where g is defined in (8.1) leads to a convex Q . Assumption 8.14 then provides a result similar to Theorem 8.9, now leading to a deformation of the interaction energy. The proof can be found in Section 8.5.4.

Theorem 8.16 (Quadratic constraint). *Let $\mu_\star \in \mathcal{P}(\mathbb{R}^d)$ be the unconstrained equilibrium measure defined in Proposition 8.5. Suppose that Assumption 8.14 holds for some $q > 1$ and $p \in (1, q)$ and let B be the set defined in (8.32). Then B is closed in the p -Wasserstein topology and*

$$\inf_{\tilde{B}} \mathcal{E} = \inf_B \mathcal{E} < +\infty. \tag{8.35}$$

Moreover

$$\mathcal{E}_B = \left\{ \mu \in \mathcal{P}(\mathbb{R}^d) : \mathcal{E}(\mu) = \inf_B \mathcal{E} \right\} = \{ \mu^\psi \}$$

where μ^ψ has compact support and is solution to, for some $\alpha \geq 0$,

$$\begin{cases} 2g * \mu^\psi + 2\alpha U_{\mu^\psi}^\psi + V = C_\psi, & \text{quasi - everywhere in } \text{supp}(\mu^\psi), \\ 2g * \mu^\psi + 2\alpha U_{\mu^\psi}^\psi + V \geq C_\psi, & \text{quasi - everywhere,} \end{cases} \tag{8.36}$$

with $C_\psi = \mathcal{E}(\mu^\psi)$. Finally, one of the two following conditions holds:

- $\mu_\star \in B$ and $\alpha = 0$;
- $\mu_\star \notin B$, in which case $Q(\mu^\psi) = 0$ and $\alpha > 0$.

The quadratic constraint leads to a change in the interaction contrarily to the linear situation, which led to a change of confinement. From Theorem 8.16, we obtain the following result.

Corollary 8.17 (From conditioning to interaction deformation). *Suppose that Assumption 8.14 holds for some $q > 1$ and $p \in (1, q)$. Let us consider a Coulomb gas $X_N = (X_{N,1}, \dots, X_{N,N}) \sim P_N$ as in (8.7), and $Y_N = (Y_{N,1}, \dots, Y_{N,N})$ with law given by*

$$Y_N \sim \text{Law} \left(X_N \left| \frac{1}{N^2} \sum_{i,j=1}^N \psi(X_{N,i} - X_{N,j}) \leq 0 \right. \right).$$

Let μ^ψ be the conditioned equilibrium measure as in Theorem 8.16. Then, almost surely, it holds

$$\lim_{N \rightarrow \infty} d_{W_p} \left(\frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}}, \mu^\psi \right) = 0, \tag{8.37}$$

regardless of the way we define the random variables Y_N in the same probability space.

Remark 8.18 (Higher order constraints, convexity, and regularity). *From the proof of Theorem 8.16, the Gibbs principle holds for a set B of the form (8.32) for Q convex and lower semicontinuous. However, we would not be able to say much on the solution in such an abstract setting. In particular, higher order statistics could be considered, leading to higher order convolutions, but checking the convexity of the associated functional may become cumbersome. By lack of applications in mind, we do not consider these higher order constraints.*

Remark 8.19 (Formula for constrained equilibrium measure under regularity assumptions). *Suppose that Assumption 8.14 holds and that V has Lipschitz continuous derivatives and that ψ is $C^2(\mathbb{R}^d)$. Then the conditioned equilibrium measure μ^ψ with appears in Theorem 8.16 and in Corollary 8.17 satisfies*

$$\begin{cases} c_d \mu^\psi - \alpha \int_{\mathbb{R}^d} \Delta \psi(\cdot, y) \mu^\psi(dy) = \frac{\Delta V}{2}, & \text{almost everywhere in } \text{supp}(\mu^\psi), \\ \mu^\psi = 0, & \text{almost everywhere outside } \text{supp}(\mu^\psi). \end{cases} \tag{8.38}$$

Indeed, (8.38) follows by applying the Laplacian on both sides of (8.36). Note that the expression (8.38) is not explicit as in the linear constraint case of Remark 8.13 because we are not able to invert the convolution associated to $\Delta \psi$ in general.

8.4 Numerical illustration

In this section we consider the problem of sampling from conditioned distributions of the form

$$\text{Law}(X_N \mid \xi_N(X_N) = 0), \quad (8.39)$$

where $\xi_N : (\mathbb{R}^d)^N \rightarrow \mathbb{R}^m$ for some $m \geq 1$, and X_N is distributed according to P_N defined in (8.7) for N fixed. We drop the index N on ξ_N in what follows to shorten the notation, and consider constraints taking values in \mathbb{R}^m for generality. Note that we consider equality rather than inequality constraints since we have seen in Sections 8.3.2 and 8.3.3 that inequality constraints are either satisfied by the equilibrium measure or saturated.

The first contribution of this section is to propose in Section 8.4.1 an algorithm for sampling from (8.39). In a second time, we present in Section 8.4.2 some numerical applications, where we illustrate the predictions of Sections 8.2 and 8.3. This is also the opportunity to explore conjectures which are not proved in the present chapter.

8.4.1 Description of the algorithm

The description of the constrained Hamiltonian Monte Carlo algorithm used for sampling, which builds upon the Hamiltonian Monte Carlo algorithm presented in Chapter 7, follows several steps. We first make precise the structure of the measure (8.39). We next introduce a constrained Langevin dynamics used for sampling, before giving the details of the numerical integration.

8.4.1.1 Dirac and Lebesgue measures on submanifolds

First, we describe more precisely the structure of the constrained measure (8.39) by introducing the submanifold \mathcal{M}_z associated with the z -level set of ξ for $z \in \mathbb{R}^m$, namely

$$\mathcal{M}_z = \{x \in (\mathbb{R}^d)^N : \xi(x) = z\}, \quad (8.40)$$

and we use the shorthand notation $\mathcal{M} = \mathcal{M}_0$. To define the conditioned measure, we use the following disintegration of (Lebesgue) measure formula: for any bounded continuous test function φ ,

$$\int_{(\mathbb{R}^d)^N} \varphi(x) dx = \int_{\mathbb{R}^m} \int_{\mathcal{M}_z} \varphi(x) \delta_{\xi(x)-z}(dx) dz. \quad (8.41)$$

Indeed, for any $z \in \mathbb{R}^m$, this defines the conditioned measure $\delta_{\xi(x)-z}(dx)$, see [302, Section 2.3.2]. Since P_N is given by (8.7), the constrained measure (8.39) can be written with the conditioned measure $\delta_{\xi(x)}(dx)$ associated with \mathcal{M} through, for any bounded continuous φ ,

$$\mathbf{E}[\varphi(X_N) \mid \xi_N(X_N) = 0] = \frac{1}{Z_N^\xi} \int_{\mathcal{M}} \varphi(x) e^{-\beta_N H_N(x)} \delta_{\xi(x)}(dx), \quad (8.42)$$

where Z_N^ξ is a normalizing constant. The measure of interest is therefore

$$P_N^\xi(dx) = \frac{e^{-\beta_N H_N(x)}}{Z_N^\xi} \delta_{\xi(x)}(dx). \quad (8.43)$$

In order to obtain a better understanding of (8.43), we relate the conditioned measure proportional to $\delta_{\xi(x)}(dx)$ to the Lebesgue measure induced on the submanifold \mathcal{M} by the canonical Euclidean scalar product, which we denote by $\sigma_{\mathcal{M}}(dx)$. We use to this end the co-area formula [7, 173, 302]. We denote by $\nabla \xi = (\nabla \xi_1, \dots, \nabla \xi_m) \in \mathbb{R}^{dN \times m}$ and introduce the Gram matrix:

$$G(x) = \nabla \xi(x)^T \nabla \xi(x) \in \mathbb{R}^{m \times m}, \quad (8.44)$$

where the superscript T denotes matrix transposition. In what follows, we assume that the Gram matrix (8.44) is non-degenerate in the sense that $G(x)$ is invertible for x in a neighborhood of \mathcal{M} (see [302, Proposition 2.1]).

Proposition 8.20. *The measures $\delta_{\xi(x)}(dx)$ and $\sigma_{\mathcal{M}}(dx)$ are related by*

$$\delta_{\xi(x)}(dx) = |\det G(x)|^{-\frac{1}{2}} \sigma_{\mathcal{M}}(dx). \quad (8.45)$$

In particular it holds

$$P_N^\xi(dx) = \frac{e^{-\beta_N H_N^\xi(x)}}{Z_N^\xi} \sigma_{\mathcal{M}}(dx), \quad (8.46)$$

where

$$H_N^\xi(x) = H_N(x) + U_N(x), \quad U_N(x) = -\frac{1}{2\beta_N} \log |\det G(x)|. \quad (8.47)$$

Remark 8.21 (Parametrization invariance). *It seems at first sight that the definition of the conditioned measure in (8.41) depends on the choice of parametrization of ξ , but it does not. To illustrate this point, we consider for simplicity that $m = 1$ and $\mathcal{M} = \{x \in (\mathbb{R}^d)^N : \xi(x) = 0\}$.*

First, the induced Lebesgue measure on \mathcal{M} does not depend on the parametrization of \mathcal{M} . Consider next a smooth function $F : \mathbb{R} \rightarrow \mathbb{R}$ such that $F(0) = 0$ and $F'(0) \neq 0$, and the change of parametrization $\mathcal{M} = \{x \in (\mathbb{R}^d)^N : F(\xi(x)) = 0\}$. The gradient of the constraint at $x \in (\mathbb{R}^d)^N$ is then $\nabla F(\xi(x)) = F'(\xi(x)) \nabla \xi(x)$. Since $\xi(x) = 0$ for $x \in \mathcal{M}$, the right hand side of (8.45) is changed only by a multiplicative factor $|F'(0)|^2 \neq 0$. Therefore, the conditioned probability measure (8.43) is left unchanged.

The aim is therefore to sample from (8.46), which is not an easy task except in very particular situations, like the one studied in Section 8.2. The attempts in [201, 200, 319] show that a naive approach is not efficient in general, since the conditioning event is rare. Actually, sampling probability distributions under constraint is a long standing problem in molecular dynamics and computational statistics. Concerning molecular simulation, one can be interested in fixing some degrees of freedom of a system like bond lengths, or the value of a so-called reaction coordinate, typically for free energy computations – we refer *e.g.* to [356, 106, 301] for more details. An example of application in computational statistics is for instance Approximate Bayesian Computations, see [403, 322].

For sampling measures on submanifolds, a naive penalization of the constraint is not a good idea in general, because it typically generates a very stiff dynamics whose numerical integration is difficult. Moreover, our problem is made harder by the singularity of the pair interaction in the Hamiltonian (8.8). We have seen in Chapter 7 that Hybrid Monte Carlo schemes (relying on a second order discretization of an underdamped Langevin dynamics with a Metropolis–Hastings acceptance rule) provide efficient methods for sampling such probability distributions. An issue when combining a Metropolis–Hastings rule with a projection on a submanifold is that reversibility may be lost, which introduces a bias. A recent strategy has been to introduce a reversibility check in addition to the standard acceptance–rejection rule, which makes the HMC scheme under constraint reversible [430, 303]. Note that [432] proposes an interesting alternative to the scheme used here, but which is not compatible with a Metropolis selection procedure in its current form. We thus present the algorithm as written in [303], with some simplifications and adaptations to our context, for which we introduce next the constrained Langevin dynamics.

8.4.1.2 Constrained Langevin dynamics

We define here an underdamped Langevin dynamics over the submanifold \mathcal{M} , whose invariant measure has a marginal in position which coincides with (8.46). We motivate using this dynamics by first considering the problem of sampling from the unconstrained measure P_N . For a given $\gamma > 0$, we define

$$\begin{cases} dX_t = Y_t dt, \\ dY_t = -\nabla H_N(X_t) dt - \gamma Y_t dt + \sqrt{\frac{2\gamma}{\beta_N}} dB_t, \end{cases} \quad (8.48)$$

where $(B_t)_{t \geq 0}$ is a dN -dimensional Wiener process. Like in Chapter 7, $(X_t)_{t \geq 0}$ stands for a position, while $(Y_t)_{t \geq 0}$ represents a momentum variable. Let us mention again that the long time convergence of the law of this process towards P_N (a difficult problem due to the singularity of the Hamiltonian) can be proved through Lyapunov function techniques [316] (see Chapter 3). In practice, the singularity of g also makes the numerical integration of (8.48) difficult, and a Metropolis–Hastings selection can be used to stabilize the numerical discretization [82]. The algorithm described below makes precise how to adapt this strategy to sample measures constrained to the submanifold \mathcal{M} .

Since we aim at sampling from (8.46), it is natural to consider the dynamics (8.48) with positions constrained to the submanifold (8.40), that is

$$\begin{cases} dX_t = Y_t dt, \\ dY_t = -\nabla H_N(X_t) dt - \gamma Y_t dt + \sqrt{\frac{2\gamma}{\beta_N}} dB_t + \nabla \xi(X_t) d\theta_t, \\ \xi(X_t) = 0, \end{cases} \quad (8.49)$$

where $(\theta_t)_{t \geq 0} \in \mathbb{R}^m$ is a Lagrange multiplier enforcing the dynamics to stay on \mathcal{M} . Let us emphasize that the position constraint induces a hidden constraint on the momenta in (8.49), which reads

$$\forall t \geq 0, \quad \nabla \xi(X_t)^T Y_t = 0.$$

The above relation is obtained by taking the derivative of $t \mapsto \xi(X_t)$ along the dynamics (8.49). This implies that momenta are orthogonal to the submanifold's zero level set, which is a natural geometric constraint [302, 303]. However, as hinted above, the dynamics (8.49) does not sample from the conditioned measure (8.39), as shown in the following proposition [302, 303].

Proposition 8.22 (Invariant measure). *The dynamics (8.49) has a unique invariant measure with marginal distribution in position given by*

$$\frac{e^{-\beta_N H_N(x)}}{\bar{Z}_N^\xi} \sigma_{\mathcal{M}}(dx),$$

where \bar{Z}_N^ξ is a normalization constant.

Although (8.49) does not sample from P_N^ξ , we have seen how to fix this problem. More precisely, Proposition 8.20 shows that the dynamics (8.49) run with the modified Hamiltonian H_N^ξ defined in (8.47) samples from P_N^ξ .

However, in practice, it may be preferable not to use the gradient of U_N , since it involves the Hessian of the constraint ξ and may be cumbersome to compute. Therefore, we will not run the dynamics (8.49) with the modified Hamiltonian H_N^ξ but with H_N , and perform some reweighting to correct for the bias arising from the factor $|\det G(x)|^{-\frac{1}{2}}$. As explained in Remark 8.27 below in the context of a HMC discretization, this ensures that we are sampling from the correct target distribution while only moderately increasing the rejection rate.

8.4.1.3 Discretization

In order to make a practical use of (8.49) combined with Proposition 8.20, we need to define a discretization scheme. We present below the strategy proposed by [303], which relies on a second order discretization of (8.49) with a Metropolis–Hastings selection and a reversibility check.

As discussed after (8.49), momenta are orthogonal to the level sets of the submanifold. We then define

$$\Pi_{\mathcal{M}^\perp} = \text{Id} - \nabla \xi(x) G^{-1}(x) \nabla \xi(x)^T \in (\mathbb{R}^d)^N, \quad (8.50)$$

whose action is to project the momentum orthogonally to the submanifold \mathcal{M} . We next define the RATTLE scheme, which is a second order discretization of the Hamiltonian part of (8.49). A particular feature of this numerical integration is to be reversible up to momentum reversal, at least for sufficiently small time steps. We assume in Algorithm 8.23 that this condition holds, and we refer to [215, Section VII.1.4] and [303] for more details.

Algorithm 8.23 (RATTLE). *Starting from a configuration (x_m, y_m) with $x_m \in \mathcal{M}$ and $\nabla \xi(x_m)^T y_m = 0$,*

1. $y_{m+\frac{1}{4}} = y_m - \frac{\Delta t}{2} \nabla H_N(x_m)$.
2. $x_{m+\frac{1}{2}} = x_m + \Delta t y_{m+\frac{1}{4}}$.
3. *Compute the Lagrange multiplier $\theta_m \in \mathbb{R}^m$ associated with $x_{m+\frac{1}{2}}$ to enforce the constraint, using Algorithm 8.25 below (if convergence has been reached).*

4. Project as $x_{m+1} = x_{m+\frac{1}{2}} + \nabla\xi(x_m)\theta_m$ and $y_{m+\frac{1}{2}} = y_{m+\frac{1}{4}} + \nabla\xi(x_m)\theta_m/\Delta t$.
5. $y_{m+\frac{3}{4}} = y_{m+\frac{1}{2}} - \frac{\Delta t}{2}\nabla H_N(x_{m+1})$.
6. $y_{m+1} = \Pi_{\mathcal{M}^\perp} y_{m+\frac{3}{4}}$ where the projector $\Pi_{\mathcal{M}^\perp}$ is defined in (8.50).

Finally, return (x_{m+1}, y_{m+1}) .

We can now present the algorithm used to sample the conditioned distribution by integrating (8.49), which runs as follows. First, the momenta y_m are updated to \tilde{y}_m according to the Ornstein–Uhlenbeck process in (8.49) projected orthogonally to the submanifold with $\Pi_{\mathcal{M}^\perp}$. Next, we evolve the configuration (x_m, \tilde{y}_m) with a RATTLE step, leading to $(\hat{x}_{m+1}, \hat{y}_{m+1})$. However, reversibility may be lost in the procedure for two reasons: either it is not possible to perform one step of RATTLE starting from $(\hat{x}_{m+1}, -\hat{y}_{m+1})$, or the image of $(\hat{x}_{m+1}, -\hat{y}_{m+1})$ differs from $(x_m, -\tilde{y}_m)$. In both cases, the RATTLE move is rejected, and the configuration is updated as $(x_m, -\tilde{y}_m)$ (mind the fact that momenta are reversed, which ensures the reversibility of the scheme with respect to momenta reversal). Finally, a Metropolis–Hastings acceptance rule corrects for the time step bias in the sampling. The full algorithm reads as follows [303].

Algorithm 8.24 (Constrained HMC with reversibility check). Fix $T > 0$, $\Delta t > 0$, $\gamma > 0$, $K_{\max} \geq 1$, $N_{\text{iter}} = \lceil T/\Delta t \rceil$ and choose an initial configuration (x_0, y_0) with $x_0 \in \mathcal{M}$ and $\nabla\xi(x_0)^T y_0 = 0$ (possibly obtained by projection). Set also thresholds $\varepsilon_{\text{rev}}, \varepsilon_{\mathcal{N}} > 0$, and define $\eta_{\Delta t} = e^{-\gamma\Delta t}$. For $m = 0, \dots, N_{\text{iter}} - 1$, run the following steps:

1. Resample the momenta as

$$\tilde{y}_m = \Pi_{\mathcal{M}^\perp} \left(\eta_{\Delta t} y_m + \sqrt{\frac{1 - \eta_{\Delta t}^2}{\beta_n}} G_m \right),$$

where G_m are independent d_n -dimensional standard Gaussian random variables.

2. Perform one RATTLE step with Algorithm 8.23 starting from (x_m, \tilde{y}_m) , providing $(\hat{x}_{m+1}, \hat{y}_{m+1})$ if the Newton algorithm with $K_{\max}, \varepsilon_{\mathcal{N}}$ has converged; otherwise set $(x_{m+1}, y_{m+1}) = (x_m, -\tilde{y}_m)$ and increment m .
3. Compute a RATTLE backward step from $(\hat{x}_{m+1}, -\hat{y}_{m+1})$, providing $(x_m^{\text{rev}}, y_m^{\text{rev}})$ if Newton with $K_{\max}, \varepsilon_{\mathcal{N}}$ has converged. If Newton algorithm has not converged or if $|x_m - x_m^{\text{rev}}| > \varepsilon_{\text{rev}}$, reject the move by setting $(x_{m+1}, y_{m+1}) = (x_m, -\tilde{y}_m)$ and increment m .
4. Compute the Metropolis–Hastings ratio

$$p_m = 1 \wedge \exp \left[-\beta_n \left(H_N^\xi(\hat{x}_{m+1}) + \frac{|\hat{y}_{m+1}|^2}{2} - H_N^\xi(x_m) - \frac{|\tilde{y}_m|^2}{2} \right) \right], \quad (8.51)$$

and set

$$(x_{m+1}, y_{m+1}) = \begin{cases} (\hat{x}_{m+1}, \hat{y}_{m+1}) & \text{with probability } p_m, \\ (x_m, -\tilde{y}_m) & \text{with probability } 1 - p_m. \end{cases}$$

A particularity of our implementation with respect to [303] is that we run the dynamics with the Hamiltonian H_N while the Metropolis–Hastings ratio (8.51) (step (4) in Algorithm 8.24) is computed with the modified Hamiltonian H_N^ξ . As pointed out in Remark 8.27 below, the modification induced by the correction term U_N in (8.47) is generally small. Therefore, considering H_N for the dynamics allows to avoid the computation of the Hessian of ξ , while the selection rule corrects for this small error.

In order for our description to be complete, we define how to project the position on \mathcal{M} (step (3) in Algorithm 8.23). We use for this a variant of Newton’s algorithm.

Algorithm 8.25 (Newton algorithm). Consider a tolerance threshold $\varepsilon_{\mathcal{N}} > 0$ and a maximal number of steps $K_{\max} \geq 1$. If one starts from an initial position $x^0 \notin \mathcal{M}$ and a Lagrange multiplier $\theta^0 = 0 \in \mathbb{R}^m$, the projection procedure reads as follows: while $k \leq K_{\max}$,

1. Compute $M_k = \nabla \xi(x^0)^T \nabla \xi(x^k) \in \mathbb{R}^{m \times m}$.
2. Set $\theta^{k+1} = \theta^k - M_k^{-1} \xi(x^k)$.
3. Define the new position $x^{k+1} = x^k + \nabla \xi(x^0)^T \theta^{k+1}$.
4. If $\max(|\theta^{k+1} - \theta^k|, |\xi(x^k)|) \leq \varepsilon_N$, the algorithm has converged, else go back to step (1).

If the algorithm has converged in $k \leq K_{\max}$ steps, return the value θ^k of the Lagrange multiplier.

As we mentioned before Algorithm 8.24, we consider a fixed direction $\nabla \xi(x^0)$ for projection. This is needed to preserve the reversibility property of the final algorithm [303]. The procedure works as soon as the matrix M_k defined in step (1) is indeed invertible at each step of the inner loop. This is the case in our situation since we take $m = 1$, and we refer to [303] for more details. We are now ready to use Algorithm 8.24 to sample from the constrained distribution, challenge the theoretical results of Sections 8.2 and 8.3 and explore conjectures.

Remark 8.26 (Rejection sources). *For a standard HMC scheme like in Chapter 7, rejection is only due to the Metropolis–Hastings selection (step (4) in Algorithm 8.24). Here, rejection can be due to the following reasons:*

- the Newton algorithm in Step (2) (forward move) has not converged;
- the Newton algorithm in Step (3) (backward or reversed move) has not converged;
- the reversibility check in Step (3) has failed;
- the Metropolis rule in Step (4) has rejected the step.

In any case, the first step resamples the momentum variable according to the Ornstein–Uhlenbeck process part in (8.48), and rejection comes with a reversal of momenta. Let us also mention that, when the ratio (8.51) is computed with H_N (up to an additive constant), the Metropolis rejection rate should scale as Δt^3 . This will be the case in our situation when ξ and φ are linear since in this case the additional term U_N in (8.47) is constant, and we numerically observe this rate of decay (see Figure 8.2).

Remark 8.27 (Correction term). *Proposition 8.20 shows that the Hamiltonian of the system must be modified in order for the constrained dynamics (8.49) to sample from the probability distribution (8.43). However, in Algorithm 8.24, we run the dynamics with H_N and perform the selection with H_N^ξ . This is motivated by the following scaling argument. Consider*

$$\xi(x) = \frac{1}{N} \sum_{i=1}^N \varphi(x_i)$$

for some real-valued smooth function φ , which corresponds to the linear constraint situation described in Section 8.3.2. In this case, the corrector term in (8.47) reads

$$U_N(x) = -\frac{1}{2\beta_N} \log \left(\sum_{i=1}^N |\nabla \varphi(x_i)| \right),$$

up to an additive constant. This means that the correction term in (8.47) scales like $O(\log(N)/N^2)$ when $\beta_N = \beta N^2$, whereas the remainder of the Hamiltonian is $O(1)$. As a result, the correction is much smaller than the Hamiltonian energy H_N , and we may neglect it in the dynamics. This allows to avoid computing the Hessian of the constraint ξ at the price of a small increase in the rejection rate.

8.4.2 Numerical results

8.4.2.1 Linear statistics with linear constraint: the influence of confinement

Since one motivation for our work was to study the trace constraint with quadratic confinement, as detailed in Section 8.2, we first consider the model presented in Theorem 8.1 with $d = 2$, $\beta_N = N^2$ and

$$\forall x \in \mathbb{R}^2, \quad \varphi(x) = c - x \cdot v, \quad v = \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \quad (8.52)$$

We run Algorithm 8.24 setting $N = 300$, $T = 10^6$, $\Delta t = 0.5$, $\gamma = 1$ and $\varepsilon_{\mathcal{N}} = \varepsilon_{\text{rev}} = 10^{-12}$ with $K_{\text{max}} = 20$. In all the simulations in dimension 2, the initial configuration is drawn uniformly over $[-1, 1]^2$. We first set $V(x) = 2|x|^2$ and $c = 1$, so that according to Theorem 8.1, the conditional law of the empirical measure μ_N under P_N with the constraint $\mu_N(\varphi) = 0$ should converge in the limit of large N towards a disk of radius $1/\sqrt{2}$, centered at $(1, 0)$ in \mathbb{R}^2 . The simulations presented in Figure 8.1 show a very good agreement with the expected result.




Figure 8.1 – Study of the quadratic confinement for $N = 300$ without constraint (left) and with the constraint (8.52) (right). We see that the constrained measure is a disk of radius $1/\sqrt{2}$ centered at $(1, 0)$.

In this simple case, the Hamiltonian in (8.47) is only modified by a constant, so we expect the Metropolis rejection rate (step (5) in Algorithm 8.24) to scale like $O(\Delta t^3)$ when $\Delta t \rightarrow 0$. In Figure 8.2, we plot this rate in log-log coordinates (setting here $N = 20$ to reduce the computation time). The slope is indeed close to 3, which confirms our expectation.




Figure 8.2 – Study of the rejection rate of the Metropolis–Hastings selection rule with $N = 20$ (step (5) in Algorithm 8.24) in log-log coordinate. The slope of the linear fit is about 2.9.

In order to show that the translation phenomenon is specific to the quadratic confinement, we first consider the case of a quartic confinement potential, namely $V(x) = |x|^4/2$ subject to the constraint (8.52) with $c = 0.5$. This choice for V together with φ defined in (8.52) satisfies Assumption 8.8, so that Theorem 8.9 applies. However, no analytic solution is a priori available because the rotational symmetry is lost. The unconstrained equilibrium measure in Figure 8.3 (left) shows a depletion of the density around $(0, 0)$. In Figure 8.3 (right), we observe that the shape of the distribution is significantly modified by the constraint, and does not possess any rotational invariance. As could have been expected, the particles close to the origin feel a weaker confinement, so the distribution is more concentrated near the outer edge.




Figure 8.3 – Study of the quartic confinement for $N = 300$ without constraint (left) and with the constraint (8.52) (right). The shape of the equilibrium measure is significantly distorted by the constraint.

Another interesting case is when the confinement is weaker than quadratic, *e.g.* $V(x) = \frac{4}{3}|x|^{\frac{3}{2}}$, for which Theorem 8.9 still applies. The results are shown in Figure 8.4, considering again the constraint (8.52) with $c = 0.5$. We observe that the shape of the distribution also significantly changes by spreading in the direction of the constraint. This can be interpreted as follows: since the confinement is stronger at the origin, the more likely way to observe a fluctuation of the barycenter (or less costly in terms of energy) is in this case to spread the distribution.

Quite interestingly, for both potentials the distribution obtained as $c \rightarrow +\infty$ seems to reach a limiting ellipsoidal shape, under an appropriate rescaling (figures not shown here). Studying more precisely these limiting shapes and the rate at which they appear is an interesting open problem.




Figure 8.4 – Study of the weak confinement for $N = 300$ without constraint (left) and with the constraint (8.52) (right). The constraint now spreads the equilibrium measure to the right.

8.4.2.2 Other constraints in dimension two

In order to illustrate the efficiency of our algorithm in situations richer than the linear constraint with a linear function φ , we now present two other cases. First, we keep a linear constraint with $V(x) = 2|x|^2$, but set

$$\forall x \in \mathbb{R}^2, \quad \varphi(x) = c - \frac{\cos(5x_1) + \cos(5x_2)}{2},$$

where x_1 and x_2 denote here the first and second coordinates of $x \in \mathbb{R}^2$. This choice is motivated by Remark 8.13: since the Laplacian of φ takes positive and negative values, we expect the particles to concentrate in some regions of \mathbb{R}^2 , possibly leading to a phase separation. Note also that, in order for the two last conditions in Assumption 8.8 to be satisfied, we need to choose $c \in (-1, 1)$. We set

again $N = 300$ but $\Delta t = 0.4$ to reduce the rejection rate. The other parameters are the same as in Section 8.4.2.1. We plot in Figure 8.5 the result of the simulation for $c = 0.2$ and $c = 0.5$. The particles concentrate in the regions where the cosines are higher, which seems to lead to a phase separation for the second value of c .




Figure 8.5 – Study of the cosine constraint for $N = 300$ with $c = 0.2$ (left) and $c = 0.5$ (right). A phase separation appears as the particles are constrained to stay in the local maxima of the cosines.

In order to illustrate the results of Section 8.3.3, we consider a quadratic constraint $\xi : (\mathbb{R}^d)^N \rightarrow \mathbb{R}$ of the form

$$\forall x \in (\mathbb{R}^d)^N, \quad \xi(x) = \frac{1}{N^2} \sum_{i,j=1}^N \psi(x_i, x_j), \quad (8.53)$$

with, for $x, y \in \mathbb{R}^2$,

$$\psi(x, y) = \phi(x - y), \quad \text{and} \quad \phi(x) = c - |x|. \quad (8.54)$$

A motivation for this choice is to modify the *rigidity* of the gas by constraining the particles to be closer or further one from another in average. In order to make this rigidity anisotropic, we also consider (8.54) with

$$\forall x \in \mathbb{R}^2, \quad \phi(x) = c - |x_1|. \quad (8.55)$$

The choice (8.55) modifies the rigidity only in one direction. For illustration we take $V(x) = |x|^4/2$, $N = 50$, $\Delta t = 0.5$, $T = 10^6$ (we take a lower number of particles because the constraint makes the dynamics quite stiff). We set $c = 1$ for (8.54) and $c = 0.5$ for (8.55), which forces the particles to move away from each other. These choices for ψ satisfy the conditions of Proposition 8.15, and the application Q defined in (8.33) can be proved to be convex, so Assumption 8.14 is satisfied and Theorem 8.16 applies. The distribution obtained for the constraint (8.54), presented in Figure 8.6 (left), shows that the more likely way for the particles to be repelled by the constraint induced by ψ is to move away from the center and concentrate on the edge, compared to Figure 8.3 (left). For the constraint (8.55), we clearly observe in Figure 8.6 (right) the effect of anisotropy.

8.4.2.3 A one dimensional example

We consider the Gaussian Unitary Ensemble (GUE), which is a degenerate two-dimensional Coulomb gas for which the particles are confined on the real axis. Its corresponds in a sense to (8.7) with $d = 1$, $V(x) = 2|x|^2$ but $g(x) = -\log|x|$, and $\beta_N = N^2$. It is known that the equilibrium measure is then the Wigner semi-circle law, and we refer for instance to [29] for a large deviations study. We can apply Theorem 8.1 for the linear constraint (8.52). In this case, the Wigner semi-circle law is indeed translated by a factor c (figure not shown here).

Next, in order to illustrate a case which is not covered by our analysis, we want to sample the spectrum of those matrices whose determinant is equal to ± 1 . In our context, this corresponds to the configurations $x \in (\mathbb{R}^d)^N$ with $\prod_{i=1}^N |x_i| = 1$. By taking the logarithm, this constraint is actually of the form (8.42) (by Remark 8.21 the conditioned probability measure (8.39) does not depend the parametrization) with

$$\forall x \in (\mathbb{R}^d)^N, \quad \xi(x) = \frac{1}{N} \sum_{i=1}^N \log|x_i|. \quad (8.56)$$




Figure 8.6 – Study of the quartic confinement for $N = 50$ with the quadratic statistics constraint (8.53)-(8.54) where $c = 1$ (left), and with the constraint (8.55) with $c = 0.5$ (right). This has to be compared to the unconstrained distribution in Figure 8.3 (left).

We plot in Figure 8.7 the distribution for $n = 300$, $T = 10^5$ and $\Delta t = 0.05$ for the unconstrained log-gas, and with the constraint (8.56) for $\Delta t = 0.01$ (starting with particles equally spaced over the interval $[-1, 1]$). We observe what looks like a symmetrized Marchenko–Pastur distribution. Actually Remark 8.13 suggests that the effective potential of the constrained distribution is $|\cdot|^2 - \alpha \log |\cdot|$ for some $\alpha > 0$, which is not that far from the Laguerre potential $|\cdot| - \alpha \log |\cdot|$.




Figure 8.7 – Study of the one dimensional log-gas for $N = 300$ without constraint (top) and with the constraint (8.52) (bottom). This corresponds to a deformation of the semi-circle distribution.

8.5 Proofs

8.5.1 Proof of Theorem 8.1

This section is devoted to the proof of Theorem 8.1. The following lemma is some sort of quantitative Wasserstein version of [55, Lemma C.1].

Lemma 8.28 (Translation). *Let $d \geq 1$, $p \geq 1$, $\mu_1, \mu_2 \in \mathcal{P}_p(\mathbb{R}^d)$ and $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}$ be a measurable function. Then for all $m_1, m_2 \in \mathbb{R}^d$,*

$$d_{W_p}^p(\mu_1 * \delta_{m_1}, \mu_2 * \delta_{m_2}) \leq 2^{p-1} |m_1 - m_2|^p + 2^{p-1} d_{W_p}^p(\mu_1, \mu_2).$$

Moreover, if $m_i = a \left(\int \varphi d\mu_i \right)$ for $i = 1, 2$ and $a \in \mathbb{R}^d$, then

$$d_{W_p}(\mu_1 * \delta_{m_1}, \mu_2 * \delta_{m_2}) \leq 2^{\frac{p-1}{p}} (1 + |a|^p \|\varphi\|_{\text{Lip}}^p)^{\frac{1}{p}} d_{W_p}(\mu_1, \mu_2).$$

Note that the right hand side is infinite if φ is not Lipschitz.

Proof. We have, by using the infimum formulation of the distance d_{W_p} ,

$$d_{W_p}^p(\mu_1 * \delta_{m_1}, \mu_2 * \delta_{m_2}) \leq 2^{p-1} |m_1 - m_2|^p + 2^{p-1} d_{W_p}^p(\mu_1, \mu_2),$$

where we used the convexity inequality $|u + v|^p \leq 2^{p-1}(|u|^p + |v|^p)$ valid for all $u, v \in \mathbb{R}^d$. Then

$$|m_1 - m_2| = \left| a \left(\int_{\mathbb{R}^d} \varphi d(\mu_1 - \mu_2) \right) \right| \leq |a| \|\varphi\|_{\text{Lip}} d_{W_1}(\mu_1, \mu_2) \leq |a| \|\varphi\|_{\text{Lip}} d_{W_p}(\mu_1, \mu_2),$$

which is the claimed estimate. □

The following lemma is a d -dimensional version of the factorization lemma in [88]. It expresses a non obvious independence between the center of mass and the shape of the cloud of particles distributed according to P_N . As noticed in [88], it reminds the structure of certain continuous spins systems such as in [81, 321].

Lemma 8.29 (Factorization). *Let us assume that the assumptions of Theorem 8.1 are satisfied, and define $u = (v, \dots, v) \in (\mathbb{R}^d)^N$. Let π and π^\perp be the orthogonal projections in $(\mathbb{R}^d)^N$ on the linear subspaces*

$$L = \mathbb{R}u \quad \text{and} \quad L^\perp = \{x \in (\mathbb{R}^d)^N : x \cdot u = 0\}.$$

Then, abridging X_N into X , the following properties hold:

- for all $x \in (\mathbb{R}^d)^N$, denoting $s(x) = \frac{x_1 + \dots + x_N}{N} \in \mathbb{R}^d$, we have

$$\begin{aligned} \pi(x) &= (s(x) \cdot v)u = ((s(x) \cdot v)v, \dots, (s(x) \cdot v)v), \\ \pi^\perp(x) &= x - \pi(x) = x - (s(x) \cdot v)u = (x_1 - (s(x) \cdot v)v, \dots, x_N - (s(x) \cdot v)v); \end{aligned}$$

- $\pi(X)$ and $\pi^\perp(X)$ are independent random vectors;
- $\pi(X)$ is Gaussian with law $\mathcal{N}\left(0, \frac{N}{2\beta_N}\right)u$, so that $s(X) \cdot v$ has law $\mathcal{N}\left(0, \frac{N}{2\beta_N}\right)$;
- $\pi^\perp(X)$ has law of density proportional to $x \in L^\perp \mapsto e^{-\beta_N H_N(x)}$ with respect to the trace of the Lebesgue measure on the linear subspace L^\perp of \mathbb{R}^{dN-1} .

Proof of Lemma 8.29. Since $|v| = 1$, we have $|u| = \sqrt{N}$, so the orthonormal projection on L reads, for $x \in (\mathbb{R}^d)^N$,

$$\pi(x) = \frac{x \cdot u}{|u|^2} u = \left(\frac{1}{N} \sum_{i=1}^N x_i \cdot v \right) u = (s(x) \cdot v)u.$$

The expression of π^\perp follows easily. For all $x \in (\mathbb{R}^d)^N$, from $x = \pi(x) + \pi^\perp(x)$ we get

$$|x|^2 = |\pi(x)|^2 + |\pi^\perp(x)|^2.$$

On the other hand, for all $i, j \in \{1, \dots, N\}$ it holds

$$\begin{aligned} x_i - x_j &= \pi(x)_i + \pi^\perp(x)_i - \pi(x)_j - \pi^\perp(x)_j \\ &= s(x) + \pi^\perp(x)_i - s(x) - \pi^\perp(x)_j \\ &= \pi^\perp(x)_i - \pi^\perp(x)_j. \end{aligned}$$

Since $V(x) = |x|^2$, it follows that for all $x = (x_1, \dots, x_N) \in (\mathbb{R}^d)^N$,

$$H_N(x) = \frac{1}{N}|x|^2 + \frac{1}{N^2} \sum_{i \neq j} g(x_i - x_j) = \frac{1}{N}|\pi(x)|^2 + H_N(\pi^\perp(x)).$$

Now, let u_1, \dots, u_{dN} be an orthogonal basis of $(\mathbb{R}^d)^N = \mathbb{R}^{dN}$ with $u_1 = u/\sqrt{N} \in L$. For all $x \in (\mathbb{R}^d)^N$ we write $x = \sum_{i=1}^{dN} t_i(x)u_i$. We have $\pi(x) = t_1(x)u_1 = (s(x) \cdot v)u$ and $\pi^\perp(x) = \sum_{i=2}^{dN} t_i(x)u_i$. Then, since $\beta_N = \beta N^2$ we have, for all bounded measurable $f : L \rightarrow \mathbb{R}$ and $g : L^\perp \rightarrow \mathbb{R}$,

$$\begin{aligned} \mathbb{E}\left[f(\pi(X))g(\pi^\perp(X))\right] &= Z^{-1} \int_{(\mathbb{R}^d)^N} f(\pi(x))g(\pi^\perp(x))e^{-\frac{\beta_N}{N}|\pi(x)|^2} e^{-\beta_N H_N(\pi^\perp(x))} dx_1 \dots dx_N \\ &= Z^{-1} \left(\int_{\mathbb{R}} f(t')e^{-\frac{\beta_N}{N}|t'|^2} dt' \right) \left(\int_{\mathbb{R}^{dN-1}} g(t'')e^{-\beta_N H_N(t'')} dt'' \right), \end{aligned}$$

where $t' = t_1 u_1$, $dt' = dt_1$, $t'' = \sum_{i=2}^{dN} t_i u_i$ and $dt'' = \prod_{i=2}^{dN} dt_i$. This concludes the proof of the last two points of the lemma. \square

We can now turn to the proof of Theorem 8.1.

Proof of Theorem 8.1. Thanks to Lemma 8.29, we have, denoting again by $u = (v, \dots, v) \in (\mathbb{R}^d)^N$,

$$\begin{aligned} \text{Law} \left(X_N \mid \frac{1}{N} \sum_{i=1}^N \varphi(X_i) = 0 \right) &= \text{Law} \left(X_N \mid \frac{1}{N} \sum_{i=1}^N X_{N,i} \cdot v = c \right) \\ &= \text{Law} \left(X_N \mid s(X_N) \cdot v = c \right) \\ &= \text{Law} \left(X_N \mid (s(X_N) \cdot v)u = cu \right) \\ &= \text{Law} \left(X_N \mid \pi(X_N) = cu \right) \\ &= \text{Law} (cu + \pi^\perp(X_N)) \\ &= \text{Law} (\tilde{X}_N), \end{aligned} \tag{8.57}$$

where $\tilde{X}_N = cu + \pi^\perp(X_N) = cu + X_N - \pi(X_N)$. We also have

$$\tilde{X}_N = \left((c - s(X_N) \cdot v)v + X_{N,1}, \dots, (c - s(X_N) \cdot v)v + X_{N,N} \right) \quad \text{where} \quad s(X_N) = \frac{X_{N,1} + \dots + X_{N,N}}{N}.$$

In other words (recall that $\varphi(x) = x \cdot v - c$)

$$\tilde{X}_N = \left(X_{N,1} - \frac{1}{N} \sum_{i=1}^N \varphi(X_{N,i})v, \dots, X_{N,N} - \frac{1}{N} \sum_{i=1}^N \varphi(X_{N,i})v \right),$$

so that

$$\tilde{\mu}_N = \frac{1}{N} \sum_{i=1}^N \delta_{\tilde{X}_{N,i}} = \mu_N * \delta_{m_N} \quad \text{where} \quad \mu_N = \frac{1}{N} \sum_{i=1}^N \delta_{X_{N,i}} \quad \text{and} \quad m_N = v \int \varphi d\mu_N.$$

Thanks to the assumptions on V and g we know that the equilibrium measure μ_\star is the uniform distribution on a ball of radius 1. Now we note that $\|\varphi\|_{\text{Lip}} \leq 1$ and $\int \varphi d\mu_\star = c$, so that by Lemma 8.28, denoting $\mu^\varphi = \delta_{cv} * \mu$, for all $p \geq 1$, there exists $c_p > 0$ with

$$d_{W_p}(\tilde{\mu}_N, \mu^\varphi) \leq c_p d_{W_p}(\mu_N, \mu_\star). \tag{8.58}$$

On the other hand, the large deviations principle [84, Proof of Theorem 1.1(4)] when $\beta_N/N \rightarrow +\infty$ gives, for all $\varepsilon > 0$,

$$\sum_N \mathbb{P}(\mathrm{d}_{\mathrm{BL}}(\mu_N, \mu_\star) \geq \varepsilon) < \infty \quad (8.59)$$

(alternatively we could use the concentration of measure [85, Theorem 1.5] and get the result for d_{W_1} as well). This summable convergence in probability towards a non-random limit, known as *complete convergence* [429], is equivalent, via Borel–Cantelli lemmas, to state that almost surely, $\lim_{N \rightarrow \infty} \mathrm{d}_{\mathrm{BL}}(\mu_N, \mu_\star) = 0$, regardless of the way we defined the random variables X_N and thus the random measures μ_N on the same probability space.

In order to upgrade the convergence from d_{BL} to d_{W_p} for all $p \geq 1$, we note that from [85, Theorem 1.12], there exists $r_0 > 0$ such that for all $r \geq r_0$,

$$\sum_N \mathbb{P}\left(\max_{1 \leq k \leq N} |X_{N,k}| \geq r\right) < \infty. \quad (8.60)$$

Now for all $p \geq 1$ and all $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ supported in the ball of \mathbb{R}^d of radius $r \geq 1$, we have

$$\mathrm{d}_{W_p}^p(\mu, \nu) \leq (2r)^{p-1} \mathrm{d}_{W_1}(\mu, \nu) \leq r(2r)^{p-1} \mathrm{d}_{\mathrm{BL}}(\mu, \nu).$$

Also, by combining (8.59) and (8.60), we obtain that for all $p \geq 1$ and all $\varepsilon > 0$,

$$\sum_N \mathbb{P}(\mathrm{d}_{W_p}(\mu_N, \mu_\star) \geq \varepsilon) < \infty. \quad (8.61)$$

By the Borel–Cantelli lemma, for all $p \geq 1$, almost surely, $\lim_{N \rightarrow \infty} \mathrm{d}_{W_p}(\mu_N, \mu_\star) = 0$, regardless of the way we define the random variables X_N on the same probability space. Finally, since $p \mapsto \mathrm{d}_{W_p}$ is monotonic in p , we can make the almost sure event valid for all p by taking the intersection of all the almost sure events obtained for integer values of p .

By combining (8.61) with (8.58), we obtain that for all $p \geq 1$ and $\varepsilon > 0$,

$$\sum_N \mathbb{P}(\mathrm{d}_{W_p}(\tilde{\mu}_N, \mu^\varphi) \geq \varepsilon) < \infty. \quad (8.62)$$

Now if Y_N is a random vector of $(\mathbb{R}^d)^N$ such that $\mathrm{Law}(Y_N) = \mathrm{Law}(X_N \mid \varphi(X_{N,1}) + \dots + \varphi(X_{N,N}) = 0)$ then, denoting by $\mu_N^Y = \frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}}$, using (8.57) and the fact that μ^φ is deterministic, we get

$$\mathrm{d}_{W_p}(\mu_N^Y, \mu^\varphi) \stackrel{\mathrm{d}}{=} \mathrm{d}_{W_p}(\tilde{\mu}_N, \mu^\varphi).$$

Therefore, from (8.62) we get, for all $p \geq 1$ and all $\varepsilon > 0$,

$$\sum_N \mathbb{P}(\mathrm{d}_{W_p}(\mu_N^Y, \mu^\varphi) \geq \varepsilon) < \infty.$$

By the Borel–Cantelli lemma, for all $p \geq 1$, almost surely, $\lim_{N \rightarrow \infty} \mathrm{d}_{W_p}(\mu_N^Y, \mu^\varphi) = 0$, regardless of the way we define the random variables Y_N on the same probability space. Finally, since d_{W_p} is monotonic in p , we can make the almost sure event valid for all p by taking the intersection of all the almost sure events obtained for integer values of p . \square

Note that the above proof relies crucially, via Lemma 8.29, on the quadratic nature of V . However the Coulomb nature of g is less crucial and the result should remain essentially valid provided that the convergence to the equilibrium measure holds, for instance at the level of generality of the assumptions of the large deviations principle in [84].

8.5.2 Proofs of Section 8.3.1

We start with the proof of the abstract Gibbs conditioning principle.

Proof of Proposition 8.3. The set \mathcal{I}_B defined in (8.21) is not empty because the infimum is finite by (8.20), B is closed and I has compact level sets (which also implies that I is lower semicontinuous),

so the infimum is attained at least for one measure. Moreover, \mathcal{I}_B is closed by lower semicontinuity of I . Now, since

$$\frac{1}{\beta_N} \log \mathbb{P} \left(Y_N \in A_\varepsilon \mid Y_N \in B \right) = \frac{1}{\beta_N} \log \mathbb{P} (Y_N \in A_\varepsilon \cap B) - \frac{1}{\beta_N} \log \mathbb{P} (Y_N \in B),$$

the result follows from an upper bound on $\mathbb{P}(A_\varepsilon \cap B)$ and a lower bound on $\mathbb{P}(B)$. The upper bound of the large deviations principle implies that

$$\overline{\lim}_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P} (Y_N \in A_\varepsilon \cap B) \leq - \inf_{A_\varepsilon \cap B} I. \tag{8.63}$$

Assume first that $\overline{A_\varepsilon \cap B} \neq \emptyset$. Since $A_\varepsilon = \{y \in \mathcal{Y}, d(y, \mathcal{I}_B) > \varepsilon\}$, the lower semi-continuity of I shows that there exists $c_\varepsilon > 0$ for which

$$\inf_{A_\varepsilon \cap B} I \geq c_\varepsilon + \inf_B I,$$

so that

$$\overline{\lim}_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P} (Y_N \in A_\varepsilon \cap B) \leq - \inf_B I - c_\varepsilon. \tag{8.64}$$

If $\overline{A_\varepsilon \cap B} = \emptyset$, the infimum in the right hand side of (8.63) is equal to $+\infty$ so that (8.64) still holds. The lower bound for the set B reads

$$\underline{\lim}_{N \rightarrow +\infty} \frac{1}{\beta_N} \log \mathbb{P} (Y_N \in B) \geq - \inf_{\dot{B}} I.$$

Since B satisfies (8.20), it holds

$$\overline{\lim}_{N \rightarrow +\infty} - \frac{1}{\beta_N} \log \mathbb{P} (Y_N \in B) \leq \inf_B I,$$

which, together with (8.64), leads to (8.22).

Finally, if $Y'_N \sim \text{Law}(Y_N \mid Y_N \in B)$ and if we define the Y'_N 's on the same probability space then, for all $\varepsilon > 0$, by the Borel–Cantelli lemma, $\sum_N \mathbb{P}(Y'_N \in A_\varepsilon) < \infty$ and thus, almost surely, $Y'_N \notin A_\varepsilon$ for large enough N . Since the set A_ε depends on $\varepsilon > 0$, by taking $\varepsilon \rightarrow 0$ with $\varepsilon \in \mathbb{Q}$, we obtain that almost surely, $\lim_{N \rightarrow \infty} d(Y'_N, \mathcal{I}_B) = 0$. \square

We next recall elements of proof for the properties of the rate function \mathcal{E} .

Proof of Proposition 8.5. Consider a probability measure $\mu \in D_{\mathcal{E}}$, so

$$\int_{\mathbb{R}^d} V(x) \mu(dx) + \iint_{\mathbb{R}^d \times \mathbb{R}^d} g(x-y) \mu(dx) \mu(dy) < +\infty.$$

Since V satisfies Assumption 8.4 and therefore beats g at infinity (in particular when $d = 2$), we have

$$\int_{\mathbb{R}^d} |x|^p \mu(dx) < +\infty,$$

for $1 < p < q$. Thus $D_{\mathcal{E}} \subset \mathcal{P}_p(\mathbb{R}^d)$.

In order to show the convexity of \mathcal{E} is it sufficient to study that of J defined in (8.17). When $d \geq 3$, this is a consequence of [84, Lemma 3.1]. In the case $d = 2$, convexity over $D_{\mathcal{E}}$ is also shown in [84, Section 3]. Note that convexity is in general due to a Bochner-type positivity of the interaction kernel.

Finally, the existence of a unique minimizer with compact support solving (8.24) follows from [84, Theorem 1.2] for $d \geq 3$ and [380, Chapter I, Theorem 1.3] for $d = 2$. Since the minimizer of \mathcal{E} over $\mathcal{P}(\mathbb{R}^d)$ has compact support, the three problems in (8.23) clearly coincide. \square

We finally present the proof of Theorem 8.7, which is a consequence of Proposition 8.3 and the Borel–Cantelli lemma.

Proof of Theorem 8.7. Under P_N , the empirical measure μ_N associated to X_N satisfies a LDP in the p -Wasserstein topology with good rate function \mathcal{E} . Since B is assumed to be a closed continuity set for the p -Wasserstein topology, the set \mathcal{E}_B defined in (8.26) is closed and non-empty by Proposition 8.3.

For simplicity we denote by

$$\mu_N^Y = \frac{1}{N} \sum_{i=1}^N \delta_{Y_{N,i}}$$

the empirical measure associated to Y_N , where $Y_N \sim \text{Law}(X_N \mid \mu_N \in B)$. For any $\varepsilon > 0$, we define the set A_ε as in Proposition 8.3. Then, there exists $c_\varepsilon > 0$ such that

$$\begin{aligned} \sum_N \mathbb{P}(\text{d}_{W_p}(\mu_N^Y, \mathcal{E}_B) > \varepsilon) &= \sum_N \mathbb{P}(\text{d}_{W_p}(\mu_N, \mathcal{E}_B) > \varepsilon \mid \mu_N \in B) \\ &= \sum_N \mathbb{P}(\mu_N \in A_\varepsilon \mid \mu_N \in B) \\ &\leq C \sum_N e^{-\beta_N c_\varepsilon} < +\infty, \end{aligned}$$

for some $C > 0$. Since $\beta_N \gg N$, the Borell–Cantelli lemma implies that

$$\lim_{N \rightarrow \infty} \text{d}_{W_p}(\mu_N^Y, \mathcal{E}_B) = 0,$$

almost surely in any probability space, which concludes the proof. \square

8.5.3 Proof of Theorem 8.9

The proof is decomposed into four steps. We first show that under Assumption 8.8, the set B is an I -continuity set for the electrostatic energy \mathcal{E} . Next, we show that any minimizer of \mathcal{E} over B has a compact support, and hence the minimizer is actually unique. The last two steps characterize the minimizer through (8.29).

Step 1: I -continuity. Let us first show that B is closed for the p -Wasserstein topology by showing that $B^c = \{\mu \in \mathcal{P}_p(\mathbb{R}^d) \mid \mu(\varphi) > 0\}$ is open. Take $\mu \in B^c$ and ν such that $\text{d}_{W_p}(\mu, \nu) \leq \varepsilon^{\frac{1}{p}}$ for some $\varepsilon > 0$. By definition of the p -Wasserstein distance it holds

$$\sup_{\substack{f \in L^1(\mu), g \in L^1(\nu) \\ f(x) \leq g(y) + |x-y|^p}} \left(\int_{\mathbb{R}^d} f d\mu - \int_{\mathbb{R}^d} g d\nu \right) \leq \varepsilon. \tag{8.65}$$

Since $\|\varphi\|_{\infty, p} < +\infty$, for any $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ it holds $\varphi \in L^1(\nu) \cap L^1(\mu)$. Moreover, $\|\varphi\|_{\text{Lip}} < +\infty$ and φ cannot be a constant function because this would contradict the existence of μ_\pm in Assumption 8.8, so $\|\varphi\|_{\text{Lip}} > 0$. As a result, for $|x - y| \geq 1$ we have

$$|\varphi(x) - \varphi(y)| \leq \|\varphi\|_{\text{Lip}} |x - y| \leq \|\varphi\|_{\text{Lip}} |x - y|^p.$$

Therefore, $\varphi/\|\varphi\|_{\text{Lip}}$ satisfies the inf-convolution condition in (8.65) and we may pick $f = g = \varphi/\|\varphi\|_{\text{Lip}}$ so that

$$\int_{\mathbb{R}^d} \frac{\varphi}{\|\varphi\|_{\text{Lip}}} d\mu - \int_{\mathbb{R}^d} \frac{\varphi}{\|\varphi\|_{\text{Lip}}} d\nu \leq \varepsilon,$$

which becomes

$$\nu(\varphi) \geq \mu(\varphi) - \varepsilon \|\varphi\|_{\text{Lip}} > 0 \quad \text{for } \varepsilon < \frac{\mu(\varphi)}{\|\varphi\|_{\text{Lip}}}.$$

As a result, B^c is open and B is closed for the p -Wasserstein topology.

We now prove that B is an I -continuity set, namely that (8.28) holds. By the same reasoning as above, the existence of $\mu_- \in D_\mathcal{E}$ such that $\mu_-(\varphi) < 0$ ensures that $\mu_- \in \dot{B}$ so

$$\inf_{\dot{B}} \mathcal{E} < +\infty.$$

Since B is closed and \mathcal{E} has compact level sets (in particular it is lower semicontinuous), there exists $\bar{\mu}$ such that

$$\mathcal{E}(\bar{\mu}) = \inf_B \mathcal{E}.$$

If $\bar{\mu}(\varphi) < 0$, it holds $\bar{\mu} \in \mathring{B}$ and the proof is complete. Thus we may assume that $\bar{\mu}(\varphi) = 0$ and, by considering a minimizing sequence, for any $\varepsilon > 0$ we may find $\mu_\varepsilon \in \mathring{B}$ such that

$$\mathcal{E}(\mu_\varepsilon) \leq \inf_{\mathring{B}} \mathcal{E} + \varepsilon. \quad (8.66)$$

For $t \in [0, 1]$, we introduce $\mu_t = t\mu_\varepsilon + (1-t)\bar{\mu} \in D_{\mathcal{E}}$. Since $\bar{\mu}(\varphi) = 0$ it holds $\mu_t \in \mathring{B}$ for any $t \in (0, 1]$. By convexity of \mathcal{E} on its domain we have

$$\mathcal{E}(\mu_t) \leq t\mathcal{E}(\mu_\varepsilon) + (1-t)\mathcal{E}(\bar{\mu}).$$

We now proceed by contradiction by assuming that $\mathcal{E}(\bar{\mu}) = \inf_{\mathring{B}} \mathcal{E} - \eta$ for some $\eta > 0$. Recalling (8.66), we have for some $\varepsilon > 0$ and any $t \in (0, 1]$

$$\mathcal{E}(\mu_t) \leq t \left(\inf_{\mathring{B}} \mathcal{E} + \varepsilon \right) + (1-t) \left(\inf_{\mathring{B}} \mathcal{E} - \eta \right) = \inf_{\mathring{B}} \mathcal{E} + t\varepsilon - (1-t)\eta.$$

Considering

$$t < \frac{\eta}{\varepsilon + \eta},$$

we obtain that $\mu_t \in \mathring{B}$ with

$$\mathcal{E}(\mu_t) < \inf_{\mathring{B}} \mathcal{E},$$

which is a contradiction. Therefore, (8.28) holds true.

Step 2: the minimizer is unique and has compact support. We now show that any minimizer μ^φ has a compact support, before turning to uniqueness. We detail the proof for $d \geq 3$ following [84] by highlighting the necessary modifications, and leave the proof for $d = 2$ to the reader (which is deduced from [380, Chapter I, Theorem 1.3]). We introduce

$$\zeta = \inf_B \mathcal{E}$$

and, for any compact K ,

$$\zeta_K = \inf_{B_K} \mathcal{E}, \quad \text{where } B_K = \{\mu \in B \mid \text{supp}(\mu) \subset K\}.$$

By Assumption 8.8, B_K is non empty for K large enough (consider μ_- restricted to K). By Assumption 8.4, for any constant C the set

$$K = \{x \in \mathbb{R}^d, V(x) \leq C\}$$

is compact. In all what follows, we assume that $V \geq 0$. Since V is lower bounded and defined up to a constant, there is no loss of generality.

Let us show that $\zeta = \zeta_K$ for C large enough. Since the minimizers on B and \mathring{B} coincide, we can consider a measure $\mu \in B$ such that $\mu(\varphi) < 0$ and $\mathcal{E}(\mu) \leq \zeta + 1$. If $\mu(K) = 1$, the measure has compact support and we are done, so we assume that $\mu(K) < 1$. The goal of the following computations is to build a measure $\mu_K \in B$ supported in K such that $\mathcal{E}(\mu_K) < \mathcal{E}(\mu)$; this contradiction will show that ζ and ζ_K are equal. Let us first show that $\mu(K) > 0$ for C large enough. Indeed,

$$\zeta + 1 \geq \mathcal{E}(\mu) = \underbrace{\int_K V d\mu}_{\geq 0} + \int_{K^c} V d\mu + \underbrace{J(\mu)}_{\geq 0} \geq C(1 - \mu(K)),$$

which shows that $\mu(K) > 0$ if $C > \zeta + 1$. We may therefore define the restriction

$$\mu_K(\cdot) = \frac{\mu(K \cap \cdot)}{\mu(K)}.$$

Since $\mu(K) < 1$, we define similarly μ_{K^c} . The measure μ then reads

$$\mu = \mu(K)\mu_K + (1 - \mu(K))\mu_{K^c}.$$

Moreover, we chose μ such that $\mu(\varphi) < 0$, so it holds $\mu_K \in B_K$ for C large enough. Using the positivity of V and J (since $d \geq 3$ it holds $g \geq 0$), and $\mu(K) < 1$, we obtain that

$$\begin{aligned} \mathcal{E}(\mu) &\geq \mu(K) \int_{\mathbb{R}^d} V d\mu_K + (1 - \mu(K)) \int_{\mathbb{R}^d} V d\mu_{K^c} + \mu(K)^2 J(\mu_K) \\ &\geq \mu(K)^2 J(\mu_K) + \mu(K)^2 \int_{\mathbb{R}^d} V d\mu_K + (1 - \mu(K))C \\ &\geq \mu(K)^2 \mathcal{E}(\mu_K) + (1 - \mu(K))C. \end{aligned}$$

Let us proceed by contradiction by assuming that $\mathcal{E}(\mu_K) \geq \mathcal{E}(\mu)$, which leads to

$$\mathcal{E}(\mu) \geq \mu(K)^2 \mathcal{E}(\mu) + (1 - \mu(K))C.$$

Since $\mathcal{E}(\mu) \leq \zeta + 1$ we obtain

$$(\zeta + 1)(1 - \mu(K)^2) \geq (1 - \mu(K))C.$$

Simplifying by $1 - \mu(K)$ we have

$$2(\zeta + 1) \geq C,$$

which is absurd for $C > 2(\zeta + 1)$. Since $\mu_K \in B$ for C large enough, this shows that ζ and ζ_K coincide and that any minimizer has compact support.

In the above proof, the only modification with respect to previous works (see for instance [84]) is to check that the restricted measure μ_K satisfies the constraint for C large enough. This is done by picking the measure μ close to the minimum and such that $\mu(\varphi) < 0$. The same strategy can be used in the situation where $d = 2$ by writing

$$\mathcal{E}(\mu) = \iint_{\mathbb{R}^d \times \mathbb{R}^d} \left(\frac{V(x) + V(y)}{2} - \log|x - y| \right) \mu(dx)\mu(dy),$$

and adapting [380, Chapter I, Theorem 1.3] since $V(x) + V(y)$ dominate $\log|x - y|$ at infinity by Assumption 8.4.

In order to show that the minimizer is unique, it suffices to notice that for two probability measures μ, ν with compact support, it holds $J(\nu - \mu) = 0$ if and only if $\mu = \nu$, see [284, Theorems 1.15 and 1.16].

Step 3: Lagrange multiplier. We now turn to a first step towards the expression of μ^φ involving a Lagrange multiplier α . We adapt the proof of [26, Theorem 3.1] by introducing the following subset of \mathbb{R}^2 :

$$R = \{(\mathcal{E}(\mu) - \mathcal{E}(\mu^\varphi) + a_0, \mu(\varphi) + a_1) : a_0 > 0, a_1 > 0, \mu \in D_{\mathcal{E}}\}. \quad (8.67)$$

Since \mathcal{E} is convex on its domain $D_{\mathcal{E}}$ (which is convex), R is a non void convex subset of \mathbb{R}^2 that does not contain $(0, 0)$ (recall also that $D_{\mathcal{E}} \subset \mathcal{P}_p(\mathbb{R}^d)$ so the constraint takes finite values). Separating R from $(0, 0)$ with a hyperplane (see [26, Corollary 1.41]), this ensures the existence of $(\alpha_0, \alpha_1) \in \mathbb{R}^2 \setminus \{(0, 0)\}$ such that, for any $\mu \in D_{\mathcal{E}}$ and $a_0, a_1 > 0$ it holds

$$\alpha_0(\mathcal{E}(\mu) - \mathcal{E}(\mu^\varphi) + a_0) + \alpha_1(\mu(\varphi) + a_1) \geq 0.$$

By taking $a_1 \rightarrow -\mu^\varphi(\varphi) \geq 0$ and $\mu = \mu^\varphi$ in the above equation, we obtain that $\alpha_0 \geq 0$. Then, choosing $\mu = \mu^\varphi$, $a_0 \rightarrow 0$ and $a_1 > -\mu^\varphi(\varphi) \geq 0$ we find $\alpha_1 \geq 0$. Taking $a_0, a_1 \rightarrow 0$, we obtain

$$\forall \mu \in D_{\mathcal{E}}, \quad \alpha_0 \mathcal{E}(\mu^\varphi) \leq \alpha_0 \mathcal{E}(\mu) + \alpha_1 \mu(\varphi). \quad (8.68)$$

We now prove that $\alpha_0 > 0$ by contradiction, using a kind of qualification of constraint argument. If $\alpha_0 = 0$, (8.68) becomes

$$\forall \mu \in D_{\mathcal{E}}, \quad 0 \leq \alpha_1 \mu(\varphi).$$

Since $\alpha_1 \neq 0$ in this case, the above equation contradicts Assumption 8.8 by taking $\mu = \mu_-$, so $\alpha_0 > 0$ and we may renormalize (8.68) into

$$\forall \mu \in D_{\mathcal{E}}, \quad \mathcal{E}(\mu^\varphi) \leq \mathcal{E}(\mu) + \alpha \mu(\varphi), \quad (8.69)$$

where we set $\alpha = \alpha_1/\alpha_0 \geq 0$.

Finally, we show that either $\mu^\varphi = \mu_\star$, in which case $\alpha = 0$, or $\mu^\varphi(\varphi) = 0$ and $\alpha > 0$. First, if $\mu_\star \in B$, μ_\star satisfies the constraint and we know from Proposition 8.5 that it solves (8.29) with $\alpha = 0$. Otherwise, it holds $\mu_\star(\varphi) > 0$. Assume that $\mu^\varphi(\varphi) < 0$ and take $\mu = \mu^\varphi$ in (8.69), so

$$\alpha\mu^\varphi(\varphi) \geq 0,$$

which implies $\alpha = 0$. Therefore, (8.69) implies that μ^φ is the global minimizer μ_\star , which is in contradiction with $\mu_\star(\varphi) = \mu^\varphi(\varphi) < 0$, so the minimizer actually saturates the constraint and $\alpha > 0$.

Step 4: potential equation. In order to derive the equation for μ^φ , we follow [84, Section 4] by introducing the modified potential and electrostatic energy, for $\mu \in \mathcal{P}(\mathbb{R}^d)$,

$$V_\alpha = V + \alpha\varphi, \quad \mathcal{E}_\alpha(\mu) = \int_{\mathbb{R}^d} V_\alpha(x)\mu(dx) + J(\mu).$$

Since the case when $\alpha = 0$ corresponds to no-conditioning and we already know the equation satisfied by the equilibrium measure in this case, we restrict our attention to the situation in which $\alpha > 0$ and $\mu^\varphi(\varphi) = 0$. We define next, for any $\mu \in D_\mathcal{E}$,

$$\forall t \in (0, 1), \quad \psi(t) = \mathcal{E}_\alpha((1-t)\mu^\varphi + t\mu).$$

Because of (8.69) and the convexity of ψ , it holds $\psi'(0) \geq 0$, so that

$$\begin{aligned} 0 \leq \psi'(0) &= \int_{\mathbb{R}^d} V_\alpha d(\mu - \mu^\varphi) + 2J(\mu^\varphi, \mu - \mu^\varphi) \\ &\leq \int_{\mathbb{R}^d} V_\alpha d\mu + 2J(\mu^\varphi, \mu) - \left(\int_{\mathbb{R}^d} V_\alpha d\mu^\varphi + 2J(\mu^\varphi, \mu^\varphi) \right) \\ &\leq \int_{\mathbb{R}^d} (V_\alpha + 2U_{\mu^\varphi})d\mu - C_\varphi, \end{aligned}$$

where we set $U_{\mu^\varphi} = \mu^\varphi * g$ and

$$C_\varphi = \int_{\mathbb{R}^d} V_\alpha d\mu^\varphi + 2J(\mu^\varphi, \mu^\varphi) = \int_{\mathbb{R}^d} (V + 2U_{\mu^\varphi}) d\mu,$$

since $\mu^\varphi(\varphi) = 0$. The above inequality may be rewritten as

$$\forall \mu \in D_\mathcal{E}, \quad \int_{\mathbb{R}^d} (V_\alpha + 2U_{\mu^\varphi} - C_\varphi) d\mu \geq 0,$$

which proves the second line of (8.29) (by definition of quasi-everywhere).

Let us now prove the first line in (8.29) by contradiction. Assume that there is $x \in \text{supp}(\mu^\varphi)$ such that $V_\alpha(x) + 2U_{\mu^\varphi}(x) > C_\varphi$. Since μ^φ has compact support, U_{μ^φ} is lower semi-continuous [284, page 59]. By lower semi-continuity of V , there exists a neighborhood \mathcal{U} of x and $\varepsilon > 0$ such that

$$\forall x \in \mathcal{U}, \quad V_\alpha(x) + 2U_{\mu^\varphi}(x) \geq C_\varphi + \varepsilon.$$

Integrating with respect to μ^φ and using $\mu^\varphi(\varphi) = 0$ leads to

$$\begin{aligned} C_\varphi &= \int_{\mathbb{R}^d} (V_\alpha + 2U_{\mu^\varphi}) d\mu^\varphi = \int_{\mathcal{U}} (V + 2U_{\mu^\varphi}) d\mu^\varphi + \int_{\mathbb{R}^d \setminus \mathcal{U}} (V + 2U_{\mu^\varphi}) d\mu^\varphi \\ &\geq (C_\varphi + \varepsilon)\mu^\varphi(\mathcal{U}) + \int_{\mathbb{R}^d \setminus \mathcal{U}} (V + 2U_{\mu^\varphi}) d\mu^\varphi. \end{aligned}$$

Since $V_\alpha + 2U_{\mu^\varphi} \geq C_\varphi$ quasi-everywhere and $\mu^\varphi \in D_\mathcal{E}$, the above inequality becomes

$$C_\varphi = \int_{\mathbb{R}^d} (V_\alpha + 2U_{\mu^\varphi}) d\mu^\varphi \geq C_\varphi + \varepsilon\mu^\varphi(\mathcal{U}).$$

We reach a contradiction by noting that $\mu^\varphi(\mathcal{U}) > 0$ since \mathcal{U} is a neighborhood of $x \in \text{supp}(\mu^\varphi)$ (and using the definition of the support), which proves the first line of (8.29).

8.5.4 Proof of Theorem 8.16

We outline the proof of Theorem 8.16, which follows the same lines as in the linear case.

Proof. We show below that B defined in (8.32) is closed for the p -Wasserstein topology under Assumption 8.14. For this, we show that B^c is open by picking $\mu \in \mathcal{P}_p(\mathbb{R}^d)$ such that $Q(\mu) > 0$ and using again that, for $\varepsilon > 0$ and $\nu \in \mathcal{P}_p(\mathbb{R}^d)$ such that $d_{W_p}(\mu, \nu) \leq \varepsilon^{\frac{1}{p}}$, it holds, by (8.15),

$$\sup_{\substack{f \in L^1(\mu), g \in L^1(\nu) \\ f(x) \leq g(y) + |x-y|^p}} \left(\int_{\mathbb{R}^d} f d\mu - \int_{\mathbb{R}^d} g d\nu \right) \leq \varepsilon. \quad (8.70)$$

First, we note that for any $\mu, \nu \in \mathcal{P}_p(\mathbb{R}^d)$ it holds $\|U_\mu^\psi\|_{\infty, p} < +\infty$ and $\|U_\nu^\psi\|_{\infty, p} < +\infty$. Therefore, $U_\mu^\psi \in L^1(\mu) \cap L^1(\nu)$ and $U_\nu^\psi \in L^1(\mu) \cap L^1(\nu)$ for any probability measures μ, ν with moments of order p . Next, by (8.34), it holds $\|U_\mu^\psi\|_{\text{Lip}} \leq C_{\text{Lip}}$ and $\|U_\nu^\psi\|_{\text{Lip}} \leq C_{\text{Lip}}$. We assume for now that these norms are non-zero, so we may first choose, $f = g = U_\mu^\psi / \|U_\mu^\psi\|_{\text{Lip}}$, which leads to

$$Q(\mu) - \iint_{\mathbb{R}^d \times \mathbb{R}^d} \psi(x, y) \nu(dx) \mu(dy) \leq \varepsilon \|U_\mu^\psi\|_{\text{Lip}}. \quad (8.71)$$

Symmetrically we take $f = g = -U_\nu^\psi / \|U_\nu^\psi\|_{\text{Lip}}$, which leads to

$$\iint_{\mathbb{R}^d \times \mathbb{R}^d} \psi(y, x) \nu(dx) \mu(dy) - Q(\nu) \leq \varepsilon \|U_\nu^\psi\|_{\text{Lip}}. \quad (8.72)$$

By summing (8.71) and (8.72) and using the symmetry of ψ , we obtain

$$Q(\nu) \geq Q(\mu) - \varepsilon (\|U_\nu^\psi\|_{\text{Lip}} + \|U_\mu^\psi\|_{\text{Lip}}) \geq Q(\mu) - 2C_{\text{Lip}}\varepsilon. \quad (8.73)$$

This shows that $Q(\nu) > 0$ for $\varepsilon < Q(\mu)/(2C_{\text{Lip}})$. To finish the argument, we consider the cases where the Lipschitz norm of the potentials generated by μ and ν may be zero. Suppose first that $\|U_\mu^\psi\|_{\text{Lip}} = 0$. This implies the existence of $c_\mu \in \mathbb{R}$ such that

$$\forall x \in \mathbb{R}^d, \quad \int_{\mathbb{R}^d} \psi(x, y) \mu(dy) = c_\mu. \quad (8.74)$$

Integrating the above equation with respect to μ shows that $c_\mu = Q(\mu) > 0$. As a result, if $\|U_\mu^\psi\|_{\text{Lip}} = 0$ and ν is such that $\|U_\nu^\psi\|_{\text{Lip}} > 0$ we can consider (8.72), which becomes (integrating (8.74) with respect to ν)

$$Q(\nu) \geq Q(\mu) - \varepsilon \|U_\nu^\psi\|_{\text{Lip}} \geq Q(\mu) - \varepsilon C_{\text{Lip}}.$$

In this case, $Q(\nu) > 0$ for $\varepsilon < Q(\mu)/C_{\text{Lip}}$. Then, if $\|U_\nu^\psi\|_{\text{Lip}} = 0$ it holds, for some $c_\nu \in \mathbb{R}$,

$$\forall x \in \mathbb{R}^d, \quad \int_{\mathbb{R}^d} \psi(x, y) \nu(dy) = c_\nu. \quad (8.75)$$

Integrating with respect to μ and using the symmetry of ψ we obtain that $c_\nu = c_\mu > 0$. Integrating next (8.75) with respect to ν shows that $Q(\nu) = c_\mu = Q(\mu) > 0$. Finally, if $\|U_\mu^\psi\|_{\text{Lip}} > 0$ but $\|U_\nu^\psi\|_{\text{Lip}} = 0$, (8.75) holds with $c_\nu = Q(\nu)$ so that (8.71) becomes

$$Q(\nu) \geq Q(\mu) - \varepsilon C_{\text{Lip}},$$

and the same conclusion follows. As a result, in any case the measures ν such that $d_{W_p}(\mu, \nu) \leq \varepsilon^{\frac{1}{p}}$ for $\varepsilon < Q(\mu)/(2C_{\text{Lip}})$ belong to B^c so that B^c is open and B is closed in the p -Wasserstein topology.

We next show that B is an I -continuity set. The existence of $\mu_- \in D_\mathcal{E}$ such that $Q(\mu_-) < 0$ ensures that

$$\inf_B \mathcal{E} < +\infty.$$

Since \mathcal{E} has compact level sets and B is closed, there exists $\bar{\mu}$ such that

$$\mathcal{E}(\bar{\mu}) = \inf_B \mathcal{E}.$$

If $Q(\bar{\mu}) < 0$, I -continuity is proven, so we may assume that $Q(\bar{\mu}) = 0$. Like in the linear case, we may take $\mu_\varepsilon \in \mathring{B}$ such that

$$\mathcal{E}(\mu_\varepsilon) \leq \inf_{\mathring{B}} \mathcal{E} + \varepsilon,$$

and consider the convex combination $\mu_t = t\mu_\varepsilon + (1-t)\bar{\mu}$ for $t \in (0, 1)$. The convexity of Q shows that, for any $t \in (0, 1)$ it holds

$$Q(\mu_t) \leq tQ(\mu_\varepsilon) + (1-t)Q(\bar{\mu}) < 0,$$

so that $\mu_t \in \mathring{B}$. Proceeding by contradiction by supposing that $\mathcal{E}(\bar{\mu}) < \inf_{\mathring{B}} \mathcal{E}$, we obtain that for $t > 0$ small enough it holds $\mu_t \in \mathring{B}$ and

$$\mathcal{E}(\mu_t) < \inf_{\mathring{B}} \mathcal{E},$$

which is a contradiction, proving that B is an I -continuity set.

One can next follow the step 2 of the proof of Theorem 8.9 to show that the minimizer μ^ψ is unique with compact support.

At this stage, the rest of Theorem 8.16 follows like the proof of Theorem 8.9. In particular, we can introduce a set similar to (8.67) by setting

$$R = \{(\mathcal{E}(\mu) - \mathcal{E}(\mu^\varphi) + a_0, Q(\mu) + a_1) : a_0 > 0, a_1 > 0, \mu \in D_\mathcal{E}\}.$$

The set R is convex by convexity of Q , so the same convex separation theorem can be used, and we can show that there exists $\alpha \geq 0$ such that

$$\forall \mu \in D_\mathcal{E}, \quad \mathcal{E}(\mu^\psi) \leq \mathcal{E}(\mu) + \alpha Q(\mu).$$

In this procedure, we use the existence of μ_\pm from Assumption 8.14 in order to reproduce the qualification of constraint argument. This leads to computations where the interaction energy J is replaced by

$$J_\alpha(\mu, \nu) = J(\mu, \nu) + \alpha \iint_{\mathbb{R}^d \times \mathbb{R}^d} \psi(x, y) \mu(dx) \nu(dy),$$

which leads to (8.36) by following the step 4 of the proof of Theorem 8.9. \square

We insist on the importance of the convexity of Q for the above proof to be valid.

- [1] A. Abdulle, D. Cohen, G. Vilmart, and K. C. Zygalakis. High weak order methods for stochastic differential equations based on modified equations. *SIAM J. Sci. Comput.*, 34(3):A1800–A1823, 2012.
- [2] A. Abdulle, G. Vilmart, and K. C. Zygalakis. High order numerical approximation of the invariant measure of ergodic SDEs. *SIAM J. Numer. Anal.*, 52(4):1600–1622, 2014.
- [3] A. Abdulle, G. Vilmart, and K. C. Zygalakis. Long time accuracy of Lie–Trotter splitting methods for Langevin dynamics. *SIAM J. Numer. Anal.*, 53(1):1–16, 2015.
- [4] S. Adelman and J. Doll. Generalized Langevin equation approach for atom/solid-surface scattering: General formulation for classical scattering off harmonic solids. *J. Chem. Phys.*, 64(6):2375–2388, 1976.
- [5] T. P. I. Ahamed, V. S. Borkar, and S. Juneja. Adaptive importance sampling technique for Markov chains using stochastic approximation. *Op. Res.*, 54(3):489–504, 2006.
- [6] R. Ahlswede and V. M. Blinovskiy. Large deviations in quantum information theory. *Probl. Inform. Transm.*, 39(4):373–379, 2003.
- [7] L. Ambrosio, N. Fusco, and D. Pallara. *Functions of Bounded Variation and Free Discontinuity Problems*, volume 254 of *Oxford Mathematical Monographs*. Clarendon Press Oxford, 2000.
- [8] A. Aminov, G. Bunin, and Y. Kafri. Singularities in large deviation functionals of bulk-driven transport models. *J. Stat. Mech.*, 2014(8):P08017, 2014.
- [9] G. W. Anderson, A. Guionnet, and O. Zeitouni. *An Introduction to Random Matrices*, volume 118 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010.
- [10] J. B. Anderson. A random-walk simulation of the Schrödinger equation: H+3. *J. Chem. Phys.*, 63(4):1499–1503, 1975.
- [11] L. Angeli, S. Grosskinsky, and A. M. Johansen. Limit theorems for cloning algorithms. *Preprint arXiv:1902.00509*, 2019.
- [12] L. Angeli, S. Grosskinsky, A. M. Johansen, and A. Pizzoferrato. Rare event simulation for stochastic dynamics in continuous time. *Preprint arXiv:1810.00693*, 2018.
- [13] D. Aristoff, T. Lelièvre, C. G. Mayne, and I. Teo. Adaptive multilevel splitting in molecular dynamics simulations. *ESAIM: Proc.*, 48:215–225, 2015.
- [14] S. N. Armstrong, S. Serfaty, and O. Zeitouni. Remarks on a constrained optimization problem for the Ginibre ensemble. *Potential Anal.*, 41(3):945–958, 2014.
- [15] S. Asmussen and P. W. Glynn. *Stochastic Simulation: Algorithms and Analysis*, volume 57 of *Stochastic Modelling and Applied Probability*. Springer Science & Business Media, 2007.

- [16] F. Augeri. Large deviations principle for the largest eigenvalue of Wigner matrices without Gaussian tails. *Electron. J. Probab.*, 21, 2016.
- [17] F. Augeri. On heavy-tail phenomena in some large deviations problems. *Preprint arXiv:1706.06184*, 2017.
- [18] Y. Baik, Y. Kafri, and V. Lecomte. Dynamical phase transitions in the current distribution of driven diffusive channels. *J. Phys. A*, 51(10):105001, 2018.
- [19] J. Baik, G. Ben Arous, and S. Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *Ann. Probab.*, 33(5):1643–1697, 2005.
- [20] D. Bakry, F. Barthe, P. Cattiaux, and A. Guillin. A simple proof of the Poincaré inequality for a large class of probability measures. *Electron. Commun. Probab.*, 13:60–66, 2008.
- [21] D. Bakry, I. Gentil, and M. Ledoux. *Analysis and Geometry of Markov Diffusion Operators*, volume 348 of *Grundlehren der mathematischen Wissenschaften*. Springer Science & Business Media, 2013.
- [22] R. Balian. *From Microphysics to Macrophysics: Methods and Applications of Statistical Physics. Volume I*, volume 2 of *Theoretical and Mathematical Physics*. Springer Science & Business Media, 2007.
- [23] V. Bansaye, B. Cloez, and P. Gabriel. Ergodic behavior of non-conservative semigroups via generalized Doeblin’s conditions. *Preprint arXiv:1710.05584*, 2017.
- [24] V. Bansaye, B. Cloez, P. Gabriel, and A. Marguet. A non-conservative Harris’ ergodic theorem. *Preprint arXiv:1903.03946*, 2019.
- [25] A. C. Barato and R. Chetrite. A formal view on level 2.5 large deviations and fluctuation relations. *J. Stat. Phys.*, 160(5):1154–1172, 2015.
- [26] V. Barbu and T. Precupanu. *Convexity and Optimization in Banach Spaces*, volume 10 of *Mathematics and its Applications*. Springer Science & Business Media, 2012.
- [27] R. Bardenet and A. Hardy. Monte Carlo with Determinantal Point Processes. *Preprint arXiv:1605.00361v1*, 2016.
- [28] A. Basu, T. Bhattacharyya, and V. S. Borkar. A learning algorithm for risk-sensitive cost. *Math. Oper. Res.*, 33(4):880–898, 2008.
- [29] G. Ben Arous and A. Guionnet. Large deviations for Wigner’s law and Voiculescu’s non-commutative entropy. *Probab. Theory Relat. Fields*, 108(4):517–542, 1997.
- [30] A. Benveniste, M. Métivier, and P. Priouret. *Adaptive Algorithms and Stochastic Approximations*, volume 22 of *Stochastic Modelling and Applied Probability*. Springer Science & Business Media, 2012.
- [31] C. Berg, J. P. R. Christensen, and P. Ressel. *Harmonic Analysis on Semigroups: Theory of Positive Definite and Related Functions*, volume 100 of *Graduate Texts in Mathematics*. Springer, 1984.
- [32] R. J. Berman. On large deviations for Gibbs measures, mean energy and Gamma-convergence. *Preprint arXiv:1610.08219v1*, 2016.
- [33] R. J. Berman. The Coulomb gas, potential theory and phase transitions. *Preprint arXiv:1811.10249*, 2018.
- [34] L. Bertini, A. De Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim. Macroscopic fluctuation theory. *Rev. Mod. Phys.*, 87(2):593, 2015.
- [35] L. Bertini, A. D. Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim. Fluctuations in stationary nonequilibrium states of irreversible processes. *Phys. Rev. Lett.*, 87(4):040601, 2001.
- [36] L. Bertini, A. D. Sole, D. Gabrielli, G. Jona-Lasinio, and C. Landim. Macroscopic fluctuation theory for stationary non-equilibrium states. *J. Stat. Phys.*, 107(3):635–675, 2002.

- [37] A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, and A. Stuart. Optimal tuning of the hybrid Monte Carlo algorithm. *Bernoulli*, 19(5A):1501–1534, 2013.
- [38] R. N. Bhattacharya. On the functional central limit theorem and the law of the iterated logarithm for Markov processes. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 60(2):185–201, 1982.
- [39] J. Bierkens, A. Bouchard-Côté, A. Doucet, A. B. Duncan, P. Fearnhead, T. Lienart, G. Roberts, and S. J. Vollmer. Piecewise deterministic Markov processes for scalable Monte Carlo on restricted domains. *Stat. Probab. Lett.*, 136:148–154, 2018.
- [40] J. Bierkens, V. Y. Chernyak, M. Chertkov, and H. J. Kappen. Linear PDEs and eigenvalue problems corresponding to ergodic stochastic optimization problems on compact manifolds. *J. Stat. Mech.*, 2016:013206, 2016.
- [41] J. Bierkens, P. Fearnhead, and G. Roberts. The zig-zag process and super-efficient sampling for Bayesian analysis of big data. *Ann. Stat.*, 47(3):1288–1320, 2019.
- [42] P. Billingsley. *Convergence of Probability Measures*, volume 9 of *Wiley Series in Probability and Statistics*. John Wiley & Sons, 2013.
- [43] J. Birrell and L. Rey-Bellet. Uncertainty quantification for Markov processes via variational principles and functional inequalities. *Preprint arXiv:1812.05174*, 2018.
- [44] I. Bjelaković, J.-D. Deuschel, T. Krüger, R. Seiler, R. Siegmund-Schultze, and A. Szkoła. A quantum version of Sanov’s theorem. *Comm. Math. Phys.*, 260(3):659–671, 2005.
- [45] X. Blanc and M. Lewin. The crystallization conjecture: a review. *EMS Surv. Math. Sci.*, 2(2):225–306, 2015.
- [46] J. Blanchet and H. Lam. State-dependent importance sampling for rare-event simulation: An overview and recent advances. *Surveys in Operations Research and Management Science*, 17(1):38–59, 2012.
- [47] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe. Variational inference: A review for statisticians. *J. Am. Stat. Assoc.*, 112(518):859–877, 2017.
- [48] T. Bodineau and B. Derrida. Distribution of current in nonequilibrium diffusive systems and phase transitions. *Phys. Rev. E*, 72(6):066110, 2005.
- [49] T. Bodineau and R. Lefevere. Large deviations of lattice Hamiltonian dynamics coupled to stochastic thermostats. *J. Stat. Phys.*, 133(1):1–27, 2008.
- [50] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler. Transition path sampling: Throwing ropes over rough mountain passes, in the dark. *Ann. Rev. Phys. Chem.*, 53(1):291–318, 2002.
- [51] F. Bolley, D. Chafaï, and J. Fontbona. Dynamics of a planar Coulomb gas. *Ann. Appl. Probab.*, 28(5):3152–3183, 2018.
- [52] F. Bolley and I. Gentil. Phi-entropy inequalities and Fokker-Planck equations. In *Progress in Analysis and its Applications*, pages 463–469. World Scientific, 2010.
- [53] F. Bolley, A. Guillin, and C. Villani. Quantitative concentration inequalities for empirical measures on non-compact spaces. *Probab. Theory Relat. Fields*, 137(3-4):541–593, 2007.
- [54] C. Bordenave and P. Caputo. A large deviations principle for Wigner matrices without Gaussian tails. *Ann. Probab.*, 42(6):2454–2496, 2014.
- [55] C. Bordenave, P. Caputo, and D. Chafaï. Circular law theorem for random Markov matrices. *Probab. Theory Relat. Fields*, 152(3-4):751–779, 2012.
- [56] C. Bordenave and D. Chafaï. Around the circular law. *Probab. Surv.*, 9, 2012.
- [57] V. S. Borkar. Learning algorithms for risk-sensitive control. In *Proc. 19th Int. Symp. Math. Theory Networks and Systems*, pages 1327–1332, 2010.

- [58] V. S. Borkar, S. Juneja, and A. A. Kherani. Performance analysis conditioned on rare events: An adaptive simulation scheme. *Commun. Info. Syst.*, 3(4):259–278, 2004.
- [59] A. A. Borovkov. Large deviation probabilities for random walks with semiexponential distributions. *Sib. Math. J.*, 41(6):1061–1093, 2000.
- [60] N. Bou-Rabee, A. Eberle, and R. Zimmer. Coupling and Convergence for Hamiltonian Monte Carlo. Preprint [arXiv:1805.00452v1](https://arxiv.org/abs/1805.00452v1), 2018.
- [61] N. Bou-Rabee and M. Hairer. Nonasymptotic mixing of the MALA algorithm. *IMA J. Numer. Anal.*, 33(1):80–110, 2012.
- [62] N. Bou-Rabee and H. Owhadi. Long-run accuracy of variational integrators in the stochastic context. *SIAM J. Numer. Anal.*, 48(1):278–297, 2010.
- [63] N. Bou-Rabee and J. M. Sanz-Serna. Geometric integrators and the Hamiltonian Monte Carlo method. Preprint [arXiv:1711.05337v1](https://arxiv.org/abs/1711.05337v1), 2017.
- [64] A. Bouchard-Côté, S. J. Vollmer, and A. Doucet. The bouncy particle sampler: A nonreversible rejection-free Markov chain Monte Carlo method. *J. Am. Stat. Assoc.*, 113(522):855–867, 2018.
- [65] J.-P. Bouchaud and M. Potters. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management*. Cambridge University Press, Cambridge, 2nd edition, 2000.
- [66] J.-P. Bouchaud and M. Potters. Financial applications of random matrix theory: a short review. Preprint [arXiv:0910.1205](https://arxiv.org/abs/0910.1205), 2009.
- [67] F. Bouchet and J. Reygner. Generalisation of the Eyring–Kramers transition rate formula to irreversible diffusion processes. *Ann. Henri Poincaré*, 17(12):3499–3532, 2016.
- [68] M. Boué and P. Dupuis. A variational representation for certain functionals of Brownian motion. *Ann. Probab.*, 26(4):1641–1659, 1998.
- [69] A. Bovier and F. Den Hollander. *Metastability: a Potential-Theoretic Approach*, volume 351 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2016.
- [70] G. E. Box and G. C. Tiao. *Bayesian Inference in Statistical Analysis*, volume 40 of *Wiley Classics Library*. John Wiley & Sons, 2011.
- [71] C.-E. Bréhier and T. Lelièvre. On a new class of score functions to estimate tail probabilities of some stochastic processes with Adaptive Multilevel Splitting. *Chaos*, 29(3):033126, 2019.
- [72] H. Brezis. *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Universitext. Springer Science & Business Media, 2010.
- [73] S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, editors. *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Handbooks of Modern Statistical Methods. CRC Press, Boca Raton, FL, 2011.
- [74] N. Brosse, A. Durmus, É. Moulines, and S. Sabanis. The tamed unadjusted Langevin algorithm. Preprint [arXiv:1710.05559v2](https://arxiv.org/abs/1710.05559v2), 2017.
- [75] J. A. Bucklew. *Introduction to Rare Event Simulation*. Springer Series in Statistics. Springer, New York, 2004.
- [76] R. Butez and D. García-Zelada. Extremal particles of two-dimensional Coulomb gases and random polynomials on a positive background. Preprint [arXiv:1811.12225](https://arxiv.org/abs/1811.12225), 2018.
- [77] D. M. Ceperley and B. Alder. Ground state of the electron gas by a stochastic method. *Phys. Rev. Lett.*, 45(7):566, 1980.
- [78] M. Ceriotti, G. Bussi, and M. Parrinello. Colored-noise thermostats à la carte. *J. Chem. Theory Comput.*, 6(4):1170–1180, 2010.

- [79] F. Cérou and A. Guyader. Adaptive multilevel splitting for rare event analysis. *Stoch. Anal. Appl.*, 25(2):417–443, 2007.
- [80] F. Cérou, A. Guyader, T. Lelièvre, and D. Pommier. A multiple replica approach to simulate reactive trajectories. *J. Chem. Phys.*, 134(5):054108, 2011.
- [81] D. Chafaï. Glauber versus Kawasaki for spectral gap and logarithmic Sobolev inequalities of some unbounded conservative spin systems. *Markov Process. Relat. Fields*, 9(3):341–362, 2003.
- [82] D. Chafaï and G. Ferré. Simulating Coulomb and log-gases with Hybrid Monte Carlo algorithms. *J. Statist. Phys.*, 174(3):692–714, 2018.
- [83] D. Chafaï, G. Ferré, and G. Stoltz. Coulomb gases under constraint: some theoretical and numerical results. *Preprint arxiv:1907.05803*, 2019.
- [84] D. Chafaï, N. Gozlan, and P.-A. Zitt. First-order global asymptotics for confined particles with singular pair repulsion. *Ann. Appl. Probab.*, 24(6):2371–2413, 2014.
- [85] D. Chafaï, A. Hardy, and M. Maïda. Concentration for Coulomb gases and Coulomb transport inequalities. *J. Func. Anal.*, 275(6):1447–1483, 2018.
- [86] D. Chafaï and S. Péché. A note on the second order universality at the edge of Coulomb gases on the plane. *J. Stat. Phys.*, 156(2):368–383, 2014.
- [87] D. Chafaï and E. Saff. Aspects of an Euclidean log-gas. work in progress, 2018.
- [88] D. Chafaï and J. Lehec. On Poincaré and logarithmic Sobolev inequalities for a class of singular Gibbs measures. *Geom. Funct. Anal.*, 2019. Lecture Notes in Mathematics 2256.
- [89] N. Champagnat and D. Villemonais. General criteria for the study of quasi-stationarity. *Preprint arXiv:1712.08092*, 2017.
- [90] N. Champagnat and D. Villemonais. Lyapunov criteria for uniform convergence of conditional distributions of absorbed Markov processes. *Preprint arXiv:1704.01928*, 2017.
- [91] N. Champagnat and D. Villemonais. Practical criteria for R-positive recurrence of unbounded semigroups. *Preprint arXiv:1904.08619*, 2019.
- [92] F. Chatelin. The spectral approximation of linear operators with applications to the computation of eigenelements of differential and integral operators. *SIAM Rev.*, 23(4):495–522, 1981.
- [93] F. Chatelin. *Spectral approximation of linear operators*. SIAM, 2011.
- [94] D. Chauveau and J. Diebolt. Estimation of the asymptotic variance in the CLT for Markov chains. *Stoch. Models*, 2003.
- [95] V. Y. Chernyak, M. Chertkov, J. Bierkens, and H. J. Kappen. Stochastic optimal control as non-equilibrium statistical mechanics: Calculus of variations over density and current. *J. Phys. A: Math. Theor.*, 47(2):022001, 2014.
- [96] R. Chetrite and H. Touchette. Nonequilibrium microcanonical and canonical ensembles and their equivalence. *Phys. Rev. Lett.*, 111:120601, 2013.
- [97] R. Chetrite and H. Touchette. Nonequilibrium Markov processes conditioned on large deviations. *Ann. Henri Poincaré*, 16(9):2005–2057, 2015.
- [98] R. Chetrite and H. Touchette. Variational and optimal control representations of conditioned and driven processes. *J. Stat. Mech. Theor. Exp.*, 2015(12):P12001, 2015.
- [99] G. Ciccotti and G. Jacucci. Direct computation of dynamical response by molecular dynamics: The mobility of a charged Lennard-Jones particle. *Phys. Rev. Lett.*, 35(12):789, 1975.
- [100] S. Ciliberto, S. Joubaud, and A. Petrosyan. Fluctuations in out-of-equilibrium systems: From theory to experiment. *J. Stat. Mech.*, 2010(12):P12003, 2010.

- [101] P. Collet, S. Martínez, and J. San Martín. *Quasi-stationary Distributions: Markov Chains, Diffusions and Dynamical Systems*. Probability and its Applications. Springer Science & Business Media, 2012.
- [102] R. Couillet and F. Benaych-Georges. Kernel spectral clustering of large dimensional data. *Electron. J. Stat.*, 10(1):1393–1454, 2016.
- [103] H. Cramér and H. Touchette. On a new limit theorem in probability theory (translation of ‘Sur un nouveau théorème-limite de la théorie des probabilités’). *Preprint arXiv:1802.05988*, 2018.
- [104] I. Csiszár. Sanov property, generalized I-projection and a conditional limit theorem. *Ann. Probab.*, 12(3):768–793, 1984.
- [105] A. Dalalyan and L. Riou-Durand. On sampling from a log-concave density using kinetic Langevin diffusions. Preprint [arXiv:1807.09382v1](https://arxiv.org/abs/1807.09382v1), 2018.
- [106] E. Darve. Thermodynamic integration using constrained and unconstrained dynamics. In C. Chipot and A. Pohorille, editors, *Free energy calculations*, volume 86 of *Springer Series in Chemical Physics*, pages 119–170. Springer, 2007.
- [107] R. Dautray and J.-L. Lions. *Mathematical Analysis and Numerical Methods for Science and Technology. Evolution Problems I*, volume 5. Springer, 2000.
- [108] D. S. Dean and S. N. Majumdar. Large deviations of extreme eigenvalues of random matrices. *Phys. Rev. Lett.*, 97(16):160201, 2006.
- [109] T. Dean and P. Dupuis. Splitting for rare event simulation: A large deviation approach to design and analysis. *Stoch. Proc. Appl.*, 119(2):562–587, 2009.
- [110] A. Debussche and E. Faou. Weak backward error analysis for SDEs. *SIAM J. Numer. Anal.*, 50(3):1735–1752, 2012.
- [111] L. Decreusefond, I. Flint, and A. Vergne. A note on the simulation of the Ginibre point process. *J. Appl. Probab.*, 52(4):1003–1012, 2015.
- [112] K. Deimling. *Nonlinear Functional Analysis*. Courier Corporation, 2010.
- [113] P. Del Moral. *Feynman–Kac Formulae*. Probability and its Applications. Springer, 2004.
- [114] P. Del Moral and A. Guionnet. On the stability of interacting processes with applications to filtering and genetic algorithms. *Annales de l’IHP Probabilités et statistiques*, 37(2):155–194, 2001.
- [115] P. Del Moral and L. Miclo. Branching and interacting particle systems approximations of Feynman–Kac formulae with applications to non-linear filtering. In *Séminaire de probabilités XXXIV*, pages 1–145. Springer, 2000.
- [116] P. Del Moral and L. Miclo. On the stability of nonlinear Feynman–Kac semigroups. In *Annales de la Faculté des Sciences Toulouse Mathématiques*, volume 11, pages 135–175. Université Paul Sabatier, 2002.
- [117] P. Del Moral and L. Miclo. Particle approximations of Lyapunov exponents connected to Schrödinger operators and Feynman–Kac semigroups. *ESAIM: Probab. Stat.*, 7:171–208, 2003.
- [118] C. Dellago, P. G. Bolhuis, and P. L. Geissler. Transition path sampling. *Adv. Chem. Phys.*, 123:1–78, 2003.
- [119] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*, volume 38 of *Stochastic Modelling and Applied Probability*. Springer-Verlag, Berlin, 2010. Corrected reprint of the second (1998) edition.
- [120] J. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [121] F. den Hollander. *Large Deviations*, volume 14 of *Fields Institute Monographs*. American Mathematical Society, 2000.

- [122] B. Derrida. Non-equilibrium steady states: Fluctuations and large deviations of the density and of the current. *J. Stat. Mech.*, 2007(07):P07023, 2007.
- [123] B. Derrida, J. L. Lebowitz, and E. R. Speer. Large deviation of the density profile in the steady state of the open symmetric simple exclusion process. *J. Stat. Phys.*, 107(3):599–634, 2002.
- [124] J.-D. Deuschel and D. W. Stroock. *Large Deviations*, volume 137 of *Pure and Applied Mathematics*. Academic Press, 1989.
- [125] G. Di Gesù, T. Lelièvre, D. Le Peutrec, and B. Nectoux. Jump Markov models and transition state theory: the quasi-stationary distribution approach. *Farad. Disc.*, 195:469–495, 2017.
- [126] I. Dinwoodie. Identifying a large deviation rate function. *Ann. Probab.*, 21(1):216–231, 1993.
- [127] R. L. Dobrushin. Central limit theorem for nonstationary Markov chains. I. *Theory Probab. Appl.*, 1(1):65–80, 1956.
- [128] R. L. Dobrushin. Central limit theorem for nonstationary Markov chains. II. *Theory Probab. Appl.*, 1(4):329–383, 1956.
- [129] J. Dolbeault, C. Mouhot, and C. Schmeiser. Hypocoercivity for linear kinetic equations conserving mass. *Trans. Amer. Math. Soc.*, 367(6):3807–3828, 2015.
- [130] M. Donsker and M. Kac. A sampling method for determining the lowest eigenvalue and the principal eigenfunction of Schrödinger’s equation. *J. Res. Natl. Bur. Stand.*, 44(50):551–557, 1950.
- [131] M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. I. *Comm. Pure Appl. Math.*, 28:1–47, 1975.
- [132] M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. II. *Comm. Pure Appl. Math.*, 28:279–301, 1975.
- [133] M. D. Donsker and S. R. S. Varadhan. On a variational formula for the principal eigenvalue for operators with maximum principle. *Proc. Natl. Acad. Sci.*, 72(3):780–783, 1975.
- [134] M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. III. *Comm. Pure Appl. Math.*, 29(4):389–461, 1976.
- [135] M. D. Donsker and S. R. S. Varadhan. Asymptotic evaluation of certain Markov process expectations for large time. IV. *Comm. Pure Appl. Math.*, 36(2):183–212, 1983.
- [136] J. L. Doob. Conditional Brownian motion and the boundary limits of harmonic functions. *Bulletin de la Société Mathématique de France*, 85:431–458, 1957.
- [137] R. Douc and E. Moulines. Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *ESAIM: Proceedings*, 19:101–107, 2007.
- [138] R. Douc, E. Moulines, and J. Olsson. Long-term stability of sequential Monte Carlo methods under verifiable conditions. *Ann. Appl. Probab.*, 24(5):1767–1802, 2014.
- [139] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer, 2001.
- [140] A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of nonlinear filtering*, 12(656-704):3, 2009.
- [141] Y. Du. *Order Structure and Topological Methods in Nonlinear Partial Differential Equations*, volume 1 of *Partial Differential Equations and Applications*. World Scientific, 2006.
- [142] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth. Hybrid Monte Carlo. *Phys. Lett. B*, 195(2):216–222, 1987.
- [143] G. Dubach. Powers of Ginibre eigenvalues. *Electron. J. Probab.*, 23, 2018.

- [144] L. Dumaz and C. Labbé. Localization of the continuous Anderson Hamiltonian in 1-D. *Probab. Theory Relat. Fields*, pages 1–67, 2017.
- [145] I. Dumitriu and A. Edelman. Matrix models for beta ensembles. *J. Math. Phys.*, 43(11):5830–5847, 2002.
- [146] A. B. Duncan, T. Lelièvre, and G. Pavliotis. Variance reduction using nonreversible Langevin samplers. *J. Stat. Phys.*, 163(3):457–491, 2016.
- [147] P. Dupuis and R. S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*, volume 902 of *Wiley Series in Probability and Statistics*. John Wiley & Sons, 2011.
- [148] P. Dupuis, V. Laschos, and K. Ramanan. Large deviations for empirical measures generated by Gibbs measures with singular energy functionals. Preprint [arXiv:1511.06928v1](https://arxiv.org/abs/1511.06928v1), 2015.
- [149] P. Dupuis and D. Lipshutz. Large deviations for the empirical measure of a diffusion via weak convergence methods. *Stoch. Process. Appl.*, 128(8):2581–2604, 2018.
- [150] P. Dupuis and H. Wang. Dynamic importance sampling for uniformly recurrent Markov chains. *Ann. Appl. Prob.*, 15(1A):1–38, 2005.
- [151] A. Durmus, A. Guillin, and P. Monmarché. Geometric ergodicity of the bouncy particle sampler. Preprint [arXiv:1807.05401](https://arxiv.org/abs/1807.05401), 2018.
- [152] A. Durmus, A. Guillin, and P. Monmarché. Piecewise deterministic Markov processes and their invariant measure. Preprint [arXiv:1807.05421](https://arxiv.org/abs/1807.05421), 2018.
- [153] A. Durmus, E. Moulines, and E. Saksman. On the convergence of Hamiltonian Monte Carlo. Preprint [arXiv:1705.00166v1](https://arxiv.org/abs/1705.00166v1), 2017.
- [154] W. E, J. Han, and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Comm. Math. Stat.*, 5(4):349–380, 2017.
- [155] W. E, W. Ren, and E. Vanden-Eijnden. String method for the study of rare events. *Phys. Rev. B*, 66(5):052301, Aug. 2002.
- [156] W. E, W. Ren, and E. Vanden-Eijnden. Minimum action method for the study of rare events. *Comm. Pure App. Math.*, 57(5):637–656, May 2004.
- [157] W. E, W. Ren, and E. Vanden-Eijnden. Finite temperature string method for the study of rare events. *J. Phys. Chem. B*, 109(14):6688–6693, 2005.
- [158] W. E, W. Ren, and E. Vanden-Eijnden. Simplified and improved string method for computing the minimum energy paths in barrier-crossing events. *J. Chem. Phys.*, 126:164103, 2007.
- [159] W. E and E. Vanden Eijnden. Towards a theory of transition paths. *J. Stat. Phys.*, 123:503–523, 2006.
- [160] A. Eberle, A. Guillin, and R. Zimmer. Quantitative Harris-type theorems for diffusions and McKean–Vlasov processes. *Trans. Amer. Math. Soc.*, 371(10):7135–7173, 2019.
- [161] J.-P. Eckmann and M. Hairer. Spectral properties of hypoelliptic operators. *Commun. Math. Phys.*, 235(2):233–253, 2003.
- [162] A. Edelman and N. R. Rao. Random matrix theory. *Acta Numer.*, 14:233–297, 2005.
- [163] N. El Karoui. The spectrum of kernel random matrices. *Ann. Stat.*, 38(1):1–50, 2010.
- [164] M. El Makrini, B. Jourdain, and T. Lelièvre. Diffusion Monte Carlo method: Numerical analysis in a simple case. *ESAIM: Math. Model. Numer. Anal.*, 41(2):189–213, 2007.
- [165] R. S. Ellis. Large deviations for a general class of random vectors. *Ann. Probab.*, 12(1):1–12, 1984.

- [166] R. S. Ellis. *Entropy, Large Deviations and Statistical Mechanics*, volume 271 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2007.
- [167] C. Enaud and B. Derrida. Large deviation functional of the weakly asymmetric exclusion process. *J. Stat. Phys.*, 114(3):537–562, 2004.
- [168] R. Erban and H. G. Othmer. From individual to collective behavior in bacterial chemotaxis. *SIAM J. Appl. Math.*, 65(2):361–391, 2004.
- [169] L. Erdős and H.-T. Yau. *A Dynamical Approach to Random Matrix Theory*, volume 28 of *Courant Lecture Notes in Mathematics*. American Mathematical Society, Providence, RI, 2017.
- [170] C. P. Espigares, P. L. Garrido, and P. I. Hurtado. Dynamical phase transition for current statistics in a simple driven diffusive system. *Phys. Rev. E*, 87:032115, 2013.
- [171] D. J. Evans and G. Morriss. *Statistical Mechanics of Nonequilibrium Liquids*. Cambridge University Press, 2008.
- [172] J. Evans. Hypocoercivity in Phi-entropy for the linear relaxation Boltzmann equation on the torus. *Preprint arXiv:1702.04168*, 2017.
- [173] L. Evans and R. Gariepy. *Measure Theory and Fine Properties of Functions*, volume 5 of *Studies in Advanced Mathematics*. CRC Press, 2018.
- [174] L. C. Evans. *Partial Differential Equations*, volume 19 of *Graduate Studies in Mathematics*. American Mathematical Society, 2010.
- [175] Z. F. Ezawa. *Quantum Hall effects*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, second edition, 2008.
- [176] M. Fathi, A.-A. Homman, and G. Stoltz. Error analysis of the transport properties of Metropolized schemes. *ESAIM: Proceedings and Surveys*, 48:341–363, 2015.
- [177] M. Fathi and G. Stoltz. Improving dynamical properties of metropolized discretizations of overdamped Langevin dynamics. *Numer. Math.*, 136(2):1–58, 2015.
- [178] J. Feng and T. G. Kurtz. *Large Deviations for Stochastic Processes*, volume 131 of *Mathematical Surveys and Monographs*. American Mathematical Soc., 2006.
- [179] G. Ferré, T. Haut, and K. Barros. Learning molecular energies using localized graph kernels. *J. Chem. Phys.*, 146(11):114107, 2017.
- [180] G. Ferré, J.-B. Maillet, and G. Stoltz. Permutation-invariant distance between atomic configurations. *J. Chem. Phys.*, 143(10):104114, 2015.
- [181] G. Ferré, M. Rousset, and G. Stoltz. More on the stability of Feynman–Kac semigroups. *Preprint arXiv:1807.00390*, 2018.
- [182] G. Ferré and G. Stoltz. Error estimates on ergodic properties of Feynman–Kac semigroups. *Numer. Math.*, 143(2):261–313, 2019.
- [183] G. Ferré and G. Stoltz. Large deviations of empirical measures of diffusions in fine topologies. *Preprint arXiv:1906.0941*, 2019.
- [184] G. Ferré and H. Touchette. Adaptive sampling of large deviations. *J. Stat. Phys.*, 172(6):1525–1544, 2018.
- [185] G. Ferré, E. Vanden-Eijnden, and T. Gräfke. Approximate controls for small temperature systems. *In preparation*, 2019.
- [186] W. H. Fleming. Exit probabilities and optimal stochastic control. *Appl. Math. Optim.*, 4(1):329–346, 1977.
- [187] W. H. Fleming and H. M. Soner. *Controlled Markov Processes and Viscosity Solutions*, volume 25 of *Stochastic Modelling and Applied Probability*. Springer, New York, 2006.

- [188] P. J. Forrester. *Log-gases and Random Matrices*, volume 34 of *London Mathematical Society Monographs Series*. Princeton University Press, Princeton, NJ, 2010.
- [189] P. J. Forrester. Analogies between random matrix ensembles and the one-component plasma in two-dimensions. *Nuclear Phys. B*, 904:253–281, 2016.
- [190] W. Foulkes, L. Mitas, R. Needs, and G. Rajagopal. Quantum Monte Carlo simulations of solids. *Rev. Mod. Phys.*, 73(1):33, 2001.
- [191] M. Freidlin and M. Weber. Random perturbations of dynamical systems and diffusion processes with conservation laws. *Probab. Theory Relat. Fields*, 128(3):441–466, 2004.
- [192] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Hamiltonian Systems*, volume 523 of *Grundlehren der mathematischen Wissenschaften*. American Mathematical Soc., 1994.
- [193] M. I. Freidlin and A. D. Wentzell. *Random Perturbations of Dynamical Systems*, volume 260 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1998.
- [194] G. Gallavotti and E. G. Cohen. Dynamical ensembles in nonequilibrium statistical mechanics. *Phys. Rev. Lett.*, 74(14):2694, 1995.
- [195] F. Gao, A. Guillin, and L. Wu. Bernstein-type concentration inequalities for symmetric Markov processes. *Theory Probab. Appl.*, 58(3):358–382, 2014.
- [196] D. García-Zelada. A large deviation principle for empirical measures on Polish spaces: Application to singular Gibbs measures on manifolds. In *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, volume 55, pages 1377–1401. Institut Henri Poincaré, 2019.
- [197] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland. Dynamical first-order phase transition in kinetically constrained models of glasses. *Phys. Rev. Lett.*, 98(19):195702, 2007.
- [198] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland. First-order dynamical phase transition in models of glasses: An approach based on ensembles of histories. *J. Phys. A*, 42(7):075007, 2009.
- [199] J. Gärtner. On large deviations from the invariant measure. *Theor. Probab. Appl.*, 22(1):24–39, 1977.
- [200] S. Ghosh and A. Nishry. Point processes, hole events, and large deviations: Random complex zeros and Coulomb gases. *Constr. Approx.*, 48(1):101–136, 2018.
- [201] S. Ghosh and A. Nishry. Gaussian complex zeros on the hole event: the emergence of a forbidden region. *Comm. Pure Appl. Math.*, 72(1):3–62, 2019.
- [202] C. Giardinà, J. Kurchan, and L. Peliti. Direct evaluation of large-deviation functions. *Phys. Rev. Lett.*, 96(12):120603, 2006.
- [203] J. Ginibre. Statistical ensembles of complex, quaternion, and real matrices. *J. Math. Phys.*, 6(3):440–449, 1965.
- [204] M. Gorissen and C. Vanderzande. Finite size scaling of current fluctuations in the totally asymmetric exclusion process. *J. Phys. A: Math. Theor.*, 44(11):115005, 2011.
- [205] F. Gosselin. Asymptotic behavior of absorbing Markov chains conditional on nonabsorption for applications in conservation biology. *Ann. Appl. Probab.*, pages 261–284, 2001.
- [206] K. Gourgoulias, M. A. Katsoulakis, L. Rey-Bellet, and J. Wang. How biased is your model? Concentration inequalities, information and model bias. *Preprint arXiv:1706.10260*, 2017.
- [207] N. Gozlan and C. Léonard. A large deviation approach to some transportation cost inequalities. *Probab. Theory Relat. Fields*, 139(1-2):235–283, 2007.
- [208] T. Grafke, R. Grauer, and T. Schäfer. The instanton method and its numerical implementation in fluid mechanics. *J. Phys. A: Math. Theor.*, 48(33):333001, 2015.

- [209] T. Grafke, R. Grauer, T. Schäfer, and E. Vanden-Eijnden. Arclength Parametrized Hamilton's Equations for the Calculation of Instantons. *Multiscale Model. Simul.*, 12(2):566–580, 2014.
- [210] T. Grafke, T. Schäfer, and E. Vanden-Eijnden. Long Term Effects of Small Random Perturbations on Dynamical Systems: Theoretical and Computational Tools. In *Recent Progress and Modern Challenges in Applied Mathematics, Modeling and Computational Science*, Fields Institute Communications, pages 17–55. Springer, New York, NY, 2017.
- [211] T. Grafke and E. Vanden-Eijnden. Numerical computation of rare events via large deviation theory. *Preprint arXiv:1812.00681*, 2018.
- [212] P. Grassberger. Go with the winners: a general Monte Carlo strategy. *Comp. Phys. Comm.*, 147(1-2):64–70, 2002.
- [213] R. Grimm and R. Storer. Monte-Carlo solution of Schrödinger's equation. *J. Comput. Phys.*, 7(1):134–156, 1971.
- [214] D. Guibourg, L. Hervé, and J. Ledoux. Quasi-compactness of Markov kernels on weighted-supremum spaces and geometrical ergodicity. *Preprint arXiv:1110.3240*, 2011.
- [215] E. Hairer, C. Lubich, and G. Wanner. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, volume 31 of *Springer Series in Computational Mathematics*. Springer Science & Business Media, 2006.
- [216] M. Hairer. Exponential mixing for a stochastic PDE driven by degenerate noise. *arXiv:math-ph/0103039*, 2001.
- [217] M. Hairer. Ergodic properties of Markov processes. *Lecture notes*, 2006.
- [218] M. Hairer. Non-Markovian processes. *Trends in Stochastic Analysis*, 353:65–102, 2009.
- [219] M. Hairer and J. C. Mattingly. Yet another look at Harris' ergodic theorem for Markov chains. In *Seminar on Stochastic Analysis, Random Fields and Applications VI*, pages 109–117. Springer, 2011.
- [220] M. Hairer and G. Pavliotis. From ballistic to diffusive behavior in periodic potentials. *J. Stat. Phys.*, 131(1):175–202, 2008.
- [221] M. Hairer and J. Weare. Improved diffusion Monte Carlo. *Comm. Pure Appl. Math.*, 67(12):1995–2021, 2014.
- [222] M. Hairer and J. Weare. The Brownian fan. *Comm. Pure Appl. Math.*, 68(1):1–60, 2015.
- [223] J. Han. Deep learning approximation for stochastic control problems. *Preprint arXiv:1611.07422*, 2016.
- [224] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proc. Natl. Acad. Sci.*, 115(34):8505–8510, 2018.
- [225] A. Hardy. A note on large deviations for 2D Coulomb gas with weakly confining potential. *Electron. Commun. Probab.*, 17, 2012.
- [226] A. Hardy. Polynomial ensembles and recurrence coefficients. *Constr. Approx.*, 48(1):137–162, 2018.
- [227] R. J. Harris and G. M. Schütz. Fluctuation theorems for stochastic dynamics. *J. Stat. Mech.*, 2007(07):P07020, 2007.
- [228] R. J. Harris and H. Touchette. Large deviation approach to nonequilibrium systems. In R. Klages, W. Just, and C. Jarzynski, editors, *Nonequilibrium Statistical Physics of Small Systems: Fluctuation Relations and Beyond*, volume 6 of *Reviews of Nonlinear Dynamics and Complexity*, pages 335–360, Weinheim, 2013. Wiley-VCH.
- [229] D. Harte. *Multifractals: Theory and Applications*. Chapman and Hall/CRC, 2001.

- [230] C. Hartmann, R. Banisch, M. Sarich, T. Badowski, and C. Schütte. Characterization of rare events in molecular dynamics. *Entropy*, 16(1):350, 2014.
- [231] C. Hartmann and C. Schütte. Efficient rare event simulation by optimal nonequilibrium forcing. *J. Stat. Mech. Theory Exp.*, 2012(11):11004, 2012.
- [232] C. Hartmann, C. Schütte, and W. Zhang. Model reduction algorithms for optimal control and importance sampling of diffusions. *Nonlinearity*, 29(8):2298, 2016.
- [233] L. O. Hedges, R. L. Jack, J. P. Garrahan, and D. Chandler. Dynamic order-disorder in atomistic models of structural glass formers. *Science*, 323(5919):1309–1313, 2009.
- [234] B. Helffer. *Semi-classical Analysis for the Schrödinger Operator and Applications*, volume 1336 of *Lecture Notes in Mathematics*. Springer, 2006.
- [235] B. Helffer. *Spectral Theory and its Applications*, volume 139 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2013.
- [236] B. Helffer and F. Nier. Quantitative analysis of metastability in reversible diffusion processes via a Witten complex approach: the case with boundary. *Matemática Contemporânea*, 26:41–85, 2004.
- [237] L. L. Helms. *Potential Theory*. Universitext. Springer, London, second edition, 2014.
- [238] F. Hérau and F. Nier. Isotropic hypoellipticity and trend to equilibrium for the Fokker-Planck equation with a high-degree potential. *Arch. Rational Mech. Anal.*, 171(2):151–218, 2004.
- [239] D. P. Herzog and J. C. Mattingly. Ergodicity and Lyapunov functions for Langevin dynamics with singular potentials. *Preprint arXiv:1711.02250*, 2017.
- [240] M. Heymann and E. Vanden-Eijnden. The geometric minimum action method: A least action principle on the space of curves. *Comm. Pure Appl. Math.*, 61(8):1052–1117, Aug. 2008.
- [241] M. Heymann and E. Vanden Eijnden. Pathways of maximum likelihood for rare events in nonequilibrium systems: Application to nucleation in the presence of shear. *Phys. Rev. Lett.*, 100(14):140601, 2008.
- [242] M. D. Hoffman and A. Gelman. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.*, 15:1593–1623, 2014.
- [243] T. A. Höft and B. K. Alpert. Fast updating multipole Coulombic potential calculation. *SIAM J. Sci. Comput.*, 39(3):A1038–A1061, 2017.
- [244] L. Hörmander. Hypoelliptic second order differential equations. *Acta Mathematica*, 119(1):147–171, 1967.
- [245] A. M. Horowitz. A generalized guided Monte Carlo algorithm. *Phys. Lett. B*, 268(CERN-TH-6172-91):247–252, 1991.
- [246] J. B. Hough, M. Krishnapur, Y. Peres, and B. Virág. Determinantal processes and independence. *Probab. Surv.*, 3:206–229, 2006.
- [247] M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. *Ann. Appl. Probab.*, 22(4):1611–1641, 2012.
- [248] C.-R. Hwang, S.-Y. Hwang-Ma, and S.-J. Sheu. Accelerating diffusions. *Ann. Appl. Probab.*, 15(2):1433–1444, 2005.
- [249] C.-R. Hwang, R. Normand, and S.-J. Wu. Variance reduction for diffusions. *Stoch. Process. Appl.*, 125(9):3522–3540, 2015.
- [250] R. L. Jack and P. Sollich. Large deviations and ensembles of trajectories in stochastic models. *Prog. Theoret. Phys. Suppl.*, 184:304–317, 2010.

- [251] C. Jarzynski. Equilibrium free-energy differences from nonequilibrium measurements: A master-equation approach. *Phys. Rev. E*, 56(5):5018–5035, 1997.
- [252] C. Jarzynski. Nonequilibrium equality for free energy differences. *Phys. Rev. Lett.*, 78(14):2690–2693, 1997.
- [253] A. Jasra. On the behaviour of the backward interpretation of Feynman–Kac formulae under verifiable conditions. *J. Appl. Probab.*, 52(02):339–359, 2015.
- [254] T. Jiang and Y. Qi. Spectral radii of large non-Hermitian random matrices. *J. Theoret. Probab.*, 30(1):326–364, 2017.
- [255] A. Jones and B. Leimkuhler. Adaptive stochastic methods for sampling driven molecular systems. *J. Chem. Phys.*, 135(8):084125, 2011.
- [256] S. Juneja and P. Shahabuddin. *Rare-Event Simulation Techniques: An Introduction and Recent Advances*, volume 13, chapter 11, pages 291–350. Elsevier, Amsterdam, 2006.
- [257] V. Jurdjevic. *Geometric Control Theory*, volume 52 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 1997.
- [258] M. Kac. On the average of a certain Wiener functional and a related limit theorem in calculus of probability. *Trans. Am. Math. Soc.*, 59(3):401–414, 1946.
- [259] O. Kallenberg. *Foundations of Modern Probability*. Probability and its Applications. Springer Science & Business Media, 2006.
- [260] S. C. Kapfer and W. Krauth. Cell-veto Monte Carlo algorithm for long-range systems. *Phys. Rev. E*, 94, 9 2016.
- [261] H. J. Kappen and H. C. Ruiz. Adaptive importance sampling for control and inference. *J. Stat. Phys.*, 162(5):1244–1266, 2016.
- [262] I. Karatzas and S. Shreve. *Brownian Motion and Stochastic Calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer Science & Business Media, 2012.
- [263] T. Kato. *Perturbation Theory for Linear Operators*, volume 132 of *Classics in Mathematics*. Springer Science & Business Media, 2013.
- [264] M. A. Katsoulakis, L. Rey-Bellet, and J. Wang. Scalable information inequalities for uncertainty quantification. *J. Comput. Phys.*, 336:513–545, 2017.
- [265] O. Kebiri, L. Neureither, and C. Hartmann. Adaptive importance sampling with forward-backward stochastic differential equations. In *International workshop on Stochastic Dynamics out of Equilibrium*, pages 265–281. Springer, 2017.
- [266] G. Keller. *Equilibrium States in Ergodic Theory*, volume 42 of *London Mathematical Society Student Texts*. Cambridge University Press, 1998.
- [267] M. Keyl. Quantum state estimation and large deviations. *Rev. Math. Phys.*, 18(01):19–60, 2006.
- [268] C. Kipnis and C. Landim. *Scaling Limits of Interacting Particle Systems*, volume 320 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1999.
- [269] W. Kliemann. Recurrence and invariant measures for degenerate diffusions. *Ann. Probab.*, 15(2):690–707, 1987.
- [270] P. E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, 1992.
- [271] P. E. Kloeden, E. Platen, and H. Schurz. *Numerical solution of SDE through computer experiments*. Universitext. Springer-Verlag, Berlin, 1994. With 1 IBM-PC floppy disk (3.5 inch; HD).
- [272] A. Koldobsky. *Fourier Analysis in Convex Geometry*. Number 116 in *Mathematical Surveys and Monographs*. American Mathematical Soc., 2005.

- [273] T. Komorowski, C. Landim, and S. Olla. *Fluctuations in Markov Processes: Time Symmetry and Martingale Approximation*, volume 345 of *Grundlehren der mathematischen Wissenschaften*. Springer Science & Business Media, 2012.
- [274] I. Kontoyiannis and S. P. Meyn. Spectral theory and limit theorems for geometrically ergodic Markov processes. *Ann. Appl. Probab.*, 13(1):304–362, 2003.
- [275] I. Kontoyiannis and S. P. Meyn. Large deviations asymptotics and the spectral theory of multiplicatively regular Markov processes. *Electron. J. Probab.*, 10(3):61–123, 2005.
- [276] M. Kopec. Weak backward error analysis for overdamped Langevin processes. *IMA J. Numer. Anal.*, 35(2):583–614, 2014.
- [277] M. Kopec. Weak backward error analysis for Langevin process. *BIT Numer. Math.*, 55(4):1057–1103, 2015.
- [278] M. Krishnapur, B. Rider, and B. Virág. Universality of the stochastic Airy operator. *Comm. Pure Appl. Math.*, 69(1):145–199, 2016.
- [279] S. Kullback. *Information Theory and Statistics*. Dover Books on Mathematics. Courier Corporation, 1997.
- [280] J. Kurchan. Fluctuation theorem for stochastic dynamics. *J. Phys. A*, 31(16):3719–3729, 1998.
- [281] B. R. La Cour and W. C. Schieve. A general conditional large deviation principle. *J. Stat. Phys.*, 161(1):123–130, 2015.
- [282] A. Lagendijk, B. Van Tiggelen, and D. S. Wiersma. Fifty years of Anderson localization. *Phys. Today*, 62(8):24–29, 2009.
- [283] P. Lancaster and L. Rodman. *Algebraic Riccati Equations*. Clarendon Press, 1995.
- [284] N. S. Landkof. *Foundations of Modern Potential Theory*. Springer-Verlag, New York-Heidelberg, 1972. Translated from the Russian by A. P. Doohovskoy, Die Grundlehren der mathematischen Wissenschaften, Band 180.
- [285] F. Lavancier, J. Møller, and E. Rubak. Determinantal point process models and statistical inference. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, 77(4):853–877, 2015.
- [286] J.-F. Le Gall. *Brownian Motion, Martingales, and Stochastic Calculus*, volume 274 of *Graduate Texts in Mathematics*. Springer, 2016.
- [287] T. Leblé and S. Serfaty. Large deviation principle for empirical fields of Log and Riesz gases. *Invent. Math.*, 210(3):645–757, 2017.
- [288] J. L. Lebowitz and H. Spohn. A Gallavotti–Cohen-type symmetry in the large deviation functional for stochastic dynamics. *J. Stat. Phys.*, 95(1-2):333–365, 1999.
- [289] V. Lecomte and J. Tailleur. A numerical approach to large deviations in continuous time. *J. Stat. Mech. Theor. Exp.*, 2007(03):P03004, 2007.
- [290] M. Ledoux. Differential operators and spectral distributions of invariant ensembles from the classical orthogonal polynomials. The continuous case. *Electron. J. Probab.*, 9:no. 7, 177–208, 2004.
- [291] M. Ledoux. On optimal matching of Gaussian samples. *Veroyatnost' i Statistika*, 457(25):226–264, 2017.
- [292] M. Ledoux. A fluctuation result in dual Sobolev norm for the optimal matching problem. 2019.
- [293] Y. T. Lee and S. S. Vempala. Convergence rate of Riemannian Hamiltonian Monte Carlo and faster polytope volume computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1115–1121. ACM, 2018.
- [294] B. Leimkuhler and C. Matthews. Rational construction of stochastic numerical methods for molecular sampling. *Appl. Math. Res. Express*, 2013(1):34–56, 2012.

- [295] B. Leimkuhler, C. Matthews, and G. Stoltz. The computation of averages from equilibrium and nonequilibrium Langevin molecular dynamics. *IMA J. Numer. Anal.*, 36(1):13–79, 2016.
- [296] B. Leimkuhler and S. Reich. *Simulating Hamiltonian Dynamics*, volume 14 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge university press, 2004.
- [297] B. Leimkuhler and M. Sachs. Ergodic properties of quasi-Markovian generalized Langevin equations with configuration dependent noise. *Preprint arXiv:1804.04029*, 2018.
- [298] T. Lelièvre, D. Le Peutrec, and B. Nectoux. Exit event from a metastable state and Eyring–Kramers law for the overdamped Langevin dynamics. In *International workshop on Stochastic Dynamics out of Equilibrium*, pages 331–363. Springer, 2017.
- [299] T. Lelièvre and F. Nier. Low temperature asymptotics for quasistationary distributions in a bounded domain. *Analysis & PDE*, 8(3):561–628, 2015.
- [300] T. Lelièvre, F. Nier, and G. A. Pavliotis. Optimal non-reversible linear drift for the convergence to equilibrium of a diffusion. *J. Stat. Phys.*, 152:237–274, 2013.
- [301] T. Lelièvre, M. Rousset, and G. Stoltz. *Free Energy Computations. A Mathematical Perspective*. Imperial College Press, London, 2010.
- [302] T. Lelièvre, M. Rousset, and G. Stoltz. Langevin dynamics with constraints and computation of free energy differences. *Math. Comp.*, 81(280):2071–2125, 2012.
- [303] T. Lelièvre, M. Rousset, and G. Stoltz. Hybrid Monte Carlo methods for sampling probability measures on submanifolds. *Numer. Math.*, 143(2):379–421, 2019.
- [304] T. Lelièvre and G. Stoltz. Partial differential equations and stochastic methods in molecular dynamics. *Acta Numer.*, 25:681–880, 2016.
- [305] C. Léonard. Entropic projections and dominating points. *ESAIM: Probab. Stat.*, 14:343–381, 2010.
- [306] C. Léonard and J. Najim. An extension of Sanov’s theorem: application to the Gibbs conditioning principle. *Bernoulli*, 8(6):721–743, 2002.
- [307] R. Leonarduzzi, P. Abry, H. Wendt, S. Jaffard, and H. Touchette. A generalized multifractal formalism for the estimation of nonconcave multifractal spectra. *IEEE Transactions on Signal Processing*, 67(1):110–119, 2018.
- [308] V. Letizia and S. Olla. Nonequilibrium isothermal transformations in a temperature gradient from a microscopic dynamics. *Ann. Probab.*, 45(6A):3987–4018, 2017.
- [309] D. Levesque and L. Verlet. On the theory of classical fluids II. *Physica*, 28(11):1124–1142, 1962.
- [310] D. Levesque and L. Verlet. Computer "Experiments" on Classical Fluids. III. Time-Dependent Self-Correlation Functions. *Phys. Rev. A*, 2:2514, 1970.
- [311] X. H. Li and G. Menon. Numerical solution of Dyson Brownian motion and a sampling scheme for invariant matrix ensembles. *J. Stat. Phys.*, 153(5):801–812, 2013.
- [312] H. Lie, C. Schütte, and C. Hartmann. Martingale-based gradient descent algorithm for estimating free energy values of diffusions. *Preprint*, 2015.
- [313] L.-H. Lim and J. Weare. Fast randomized iteration: diffusion Monte Carlo through the lens of numerical linear algebra. *SIAM Rev.*, 59(3):547–587, 2017.
- [314] W. Liu and L. Wu. Large deviations for empirical measures of mean-field Gibbs measures. *Stoch. Process. Appl.*, 2019.
- [315] S. Livingstone, M. F. Faulkner, and G. O. Roberts. Kinetic energy choice in Hamiltonian/hybrid Monte Carlo. *Biometrika*, 106(2):303–319, 2019.
- [316] Y. Lu and J. C. Mattingly. Geometric ergodicity of Langevin dynamics with Coulomb interactions. *Preprint arXiv:1902.00602*, 2019.

- [317] M. Maïda and É. Maurel-Segala. Free transport-entropy inequalities for non-convex potentials and application to concentration for random matrices. *Prob. Theory Relat. Fields*, 159(1-2):329–356, 2014.
- [318] S. N. Majumdar, C. Nadal, A. Scardicchio, and P. Vivo. Index distribution of Gaussian random matrices. *Phys. Rev. Lett.*, 103:220603, Nov 2009.
- [319] S. N. Majumdar, C. Nadal, A. Scardicchio, and P. Vivo. How many eigenvalues of a Gaussian random matrix are positive? *Phys. Rev. E*, 83:041105, Apr 2011.
- [320] N. Makri and W. H. Miller. Exponential power series expansion for the quantum time evolution operator. *J. Chem. Phys.*, 90(2):904–911, 1989.
- [321] F. Malrieu. Logarithmic Sobolev inequalities for some nonlinear PDE's. *Stoch. Process. Appl.*, 95(1):109–132, 2001.
- [322] J.-M. Marin, P. Pudlo, C. P. Robert, and R. J. Ryder. Approximate Bayesian computational methods. *Stat. Comput.*, 22(6):1167–1180, 2012.
- [323] G. Maruyama and H. Tanaka. Ergodic property of n -dimensional recurrent Markov processes. *Memoirs of the Faculty of Science, Kyushu University. Series A, Mathematics*, 13(2):157–172, 1959.
- [324] J. C. Mattingly, A. M. Stuart, and D. J. Higham. Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise. *Stoch. Proc. Appl.*, 101(2):185–232, 2002.
- [325] J. C. Mattingly, A. M. Stuart, and M. V. Tretyakov. Convergence of numerical time-averaging and stationary measures via Poisson equations. *SIAM J. Numer. Anal.*, 48(2):552–577, 2010.
- [326] J. Mehl, T. Speck, and U. Seifert. Large deviation function for entropy production in driven one-dimensional systems. *Phys. Rev. E*, 78:011123, 2008.
- [327] M. L. Mehta. *Random Matrices*, volume 142 of *Pure and Applied Mathematics (Amsterdam)*. Elsevier/Academic Press, Amsterdam, third edition, 2004.
- [328] M. Mella, G. Morosi, and D. Bressanini. Time step bias improvement in diffusion Monte Carlo simulations. *Phys. Rev. E*, 61(2):2050, 2000.
- [329] C. D. Meyer. *Matrix Analysis and Applied Linear Algebra*, volume 71 of *Other Titles in Applied Mathematics*. SIAM, 2000.
- [330] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer Science & Business Media, 2012.
- [331] L. Michel. About small eigenvalues of Witten Laplacian. *Pure Appl. Anal.*, 1(2):149–206, 2017.
- [332] G. N. Milstein and M. V. Tretyakov. *Stochastic Numerics for Mathematical Physics*. Scientific Computation. Springer Science & Business Media, 2013.
- [333] A. V. Nagaev. An integral limit theorem for large deviations. *Izv. Akad. Nauk UzSSR Ser. Fiz.-Mat. Nauk*, 1962(6):37–43, 1962.
- [334] A. V. Nagaev. Integral theorems with allowance for large deviations when Cramér's condition is violated. *Dokl. Akad. Nauk SSSR*, 180:279–281, 1968.
- [335] A. V. Nagaev. Integral limit theorems taking large deviations into account when Cramér's condition does not hold. I. *Theory Probab. Appl.*, 14(1):51–64, 1969.
- [336] A. V. Nagaev. Integral limit theorems taking large deviations into account when Cramér's condition does not hold. II. *Theory Probab. Appl.*, 14(2):193–208, 1969.
- [337] T. Nemoto, F. Bouchet, R. L. Jack, and V. Lecomte. Population-dynamics method with a multicanonical feedback control. *Phys. Rev. E*, 93(6):062123, 2016.
- [338] T. Nemoto, E. G. Hidalgo, and V. Lecomte. Finite-time and finite-size scalings in the evaluation of large-deviation functions: Analytical study using a birth-death process. *Phys. Rev. E*, 95(1):012102, 2017.

- [339] T. Nemoto, R. L. Jack, and V. Lecomte. Finite-size scaling of a first-order dynamical phase transition: Adaptive population dynamics and an effective model. *Phys. Rev. Lett.*, 118(11):115702, 2017.
- [340] T. Nemoto and S.-I. Sasa. Variational formula for experimental determination of high-order correlations of current fluctuations in driven systems. *Phys. Rev. E*, 83:030105, 2011.
- [341] T. Nemoto and S.-I. Sasa. Computation of large deviation statistics via iterative measurement-and-feedback procedure. *Phys. Rev. Lett.*, 112:090602, 2014.
- [342] D. Nickelsen and H. Touchette. Anomalous scaling of dynamical large deviations. *Phys. Rev. Lett.*, 121(9):090602, 2018.
- [343] I. Nourdin and G. Peccati. *Normal Approximations with Malliavin Calculus: from Stein's Method to Universality*, volume 192 of *Cambridge Tracks in Mathematics*. Cambridge University Press, 2012.
- [344] R. D. Nussbaum. The radius of the essential spectrum. *Duke Math. J.*, 37(3):473–478, 1970.
- [345] R. D. Nussbaum. Eigenvectors of order-preserving linear operators. *J. London Math. Soc.*, 58(2):480–496, 1998.
- [346] P. T. Nyawo and H. Touchette. A minimal model of dynamical phase transition. *Europhys. Lett.*, 116(5):50009, 2017.
- [347] S. Olver, R. R. Nadakuditi, and T. Trogdon. Sampling unitary ensembles. *Random Matrices Theory Appl.*, 4(1):1550002, 22, 2015.
- [348] E. Pardoux and D. Talay. Discretization and simulation of stochastic differential equations. *Acta Applicandae Mathematica*, 3(1):23–47, 1985.
- [349] É. Pardoux and Y. Veretennikov. On the Poisson equation and diffusion approximation. I. *Ann. Probab.*, 29(3):1061–1085, 2001.
- [350] G. A. Pavliotis. *Stochastic Processes and Applications: Diffusion Processes, the Fokker-Planck and Langevin Equations*, volume 60 of *Texts in Applied Mathematics*. Springer, 2014.
- [351] A. Pazy. *Semigroups of Linear Operators and Applications to Partial Differential Equations*, volume 44 of *Applied Mathematical Sciences*. Springer Science & Business Media, 2012.
- [352] H. Pham. *Continuous-time Stochastic Control and Optimization with Financial Applications*, volume 61 of *Stochastic Modelling and Applied Probability*. Springer Science & Business Media, 2009.
- [353] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM J. Cont. Opt.*, 30(4):838–855, 1992.
- [354] S. T. Rachev. *Probability Metrics and the Stability of Stochastic Models*, volume 269 of *Wiley Series in Probability and Statistics*. John Wiley & Son Ltd, 1991.
- [355] F. Ragone, J. Wouters, and F. Bouchet. Computation of extreme heat waves in climate models using a large deviation algorithm. *Proc. Natl. Acad. Sci.*, 115(1):24–29, 2018.
- [356] D. C. Rapaport. *The Art of Molecular Dynamics Simulation*. Cambridge University Press, 2004.
- [357] F. Rassoul-Agha and T. Seppäläinen. *A Course on Large Deviations with an Introduction to Gibbs Measures*, volume 162 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2015.
- [358] M. Reed and B. Simon. *Methods of Modern Mathematical Physics IV: Analysis of Operators*. Elsevier, 1978.
- [359] M. Reed and B. Simon. *Methods of Modern Mathematical Physics I: Functional Analysis*. Academic Press, San Diego, 1980.

- [360] P. Reimann. Brownian motors: Noisy transport far from equilibrium. *Phys. Rep.*, 361:57–265, 2002.
- [361] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*, volume 293 of *Grundlehren der mathematischen Wissenschaften*. Springer Science & Business Media, 2013.
- [362] L. Rey-Bellet. Ergodic properties of Markov processes. In *Open Quantum Systems II*, pages 1–39. Springer, 2006.
- [363] L. Rey-Bellet and K. Spiliopoulos. Irreversible Langevin samplers and variance reduction: A large deviations approach. *Nonlinearity*, 28(7):2081, 2015.
- [364] L. Rey-Bellet and K. Spiliopoulos. Improving the convergence of reversible samplers. *J. Stat. Phys.*, 164(3):472–494, 2016.
- [365] B. Rider. A limit theorem at the edge of a non-Hermitian random matrix ensemble. *J. Phys. A*, 36(12):3401–3409, 2003.
- [366] H. Risken. *The Fokker-Planck Equation: Methods of Solution and Applications*, volume 18 of *Springer Series in Synergetics*. Springer, Berlin, 3rd edition, 1996.
- [367] C. Robert. *The Bayesian Choice: from Decision-Theoretic Foundations to Computational Implementation*. Springer Texts in Statistics. Springer Science & Business Media, 2007.
- [368] C. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer Science & Business Media, 2013.
- [369] C. P. Robert. *Monte Carlo Methods*. Wiley Online Library, 2004.
- [370] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer-Verlag, New York, second edition, 2004.
- [371] G. O. Roberts and J. S. Rosenthal. Optimal scaling for various Metropolis-Hastings algorithms. *Stat. Sci.*, 16(4):351–367, 2001.
- [372] G. O. Roberts and J. S. Rosenthal. General state space Markov chains and MCMC algorithms. *Probab. Surv.*, 1:20–71, 2004.
- [373] G. O. Roberts and R. Tweedie. Exponential convergence of Langevin distributions and their discrete approximations. *Bernoulli*, 2(4):341–363, 1996.
- [374] C. M. Rohwer, F. Angeletti, and H. Touchette. Convergence of large deviation estimators. *Phys. Rev. E*, 92:052104, 2015.
- [375] P. J. Rossky, J. Doll, and H. Friedman. Brownian dynamics as smart Monte Carlo simulation. *J. Chem. Phys.*, 69(10):4628–4633, 1978.
- [376] M. Rousset. *Continuous time population Monte Carlo and computational physics*. PhD thesis, Université Paul Sabatier Toulouse, 2006.
- [377] M. Rousset. On the control of an interacting particle estimation of Schrödinger ground states. *SIAM J. Math. Anal.*, 38(3):824–844, 2006.
- [378] M. Rousset and G. Stoltz. Equilibrium sampling from nonequilibrium dynamics. *J. Stat. Phys.*, 123(6):1251–1272, 2006.
- [379] W. Rudin. *Real and Complex Analysis*. Tata McGraw-Hill Education, 2006.
- [380] E. B. Saff and V. Totik. *Logarithmic Potentials with External Fields*, volume 316 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 1997. Appendix B by Thomas Bloom.
- [381] A. Sarsa, J. Boronat, and J. Casulleras. Quadratic diffusion Monte Carlo and pure estimators for atoms. *J. Chem. Phys.*, 116(14):5956–5962, 2002.

- [382] A. Scardicchio, C. E. Zachary, and S. Torquato. Statistical properties of determinantal point processes in high-dimensional Euclidean spaces. *Phys. Rev. E* (3), 79(4):041108, 19, 2009.
- [383] H. H. Schaefer. *Banach Lattices and Positive Operators*, volume 215 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1974.
- [384] L. Schwartz. *Analyse I*. Hermann, 1991.
- [385] K. Sekimoto. *Stochastic Energetics*, volume 799 of *Lect. Notes. Phys.* Springer, New York, 2010.
- [386] S. Serfaty. *Coulomb gases and Ginzburg-Landau vortices*. Zurich Lectures in Advanced Mathematics. Euro. Math. Soc. (EMS), Zürich, 2015.
- [387] S. Serfaty. Systems of points with Coulomb interactions. Preprint [arXiv:1712.04095v1](https://arxiv.org/abs/1712.04095v1), 2017.
- [388] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis*. Stochastic Modeling Series. Chapman and Hall, London, 1995.
- [389] B. Simon. Semiclassical analysis of low lying eigenvalues I. Non-degenerate minima: Asymptotic expansions. *Ann. Inst. Henri Poincaré: Phys. Th.*, 38(3):295–308, 1983.
- [390] S. Smale. Mathematical problems for the next century. *Math. Intelligencer*, 20(2):7–15, 1998.
- [391] S. Smale. Mathematical problems for the next century. In *Mathematics: frontiers and perspectives*, pages 271–294. Amer. Math. Soc., Providence, RI, 2000.
- [392] H. Spohn. *Large Scale Dynamics of Interacting Particles*. Texts and Monographs in Physics. Springer Science & Business Media, 2012.
- [393] G. Stoltz and Z. Trstanova. Langevin dynamics with general kinetic energies. *Multiscale Model. Simul.*, 16(2):777–806, 2018.
- [394] D. W. Stroock and S. Karmakar. *Lectures on Topics in Stochastic Differential Equations*. Tata Institute of Fundamental Research, Bombay, 1982.
- [395] D. W. Stroock and S. R. Varadhan. On the support of diffusion processes with applications to the strong maximum principle. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971)*, volume 3, pages 333–359, 1972.
- [396] D. W. Stroock and S. R. S. Varadhan. *Multidimensional Diffusion Processes*. Springer, New York, 1979.
- [397] J. Tailleur and V. Lecomte. Simulation of large deviation functions using population dynamics. *AIP Conf. Proc.*, 1091(1):212–219, 2009.
- [398] D. Talay. Second-order discretization schemes of stochastic differential systems for the computation of the invariant law. *Stochastics*, 29(1):13–36, 1990.
- [399] D. Talay. Stochastic Hamiltonian systems: exponential convergence to the invariant measure, and discretization by the implicit Euler scheme. *Markov Process. Relat. Fields*, 8(2):163–198, 2002.
- [400] D. Talay and L. Tubaro. Expansion of the global error for numerical schemes solving stochastic differential equations. *Stoch. Anal. Appl.*, 8(4):483–509, 1990.
- [401] T. Tao. *Topics in Random Matrix Theory*, volume 132 of *Graduate Studies in Mathematics*. American Mathematical Soc., 2012.
- [402] T. Tao and V. Vu. Random matrices: universality of ESDs and the circular law. *Ann. Probab.*, 38(5):2023–2065, 2010. With an appendix by Manjunath Krishnapur.
- [403] S. Tavaré, D. J. Balding, R. C. Griffiths, and P. Donnelly. Inferring coalescence times from DNA sequence data. *Genetics*, 145(2):505–518, 1997.

- [404] H. Touchette. Introduction to dynamical large deviations of Markov processes. *Physica A*, 504:5–19.
- [405] H. Touchette. *Equivalence and Nonequivalence of the Microcanonical and Canonical Ensembles: A Large Deviations Study*. PhD thesis, Department of Physics, McGill University, December 2003.
- [406] H. Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478(1):1–69, 2009.
- [407] H. Touchette. Methods for calculating nonconcave entropies. *J. Stat. Mech.: Theory Exp.*, 2010(05):P05008, 2010.
- [408] H. Touchette. A basic introduction to large deviations: Theory, applications, simulations. *Preprint arXiv:1106.4146*, 2011.
- [409] Z. Trstanova. *Mathematical and Algorithmic Analysis of Modified Langevin Dynamics*. PhD thesis, Université Grenoble Alpes, 2016.
- [410] P. Tsoigni Nyawo and H. Touchette. Large deviations of the current for driven periodic diffusions. *Phys. Rev. E*, 94:032101, 2016.
- [411] P. Tsoigni Nyawo and H. Touchette. A minimal model of dynamical phase transition. *Europhys. Lett.*, 116:50009, 2016.
- [412] C. Umrigar, M. Nightingale, and K. Runge. A diffusion Monte Carlo algorithm with very small time-step errors. *J. Chem. Phys.*, 99(4):2865–2890, 1993.
- [413] E. Vanden-Eijnden. Transition path theory. In M. Ferrario, G. Ciccotti, and K. Binder, editors, *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*, volume 703 of *Lecture Notes in Physics*, pages 453–493. Springer, 2006.
- [414] E. Vanden-Eijnden and J. Weare. Rare Event Simulation of Small Noise Diffusions. *Comm. Pure Appl. Math.*, 65(12):1770–1803, Dec. 2012.
- [415] P. Vanetti, A. Bouchard-Côté, G. Deligiannidis, and A. Doucet. Piecewise-Deterministic Markov Chain Monte Carlo. Preprint [arXiv:1707.05296v1](https://arxiv.org/abs/1707.05296v1), 2018.
- [416] S. Varadhan. *Large Deviations and Applications*. Number 46 in Regional Conference Series in Applied Mathematics. SIAM, 1984.
- [417] D. Veneziano. Large deviations of multifractal measures. *Fractals*, 10(01):117–129, 2002.
- [418] L. Verlet. Computer "experiments" on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules. *Phys. Rev.*, 159(1):98, 1967.
- [419] C. Villani. *Topics in Optimal Transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, 2003.
- [420] C. Villani. *Hypocoercivity*, volume 950 of *Memoirs of the American Mathematical Society*. American Mathematical Soc., 2009.
- [421] P. Vivo, S. N. Majumdar, and O. Bohigas. Large deviations of the maximum eigenvalue in Wishart random matrices. *J. Phys. A: Math. Theor.*, 40(16):4317, 2007.
- [422] A. F. Voter. A method for accelerating the molecular dynamics simulation of infrequent events. *J. Chem. Phys.*, 106(11):4665–4677, 1997.
- [423] H. Vroylandt and G. Verley. Non-equivalence of dynamical ensembles and emergent non-ergodicity. *J. Stat. Phys.*, 174(2):404–432, 2019.
- [424] R. Wang, X. Wang, and L. Wu. Sanov's theorem in the Wasserstein distance: A necessary and sufficient condition. *Stat. Probab. Lett.*, 80(5-6):505–512, 2010.
- [425] E. Witten. Supersymmetry and Morse theory. *J. Diff. Geom.*, 17(4):661–692, 1982.

- [426] L. Wu. A deviation inequality for non-reversible Markov processes. *Ann. Henri Poincaré: Probab. Stat.*, 36(4):435–445, 2000.
- [427] L. Wu. Large and moderate deviations and exponential convergence for stochastic damping Hamiltonian systems. *Stoch. Process. Appl.*, 91(2):205–238, 2001.
- [428] L. Wu. Essential spectral radius for Markov semigroups (I): discrete time case. *Probab. Theory Relat. Fields*, 128(2):255–321, 2004.
- [429] J. E. Yukich. *Probability Theory of Classical Euclidean Optimization Problems*, volume 1675 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1998.
- [430] E. Zappa, M. Holmes-Cerfon, and J. Goodman. Monte Carlo on manifolds: sampling densities and integrating functions. *Comm. Pure Appl. Math.*, 71(12):2609–2647, 2018.
- [431] I. Zentner, G. Ferré, F. Poirion, and M. Benoit. A biorthogonal decomposition for the identification and simulation of non-stationary and non-Gaussian random fields. *J. Comp. Phys.*, 314:1–13, 2016.
- [432] W. Zhang. Ergodic SDEs on submanifolds and related numerical sampling schemes. *Preprint arXiv:1702.08064*, 2018.
- [433] W. Zhang, H. Wang, C. Hartmann, M. Weber, and C. Schütte. Applications of the cross-entropy method to importance sampling and optimal control of diffusions. *SIAM J. Sci. Comput.*, 36(6):A2654–A2672, 2014.
- [434] V. M. Zolotarev. *Modern Theory of Summation of Random Variables*. Modern Probability and Statistics. VSP, 1997.
- [435] R. Zwanzig. Nonlinear generalized Langevin equations. *J. Stat. Phys.*, 9(3):215–220, 1973.
- [436] K. Zygalakis. On the existence and the applications of modified equations for stochastic differential equations. *SIAM J. Sci. Comput.*, 33(1):102–130, 2011.

- I*-continuity set, 45, 51, 235, 238
- h*-transform, 79, 130, 182
- Adaptive algorithm, 186
- Approximate eigenvalue problem, 155
- Approximate inverse operator, 175
- Asymptotic variance, 35
- Central limit theorem, 35
- Co-area formula, 240
- Collatz–Wielandt formula, 53
- Compact level sets, 44, 108
- Compact operator, 95, 129
- Cone in a Banach space, 95
- Contraction principle, 50, 113
- Controllability, 109
- Coulomb gas, 228
- Distance
 - bounded-Lipschitz, 230
 - Fortet–Mourier, 55
 - total variation, 96
 - Wasserstein, 107, 230
 - weighted total variation, 107
- Dobrushin ergodic coefficient, 177
- Donsker–Varadhan formula, 53, 113
- Electrostatic energy, 211, 228
- Entropy, 48
- Ergodicity, 32, 95, 183
- Error on the invariant measure, 58
- Essential spectral radius, 79, 95
- Euler–Maruyama, 57, 161, 215
- Exponential tightness, 45, 133
- Feller, 78
- Fenchel–Legendre transform, 46, 137
- Feynman–Kac
 - ergodicity, 80, 150
 - evolution operator, 53, 124, 182
 - invariant measure, 150
 - operator, 53, 114, 128, 146
 - semigroup, 76
- Fisher information, 114
- Gärtner–Ellis, 111
- Galerkin discretization, 162, 186
- Gateau-differentiability, 132
- Gaussian Unitary Ensemble, 212
- Generalized Hamiltonian Monte Carlo, 242
- Gibbs conditioning principle, 51, 233
- Ginibre Ensemble, 41, 213, 232
- Hamilton–Jacobi equation, 62, 200
- Hamiltonian, 229
- Hamiltonian dynamics, 30, 118
- Hamiltonian Monte Carlo, 215
- Hypoellipticity, 84, 109
- Instanton, 39, 199
- Interacting particles systems, 60, 161
- Invariant measure, 28, 108
- Irreducibility, 78, 128
- Kantorovitch–Rubinstein duality, 231
- Lagrange multiplier, 254
- Langevin
 - constrained dynamics, 240
 - overdamped dynamics, 29, 116
 - underdamped dynamics, 29, 118
- Large deviations principle, 44, 111
- Linear constraint, 235
- Local martingale, 85, 123
- Lyapunov
 - condition, 33, 78, 84, 95, 109
 - function, 34, 216
- Measures on submanifolds, 239
- Metropolis adjusted Langevin algorithm, 215
- Minorization condition, 78, 95, 150
- Numerical
 - analysis, 153
 - discretization, 148, 184
 - expansion, 201
 - scheme, 161
- Optimal control, 61, 197
- Poincaré inequality, 90
- Power method, 185

-
- Quadratic constraint, [237](#)
 - Quasi-everywhere property, [231](#)

 - Rate function, [44](#), [114](#)
 - RATTLE integrator, [241](#)
 - Regularization property, [126](#)
 - Reversibility, [28](#)
 - Riccati equation, [200](#)

 - Scaled Cumulant Generating Function, [46](#), [82](#), [111](#),
[156](#)
 - Separation of convex sets, [254](#)
 - Sobolev spaces, [108](#)
 - Spectral radius, [79](#), [95](#)
 - Support of a measure, [231](#)

 - Theorem
 - Ascoli, [96](#)
 - Cramer, [47](#), [112](#)
 - Gärtner–Ellis, [49](#), [134](#)
 - Girsanov, [185](#), [197](#)
 - Heine–Cantor, [96](#)
 - Krein–Rutman, [96](#), [129](#)
 - Sanov, [48](#)

 - Topology
 - τ , [107](#), [111](#)
 - strong, [48](#)
 - Wasserstein, [48](#)
 - weak, [48](#)

 - TU-lemma, [157](#)

 - Ultra-Feller, [96](#)

 - Witten Laplacian, [112](#)
 - Witten transform, [137](#)