



HAL
open science

Annealed Flow Transport Monte Carlo

Michael Arbel, Alexander G D G Matthews, Arnaud Doucet

► **To cite this version:**

Michael Arbel, Alexander G D G Matthews, Arnaud Doucet. Annealed Flow Transport Monte Carlo. ICML 2021 - 38th International Conference on Machine Learning, Jul 2021, Online, France. pp.1-70, 10.48550/arXiv.2102.07501 . hal-03455478

HAL Id: hal-03455478

<https://hal.science/hal-03455478>

Submitted on 29 Nov 2021

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.

Annealed Flow Transport Monte Carlo

Michael Arbel^{*1} Alexander G. D. G. Matthews^{*2} Arnaud Doucet²

Abstract

Annealed Importance Sampling (AIS) and its Sequential Monte Carlo (SMC) extensions are state-of-the-art methods for estimating normalizing constants of probability distributions. We propose here a novel Monte Carlo algorithm, Annealed Flow Transport (AFT), that builds upon AIS and SMC and combines them with normalizing flows (NFs) for improved performance. This method transports a set of particles using not only importance sampling (IS), Markov chain Monte Carlo (MCMC) and resampling steps - as in SMC, but also relies on NFs which are learned sequentially to push particles towards the successive annealed targets. We provide limit theorems for the resulting Monte Carlo estimates of the normalizing constant and expectations with respect to the target distribution. Additionally, we show that a continuous-time scaling limit of the population version of AFT is given by a Feynman–Kac measure which simplifies to the law of a controlled diffusion for expressive NFs. We demonstrate experimentally the benefits and limitations of our methodology on a variety of applications.

1. Introduction

Let π be a target density on $\mathcal{X} \subseteq \mathbb{R}^d$ w.r.t. the Lebesgue measure known up to a normalizing constant Z . We want to estimate Z and approximate expectations with respect to π . This has applications in Bayesian statistics but also variational inference (VI) (Mnih and Rezende, 2016) and compression (Li and Chen, 2019; Huang et al., 2020) among others. AIS (Neal, 2001) and its SMC extensions (Del Moral et al., 2006) are state-of-the-art Monte Carlo methods addressing this problem which rely on a

sequence of annealed targets $\pi_k \propto \pi_0^{1-\beta_k} \pi_K^{\beta_k}$ bridging smoothly an easy-to-sample distribution π_0 to $\pi_K := \pi$ for $0 = \beta_0 < \beta_1 < \dots < \beta_K = 1$ and MCMC kernels of invariant distributions π_k (Zhou et al., 2016; Lorente et al., 2020). In their simplest instance, SMC samplers propagate N particles approximating π_k at time k . These particles are reweighted according to weights proportional to π_{k+1}/π_k at time $k+1$ to build an IS approximation of π_{k+1} , then one resamples N times from this approximation and finally mutate the resampled particles according to MCMC steps of invariant distribution π_{k+1} . This procedure can provide high-variance estimators if the discrepancy between π_k and π_{k+1} is significant as the resulting IS weights then have a large variance and/or if the MCMC kernels mix poorly. This can be reduced by increasing K and the number of MCMC steps at each temperature but comes at an increasing computational cost.

An alternative approach is to build a transport map $T : \mathcal{X} \rightarrow \mathcal{X}$ to ensure that if $X \sim \pi_0$ then the distribution of $X' = T(X)$ denoted $T_{\#}\pi_0$ is approximately equal to π . In (El Moselhy and Marzouk, 2012), this map is parameterized using a polynomial chaos expansion and learned by minimizing a regularized Kullback-Leibler (KL) divergence between $T_{\#}\pi_0$ and π ; see also (Marzouk et al., 2016). Taghvaei et al. (2020) and Olmez et al. (2020) obtain transport maps by solving a Poisson equation. However, they do not correct for the discrepancy between $T_{\#}\pi_0$ and π using IS. Doing so would incur a $O(d^3)$ cost when computing the Jacobian. Normalizing Flows (NFs) are an alternative flexible class of diffeomorphisms with easy-to-compute Jacobians (Rezende and Mohamed, 2015). These can be used to parameterize T and are also typically learned by minimizing $\text{KL}(T_{\#}\pi_0 || \pi)$ or a regularized version of it. This approach has been investigated in many recent work; see e.g. (Gao et al., 2020; Nicoli et al., 2020; Noé et al., 2019; Wirmsberger et al., 2020). Although it is attractive, it is also well-known that optimizing this ‘mode-seeking’ KL can lead to an approximation of the target $T_{\#}\pi_0$ which has thinner tails than the target π and ignore some of its modes; see e.g. (Domke and Sheldon, 2018).

In this paper, our contributions are as follows.

- We propose Annealed Flow Transport (AFT), a methodology that takes advantages of the strengths of both

^{*}Equal contribution ¹Gatsby Computational Neuroscience Unit, University College London ²DeepMind. Correspondence to: Michael Arbel <michael.n.arbel@gmail.com>, Alexander G. D. G. Matthews <alexmatthews@google.com>, Arnaud Doucet <arnauddoucet@google.com>.

SMC and NFs. Given particles approximating π_k at time k , we learn a NF T_{k+1} minimizing the KL between $(T_{k+1})_{\#}\pi_k$ and π_{k+1} . As π_k is closer to π_{k+1} than π_0 is from $\pi_K = \pi$, learning such a NF is easier and less prone to mode collapse. Additionally the use of MCMC steps in SMC samplers allows the particles to diffuse and further prevent such collapse. Having obtained T_{k+1} , we then apply this mapping to the particles before building an IS approximation of π_{k+1} and then use resampling and MCMC steps.

- We establish a weak law of large numbers and a Central Limit Theorem (CLT) for the resulting Monte Carlo estimates of Z and expectations w.r.t. π . Available CLT results for SMC (Chopin, 2004; Del Moral, 2004; Künsch, 2005; Beskos et al., 2016) do not apply here as the transport maps are learned from particles.
- When one relies on Unadjusted Langevin algorithm (ULA) kernels to mutate particles, a time-rescaled population version of AFT without resampling is shown to converge as $K \rightarrow \infty$ towards a Feynman–Kac measure. For NFs expressive enough to include exact transport maps between successive distributions, this measure corresponds to the measure induced by a controlled Langevin diffusion.
- We demonstrate the performance of AFT on a variety of benchmarks, showing that it can improve over SMC for a given number of temperatures.

Related Work. The use of deterministic maps with AIS (Vaikuntanathan and Jarzynski, 2011) and SMC (Akyildiz and Míguez, 2020; Everitt et al., 2020; Heng et al., 2021) has already been explored. However, Everitt et al. (2020) and Vaikuntanathan and Jarzynski (2011) do not propose a generic methodology to build such maps while Akyildiz and Míguez (2020) introduce mode-seeking maps and do not correct for the incurred bias. Heng et al. (2021) rely on quadrature and a system of time-discretized nonlinear ordinary differential equations: this can be computationally cheaper than learning NFs but is application specific. NFs benefit from easy-to-compute Jacobians and a large and quickly expanding literature (Papamakarios et al., 2019); e.g., as both MCMC and NFs on manifolds have been developed, our algorithm can be directly extended to such settings.

Evidence Lower Bounds (ELBOs) based on unbiased estimators of Z have also been mentioned in (Salimans et al., 2015; Goyal et al., 2017; Caterini et al., 2018; Huang et al., 2018; Wu et al., 2020; Thin et al., 2021). These estimators generalize AIS, and are obtained using sequential IS, transport maps and MCMC. However, when MCMC kernels such as Metropolis–Hastings (MH) or Hamiltonian Monte Carlo (HMC) are used, accept/reject steps lead to high variance estimates of ELBO gradients (Thin et al., 2021). Moreover, while SMC (i.e. combining sequential

IS and resampling) can also be used to define an ELBO, resampling steps correspond to sampling discrete distributions and lead to high variance gradient estimates; see e.g. (Maddison et al., 2017; Le et al., 2018; Naesseth et al., 2018) in the context of state-space models. The algorithm proposed here does not rely on the ELBO, so it can use arbitrary MCMC kernels and exploit the benefits of resampling. Moreover, it only requires a single pass through the $K + 1$ annealed distributions: there is no need to iteratively run sequential IS or SMC for estimating Z and an ELBO gradient estimate.

Optimal control ideas have also been proposed to improve SMC by introducing an additive drift to a time-inhomogeneous ULA to improve sampling; see Richard and Zhang (2007); Kappen and Ruiz (2016); Guarniero et al. (2017); Heng et al. (2020). The proposed iterative algorithms require estimating value functions but, to be implementable, the approximating function class has to be severely restricted. The algorithm proposed here is much more widely applicable and can use sophisticated MCMC kernels.

Finally, alternative particle methods based on gradient flows in the space of probability measures have been proposed to provide an approximation of π , such as Stein Variational Gradient Descent (SVGD) (Liu and Wang, 2016; Liu et al., 2019; Wang and Li, 2019; Zhu et al., 2020; Reich and Weissmann, 2021). However, their consistency results require both K , the number of time steps, and N , the number of particles, to go to infinity. In contrast, AFT only needs $N \rightarrow \infty$. Moreover, they require specifying a suitable Reproducing Kernel Hilbert Space or performing kernel density estimation, which can be challenging in high dimension. Additionally, contrary to AFT, these methods do not provide an estimate of Z . One recent exception is the work of Han and Liu (2017) which combines SVGD with IS to estimate Z but this requires computing Jacobians of computational cost $O(d^3)$.

2. Sequential Monte Carlo samplers

We provide here a brief overview of SMC samplers and their connections to AIS. More details can be found in (Del Moral et al., 2006; Dai et al., 2020).

We will rely on the following notation for the annealed densities $(\pi_k)_{0 \leq k \leq K}$ targeted by SMC:

$$\pi_k(x) = \frac{\gamma_k(x)}{Z_k} = \frac{\exp(-V_k(x))}{Z_k},$$

where $Z_0 = 1$ so $\pi_0(x) = \gamma_0(x)$ and $V_k(x) = (1 - \beta_k)V_0 + \beta_k V_K$ for $0 = \beta_0 < \beta_1 < \dots < \beta_K = 1$. However, we could use more generally any sequence of distributions bridging smoothly π_0 to $\pi_K = \pi$.

2.1. Sequential importance sampling

Let us first ignore the key resampling steps used by SMC. In this case, SMC boils down to a sequential IS technique where one approximates π_k at time k . We first sample $X_0 \sim \pi_0$ at time $k = 0$, then at time $k \geq 1$, obtain a new sample $X_k \sim M_k(X_{k-1}, \cdot)$ using a Markov kernel M_k . For the distribution of X_k to be closer to π_k than the one of X_{k-1} , M_k is typically selected as a MCMC kernel of invariant density π_k such as MH or HMC, or of approximate invariant density π_k such as ULA. Hence, by construction, the joint density of $X_{0:k}$ is

$$\bar{\eta}_k(x_{0:k}) = \pi_0(x_0) \prod_{l=1}^k M_l(x_{l-1}, x_l). \quad (1)$$

The resulting marginal η_k of X_k under $\bar{\eta}_k$ usually differs from π_k . If one could evaluate η_k pointwise, then IS could be used to correct for the discrepancy between η_k and π_k using the IS weight $w_k(x_k) = \gamma_k(x_k)/\eta_k(x_k)$. Unfortunately, η_k is intractable in all but toy scenarios. Instead, SMC samplers introduce joint target densities $\bar{\pi}_k(x_{0:k})$ to compute tractable IS weights $w_k(x_{0:k})$ over the whole path $X_{0:k}$ defined by

$$\bar{\pi}_k(x_{0:k}) = \pi_k(x_k) \prod_{l=0}^{k-1} L_l(x_{l+1}, x_l), \quad (2)$$

here L_l are ‘‘backward’’ Markov kernels moving each sample X_{l+1} into a sample X_l starting from a *virtual* sample X_k from π_k ¹. Hence by construction π_k is the marginal of $\bar{\pi}_k$ at time k . The backward kernels L_{k-1} are chosen so that the following *incremental IS weights* are well-defined

$$G_k(x_{k-1}, x_k) = \frac{\gamma_k(x_k) L_{k-1}(x_k, x_{k-1})}{\gamma_{k-1}(x_{k-1}) M_k(x_{k-1}, x_k)}, \quad (3)$$

and, from (1) and (2), one obtains

$$w_k(x_{0:k}) := \frac{\bar{\gamma}_k(x_{0:k})}{\bar{\eta}_k(x_{0:k})} = \prod_{l=1}^k G_l(x_{l-1}, x_l), \quad (4)$$

where $\bar{\gamma}_k(x_{0:k}) = Z_k \bar{\pi}_k(x_{0:k})$ is the unnormalized joint target. Using IS, it is thus straightforward to check that

$$Z_k = \bar{\eta}_k[w_k], \quad \bar{\pi}_k[f] = \frac{\bar{\eta}_k[w_k f]}{\bar{\eta}_k[w_k]}, \quad (5)$$

where $f(x_{0:k})$ is a function of the whole trajectory $x_{0:k}$ and $\mu[g]$ is a shorthand notation for the expectation $\mathbb{E}_{X \sim \mu}[g(X)]$. As π_k is a marginal of $\bar{\pi}_k$, we can also estimate expectations w.r.t. to π_k using $\bar{\pi}_k[f] = \pi_k[f]$ for

¹As in (Crooks, 1998; Neal, 2001; Del Moral et al., 2006; Dai et al., 2020), we do not use measure-theoretic notation here but it should be kept in mind that the kernels M_l do not necessarily admit a density w.r.t. Lebesgue measure; e.g. a MH kernel admits an atomic component. For completeness, a formal measure-theoretic presentation of the results of this section is given in Appendix A.

$f(x_{0:k}) = f(x_k)$. From (5), it is thus possible to derive consistent estimators of Z_k and $\pi_k[f]$ by sampling N ‘particles’ $X_{0:k}^i \sim \bar{\eta}_k$ where $i = 1, \dots, N$ and using

$$Z_k^N = \frac{1}{N} \sum_{i=1}^N w_k(X_{0:k}^i), \quad \pi_k^N[f] = \sum_{i=1}^N W_k^i f(X_k^i), \quad (6)$$

where $W_k^i \propto w_k(X_{0:k}^i)$, $\sum_{i=1}^N W_k^i = 1$.

When the kernels M_k are π_k -invariant and we select L_{k-1} as the reversal of M_k , i.e. $\pi_k(x) M_k(x, x') = \pi_k(x') L_{k-1}(x', x)$, it is easy to check that $G_l(x_{l-1}, x_l) = \gamma_l(x_{l-1})/\gamma_{l-1}(x_{l-1})$. In that case, (5) corresponds to AIS (Neal, 2001) and is also known as the Jarzynski–Crooks identity (Jarzynski, 1997; Crooks, 1998). When π_k is a sequence of posterior densities, a similar construction was also used in (MacEachern et al., 1999; Gilks and Berzuini, 2001; Chopin, 2002). The generalized identity (5) allows the use of more general dynamics, including deterministic maps which will be exploited by our algorithm.

In practice, the choice of the backward transition kernels has a large impact on the variance of the estimates (6). (Del Moral et al., 2006) identified the backward kernels minimizing the variance of the IS weights (3)–(4) and proposed various approximations to them.

2.2. Sequential Monte Carlo

To reduce the variance of the IS estimators (6), SMC samplers combine sequential IS steps with resampling steps. Given an IS approximation $\pi_{k-1}^N = \sum_{i=1}^N W_{k-1}^i \delta_{X_{k-1}^i}$ of π_{k-1} at time $k-1$, one resamples N times from π_{k-1}^N to obtain particles approximately distributed according to π_{k-1} . This has for effect of discarding particles with low weights and replicating particles with high weights, this helps focusing subsequent computation on ‘‘promising’’ regions of the space. Empirically, resampling usually provides lower variance unbiased estimates of normalizing constants and is computationally very cheap; see e.g. (Chopin, 2002; Hukushima and Iba, 2003; Del Moral et al., 2006; Rousset and Stoltz, 2006; Zhou et al., 2016; Barash et al., 2017). The resampled particles are then evolved according to M_k , weighted according to G_k and resampled again.

3. Annealed Flow Transport Monte Carlo

We now introduce AFT, a new flexible adaptive Monte Carlo method that leverages NFs. Given the particle approximations $\pi_{k-1}^N := \sum_{i=1}^N W_{k-1}^i \delta_{X_{k-1}^i}$ and Z_{k-1}^N at time $k-1$, AFT computes an approximation π_k^N and Z_k^N by performing four main sub-steps: *Transport, Importance Sampling, Resampling and Mutation*, as summarized in Algorithm 1. Whenever the index i is used in the algorithm, we mean ‘for all $i \in \{1, \dots, N\}$ ’. These four sub-steps are now detailed below.

Algorithm 1 Annealed Flow Transport

- 1: **Input:** number of particles N , unnormalized annealed targets $\{\gamma_k\}_{k=0}^K$ such that $\gamma_0 = \pi_0$ and $\gamma_K = \gamma$, resampling threshold $A \in [1/N, 1)$.
- 2: **Output:** Approximations π_K^N and Z_K^N of π and Z .
- 3: Sample $X_0^i \sim \pi_0$ and set $W_0^i = \frac{1}{N}$ and $Z_0^N = 1$.
- 4: **for** $k = 1, \dots, K$ **do**
- 5: Compute $\mathcal{L}_k^N(T)$ using (8).
- 6: Solve $T_k \leftarrow \operatorname{argmin}_{T \in \mathcal{T}} \mathcal{L}_k^N(T)$ using e.g. SGD.
- 7: Transport particles: $\tilde{X}_k^i = T_k(X_{k-1}^i)$.
- 8: Estimate normalizing constant Z_k :

$$Z_k^N \leftarrow Z_{k-1}^N \left(\sum_{i=1}^N W_{k-1}^i G_{k, T_k}(X_{k-1}^i) \right).$$
- 9: Compute IS weights:

$$w_k^i \leftarrow W_{k-1}^i G_{k, T_k}(X_{k-1}^i) \text{ // unnormalized}$$

$$W_k^i \leftarrow \frac{w_k^i}{\sum_{j=1}^N w_k^j} \text{ // normalized}$$
- 10: Compute effective sample size ESS_k^N using (10).
- 11: **if** $\text{ESS}_k^N/N \leq A$ **then**
- 12: Resample N particles denoted abusively also \tilde{X}_k^i according to the weights W_k^i , then set $W_k^i = \frac{1}{N}$.
- 13: **end if**
- 14: Sample $X_k^i \sim K_k(\tilde{X}_k^i, \cdot)$. // MCMC
- 15: **end for**

3.1. Transport map estimation

In this step, we learn a NF T_k that moves each sample X_{k-1} from π_{k-1} to a sample $\tilde{X}_k = T_k(X_{k-1})$ as close as possible to π_k by minimizing an estimate of $\text{KL}(T_{\#}\pi_{k-1} || \pi_k)$ over a set \mathcal{T} of NFs. This KL can be decomposed as a sum of a loss term $\mathcal{L}_k(T)$ and a term $\log \frac{Z_k}{Z_{k-1}}$ that can be ignored as it is independent of the NF T . A simple change of variables allows us to express the loss term $\mathcal{L}_k(T)$ as an expectation under π_{k-1} of some tractable function $x \mapsto h_T(x)$:

$$\begin{aligned} \mathcal{L}_k(T) &:= \pi_{k-1}[h_T], \\ h_T(x) &:= V_k(T(x)) - V_{k-1}(x) - \log |\nabla T(x)|. \end{aligned} \quad (7)$$

The Jacobian determinant of T in (7) can be evaluated efficiently for NFs while the expectation under π_{k-1} can be estimated using π_{k-1}^N thus yielding the empirical loss:

$$\mathcal{L}_k^N(T) := \sum_{i=1}^N W_{k-1}^i h_T(X_{k-1}^i). \quad (8)$$

In practice, (8) is optimized over the NF parameters using gradient descent. The resulting NF T_k is then used to transport each particle X_{k-1}^i to $\tilde{X}_k^i = T_k(X_{k-1}^i)$ ². However, the loss (8) being not necessarily convex, the solution T_k is likely to be sub-optimal. This is not an issue, since IS is used to correct for such approximation error as we will see

²We should write T_k^N to indicate the dependence of our estimate of N but do not to simplify notation.

next. We also emphasize that the convergence results for this scheme presented in Section 4 do not require finding a global minimizer of this non-convex optimization problem.

3.2. Importance Sampling, Resampling and Mutation

Importance Sampling. This step corrects for the NF T_k being only an approximate transport between π_{k-1} and π_k . In this case, we have $M_k^{\text{trans}}(x, x') = \delta_{T_k(x)}(x')$ and by selecting $L_{k-1}^{\text{trans}}(x, x') = \delta_{T_k^{-1}(x')}(x)$ then the incremental IS weight (3) is given by a simple change-of-variables formula

$$G_{k, T_k}(x_{k-1}) = \frac{\gamma_k(T_k(x_{k-1})) |\nabla T_k(x_{k-1})|}{\gamma_{k-1}(x_{k-1})}. \quad (9)$$

Using (9), we can update the weights $w_k^i = W_{k-1}^i G_{k, T_k}(X_{k-1}^i)$ to account for the errors introduced by T_k . When T_k are exact transport maps from π_{k-1} to π_k , the incremental weight in (9) becomes constant and equal to the ratio Z_k/Z_{k-1} . Thus, introducing the NF T_k can be seen as a way to reduce the variance of the IS weights in the SMC sampler.

Resampling. As discussed in Section 2.2, resampling can be very beneficial but it should only be performed when the variance of the IS weights is too high (Liu and Chen, 1995) as measured by the Effective Sample Size (ESS)

$$\text{ESS}_k^N = \left(\sum_{i=1}^N (W_k^i)^2 \right)^{-1}, \quad (10)$$

which is such that $\text{ESS}_k^N \in [1, N]$. When ESS_k^N/N is smaller than some prescribed threshold $A \in [1/N, 1)$ (we use $A = 0.3$), resampling is triggered and each particle \tilde{X}_k^i is then resampled without replacement from the set of N available particles $\{\tilde{X}_k^i\}_{i \in [1:N]}$ according to a multinomial distribution with weights $\{W_k^i\}_{i \in [1:N]}$. The weights are then reset to uniform ones; i.e. $W_k^i = \frac{1}{N}$. More sophisticated lower variance resampling schemes have also been proposed; see e.g. (Kitagawa, 1996; Chopin, 2004).

Mutation. The final step consists in mutating the particles using a π_k -invariant MCMC kernel K_k , i.e. using $X_k^i \sim K_k(\tilde{X}_k^i, \cdot)$. This allows particles to better explore the space.

Note that if the transport maps T_k were known, Algorithm 1 could be reinterpreted as a specific instance of a SMC as detailed in Section 2 where at each time $k \geq 1$ we perform two time steps of a standard SMC sampler by applying first a transport step $M_k^{\text{trans}}(x, x') = \delta_{T_k(x)}(x')$ then a mutation step $M_k^{\text{mut}}(x, x') = K_k(x, x')$; see Appendix B.1 for details.

3.3. Variants and Extensions

Contrary to standard SMC, the estimates Z_k^N returned by Algorithm 1 are biased because of the dependence of the NF T_k on the particles. To obtain unbiased estimates of Z_k and to avoid over-fitting of the NF to the N particles, a variant of Algorithm 1 described in Algorithm 2 (see Appendix F) is used in the experimental evaluation. This variant employs three sets of particles: the *training set* is used to evaluate the loss (8), the *validation set* is used in a stopping criterion when learning the NF and the *test set* is independent from the rest and is computed sequentially using the learned NFs. It would also be possible to combine AFT with various extensions to SMC that were already proposed in the literature. For example, we can select adaptively the annealing parameters β_k to ensure the ESS only decreases by a pre-determined percentage (Jasra et al., 2011; Schäfer and Chopin, 2013; Beskos et al., 2016; Zhou et al., 2016) or use the approximation of π_k obtained at step 13 of Algorithm 1 to determine the parameters of the MCMC kernel K_k (Del Moral et al., 2012a; Buchholz et al., 2021).

4. Asymptotic analysis

We establish here a law of large numbers and a CLT for the particle estimates $\pi_k^N[f]$ and Z_k^N of $\pi_k[f]$ and Z_k . We denote by \xrightarrow{P} convergence in probability and by \xrightarrow{D} convergence in distribution.

4.1. Weak law of large numbers

Theorem 1 shows that $\pi_k^N[f]$ and Z_k^N are consistent estimators of $\pi_k[f]$ and Z_k , hence of $\pi[f]$ and Z at time $k = K$.

Theorem 1 (weak law of large numbers). *Let f be a function s.t. $|f(x)| \leq C(1 + \|x\|^4)$ for all $x \in \mathcal{X}$ and for some $C > 0$. Under Assumptions (A) to (D) and for any $k \in 0, \dots, K$:*

$$(\mathcal{R}_k) : \quad \pi_k^N[f] \xrightarrow{P} \pi_k[f], \quad Z_k^N \xrightarrow{P} Z_k.$$

The result is proven in Appendix C.3 and relies on four assumptions stated in Appendix C.1: (A) on the smoothness of the Markov kernels K_k , (B) on the moments of π_k , (C) on the smoothness of the family of NFs and (D) on the boundedness of the incremental IS weight $G_{k,T}(x)$. Perhaps surprisingly, Theorem 1 does not require the NFs to converge as $N \rightarrow \infty$. This is a consequence of Proposition 9 in Appendix C.3 which ensures uniform consistency of the particle approximation regardless of the choice of the NFs. However, convergence of the NFs is required to obtain a CLT result as we see next. Theorem 4 of Appendix C.3 states a similar result for Algorithm 2 of Appendix F.

4.2. Central Limit theorem

Besides assumptions (A) to (D), we make five assumptions stated in Appendix C.1: (E) on the Markov kernels K_k strengthens (A) and is satisfied by many commonly used Markov kernels as shown in C.2. The smoothness assumptions (F) and (G) on the family \mathcal{T} of NFs and potentials V_k are also standard. Finally, (H) and (I) describe the asymptotic behavior of T_k . We do not require T_k to be a global minimizer of the loss \mathcal{L}_k^N , neither do we assume it to be an exact local minimum of \mathcal{L}_k^N . Instead, (H) only needs T_k to be an approximate local minimum of \mathcal{L}_k^N and (I) implies that T_k converges in probability towards a *strict local* minimizer T_k^* of \mathcal{L}_k as $N \rightarrow \infty$.

Before stating the CLT result, we need to introduce the asymptotic incremental variance $\mathbb{V}_k^{\text{inc}}[f]$ at iteration k . To this end, consider the set of limiting re-sampling times $\mathcal{K}_{\text{opt}} := \{k_0, \dots, k_P\} \subset \{0, \dots, K\}$ defined recursively by $k_{p+1} := \inf\{k_p < k : \text{nESS}_k \leq A\}$ and $k_{P+1} := K + 1$ where $\text{ESS}_k^N / N \xrightarrow{N \rightarrow \infty} \text{nESS}_k$ with

$$\text{nESS}_k = \frac{\pi_{k_p} [\mathbb{E}[w_k^* | X_{k_p}]]^2}{\pi_{k_p} [\mathbb{E}[(w_k^*)^2 | X_{k_p}]]},$$

the expectation being w.r.t. to $X_s \sim K_s(T_s^*(X_{s-1}), \cdot)$ for $k_p + 1 \leq s \leq k$, while $X_{k_p} \sim \pi_{k_p}$ and $w_k^* = \prod_{s=k_p+1}^k G_{s,T_s^*}(X_{s-1})$ is the product of the incremental IS weights using the locally optimal NFs T_s^* . The variance $\mathbb{V}_k^{\text{inc}}[f]$ at time k is given by:

$$\mathbb{V}_k^{\text{inc}}[f] = \begin{cases} Z_k^2 \text{Var}_{\pi_k}[f], & k \in \mathcal{K}, \\ Z_{k_p}^2 \pi_{k_p} [\mathbb{E}[(w_k^*)^2 \mathcal{G}_k[f] | X_{k_p}]], & k_p < k < k_{p+1}, \end{cases}$$

with $\mathcal{G}_k[f] := K_k[f^2](T_k^*(X_{k-1})) - K_k[f]^2(T_k^*(X_{k-1}))$.

Theorem 2 (Central limit theorem). *Let f be a real valued function s.t., for some $C > 0$, $f(x) \leq C(1 + \|x\|^2)$ and*

$$\|f(x) - f(x')\| \leq C(1 + \|x\|^3 + \|x'\|^3)\|x - x'\|.$$

Then, under Assumptions (A) to (I) and for $0 \leq k \leq K$:

$$(CLT_k) : \quad \begin{cases} \sqrt{N}(\pi_k^N[f] - \pi_k[f]) & \xrightarrow{D} \mathcal{N}(0, \mathbb{V}_k^\pi[f]), \\ \sqrt{N}(Z_k^N - Z_k) & \xrightarrow{D} \mathcal{N}(0, \mathbb{V}_k^\gamma[1]). \end{cases}$$

$\mathbb{V}_k^\gamma[f]$ and $\mathbb{V}_k^\pi[f]$ are defined recursively with $\mathbb{V}_0^\gamma[f] = \text{Var}_{\pi_0}[f]$ and

$$\begin{aligned} \mathbb{V}_k^\gamma[f] &= \mathbb{V}_k^{\text{inc}}[f] + \mathbb{V}_{k-1}^\gamma[Q_{k,T_k^*}[f]], \\ \mathbb{V}_k^\pi[f] &= Z_k^{-2} \mathbb{V}_k^\gamma[f - \pi_k[f]], \end{aligned}$$

where $Q_{k,T}(x, dy) := G_{k,T}(x)K_k(T(x), dy)$.

The asymptotic variances \mathbb{V}_k^γ and \mathbb{V}_k^π depend only on the maps T_k^* and not on the local variations of the family \mathcal{T} around T_k^* . This is a consequence of the particular form of the IS weights which provide an exact correction regardless of the NF selected as summarized by the following identity:

$$\pi_k[f] = \frac{\pi_{k-1}[Q_{k,T}[f]]}{\pi_{k-1}[G_{k,T}]}, \quad \forall T \in \mathcal{T}.$$

In the ideal case when T_k^* are exact transport maps from π_{k-1} to π_k , the ESS resampling criterion ESS_k^N/N is always equal to 1 and thus resampling is never triggered. Moreover, a direct computation shows that the asymptotic variance $\mathbb{V}_k^\pi[f]$ is exactly equal to $\text{Var}_{\pi_k}[f]$. This illustrates the benefit of introducing NFs to improve SMC. A proof is provided in Appendix C.4 along with a similar result (Theorem 5) for Algorithm 2.

5. Continuous-time scaling limit

We consider the setting where π_k arise from the time-discretization of a continuous-time path $(\Pi_t)_{[0,1]}$ of densities connecting π_0 to π ; i.e. π_k is of the form $\pi_k = \Pi_{t_k}$ with $t_k = k\lambda$ and $\lambda = \frac{1}{K}$. We write $V_t(x)$ and Z_t to denote the potential and unknown normalizing constant of Π_t and $\Gamma_t(x) = \exp(-V_t(x))$. We are here interested in identifying the ‘‘population’’ behavior of AFT (i.e. $N \rightarrow \infty$) as $\lambda \rightarrow 0$ when ULA kernels are used and no resampling is performed as in AIS. To simplify the analysis, we further consider in this Section the ideal situation where T_k is an exact minimizer of the population loss \mathcal{L}_k . Rigorous proofs of the results discussed here can be found in Appendix E.

5.1. Settings

Without resampling, the population version of AFT behaves as a sequential IS algorithm as defined in Section 2.1 where it is possible to *collapse* the transport step and mutation step into one single Markov kernel $M_k(x, x') = K_k(T_k(x), x')$. Similarly we can collapse the corresponding backward kernels and the resulting extended target distributions $\bar{\pi}_k$ are still given by (5) with modified IS weights

$$w_k(x_{0:k}) = \underbrace{\prod_{l=1}^k \frac{\gamma_l(x_l)}{\gamma_l K_l(x_l)}}_{r_k(x_{1:k})} \prod_{l=1}^k G_{l,T_l}(x_{l-1}), \quad (11)$$

where $r_k(x_{1:k}) = 1$ for π_l -invariant MCMC kernels K_l as used in Algorithm 1; see Appendix B.2 for a derivation. To ensure that the laws $\bar{\eta}_k$ and $\bar{\pi}_k$ of the Markov chain $X_{0:k}$ converge to some continuous-time limits, K_k are chosen to be ULA kernels³; i.e. $K_k(x, x')$ is a Gaussian density

in x' with mean $x - \lambda V_k(x)$ and covariance $2\lambda I$. In this case, $\gamma_k K_k(x) = \int \gamma_k(y) K_k(y, x) dy$ is intractable and so is $r_k(x_{1:k})$. This is not an issue as we are only interested here in identifying the theoretical scaling limit. To ensure $\bar{\eta}_k$ and $\bar{\pi}_k$ admit a limit, we also consider NFs of the form:

$$T(x) = x + \lambda A_\theta(x),$$

where $(\theta, x) \mapsto A_\theta(x)$ is from $\Theta \times \mathcal{X}$ to \mathcal{X} and Θ is a compact parameter space. The continuous-time analogues of NFs sequences $(T_k)_{k \in \{1, \dots, K\}}$ are represented by a set \mathcal{A} of time-dependent controls of the form $\alpha_t(x) = A_{\theta_t}(x)$, where $t \mapsto \theta_t$ is a 1-Lipschitz trajectory in Θ . To any control α corresponds an NFs sequence $(T_k)_{k \in \{1, \dots, K\}}$ defined by $T_k(x) = x + \lambda \alpha_{t_k}(x)$.

5.2. Continuous-time limits

Limiting forward process. Using a similar approach to (Dalalyan, 2017), the Markov chain $X_{0:K}$ under $\bar{\eta}_K$ converges towards a stochastic process $X_{[0,1]}$ defined by the following Stochastic Differential Equation (SDE)

$$dX_t = (\alpha_t(X_t) - \nabla V_t(X_t)) dt + \sqrt{2} dB_t, \quad (12)$$

where $X_0 \sim \pi_0$ and $(B_t)_{t \geq 0}$ is a standard Brownian motion. We denote by $\bar{\Lambda}_t^\alpha$ the joint distribution of this process up to time t and by Λ_t^α its marginal at time t .

Limiting weights. The weight $w_K(X_{0:K})$ in (11) is such that $r_K(X_{1:K}) \rightarrow 1$ as the invariant distribution of the ULA kernel K_k converges to π_k when $\lambda \rightarrow 0$ while the logarithm of the product of $G_{l,T_l}(X_{l-1})$ is a Riemann sum whose limiting value is the following integral:

$$\sum_{l=1}^K \log(G_{l,T_l}(X_{l-1})) \xrightarrow{\lambda \rightarrow 0} \int_0^1 g_s^\alpha(X_s) ds,$$

with $X_{[0,1]}$ defined in (12) and $g_t^\alpha(x)$ being the dominating term in the Taylor expansion of $\log(G_{l,T_l}(x))$ w.r.t. time:

$$g_t^\alpha(x) = \nabla \cdot \alpha_t(x) - \nabla_x V_t(x)^\top \alpha_t(x) - \partial_t V_t(x).$$

The limit of IS weights $w_k(X_{0:k})$ is thus identified as

$$w_t^\alpha(X_{[0,t]}) = \exp\left(\int_0^t g_s^\alpha(X_s) ds\right).$$

In the context of *non-equilibrium dynamics*, $g_t^\alpha(x)$ is known as *instantaneous work* (Rousset and Stoltz, 2006) and is constant in the ideal case where $\Pi_t = \Lambda_t^\alpha$.

Limiting objective. To identify a non-trivial limiting loss, we consider the following aggregation of all $\mathcal{L}_k(T_k)$

$$\mathcal{L}_\lambda^{tot}(\alpha) := \lambda^{-1} \sum_{k=1}^K \mathcal{L}_k(T_k). \quad (13)$$

³The random walk MH algorithm also admits a Langevin diffusion as scaling limit when $\lambda \rightarrow 0$ (Gelfand and Mitter, 1991; Choi, 2019) but the technical analysis is much more involved.

The next result shows that (13) converges towards a non-trivial loss $\mathcal{M}(\alpha)$ as $\lambda \rightarrow 0$ under three assumptions stated in Appendix E.2: (a) and (b) on the smoothness of $V_t(x)$ and $A_\theta(x)$ and (c) on the moments of Π_t .

Proposition 1. *Under Assumptions (a) to (c), for λ small enough, it holds that for all $\alpha \in \mathcal{A}$*

$$|\mathcal{L}_\lambda^{\text{tot}}(\alpha) - \mathcal{M}(\alpha)| \leq \lambda C,$$

where C is independent of λ and

$$\mathcal{M}(\alpha) = \frac{1}{2} \int_0^1 \left(\Pi_t \left[(g_t^\alpha)^2 \right] - \Pi_t [g_t^\alpha]^2 \right) dt. \quad (14)$$

The optimal NFs $(T_k)_{1:K}$ are thus expected to converge towards some α^* minimizing $\mathcal{M}(\alpha)$ over \mathcal{A} as made precise in Proposition 29 of Appendix E.6. Moreover when the class of NFs is expressive, i.e. \mathcal{A} is rich enough, then $\mathcal{M}(\alpha^*) = 0$ and thus g_t^α are constant and α^* satisfies the Partial Differential Equation (PDE)

$$0 = \nabla \cdot \alpha^*(x) - \nabla_x V_t(x)^\top \alpha^*(x) - \partial_t V_t(x) + \Pi_t [\partial_t V_t].$$

This PDE has appeared, among others, in Lelièvre et al. (2010, pp. 273–275) and (Vaikuntanathan and Jarzynski, 2008; Reich, 2011; Heng et al., 2021). Its solution defines a deterministic flow α_t^* that transports mass along the path $(\Pi_t^\alpha)_{[0,1]}$; i.e. if X_t is a solution to an ODE of the form $\dot{X}_t = \alpha_t^*(X_t)$ with initial values $X_0 \sim \Pi_0$, then $X_t \sim \Pi_t$.

Feynman–Kac measure. Given a control α , we consider the Feynman–Kac measure $\bar{\Pi}_t^\alpha$ defined for any bounded continuous functional f of the process $X_{[0,t]}$ in (12)

$$\bar{\Pi}_t^\alpha[f] = \frac{\bar{\Lambda}_t^\alpha[w_t^\alpha f]}{\bar{\Lambda}_t^\alpha[w_t^\alpha]}. \quad (15)$$

By a similar argument as in (Rousset and Stoltz, 2006), we show in Proposition 22 of Appendix E.3 that $\bar{\Pi}_t^\alpha$ admits Π_t as a marginal at time t regardless of the choice of α . Using the optimal control α^* in (12) and (15) gives rise to $\bar{\Lambda}^*$ and $\bar{\Pi}_t^*$ which are equal when $\mathcal{M}(\alpha^*) = 0$. Next, we show that $\bar{\Pi}_t^*$ is the scaling limit of $\bar{\pi}_k$.

5.3. Convergence to the continuous-time limit

As the measures $\bar{\pi}_k$ and $\bar{\Pi}_t^*$ are defined on different spaces, we construct a sequence of interpolating measures $\bar{\Pi}_t^\lambda$ defined over the same space as $\bar{\Pi}_t^*$ and whose marginal at the joint times $\{t_0, \dots, t_K\}$ is exactly equal to $\bar{\pi}_k$; see Appendix E.1 for details. Theorem 3 provides a convergence rate for the interpolating measures $\bar{\Pi}_t^\lambda$ towards $\bar{\Pi}_t^*$ as $\lambda \rightarrow 0$, thus establishing $\bar{\Pi}_t^*$ as the scaling limit of $\bar{\pi}_k$; see Appendix E.6 for the proof.

Theorem 3. *Under Assumptions (a) to (g), then for λ small enough there exists a finite C such that for any $t \in [0, 1]$:*

$$\text{KL}(\bar{\Pi}_t^* || \bar{\Pi}_t^\lambda) \leq C\sqrt{\lambda}.$$

This result relies on Assumptions (d) to (g) in addition to Assumptions (a) to (c) which are also stated in Appendix E.2. (d) strengthens assumption (c) on the moments of Π_t . (e) guarantees the existence of a solution α^* in \mathcal{A} minimizing \mathcal{M} and controls the local behavior of \mathcal{M} near α^* . (f) guarantees the existence of solutions α^λ in \mathcal{A} minimizing $\mathcal{L}_\lambda^{\text{tot}}(\alpha)$ for any $\lambda = \frac{1}{K}$. Finally, (g) ensures the optimal control α^* induces bounded IS weights.

6. Applications

In this section we detail the practical implementation of AFT and empirically investigate performance against relevant baselines.

As discussed in Section 3.3, we use three sets of particles—‘train, test and validation’ which improves robustness, avoids overfitting the flow to the particles and gives unbiased estimates of Z when using the test set. We initialize our flows to the identity for the optimization at each time step. Algorithm 2, in the supplement gives a summary.

We concentrate our empirical value evaluation on the learnt flow, which is equivalent to using the test set particles. The learnt flow is of interest in deploying an efficient sampler on large scale distributed parallel compute resources. It is also of interest for inclusion as a subroutine in a larger system. Since modern hardware enables us to do large computations in parallel, the computation is dominated by algorithmic steps that are necessarily done serially, particularly repeat applications of the Markov kernel (Lee et al., 2010).

As our primary, strong, baseline for AFT, we use a standard instance of SMC samplers (Del Moral et al., 2006; Zhou et al., 2016) which corresponds to AIS with adaptive resampling and is also known as population annealing in physics (Hukushima and Iba, 2003; Barash et al., 2017). As observed many times in the literature and in our experiments, SMC estimates are of lower variance than AIS estimates. This SMC baseline is closely related to AFT since it corresponds to using AFT with an identity transformation $T_k(x) = x$ instead of a learnt flow.

We largely use the number of transitions K as a proxy for compute time. This is valid when the cost of evaluating the flow is modest relative to that of the other algorithmic steps, as it is for the trained flows in all non-trivial cases we consider. We only consider flows of no more than a few layers per transition, but deeper flows could start to form an appreciable part of the serial computation. In some cases, we use variational inference (VI) as a measure of behaviour

without MCMC. In this case, evaluation time is not comparable and faster. Since we concentrate on trained flows, we do not evaluate training time in the benchmarks considered, though fast training of AFT could be of interest in further work. Both SMC and AFT use the same Markov kernels K_k , using HMC except where otherwise stated. We tune the step size to have a reasonable acceptance probability based on preliminary runs of SMC using a modest K . Then for larger K experiments, we linearly interpolate the step sizes chosen on the preliminary runs. We always use a linearly spaced geometric schedule and the initial distribution is always a multivariate standard normal. We repeat experiments 100 times. Further experimental details may be found in Appendix G. We plan to make the code available within <https://github.com/deepmind>.

6.1. Illustrative example

We start with an easily visualized two dimensional target density as shown in Figure 1. All sensible methods should work in such a low dimensional case but it can still be informative. We investigate two families of flows based on rational quadratic splines (Durkan et al., 2019). The first (termed AFTmf for mean field) operates on the two dimensions separately. The second family (denoted AFT in Figure 1) adds dependence to the splines using inverse autoregressive flows (Kingma et al., 2016). Figure 1 shows weighted samples from AFT as we anneal from a standard normal distribution. Figure 2 (a) shows that AFT reduces the variance of the normalizing constant estimator relative to SMC. Conversely, we see that AFTmf actually *increases* the variance relative to SMC for small numbers of transitions. Since the factorized approximation cannot model the dependence of variables the optimum of the KL underestimates the variance of the target. Later, in Sections 6.3 and 6.4, we discuss examples where even a simple NF leads to an improvement for a modest number of transitions.




Figure 1: Weighted samples for a 2-D target density with AFT. The colours show the normalized weights which are clipped at the 95th percentile for clarity. The final samples are visually indistinguishable from the target.

6.2. Funnel distribution

We next evaluate the performance of the method on Neal’s ten-dimensional ‘funnel’ distribution (Neal, 2003):

$$x_0 \sim \mathcal{N}(0, \sigma_f^2), \quad x_{1:9}|x_0 \sim \mathcal{N}(\mathbf{0}, \exp(x_0)\mathbf{I}).$$




Figure 2: Results from the four different examples. Cyan lines denote gold standard values of the log normalizing constant. In (c) and (d) green horizontal lines denote the median value for an importance sampling estimate based on variational inference. Note that in (d) the small AFT error bars can make it difficult to see - it can be found next to the gold standard value in each case.

Here, $\sigma_f^2 = 9$. Many MCMC methods find this example challenging because there is a variety of length scales depending on the value of x_0 and because marginally $x_{1:9}$ has heavy tails. We use here slice sampling instead of HMC for the Markov kernels as recommended in (Neal, 2003). For each flow we use an affine inverse autoregressive flow (Kingma et al., 2016). In this example, we also compare against VI (Rezende et al., 2014) which uses the same number of flows. We then apply a simple importance correction to the VI samples to give an unbiased estimate of the normalizing constant. Figure 2 (b) shows the results. We see that for small number of flows/transitions VI performs best, followed by AFT. However, VI shows little further improvement with additional flows and in this regime AFT, SMC and VI perform similarly.

6.3. Variational Autoencoder latent space

For our next example, we trained a variational autoencoder (Kingma and Welling, 2014; Rezende et al., 2014) with convolution on the binarized MNIST dataset (Salakhutdinov and Murray, 2008) and a normal encoder distribution with diagonal covariance. Using the fixed, trained, generative decoder network we investigated the quality of normalizing constant estimation which in this case corresponds to the likelihood of a data point with the distribution over the 30 latent variables marginalized out (Wu et al., 2017).

Using long run SMC on the 10000 point test set we estimate that the hold out log-likelihood per data point for the network is -86.3. For each data point we also found the optimal variational normal approximation with diagonal covariance rather than using the amortized variational approximation. Using this optimal normal approximation we investigated its variance when used as an importance proposal for the likelihood. We estimate the mean absolute error for the estimator across the test set was 0.6 nats per data point which indicates that the VI is often performing well. There was a tail of digits where VI performed relatively worse. Since these ‘difficult’ digits constituted a more challenging inference problem, we used one of these, with a VI/SMC error of 1.5 nats, to comparatively benchmark AFT in the detailed manner used in our other examples.

For the AFT flow we used an affine transformation with diagonal linear transformation matrix. The baseline VI approximation can be thought of the pushforward of a standard normal through this ‘diagonal affine’ flow. Note that since diagonal affine transformations are closed under composition there would obtain no additional expressiveness in the baseline VI approximation from adding more of them.

Figure 2 (c) shows the results for this example. Both AFT and SMC reduce in variance as the number of temperatures increases and exceed the performance of the variational

baseline. AFT has a notably lower variance than SMC for 10 and 30 temperatures- which shows the incorporation of the flows is beneficial in this case. Results for other difficult digits are shown in the appendix where the qualitative trend is similar.

6.4. Log Gaussian Cox process

We evaluate here the performance of AFT for estimating the normalizing constant of a log Gaussian Cox process applied to modelling the positions of pine saplings in Finland (Møller et al., 1998). We consider points on a discretized $d = M \times M = 1600$ grid. This results in the target density

$$\gamma(x) = \mathcal{N}(x; \mu, K) \prod_{i \in [1:M]^2} \exp(x_i y_i - a \exp(x_i)).$$

This challenging high-dimensional problem is a commonly used benchmark in the SMC literature (Heng et al., 2020; Buchholz et al., 2021). The mean and covariance function match those estimated by (Møller et al., 1998) and are detailed in the Appendix. The supplement also discusses the effect of pre-conditioners on the mixing of the Markov kernel. For the NF we again used the diagonal affine transformation. The approximating family is the push forward of the previous target distribution and thus even a simple flow can result in a good approximation. It is also fast to evaluate. Figure 2 (d) shows that the baseline VI approximation is unable to capture the posterior correlation and that AFT gives significantly more accurate results than SMC for a given number of transitions. As such, the Markov kernel and flow complement each other in this case.

7. Conclusion

We proposed Annealed Flow Transport which combines SMC samplers and normalizing flows. We studied its asymptotic behavior and showed the benefit of introducing learned flows to reduce the asymptotic variance. We identified the scaling limit of AFT as a controlled Feynman–Kac measure whose optimal control solved a flow transport problem in an idealized setting. Empirically we found multiple cases where trained AFT gave lower variance estimates than SMC for the same number of transitions, showing that we can combine the advantages of both SMC and normalizing flows. We believe AFT will be particularly useful in scenarios where it is both difficult to design fast mixing MCMC kernels and very good flows so that neither SMC nor VI provide low variance estimates.

8. Acknowledgements

The authors would like to thank Danilo Rezende, Arthur Gretton and Taylan Cemgil.

References

- Akyildiz, Ö. D. and Míguez, J. (2020). Nudging the particle filter. *Statistics and Computing*, 30(2):305–330.
- Barash, L. Y., Weigel, M., Borovskỳ, M., Janke, W., and Shchur, L. N. (2017). GPU accelerated population annealing algorithm. *Computer Physics Communications*, 220:341–350.
- Beskos, A., Jasra, A., Kantas, N., and Thiery, A. (2016). On the convergence of adaptive sequential Monte Carlo methods. *The Annals of Applied Probability*, 26(2):1111–1146.
- Beskos, A., Pinski, F., Sanz-Serna, J., and Stuart, A. (2011). Hybrid Monte Carlo on Hilbert spaces. *Stochastic Processes and their Applications*, 121(10):2201 – 2230.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations of Python+NumPy programs.
- Buchholz, A., Chopin, N., and Jacob, P. E. (2021). Adaptive tuning of Hamiltonian Monte Carlo within sequential Monte Carlo. *Bayesian Analysis to appear - arXiv preprint arXiv:1808.07730*.
- Caterini, A. L., Doucet, A., and Sejdinovic, D. (2018). Hamiltonian variational auto-encoder. In *Advances in Neural Information Processing Systems*, pages 8167–8177.
- Choi, M. C. (2019). Universality of the Langevin diffusion as scaling limit of a family of Metropolis–Hastings processes i: fixed dimension. *arXiv preprint arXiv:1907.10318*.
- Chopin, N. (2002). A sequential particle filter method for static models. *Biometrika*, 89(3):539–552.
- Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *The Annals of Statistics*, 32(6):2385–2411.
- Crooks, G. E. (1998). Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems. *Journal of Statistical Physics*, 90(5–6):1481–1487.
- Dai, C., Heng, J., Jacob, P. E., and Whiteley, N. (2020). An invitation to sequential Monte Carlo samplers. *arXiv preprint arXiv:2007.11936*.
- Dalalyan, A. S. (2017). Theoretical guarantees for approximate sampling from smooth and log-concave densities. *Journal of the Royal Statistical Society: Series B*, 3(79):651–676.
- Del Moral, P. (2004). *Feynman-Kac Formulae: Genealogical and Interacting Particle Approximations*. Springer.
- Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B*, 68(3):411–436.
- Del Moral, P., Doucet, A., and Jasra, A. (2012a). An adaptive sequential Monte Carlo method for approximate Bayesian computation. *Statistics and Computing*, 22(5):1009–1020.
- Del Moral, P., Doucet, A., and Jasra, A. (2012b). On adaptive resampling strategies for sequential Monte Carlo methods. *Bernoulli*, 18(1):252–278.
- Dillon, J. V., Langmore, I., Tran, D., Brevdo, E., Vasudevan, S., Moore, D., Patton, B., Alemi, A., Hoffman, M., and Saurous, R. A. (2017). TensorFlow Distributions. *arXiv preprint arXiv:1711.10604*.
- Domke, J. and Sheldon, D. R. (2018). Importance weighting and variational inference. In *Advances in Neural Information Processing Systems*, pages 4470–4479.
- Douc, R. and Moulines, E. (2008). Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *The Annals of Statistics*, 36(5):2344–2376.
- Dudley, R. M. (2018). *Real analysis and Probability*. CRC Press.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. (2019). Neural spline flows. In *Advances in Neural Information Processing Systems*.
- El Moselhy, T. A. and Marzouk, Y. M. (2012). Bayesian inference with optimal maps. *Journal of Computational Physics*, 231(23):7815–7850.
- Everitt, R. G., Culliford, R., Medina-Aguayo, F., and Wilson, D. J. (2020). Sequential Monte Carlo with transformations. *Statistics and Computing*, 30(3):663–676.
- Gao, C., Isaacson, J., and Krause, C. (2020). i-flow: High-dimensional Integration and Sampling with normalizing flows. *arXiv preprint arXiv:2001.05486*.
- Gelfand, S. B. and Mitter, S. K. (1991). Weak convergence of Markov chain sampling methods and annealing algorithms to diffusions. *Journal of Optimization Theory and Applications*, 68(3):483–498.

- Germain, M., Gregor, K., Murray, I., and Larochelle, H. (2015). Made: Masked autoencoder for distribution estimation. In Bach, F. and Blei, D., editors, *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 881–889, Lille, France. PMLR.
- Gilks, W. R. and Berzuini, C. (2001). Following a moving target - Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society: Series B*, 63(1):127–146.
- Goyal, A. G. A. P., Ke, N. R., Ganguli, S., and Bengio, Y. (2017). Variational walkback: Learning a transition operator as a stochastic recurrent net. In *Advances in Neural Information Processing Systems*, pages 4392–4402.
- Guarniero, P., Johansen, A. M., and Lee, A. (2017). The iterated auxiliary particle filter. *Journal of the American Statistical Association*, 112(520):1636–1647.
- Han, J. and Liu, Q. (2017). Stein variational adaptive importance sampling. *Uncertainty in Artificial Intelligence*.
- Heng, J., Bishop, A. N., Deligiannidis, G., and Doucet, A. (2020). Controlled sequential Monte Carlo. *The Annals of Statistics*, 48(5):2904–2929.
- Heng, J., Doucet, A., and Pokern, Y. (2021). Gibbs flow for approximate transport with applications to Bayesian computation. *Journal of the Royal Statistical Society Series B*, 83(1):156–187.
- Hennigan, T., Cai, T., Norman, T., and Babuschkin, I. (2020). Haiku: Sonnet for JAX.
- Hessel, M., Budden, D., Viola, F., Rosca, M., Sezener, E., and Hennigan, T. (2020). Optax: composable gradient transformation and optimisation, in jax.
- Hoffman, M. D. (2017). Learning deep latent Gaussian models with Markov chain Monte Carlo. In Precup, D. and Teh, Y. W., editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1510–1519. PMLR.
- Huang, C.-W., Tan, S., Lacoste, A., and Courville, A. C. (2018). Improving explorability in variational inference with annealed variational objectives. In *Advances in Neural Information Processing Systems*, pages 9701–9711.
- Huang, S., Makhzani, A., Cao, Y., and Grosse, R. (2020). Evaluating lossy compression rates of deep generative models. In *International Conference on Machine Learning*, pages 4444–4454. PMLR.
- Hukushima, K. and Iba, Y. (2003). Population annealing and its application to a spin glass. In *AIP Conference Proceedings*, volume 690, pages 200–206. American Institute of Physics.
- Jarzynski, C. (1997). Nonequilibrium equality for free energy differences. *Physical Review Letters*, 78(14):2690–2963.
- Jasra, A., Stephens, D. A., Doucet, A., and Tsagaris, T. (2011). Inference for Lévy-driven stochastic volatility models via adaptive sequential Monte Carlo. *Scandinavian Journal of Statistics*, 38(1):1–22.
- Kappen, H. J. and Ruiz, H. C. (2016). Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, 162(5):1244–1266.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. (2016). Improved variational inference with inverse autoregressive flow. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational Bayes. *ICLR*.
- Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5(1):1–25.
- Künsch, H. R. (2005). Recursive Monte Carlo filters: algorithms and theoretical analysis. *The Annals of Statistics*, 33(5):1983–2021.
- Le, T. A., Igl, M., Rainforth, T., Jin, T., and Wood, F. (2018). Auto-encoding sequential Monte Carlo. In *ICLR*.
- Lee, A., Yau, C., Giles, M. B., Doucet, A., and Holmes, C. C. (2010). On the utility of graphics cards to perform massively parallel simulation of advanced monte carlo methods. *Journal of Computational and Graphical Statistics*, 19(4):769–789.
- Lei Ba, J., Kiros, J. R., and Hinton, G. E. (2016). Layer Normalization. *arXiv e-prints*.
- Lelièvre, T., Rousset, M., and Stoltz, G. (2010). *Free Energy Computations: A Mathematical Perspective*. World Scientific.

- Li, Q. and Chen, Y. (2019). Rate distortion via deep learning. *IEEE Transactions on Communications*, 68(1):456–465.
- Liu, C., Zhuo, J., Cheng, P., Zhang, R., and Zhu, J. (2019). Understanding and accelerating particle-based variational inference. In *International Conference on Machine Learning*, pages 4082–4092.
- Liu, J. S. and Chen, R. (1995). Blind deconvolution via sequential imputations. *Journal of the American Statistical Association*, 90(430):567–576.
- Liu, Q. and Wang, D. (2016). Stein variational gradient descent: a general purpose Bayesian inference algorithm. In *Advances in Neural Information Processing Systems*.
- Llorente, F., Martino, L., Delgado, D., and Lopez-Santiago, J. (2020). Marginal likelihood computation for model selection and hypothesis testing: an extensive review. *arXiv preprint arXiv:2005.08334*.
- MacEachern, S. N., Clyde, M., and Liu, J. S. (1999). Sequential importance sampling for nonparametric Bayes models: The next generation. *Canadian Journal of Statistics*, 27(2):251–267.
- Maddison, C. J., Lawson, J., Tucker, G., Heess, N., Norouzi, M., Mnih, A., Doucet, A., and Teh, Y. (2017). Filtering variational objectives. In *Advances in Neural Information Processing Systems*, pages 6573–6583.
- Marzouk, Y., Moselhy, T., Parno, M., and Spantini, A. (2016). Sampling via measure transport: An introduction. *Handbook of Uncertainty Quantification*, pages 1–41.
- Mnih, A. and Rezende, D. (2016). Variational inference for Monte Carlo objectives. In *International Conference on Machine Learning*, pages 2188–2196. PMLR.
- Møller, J., Syversveen, A. R., and Waagepetersen, R. P. (1998). Log Gaussian Cox processes. *Scandinavian Journal of Statistics*, 25(3):451–482.
- Naesseth, C. A., Linderman, S. W., Ranganath, R., and Blei, D. M. (2018). Variational sequential Monte Carlo. In *AISTATS*.
- Neal, R. (2011). MCMC using Hamiltonian dynamics. *Handbook of Markov chain Monte Carlo*.
- Neal, R. M. (2001). Annealed importance sampling. *Statistics and Computing*, 11(2):125–139.
- Neal, R. M. (2003). Slice sampling. *The Annals of Statistics*, 31(3):705–767.
- Nicoli, K. A., Nakajima, S., Strodthoff, N., Samek, W., Müller, K.-R., and Kessel, P. (2020). Asymptotically unbiased estimation of physical observables with neural samplers. *Physical Review E*, 101(2):023304.
- Noé, F., Olsson, S., Köhler, J., and Wu, H. (2019). Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457):eaaw1147.
- Olmez, S. Y., Taghvaei, A., and Mehta, P. G. (2020). Deep fpf: Gain function approximation in high-dimensional setting. In *59th IEEE Conference on Decision and Control (CDC)*, pages 4790–4795. IEEE.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2019). Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*.
- Reich, S. (2011). A dynamical systems framework for intermittent data assimilation. *BIT Numerical Mathematics*, 51(1):235–249.
- Reich, S. and Weissmann, S. (2021). Fokker–Planck particle systems for Bayesian inference: Computational approaches. *SIAM/ASA Journal on Uncertainty Quantification*, 9(2):446–482.
- Rezende, D. J. and Mohamed, S. (2015). Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning - Volume 37, ICML’15*, pages 1530–1538. JMLR.org.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *ICML*, pages 1278–1286.
- Richard, J.-F. and Zhang, W. (2007). Efficient high-dimensional importance sampling. *Journal of Econometrics*, 141(2):1385–1411.
- Rousset, M. and Stoltz, G. (2006). Equilibrium sampling from nonequilibrium dynamics. *Journal of Statistical Physics*, 123(6):1251–1272.
- Salakhutdinov, R. and Murray, I. (2008). On the quantitative analysis of Deep Belief Networks. In *Proceedings of the 25th Annual International Conference on Machine Learning (ICML 2008)*, pages 872–879.
- Salimans, T., Kingma, D., and Welling, M. (2015). Markov chain Monte Carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226.
- Schäfer, C. and Chopin, N. (2013). Sequential Monte Carlo on large binary sampling spaces. *Statistics and Computing*, 23(2):163–184.

- Sen, B. (2018). A gentle introduction to empirical process theory and applications. *Lecture Notes, Columbia University*.
- Taghvaei, A., Mehta, P. G., and Meyn, S. P. (2020). Diffusion map-based algorithm for gain function approximation in the feedback particle filter. *SIAM/ASA Journal on Uncertainty Quantification*, 8(3):1090–1117.
- Thin, A., Kotelevskii, N., Durmus, A., Panov, M., Moulines, E., and Doucet, A. (2021). Monte Carlo variational auto-encoders. *International Conference on Machine Learning*.
- Turner, R. E. and Sahani, M. (2011). Two problems with variational expectation maximisation for time-series models. In Barber, D., Cemgil, T., and Chiappa, S., editors, *Bayesian Time Series Models*, chapter 5, pages 109–130. Cambridge University Press.
- Vaikuntanathan, S. and Jarzynski, C. (2008). Escorted free energy simulations: Improving convergence by reducing dissipation. *Physical Review Letters*, 100(19):190601.
- Vaikuntanathan, S. and Jarzynski, C. (2011). Escorted free energy simulations. *The Journal of Chemical Physics*, 134(5):054107.
- Van der Vaart, A. W. (2000). *Asymptotic Statistics*. Cambridge University Press.
- Wang, Y. and Li, W. (2019). Accelerated information gradient flow. *arXiv preprint arXiv:1909.02102*.
- Wirnsberger, P., Ballard, A. J., Papamakarios, G., Abercrombie, S., Racanière, S., Pritzel, A., Rezende, D., and Blundell, C. (2020). Targeted free energy estimation via learned mappings. *The Journal of Chemical Physics*, 153(14):144112.
- Wu, H., Köhler, J., and Noé, F. (2020). Stochastic normalizing flows. In *Advances in Neural Information Processing Systems*.
- Wu, Y., Burda, Y., Salakhutdinov, R., and Grosse, R. B. (2017). On the quantitative analysis of decoder-based generative models. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.
- Zhou, Y., Johansen, A. M., and Aston, J. A. (2016). Toward automatic model comparison: An adaptive sequential Monte Carlo approach. *Journal of Computational and Graphical Statistics*, 25(3):701–726.
- Zhu, M., Liu, C., and Zhu, J. (2020). Variance reduction and quasi-Newton for particle-based variational inference. In *International Conference on Machine Learning*, pages 11576–11587.

A. Using measure-theoretic notation

The Markov transition kernel M_k is defined as a map $M_k : \mathcal{X} \times \mathcal{B}(\mathcal{X}) \rightarrow [0, 1]$ where $\mathcal{B}(\mathcal{X})$ are the Borel sets, L_{k-1} is defined similarly. The joint distribution of the non-homogeneous Markov chain of initial distribution $X_0 \sim \pi_0$ and Markov transition kernel $X_l \sim M_l(X_{l-1}, \cdot)$ at time l is given at time k by

$$\bar{\eta}_k(\mathrm{d}x_{0:k}) = \pi_0(\mathrm{d}x_0) \prod_{l=1}^k M_l(x_{l-1}, \mathrm{d}x_l).$$

SMC samplers rely on the following target distribution of the form

$$\bar{\pi}_k(\mathrm{d}x_{0:k}) = \pi_k(\mathrm{d}x_k) \prod_{l=0}^{k-1} L_l(x_{l+1}, \mathrm{d}x_l),$$

and $\gamma_k(\mathrm{d}x_{0:k}) = Z_k \bar{\pi}_k(\mathrm{d}x_{0:k})$. When $\pi_l(\mathrm{d}x') L_{l-1}(x', \mathrm{d}x)$ is absolutely continuous w.r.t. $\pi_l(\mathrm{d}x) M_l(x, \mathrm{d}x')$, then we can define its Radon-Nikodym derivative and the incremental importance weight G_l through

$$G_l(x_{l-1}, x_l) = \frac{Z_l}{Z_{l-1}} \frac{\pi_l(\mathrm{d}x_l) L_{l-1}(x_l, \mathrm{d}x_{l-1})}{\pi_{l-1}(\mathrm{d}x_{l-1}) M_l(x_{l-1}, \mathrm{d}x_l)} = \frac{\gamma_l(\mathrm{d}x_l) L_{l-1}(x_l, \mathrm{d}x_{l-1})}{\gamma_{l-1}(\mathrm{d}x_{l-1}) M_l(x_{l-1}, \mathrm{d}x_l)}.$$

If G_l is defined for $l = 1, \dots, k$, then $\bar{\pi}_k$ is absolutely continuous w.r.t. $\bar{\eta}_k$ so we can write

$$\bar{\pi}_k[f] = \frac{\bar{\eta}_k[w_k f]}{\bar{\eta}_k[w_k]}, \quad \text{where} \quad w_k(x_{0:k}) = Z_k \frac{\mathrm{d}\bar{\pi}_k}{\mathrm{d}\bar{\eta}_k}(x_{0:k}) = \frac{\mathrm{d}\gamma_k}{\mathrm{d}\bar{\eta}_k}(x_{0:k}) = \prod_{l=1}^k G_l(x_{l-1}, x_l).$$

If M_k is π_k -invariant then (Crooks, 1998; Neal, 2001) select L_{k-1} at the reversal of M_k , that is the kernel satisfying $\pi_k(\mathrm{d}x) M_k(x, \mathrm{d}x') = \pi_k(\mathrm{d}x') L_{k-1}(x', \mathrm{d}x)$ and in this case

$$G_l(x_{l-1}, x_l) = \frac{Z_l}{Z_{l-1}} \frac{\mathrm{d}\pi_l}{\mathrm{d}\pi_{l-1}}(x_{l-1}) = \frac{\mathrm{d}\gamma_l}{\mathrm{d}\gamma_{l-1}}(x_{l-1}).$$

B. Extended proposal and target of AFT algorithm

In this section, assuming the transport maps T_k are here fixed, we write explicitly the extended proposal and target distributions used implicitly by the AFT algorithm if no resampling was used.

B.1. Non-collapsed version

In this case, we sample $X_0 \sim \pi_0(\cdot)$ at $k = 0$ then use $\tilde{X}_k = T_k(X_{k-1})$ followed by $X_k = K_k(\tilde{X}_k, \cdot)$ at time $k \geq 1$. Hence, using the notation $M_l^{\text{trans}}(x, x') = \delta_{T_l(x)}(x')$ and $M_l^{\text{mut}}(x, x') = K_l(x, x')$, the proposal at time k after the transport step is of the form

$$\bar{\eta}_k(x_{0:k-1}, \tilde{x}_{1:k}) = \pi_0(x_0) \left(\prod_{l=1}^{k-1} M_l^{\text{trans}}(x_{l-1}, \tilde{x}_l) M_l^{\text{mut}}(\tilde{x}_l, x_l) \right) M_k^{\text{trans}}(x_{k-1}, \tilde{x}_k),$$

and the target is

$$\bar{\pi}_k(x_{0:k-1}, \tilde{x}_{1:k}) = \pi_k(\tilde{x}_k) L_{k-1}^{\text{trans}}(\tilde{x}_k, x_{k-1}) \left(\prod_{l=0}^{k-2} L_l^{\text{mut}}(x_{l+1}, \tilde{x}_{l+1}) L_l^{\text{trans}}(\tilde{x}_{l+1}, x_l) \right),$$

where $L_{l-1}^{\text{trans}}(x, x') = \delta_{T_{l-1}^{-1}(x)}(x')$ and $L_{l-1}^{\text{mul}}(x, x') = \pi_l(x') M_l^{\text{mul}}(x', x) / \pi_l(x)$. After the mutation step at time k , the proposal is

$$\bar{\eta}_k(x_{0:k}, \tilde{x}_{1:k}) = \bar{\eta}_k(x_{0:k-1}, \tilde{x}_{1:k}) M_k^{\text{mut}}(\tilde{x}_k, x_k) = \pi_0(x_0) \left(\prod_{l=1}^k M_l^{\text{trans}}(x_{l-1}, \tilde{x}_l) M_l^{\text{mut}}(\tilde{x}_l, x_l) \right).$$

and the target is

$$\bar{\pi}_k(x_{0:k}, \tilde{x}_{1:k}) = \pi_k(x_k) \left(\prod_{l=0}^{k-1} L_l^{\text{mut}}(x_{l+1}, \tilde{x}_{l+1}) L_l^{\text{trans}}(\tilde{x}_{l+1}, x_l) \right).$$

Hence the incremental weight after a transport term at time k is of the form

$$\frac{\bar{\pi}_k(x_{0:k-1}, \tilde{x}_{1:k})}{\bar{\eta}_k(x_{0:k-1}, \tilde{x}_{1:k})} = \frac{\bar{\pi}_{k-1}(x_{0:k-1}, \tilde{x}_{1:k-1})}{\bar{\eta}_{k-1}(x_{0:k-1}, \tilde{x}_{1:k-1})} \underbrace{\frac{\pi_k(\tilde{x}_k) L_{k-1}^{\text{trans}}(\tilde{x}_k, x_{k-1})}{\pi_{k-1}(\tilde{x}_{k-1}) M_k^{\text{trans}}(x_{k-1}, \tilde{x}_k)}}_{\text{incremental weight} = \frac{Z_{k-1}}{Z_k} G_{k, T_k}(x_{k-1})},$$

while after the mutation step it is of the form

$$\frac{\bar{\pi}_k(x_{0:k}, \tilde{x}_{1:k})}{\bar{\eta}_k(x_{0:k}, \tilde{x}_{1:k})} = \frac{\bar{\pi}_k(x_{0:k-1}, \tilde{x}_{1:k})}{\bar{\eta}_k(x_{0:k-1}, \tilde{x}_{1:k})} \underbrace{\frac{\pi_k(x_k) L_{k-1}^{\text{mul}}(x_k, \tilde{x}_k)}{\pi_k(\tilde{x}_k) M_k^{\text{mul}}(\tilde{x}_k, x_k)}}_{\text{incremental weight}=1}.$$

B.2. Collapsed version

When no resampling is used, there is no use for the introduction of the random variables $\tilde{X}_{1:k}$ in the previous derivation and they can be integrated out. In this case, we collapse the transport step and mutation step into one single Markov kernel

$$\begin{aligned} M_k^{\text{col}}(x, x') &= \int M_k^{\text{trans}}(x, \tilde{x}) M_k^{\text{mul}}(\tilde{x}, x') d\tilde{x} \\ &= \int \delta_{T_k(x)}(\tilde{x}) K_k(\tilde{x}, x') d\tilde{x} \\ &= K_k(T_k(x), x'). \end{aligned} \tag{16}$$

Similarly we collapse the backward kernels used to defined the extended target distributions $\bar{\pi}_k$

$$\begin{aligned} L_{k-1}^{\text{col}}(x, x') &= \int \frac{\pi_k(\tilde{x}) K_k(\tilde{x}, x)}{\pi_k K_k(x)} \delta_{T_k^{-1}(\tilde{x})}(x') d\tilde{x} \\ &= \frac{\pi_k(T_k(x')) |\nabla T_k(x')| K_k(T_k(x'), x)}{\pi_k K_k(x)}. \end{aligned} \tag{17}$$

Contrary to Appendix B.1, we consider the more general scenario here where K_k might not be π_k invariant discussed in Section 5. From (16) and (17), $\bar{\pi}_k$ is thus given by (5) for

$$w_k(x_{0:k}) = \underbrace{\prod_{l=1}^k \frac{\gamma_l(x_l)}{\gamma_l K_l(x_l)}}_{r_k(x_{1:k})} \prod_{l=1}^k G_{l, T_l}(x_{l-1}),$$

where $r_k(x_{1:k}) = 1$ for π_l -invariant MCMC kernels K_l as used in Algorithm 1.

C. Proof of the asymptotic results

We consider the unnormalized empirical measure γ_k^N defined as:

$$\gamma_k^N = Z_k^N \pi_k^N.$$

We will provide the consistency and CLT results for both $\gamma_k^N[f]$ and $\pi_k^N[f]$ which imply the results on the normalizing constant Z_k^N as $Z_k^N = \gamma_k^N[1]$. We denote by \mathcal{F}_k^N the filtration generated by the particles X_k^i and the NFs T_{k+1} up to time

k and write $\mathbb{E}_k[f] = \mathbb{E}[f|\mathcal{F}_k^N]$. This accounts for possible randomness coming from the optimization of the NFs. We also consider \mathcal{C}_p the class of continuous functions f on \mathcal{X} with growth in x of at most p , for some non-negative integer p , i.e.

$$\mathcal{C}_p(C) = \{f : \mathcal{X} \rightarrow \mathbb{R} \mid f \text{ is continuous and } \forall x \in \mathcal{X} : \|f(x)\| \leq C(1 + \|x\|^p)\}.$$

In addition, we denote by \mathcal{LC}_p the class of functions in \mathcal{C}_p that are locally Lipschitz and with local Lipschitz constant satisfying a growth condition:

$$\mathcal{LC}_p(C) = \left\{f \in \mathcal{C}_p(C) \mid \forall x, x' \in \mathcal{X} : \|f(x) - f(x')\| \leq C \left(1 + \|x\|^{p+1} + \|x'\|^{p+1}\right) \|x - x'\|\right\}.$$

For ease of notation we also introduce the unnormalized transition kernel $Q_{k,T}$ which acts on functions f by:

$$Q_{k,T}[f](x) := G_{k,T}(x) \int f(y) K_k(T(x), dy).$$

Moreover, we overload the notation and write $\mathcal{L}_k(\theta) := \mathcal{L}_k(\tau_\theta)$ and $\mathcal{L}_k^N(\theta) := \mathcal{L}_k^N(\tau_\theta)$.

C.1. Assumptions

The following assumptions are needed for both Theorems 1 and 2.

- (A) The Markov kernel K_k preserves the class \mathcal{C}_p for any p , meaning that $K_k(f) \in \mathcal{C}_p$ whenever f in \mathcal{C}_p .
- (B) π_k admits 8th order moments.
- (C) The normalizing flows in \mathcal{T} are of the form $T(x) = \tau_\theta(x)$ where θ is a finite dimensional vector in a compact convex set Θ . Moreover, the maps $x \mapsto \tau_\theta(x)$ are L -Lipschitz and jointly continuous in x and θ .
- (D) The importance weights $G_{k,T}(x)$ are uniformly bounded over x and T .

In addition to the previous assumptions, we will need additional assumptions for the CLT result in Theorem 2. First, we strengthen Assumption (A)

- (E) The Markov kernel K_k preserves the class \mathcal{LC}_p for any p , i.e. $K_k(f) \in \mathcal{C}_p$ for any f in \mathcal{LC}_p .

We then make additional assumptions on the smoothness of the potentials V_k and the parameterization of the normalizing flows τ_θ :

- (F) The flow $(\theta, x) \mapsto \tau_\theta(x)$ admits derivatives $\nabla_\theta \tau_\theta(x)$, $\partial_{\theta_i} \partial_{x_l} \tau_\theta(x)$ and $\partial_{\theta_i} \partial_{\theta_j} \partial_{x_l} \tau_\theta(x)$, $H_x \tau_\theta(x)$ with at most linear growth in x uniformly in θ . Moreover, all singular values of $\nabla_x \tau_\theta(x)$ are lower-bounded by a positive constant $c > 0$ uniformly in x and θ .
- (G) The potentials are twice continuously differentiable and their gradients are L -Lipchitz, i.e. $\|\nabla V_k(x) - \nabla V_k(x')\| \leq L\|x - x'\|$.

Finally, we make two assumptions on the algorithm used to find θ_k^N . We denote by Θ_k^* the set of local minimizers of the population loss $\theta \mapsto \mathcal{L}_k(\theta)$.

- (H) The estimator θ_k^N satisfies:

$$\begin{aligned} \nabla \mathcal{L}_k^N(\theta_k^N) &= o_{\mathbb{P}}(1), \\ H \mathcal{L}_k^N(\theta_k^N) &\geq o_{\mathbb{P}}(1). \end{aligned}$$

- (I) There exists a local minimizer $\theta_k^* \in \Theta_k^*$ of the population loss $\mathcal{L}_k(\theta)$ such that

$$\mathbb{P} \left[\theta_k^* \in \arg \min_{\theta \in \Theta_k^*} \|\theta_k^N - \theta\| \right] \rightarrow 1.$$

Assumption **(H)** states that the algorithm finds an approximate local minimizer of the empirical loss $\mathcal{L}_k^N(\theta)$. This condition depends only on how well the algorithm is able to find a local optimum of the empirical loss accurately. In the ideal case where θ_k^N is an exact local minimizer of $\mathcal{L}_k^N(\theta)$, then the condition holds by definition. Assumption **(I)** states that as N increases θ_k^N remains within the basin of attraction of a single local optimum θ_k^* and does not jump between different solutions. For instance, in the case of gradient descent, this assumption can be satisfied if the algorithm starts from the same initial $(\theta_k)_0$ for all values of N and is iterated to obtain an estimator θ_k^N . Hence, as N increases the empirical loss will have the same basins of attraction as the population loss and the choice of the solution θ_k^* is determined only by the initial condition $(\theta_k)_0$.

C.2. Kernels satisfying Assumptions **(A)** and **(E)**

Here we provide examples of generic transition kernels K_k that satisfy Assumptions **(A)** and **(E)**. In Appendix C.2.1, we show that the kernel used in the Unadjusted Langevin Algorithm (ULA kernel) satisfies Assumptions **(A)** and **(E)** under mild assumptions on π_k . While this kernel is not exactly invariant w.r.t. π_k , we will use it in Appendix C.2.2 to construct a kernel invariant w.r.t. π_k and satisfying Assumptions **(A)** and **(E)**.

C.2.1. UNADJUSTED LANGEVIN KERNEL

We consider a slightly generalized version of the ULA kernel whose density $g(x, y)$ is given by:

$$g(x, y) \propto \exp\left(-\frac{1}{4\tau}\|y - x - \alpha\tau\nabla\log\pi_k(x)\|^2\right), \quad (18)$$

with $0 \leq \alpha \leq 1$ and $\tau > 0$. When $\alpha = 0$, one recovers the random walk kernel, while setting $\alpha = 1$ gives back the ULA kernel with discretization step-size τ .

Proposition 2. *Under Assumption **(G)**, the density $g(x, y)$ in (18) satisfies*

$$\|\nabla_x \log g(x, y)\| \lesssim (1 + \|y\| + \|x\|), \quad \|\nabla_y \log g(x, y)\| \lesssim (1 + \|y\| + \|x\|).$$

Moreover, the ULA kernel with density $g(x, y)$ satisfies Assumptions **(A)** and **(E)**.

Proof. The estimate in (21) is obtained by direct computation of the gradient of the logarithm of $g(x, y)$

$$\begin{aligned} \|\nabla_x g(x, y)\| &= \frac{1}{\tau} \left\| (y - x + \tau\alpha\nabla V_k(x))^\top (-I + \tau\alpha H_x V_k(x)) \right\| \\ &\lesssim (1 + \|x\| + \|y\|), \end{aligned}$$

where we used that $\nabla V_k(x)$ has at most a linear growth in x and $H_x V_k(x)$ is bounded by Assumption **(G)**. The second assertion is obtained similarly by directly computing the gradient w.r.t y .

To show that the ULA kernel satisfies Assumption **(A)**, consider a function f in \mathcal{C}_p , we can then write after a change of variables:

$$\begin{aligned} K_k(f)(x) &= \int f(y + x - \tau\alpha\nabla V_k(x)) \mathcal{N}(y, 0, 2\tau I) dy \\ &\lesssim \int \|y\|^p \mathcal{N}(y, 0, 2\tau I) dy + \|x\|^p + (\tau\alpha)^p \|\nabla V_k(x)\| \lesssim (1 + \|x\|^p), \end{aligned}$$

where we get the last inequality by Assumption **(G)**. It is easy to see that $x \mapsto K(f)(x)$ is continuous by smoothing with a Gaussian and recalling that $\nabla V_k(x)$ is continuous. Hence, we can conclude that $K(f) \in \mathcal{C}_p$. To show that Assumption **(E)** holds, we consider a function f in $\mathcal{L}\mathcal{C}_p$ and control the difference $|K_k(f)(x) - K_k(f)(x')|$. For concision, we introduce $r(x, y) = y + x - \tau\alpha\nabla V_k(x)$ and write:

$$\begin{aligned} |K_k(f)(x) - K_k(f)(x')| &\leq \int |f(r(x, y)) - f(r(x', y))| \mathcal{N}(y, 0, 2\tau I) dy \\ &\lesssim \int \|r(x, y) - r(x', y)\| \left(\|r(x, y)\|^{p+1} + \|r(x', y)\|^{p+1} \right) \mathcal{N}(y, 0, 2\tau I) dy. \end{aligned} \quad (19)$$

Using Assumption **(G)**, we clearly have:

$$\|r(x, y)\| \lesssim (1 + \|x\| + \|y\|), \quad \|r(x, y) - r(x', y)\| \lesssim \|x - x'\|.$$

We get the desired result by using the previous bounds in (19) and using the convexity of the power function. \square

C.2.2. METROPOLIS–HASTINGS KERNEL

For a target density π_k , we consider a Metropolis–Hasting kernel $K_k(x, dy)$ of the form:

$$K_k(x, dy) = g(x, y)\alpha(x, y) dy + \delta_x(dy) \int (1 - \alpha(x, u))g(x, u) du, \quad (20)$$

where $g(x, y)$ is the density of a proposal kernel and $\alpha(x, y)$ is the acceptance ratio:

$$\alpha(x, y) = \min \left(1, \frac{\pi_k(y)g(y, x)}{\pi_k(x)g(x, y)} \right).$$

We are in particular interested in proposals g that satisfy the growth condition:

$$\|\nabla_x \log g(x, y)\| \lesssim (1 + \|x\| + \|y\|). \quad (21)$$

By Proposition 2, the above condition is satisfied if g is a ULA kernel and if the potential V_k satisfies Assumption **(G)**.

In the next proposition, we show that Assumptions **(A)** and **(E)** hold under mild assumptions on π_k and when the proposal g satisfies (21).

Proposition 3. *Assume that Assumptions **(B)** and **(G)** hold for π_k and that g satisfies Assumption **(A)** then the MH kernel in (20) satisfies Assumption **(A)**.*

*If, in addition, g satisfies the growth condition in (21), then the MH kernel in (20) satisfies Assumption **(E)**.*

Proof. For the first part of the proof, we consider a function f in \mathcal{C}_p and write:

$$\begin{aligned} |K_k[f](x)| &= \left| f(x) \int (1 - \alpha(x, u))g(x, u) du + \int f(y)\alpha(x, y)g(x, y) dy \right| \\ &\leq |f(x)| + \left| \int f(y)g(x, y) dy \right|, \end{aligned}$$

where we used that $0 \leq \alpha(x, y) \leq 1$ to get the inequality. Since g satisfies Assumption **(A)** and $f \in \mathcal{C}_p$ we directly conclude that:

$$|K_k[f]| \lesssim 1 + \|x\|^p.$$

To prove the second part, consider a function f in \mathcal{LC}_p . We need to control the difference $|K_k[f](x) - K_k[f](x')|$:

$$\begin{aligned} K_k[f](x) - K_k[f](x') &= f(x) - f(x') + \int (f(y) - f(x))g(x, y)\alpha(x, y) dy - \int (f(y) - f(x'))g(x', y)\alpha(x', y) dy \\ &= \underbrace{(f(x) - f(x')) \int g(x, y)\alpha(x, y) dy}_A \\ &\quad + \underbrace{\int (f(y) - f(x'))(g(x, y) - g(x', y))\alpha(x, y) dy}_B \\ &\quad + \underbrace{\int (f(y) - f(x'))g(x', y)(\alpha(x, y) - \alpha(x', y)) dy}_C. \end{aligned}$$

We will control each term A , B and C independently. Since $0 \leq \alpha \leq 1$, and $f \in \mathcal{LC}_p$, we directly have $|A| \lesssim (1 + \|x\|^{p+1} + \|x'\|^{p+1})\|x - x'\|$. To control the second term B , we use the fundamental theorem of calculus which yields

$$|g(x, y) - g(x', y)| = \left| \int_0^1 g(x_t, y) \nabla_x \log g(x_t, y)^\top (x - x') dt \right|,$$

where $x_t := (1 - t)x + tx'$. Since g satisfies (21) by assumption, we can directly write:

$$|g(x, y) - g(x', y)| \leq \|x - x'\| \int_0^1 g(x_t, y) (1 + \|y\| + \|x_t\|) dt.$$

Plugging the above inequality in B and using that $f \in \mathcal{C}_p$ yields:

$$|B| \lesssim \|x - x'\| \int_0^1 \int (1 + \|y\|^p + \|x'\|^p) (1 + \|y\| + \|x_t\|) g(x_t, y) dt.$$

Since g satisfies Assumption (A), we can directly conclude that $|B| \lesssim (1 + \|x\|^{p+1} + \|x'\|^{p+1})\|x - x'\|$. Finally, to control C , we first define the function $b(x, y) = V_k(x) - V_k(y) + \log \frac{g(y, x)}{g(x, y)}$ so that the acceptance ratio can be written as $\alpha(x, y) = \min(1, e^{b(x, y)})$. Using Lemma 1, we directly have:

$$\begin{aligned} |\alpha(x, y) - \alpha(x', y)| &\leq |b(x, y) - b(x', y)| \\ &\leq |V_k(x) - V_k(x') + \log g(y, x) - \log g(y, x') + \log g(x', y) - \log g(x, y)| \\ &\leq \left| \int_0^1 (\nabla V_k(x_t) + \nabla_y \log g(y, x_t) - \nabla_x \log g(x_t, y))^\top (x - x') dt \right| \\ &\lesssim \|x - x'\| \int_0^1 (1 + \|x_t\| + \|y\|) dt \lesssim \|x - x'\| (1 + \|x\| + \|x'\| + \|y\|). \end{aligned}$$

We can therefore use the above inequality to upper-bound $|C|$ as follows

$$\begin{aligned} |C| &\lesssim \|x - x'\| \int (|f(x')| + |f(y)|) (1 + \|x\| + \|x'\| + \|y\|) g(x', y) dy \\ &\lesssim \|x - x'\| (1 + \|x\| + \|x'\|), \end{aligned}$$

where we used that f belongs to \mathcal{LC}_p and thus to \mathcal{C}_p . □

Lemma 1. *The following holds for any a, a' in \mathbb{R} :*

$$\left| \min(1, e^a) - \min(1, e^{a'}) \right| \leq |a - a'|$$

Proof. Fix a and a' in \mathbb{R} . We distinguish 4 cases:

Case 1: $a \leq 0$ and $a' \leq 0$.

$$\begin{aligned} \left| \min(1, e^a) - \min(1, e^{a'}) \right| &= |e^a - e^{a'}| = |a - a'| \int_0^1 e^{(1-t)a + ta'} dt \\ &\leq \max(e^a, e^{a'}) |a - a'| \leq |a - a'| \end{aligned}$$

where we used that $e^a \leq 1$ and $e^{a'} \leq 1$.

- Case 2: $a \geq 0$ and $a' \geq 0$

We directly have $\left| \min(1, e^a) - \min(1, e^{a'}) \right| = |1 - 1| = 0 \leq |a - a'|$.

Case 3: $a \leq 0$ and $a' \geq 0$.

$$\left| \min(1, e^a) - \min(1, e^{a'}) \right| = |e^a - 1| = |a| \int_0^1 e^{ta} dt$$

Recalling that $a \leq 0$ we have $e^{ta} \leq 1$ and $|a| = -a$. Moreover, since $a' \geq 0$ we can write

$$\left| \min(1, e^a) - \min(1, e^{a'}) \right| \leq -a \leq a' - a = |a' - a|$$

Case 4: $a' \leq 0$ and $a \geq 0$. This case is the same as case 3 by switching the roles of a and a' . □

C.3. Weak law of large numbers

For simplicity, we provide a proof of Appendix C.3 when resampling is performed at each step. This can easily be extended to adaptive resampling using techniques from (Douc and Moulines, 2008; Del Moral et al., 2012b). We denote by \xrightarrow{P} convergence in probability.

Proof. of Theorem 1. We proceed by induction. The result clearly holds for $k = 0$ by the regular law of large numbers. By induction, we assume \mathcal{R}_l holds for $0 \leq l \leq k - 1$ and we will prove that \mathcal{R}_k holds as well. Let f be a function in \mathcal{C}_4 . We use the decomposition $\pi_k^N[f] - \pi_k[f] = A_N + B_N$ with:

$$\begin{aligned} A_N &= \pi_k^N[f] - \mathbb{E}_{k-1}[\pi_k^N[f]], \\ B_N &= \mathbb{E}_{k-1}[\pi_k^N[f]] - \pi_k[f]. \end{aligned}$$

Propositions 4 and 5 show that both A_N and B_N converge in probability to 0 and imply that $\pi_k^N[f] - \pi_k[f] \xrightarrow{P} 0$. It remains to show that $\gamma_k^N[f] - \gamma_k[f] \xrightarrow{P} 0$. We recall that $\gamma_k^N[f] = \gamma_k^N[1]\pi_k^N[f]$ and $\gamma_k[f] = \gamma_k[1]\pi_k[f]$. Thus we only need to show that $\gamma_k^N[1] - \gamma_k[1] \xrightarrow{P} 0$. Recall that $\gamma_k^N[1] = \prod_{l=1}^k \pi_{l-1}^N[G_l, T_l]$ and by Proposition 4 we know that $\pi_{l-1}^N[G_l, T_l] \xrightarrow{P} \frac{Z_l}{Z_{l-1}}$, thus we directly have $\gamma_k^N[1] \xrightarrow{P} \prod_{l=1}^k \frac{Z_l}{Z_{l-1}} = \frac{Z_k}{Z_0}$. This directly implies $\gamma_k^N[1] \xrightarrow{P} \frac{Z_k}{Z_0} = \gamma_k[1]$ since $Z_0 = 1$ by construction. Finally, we conclude that $\gamma_k^N[f] - \gamma_k[f] \xrightarrow{P} 0$ using Slutsky's lemma. □

Theorem 4 (Weak law of large numbers for Algorithm 2). *Let f be a function s.t. $|f(x)| \leq C(1 + \|x\|^4)$ for all $x \in \mathcal{X}$ and for some $C > 0$. Under Assumptions (A) to (D) and for any $k \in 0, \dots, K$:*

$$(\mathcal{R}_k) : \quad \pi_k^{N_{\text{test}}}[f] \xrightarrow{P} \pi_k[f], \quad Z_k^{N_{\text{test}}, \text{test}} \xrightarrow{P} Z_k,$$

where $\pi_k^{N_{\text{test}}}$ and $Z_k^{N_{\text{test}}, \text{test}}$ are given by Algorithm 2.

Proof. The proof is a direct consequence of consistency of SMC samplers (Chopin, 2004; Del Moral, 2004). Indeed, the test particles $\left\{ X_k^{i, \text{test}} \right\}_{1:N}$ are independent from the train and validation particles $\left\{ X_k^{i, \text{train}} \right\}_{1:N}$ and $\left\{ X_k^{i, \text{val}} \right\}_{1:N}$ used to learn the flows T_k . Moreover, by Proposition 6, the importance weights $w_k^{i, \text{test}}$ correct exactly for the discrepancy between π_k and $(T_k)_{\#} \pi_{k-1}$. Hence, knowing the train and validation particles, Algorithm 2 is a standard SMC sampler with Markov transition kernel given by $M_k(x, \cdot) = (T_k)_{\#} K_k(x, \cdot)$. Therefore consistency holds and $Z_k^{N_{\text{test}}, \text{test}}$ is an unbiased estimator of Z_k . □

Proposition 4. *Under Assumptions (A) to (D) and whenever the recursion assumption \mathcal{R}_{k-1} holds for a given $k > 0$, it also holds that*

$$B_N := \mathbb{E}_{k-1}[\pi_k^N[f]] - \pi_k[f] \xrightarrow{P} 0, \quad \pi_{k-1}^N[G_{k, T_k}] \xrightarrow{P} \frac{Z_k}{Z_{k-1}},$$

for all functions f in \mathcal{C}_4 .

Proof. We use the following decomposition for $B_N = B_N^{(1)} + B_N^{(2)} + B_N^{(3)} + B_N^{(4)}$:

$$\begin{aligned} B_N^{(1)} &:= \frac{1}{\pi_{k-1}^N[G_{k,T_k}]} (\pi_{k-1}[Q_{k,T_k}[f]] - \pi_{k-1}[Q_{k,T_k^*}[f]]), \\ B_N^{(2)} &:= \frac{\pi_{k-1}[Q_{k,T_k^*}[f]]}{\pi_{k-1}^N[G_{k,T_k}]\pi_{k-1}[G_{k,T_k^*}]} (\pi_{k-1}[G_{k,T_k^*}] - \pi_{k-1}[G_{k,T_k}]), \\ B_N^{(3)} &:= \frac{1}{\pi_{k-1}^N[G_{k,T_k}]} (\pi_{k-1}^N[Q_{k,T_k}[f]] - \pi_{k-1}[Q_{k,T_k}[f]]), \\ B_N^{(4)} &:= \frac{\pi_{k-1}[Q_{k,T_k^*}[f]]}{\pi_{k-1}^N[G_{k,T_k}]\pi_{k-1}[G_{k,T_k^*}]} (\pi_{k-1}[G_{k,T_k}] - \pi_{k-1}^N[G_{k,T_k}]). \end{aligned}$$

Proposition 6 states that $\pi_{k-1}[Q_{k,T}[f]]$ is independent of the choice of T . Thus, the first two terms $B_N^{(1)}$ and $B_N^{(2)}$ are exactly 0.

We know by Proposition 7 that $Q_{k,T}[f](x)$ belongs to \mathcal{C}_4 , uniformly over \mathcal{T} . Moreover, the family $\mathcal{F} = (Q_{k,\tau_\theta}[f](x))_{\theta \in \Theta}$ is continuously indexed by the compact set Θ . We can therefore apply Proposition 9 under the recursion assumption \mathcal{R}_{k-1} to the family \mathcal{F} . This ensures that

$$\sup_{T \in \mathcal{T}} |\pi_{k-1}^N[Q_{k,T}[f]] - \pi_{k-1}[Q_{k,T}[f]]| \xrightarrow{P} 0.$$

In particular, we have

$$\begin{aligned} \pi_{k-1}^N[Q_{k,T_k}[f]] - \pi_{k-1}[Q_{k,T_k}[f]] &\xrightarrow{P} 0, \\ \pi_{k-1}^N[G_{k,T_k}] - \pi_{k-1}[G_{k,T_k}] &\xrightarrow{P} 0, \end{aligned}$$

where the last equation is obtained simply by choosing $f = 1$. This directly implies that $B_N^{(3)}$ and $B_N^{(4)}$ converge to 0 in probability using Slutsky's lemma. \square

Proposition 5. *Under Assumptions (A) to (D) and whenever the recursion assumption \mathcal{R}_{k-1} holds for a given $k > 0$, it holds that:*

$$A_N := \pi_k^N[f] - \mathbb{E}_{k-1}[\pi_k^N[f]] \xrightarrow{P} 0,$$

for all measurable functions f in \mathcal{C}_p .

Proof. We will show that the characteristic function of A_N denoted $\phi_{A_N}(t)$ converges towards 1 for all $t \in \mathbb{R}$. It is easy to see that A_N can be expressed as:

$$A_N = \frac{1}{N} \left(\sum_{i=1}^N U_{N,i} \right), \quad U_{N,i} = f(X_k^i) - \mathbb{E}_{k-1}[f(X_k^i)],$$

where, conditionally on \mathcal{F}_{k-1}^N , the variables X_k^i are independent and identically distributed according to:

$$\mathbb{P}(X_k \in dx | \mathcal{F}_{k-1}^N) = \sum_{i=1}^N \frac{G_{k,T_k}(X_{k-1}^i)}{\sum_{j=1}^N G_{k,T_k}(X_{k-1}^j)} K_k(T_k(X_{k-1}^i), dx).$$

Let us introduce the conditional characteristic function $\hat{\phi}_{U_{N,1}}(t)$ knowing \mathcal{F}_{k-1}^N defined by:

$$\hat{\phi}_{U_{N,1}}(t) = \mathbb{E}_{k-1}[e^{itU_{N,1}}].$$

This allows to express $\phi_{A_N}(t)$ in terms of $\hat{\phi}_{U_{N,1}}(t)$ as $\phi_{A_N}(t) = \mathbb{E}[\hat{\phi}_{U_{N,1}}(\frac{t}{N})^N]$. Thus, we only need to prove that $\hat{\phi}_{U_{N,1}}(\frac{t}{N})^N \xrightarrow{P} 1$.

We will rely on the following expression $\hat{\phi}_{U_{N,1}}(\frac{t}{N}) = 1 + \frac{t}{N} E_N(\frac{t}{N})$ with $E_N(u) := \int_0^1 \hat{\phi}'_{U_{N,1}}(su) ds$. We only need to prove that $E_N(\frac{t}{N}) \xrightarrow{P} 0$ as $N \rightarrow \infty$. Let us introduce the families of functions indexed by $u \in [-|t|, |t|]$:

$$g_u(x) = \int_0^1 e^{isu f(x)} ds, \quad \hat{g}_u = \int_0^1 e^{-isu \mathbb{E}_{k-1}[f(X_k)]} ds, \quad f_u(x) = f(x) g_u(x).$$

Using the dominated convergence theorem, we can further express $E_N(u)$ as follows:

$$\begin{aligned} E_N(u) &= i \mathbb{E}_{k-1} \left[U_{N,1} \int_0^1 e^{isu U_{N,1}} ds \right] \\ &= i \mathbb{E}_{k-1} [f_u(X_k)] \hat{g}_u - i \mathbb{E}_{k-1} [f(X_k)] \mathbb{E}_{k-1} [g_u(X_k)] \hat{g}_u. \end{aligned}$$

Each expectation is of the form $\mathbb{E}_{k-1}[h(X_k)] = \frac{\pi_{k-1}^{N-1}[Q_{k,T}[h]]}{\pi_{k-1}^{N-1}[G_{k,T}]}$. Using the induction hypothesis \mathcal{R}_{k-1} , and recalling that each function f_u , f and g_u belongs to \mathcal{C}_4 , it follows that each conditional expectation converges in probability towards $\pi_k[f_u]$, $\pi_k[f]$ and $\pi_k[g_u]$ while $\hat{g}_u \xrightarrow{P} \int_0^1 e^{-isu \pi_k[f]} ds$. Moreover, using the fact that the functions g_u and f_u are continuously indexed by u over the compact set $[-|t|, |t|]$, we can apply Proposition 9 to ensure that convergence is uniform over this set. This allows us to prove in particular that $E_N(\frac{t}{N}) \xrightarrow{P} 0$. We have shown so far that $\hat{\phi}_{U_{N,1}}(\frac{t}{N}) = 1 + o_{\mathbb{P}}(\frac{t}{N})$ which allows to conclude that $\hat{\phi}_{U_{N,1}}(\frac{t}{N})^N \xrightarrow{P} 1$ as $N \rightarrow \infty$ and thus $A_N \xrightarrow{P} 0$. \square

Proposition 6. *The following holds for any admissible T in \mathcal{T} and function f such that $K_k[f](x) < \infty$ and $\pi_k[f] < \infty$:*

$$\pi_{k-1}[Q_{k,T}[f]] = \frac{Z_k}{Z_{k-1}} \pi_k[f].$$

In particular, we have $\pi_{k-1}[G_{k,T}] = \frac{Z_k}{Z_{k-1}}$.

Proof. For any admissible map T we have that:

$$\begin{aligned} \pi_{k-1}[Q_{k-1,T}[f]] &= \frac{Z_k}{Z_{k-1}} \int \frac{\pi_k(T(x))}{\pi_{k-1}(x)} |\nabla T(x)| \left(\int f(y) K_k(T(x), dy) \right) \pi_{k-1}(x) dx \\ &= \frac{Z_k}{Z_{k-1}} \int \int \pi_k(z) f(y) K_k(z, dy) dz \\ &= \frac{Z_k}{Z_{k-1}} \pi_k[f]. \end{aligned}$$

The second line is obtained by a change of variables $z = T(x)$ and using that K_k is invariant w.r.t π_k . The last inequality is obtained by choosing $f = 1$. \square

Proposition 7. *Let f be a measurable function in \mathcal{C}_p for $0 \leq p \leq 4$. Then, under Assumptions (A), (C) and (D), the function $x \mapsto Q_{k,T}[f](x)$ belongs to \mathcal{C}_p uniformly over \mathcal{T} . In other words, there exists a positive constant $C > 0$ such that:*

$$|Q_{k,T}[f](x)| \leq C(1 + \|x\|^p), \quad \forall x \in \mathcal{X}, \forall T \in \mathcal{T}.$$

Proof. By Assumption (A), we have that $|K_k[f](x)| \leq C(1 + \|x\|^p)$. Moreover, using Assumption (C) we know that $T(x)$ has a linear growth in x , $\|T(x)\| \leq C'(1 + \|x\|)$ with the same constant C' for all $T \in \mathcal{T}$. Therefore, there exists a positive constant $M > 0$, such that $|K_k[f](T(x))| \leq M(1 + \|x\|^p)$ for any $T \in \mathcal{T}$ and $x \in \mathcal{X}$. Finally, we know by Assumption (D) that $G_{k,T}(x)$ is bounded uniformly over x and T . This allows us to conclude that $Q_{k,T}[f]$ has the desired growth in x which is uniform over $T \in \mathcal{T}$. \square

Proposition 8. *Let $1 \leq p \leq 4$, $C > 0$ and \mathcal{F} be a class of measurable functions in $\mathcal{C}_p(C)$ such that the bracketing number $\mathcal{N}_{[]}(\epsilon, \mathcal{F}, L_2(\pi_k))$ is finite for any $\epsilon > 0$. Then under Assumption (B) and the recursion assumption \mathcal{R}_k the following uniform convergence holds in probability*

$$\sup_{f \in \mathcal{F}} |\pi_k^N[f] - \pi_k[f]| \xrightarrow{P} 0.$$

Proof. First consider the envelope $F(x) := \sup_{s \in S} |f(x)|$ which has a growth of at most p in x by assumption on \mathcal{F} . Moreover, F is π_k -integrable by Assumption **(B)**. Fix $\epsilon > 0$. Since the bracketing number $N_{[]}(\epsilon, \mathcal{F}, L_1(\pi_k))$ is finite, there exists finitely many ϵ -brackets $([l_i, u_i])_{1 \leq i \leq I}$ whose union contains \mathcal{F} and such that $\pi_k(u_i - l_i) < \epsilon$ for every $i \in \{1, I\}$. Moreover, the functions l_i and u_i can be chosen to have a growth of at most p in x , since \mathcal{F} belongs to $\mathcal{C}_p(C)$. Hence, for every $f \in \mathcal{F}$, there is a bracket $[l_i, u_i]$ such that:

$$(\pi_k^N - \pi_k)[f] \leq (\pi_k^N - \pi_k)[u_i] + \pi_k[u_i - f] \leq (\pi_k^N - \pi_k)[u_i] + \epsilon.$$

Hence, we have:

$$\sup_{f \in \mathcal{F}} (\pi_k^N - \pi_k)[f] \leq \max_i (\pi_k^N - \pi_k)[u_i] + \epsilon.$$

Since \mathcal{R}_k holds and $u_i \in \mathcal{C}_p(C)$, the right hand side converges in probability towards ϵ . Similarly, it is possible to show that:

$$\inf_{f \in \mathcal{F}} (\pi_k^N - \pi_k)[f] \geq \min_i (\pi_k^N - \pi_k)[l_i] - \epsilon.$$

with r.h.s. converging towards $-\epsilon$ in probability. This allows us to conclude. \square

Proposition 9. *Let \mathcal{F} be a class of measurable functions in $\mathcal{C}_4(C)$ for some $C > 0$:*

$$|f(x)| \leq C(1 + \|x\|^4), \quad \forall x \in \mathcal{X}, \forall f \in \mathcal{F}.$$

*Assume that \mathcal{F} is continuously indexed by a compact set S , i.e. $s \mapsto f_s(x)$ is continuous for any $x \in \mathcal{X}$, where f_s is an element in \mathcal{F} indexed by $s \in S$. Then under Assumption **(B)** and the recursion assumption \mathcal{R}_k the following uniform convergence holds in probability*

$$\sup_{f \in \mathcal{F}} |\pi_k^N[f] - \pi_k[f]| \xrightarrow{P} 0.$$

Proof. First consider the envelope $F(x) := \sup_{s \in S} |f_s(x)|$ which has a growth of at most 4 in x by assumption over the class \mathcal{F} . Moreover, F is π_k -integrable by Assumption **(B)**. Since \mathcal{F} is continuously indexed by a compact set and has an integrable envelope F w.r.t. π_k , this implies that its bracketing number $N_{[]}(\epsilon, \mathcal{F}, L_1(\pi_k))$ is finite for every $\epsilon > 0$ (Van der Vaart, 2000, Example 19.8). We can directly apply Proposition 8 to get the desired result. \square

C.4. Proof of the central limit theorem

As shown in (Douc and Moulines, 2008, Theorem 10) and (Del Moral et al., 2012b, Section 6), the fluctuations of the SMC sampler with adaptive resampling admit the same asymptotic variance as the ideal SMC sampler with resampling at the optimal times in \mathcal{K}_{opt} . Therefore, it is enough to prove this result for the case when resampling is triggered exactly at times k in \mathcal{K} .

Proof. of Theorem 2. We will proceed by induction. For $k = 0$, the samples $(X_0^i)_{1 \leq i \leq N}$ are i.i.d. thus one can directly apply the standard central limit theorem to show that the result holds at $k = 0$, i.e. we write CLT_0 holds. By induction, let us assume that CLT_{k-1} holds for some $k > 0$, we will then show that CLT_k holds as well. Let f be a measurable real-valued function over \mathcal{X} with at most quadratic growth in x . For conciseness, we first define E_N and F_N

$$E_N = \sqrt{N}(\gamma_k^N[f] - \gamma_k[f]), \quad F_N = \sqrt{N}(\pi_k^N[f] - \pi_k[f]).$$

We need to show that E_N and F_N converge to centered Gaussians with variances $\mathbb{V}_k^\gamma[f]$ and $\mathbb{V}_k^\pi[f]$. Starting with E_N , we use the decomposition $E_N = R_N + P_N$ with

$$P_N := \mathbb{E}_{k-1}[E_N], \quad R_N := E_N - \mathbb{E}_{k-1}[E_N]. \quad (22)$$

Using the recursion assumption and Proposition 10, we can show that P_N converges in distribution towards a centered Gaussian with variance $\mathbb{V}_{k-1}^\gamma[Q_{k,T^*}[f]]$. Moreover, by Propositions 12 and 13, we know that $\mathbb{E}_{k-1}[e^{itR_N}] \xrightarrow{P}$

$\exp(-\frac{t^2}{2}\mathbb{V}_k^{inc}[f])$. This allows us to conclude that the characteristic function of E_N converges pointwise to the characteristic function of a centered Gaussian distribution with variance $\mathbb{V}_k^\gamma[f]$:

$$\phi_{E_N}(t) = \mathbb{E}[e^{itP_N}\mathbb{E}_{k-1}[e^{itR_N}]] \rightarrow \exp\left(-\frac{t^2}{2}\mathbb{V}_k^\gamma[f]\right),$$

where by definition $\mathbb{V}_k^\gamma[f] = \mathbb{V}_k^{inc}[f] + \mathbb{V}_{k-1}^\gamma[Q_{k,T^*}[f]]$. We can then conclude using Lévy continuity theorem that E_N converges in distribution towards a centered Gaussian with variance $\mathbb{V}_k^\gamma[f]$ for any f with at most quadratic growth. For F_N , we can use the following identity:

$$F_N = \frac{\sqrt{N}}{\gamma_k^N[1]}(\gamma_k^N - \gamma_k)[f - \pi_k[f]].$$

Recalling that $\gamma_k^N[1] \xrightarrow{P} \gamma_k[1]$ by Theorem 1, we can directly conclude using Slutsky's lemma that

$$F_N \xrightarrow{\mathcal{D}} \mathcal{N}(0, \mathbb{V}_k^\pi[f])$$

where convergence is in distribution and where, by definition, $\mathbb{V}_k^\pi[f] = \frac{\mathbb{V}_k^\gamma[f - \pi_k[f]]}{\gamma_k[1]^2}$. This concludes the proof. \square

Theorem 5 (Central limit theorem for Algorithm 2). *Let f be a real valued function s.t. $f(x) \leq C(1 + \|x\|^2)$ for some $C > 0$. Then, under Assumptions (A) to (G) and for any $k \in 0, \dots, K$ the same CLT result as in Theorem 2 holds when using the particles produced by Algorithm 2 instead of Algorithm 1.*

Proof. The proof proceeds by recursion exactly as in Theorem 2. The only difference is that the flow is estimated using the training and validation particles instead of the test ones. This does not affect the proof by recursion since we condition w.r.t. the sigma algebra \mathcal{F}_{k-1} generated by both test particles and by the flow at time k . We only need to make sure that θ_k^N produced by Algorithm 3 satisfies Assumptions (H) and (I).

First, the validation criterion $\mathcal{L}_k^{N_{val}}(\theta)$ converges uniformly in θ in probability towards \mathcal{L}_k and so does the training criterion $\mathcal{L}_k^{N_{train}}(\theta)$ by Proposition 9. Hence, returning the flow with smallest validation error is asymptotically equivalent to returning the flow with smallest training error as N_{val} and N_{train} increase. Moreover, since θ_k^N is obtained by performing gradient descent over $\mathcal{L}_k^{N_{train}}$, the final iterate will be the one that minimizes $\mathcal{L}_k^{N_{train}}$. Recalling now that gradient descent converges to a local minimizer, it follows that Assumption (H) holds provided that the number of iterations J is large enough as N_{train} and N_{val} increase.

Second, since the flows are all initialized to the identity for any number of particles and since both the training loss and its gradient are uniformly converging in probability towards $\mathcal{L}_k(\theta)$ and $\nabla_\theta \mathcal{L}_k(\theta)$ then the optimization trajectories obtained using $\nabla_\theta \mathcal{L}_k^N(\theta)$ also converge uniformly to the one obtained using $\nabla_\theta \mathcal{L}_k(\theta)$. Hence, for N large enough, θ^N is approaching a single local minimizer θ^* . Therefore Assumption (I) also holds. \square

Proposition 10. *Let f be a function in \mathcal{C}_2 . Under the induction assumption CLT_{k-1} , we have that:*

$$P_N \xrightarrow{\mathcal{D}} \mathcal{N}(0, \mathbb{V}_{k-1}^\gamma[Q_{k,T^*}[f]]),$$

where P_N is defined in (22).

We defer the proof of Proposition 10 to Appendix D.3 as it relies on asymptotic stochastic equicontinuity of a suitable process which will be proven later in Appendix D.

Proposition 11. *Assume that resampling is only performed at the ideal resampling times in \mathcal{K} . For $k \notin \mathcal{K}$, denote by k_p the largest integer in \mathcal{K} such that $k_p < k$. Then, the importance weights W_k^i are given by:*

$$W_k^i = \frac{1}{N} \frac{Z_{k_p}^N}{Z_k^N} w_k^i,$$

with w_k^i given by:

$$w_k^i = \prod_{s=k_p+1}^k G_{s,T_s}(X_{s-1}^i).$$

Proof. This is a simple consequence of the recursion expression of the IS weights and normalizing constants for $k > k_p$:

$$\begin{aligned} W_k^i &= \frac{W_{k-1}^i}{\pi_{k-1}^N[G_{k,T_k}]} G_{k,T_k}(X_{k-1}^i), \\ Z_k^N &= Z_{k-1}^N \pi_{k-1}^N[G_{k,T_k}], \end{aligned}$$

with $W_{k_p}^i = \frac{1}{N}$. □

Proposition 12. *Let f be a function in \mathcal{C}_2 and R_N is defined in (22) and consider the conditional characteristic function*

$$\hat{\phi}_{R_N}(t) = \mathbb{E}_{k-1}[e^{itR_N}].$$

Assume that resampling is not performed at iteration k and let k_p be the largest integer in \mathcal{K} such that resampling is performed at time k_p and $k_p < k$. Recall the expression of the asymptotic incremental variance $\mathbb{V}_k^{\text{inc}}$ when $k \notin \mathcal{K}$:

$$\mathbb{V}_k^{\text{inc}}[f] = Z_{k_p}^2 \pi_{k_p}[\mathbb{E}[G_k[f]|X_{k_p}]],$$

with $G_k[f] := K_k[f^2](T_k^(X_{k-1})) - K_k[f]^2(T_k^*(X_{k-1}))$. Then, under the recursion assumption CLT_{k-1} , we have:*

$$\hat{\phi}_{R_N}(t) \xrightarrow{P} \exp\left(-\frac{t^2}{2} \mathbb{V}_k^{\text{inc}}[f]\right).$$

Proof. First, define $U_{N,i} := \sqrt{N} Z_k^N W_k^i f(X_k^i)$ and note that R_N is expressed in term of $U_{N,i}$ as:

$$R_N = \sum_{i=1}^N U_{N,i} - \mathbb{E}_{k-1}[U_{N,i}].$$

We will use the same approach as in the proof of (Douc and Moulines, 2008, Theorem 2). For that purpose, we will show the following equations hold

$$\sum_{i=1}^N \mathbb{E}_{k-1}[U_{N,i}^2] - (\mathbb{E}_{k-1}[U_{N,i}])^2 \xrightarrow{P} \mathbb{V}_k^{\text{inc}}[f], \quad (23)$$

$$\sum_{i=1}^N \mathbb{E}_{k-1}[U_{N,i}^2 \mathbf{1}_{\{|U_{N,i}| \geq \epsilon\}}] \xrightarrow{P} 0, \quad \text{for any } \epsilon > 0. \quad (24)$$

The result will follow directly by application of (Douc and Moulines, 2008, Theorem A.3). Using the expression of W_k^i given by Proposition 11, we have that:

$$\begin{aligned} \sum_{i=1}^N \mathbb{E}_{k-1}[U_{N,i}^2] &= \left(Z_{k_p}^N\right)^2 \frac{1}{N} \sum_{i=1}^N (w_k^i)^2 K_k[f^2](T_k(X_{k-1}^i)) \xrightarrow{P} Z_{k_p}^2 \pi_{k_p}[\mathbb{E}[(w_k^*)^2 K_k[f^2] \circ T_k^*]] \\ \sum_{i=1}^N \mathbb{E}_{k-1}[U_{N,i}]^2 &= \left(Z_{k_p}^N\right)^2 \frac{1}{N} \sum_{i=1}^N (w_k^i)^2 (K_k[f](T_k(X_{k-1}^i)))^2 \xrightarrow{P} Z_{k_p}^2 \pi_{k_p}[\mathbb{E}[(w_k^*)^2 (K_k[f] \circ T_k^*)^2]]. \end{aligned}$$

The above expressions are a result of the consistency of the particles trajectories $X_{k_p:k-1}^i$ by the recursion assumption. This shows (23). The proof of (24) is the same as in (Douc and Moulines, 2008, Theorem 2). □

Proposition 13. *Let f be a function in \mathcal{C}_2 and R_N is defined in (22) and consider the conditional characteristic function*

$$\hat{\phi}_{R_N}(t) = \mathbb{E}_{k-1}[e^{itR_N}].$$

Under the recursion assumption CLT_{k-1} , and if resampling is performed at iteration k we have:

$$\hat{\phi}_{R_N}(t) \xrightarrow{P} \exp\left(-\frac{t^2}{2} \gamma_k[1]^2 \text{Var}_{\pi_k}[f]\right).$$

Proof. First note that R_N is expressed as a sum of the form:

$$R_N = \frac{1}{\sqrt{N}} \sum_{i=1}^N B_{N,i}, \quad B_{N,i} = \gamma_k^N [1](U_{N,i} - \mathbb{E}_{k-1}[U_{N,i}]),$$

where, conditionally on \mathcal{F}_{k-1}^N , the variables $(U_{N,i})_{1 \leq i \leq N}$ are independent and identically distributed as $f(X_k)$ with:

$$P(X_k \in dx | \mathcal{F}_{k-1}^N) = \sum_{i=1}^N \frac{G_{k,T_k}(X_{k-1}^i)}{\sum_{j=1}^N G_{k,T_k}(X_{k-1}^j)} K_k(T_k(X_{k-1}^i), dx).$$

Hence, we have that $\hat{\phi}_{R_N}(t) = (\hat{\phi}_{B_{N,1}}(\frac{t}{\sqrt{N}}))^N$. We will start by proving the following asymptotic decomposition for $\hat{\phi}_{B_{N,1}}$:

$$\hat{\phi}_{B_{N,1}}\left(\frac{t}{\sqrt{N}}\right) = 1 - \frac{t^2}{2N} \gamma_k^2 [1] \text{Var}_{\pi_k}[f] + \frac{t^2}{N} o_{\mathbb{P}}(1).$$

Since $B_{N,1}$ is centered conditionally on \mathcal{F}_{k-1}^N and admits a finite second-order moment, the function $s \mapsto \hat{\phi}_{B_{N,1}}(s)$ is twice differentiable and satisfies $\hat{\phi}_{B_{N,1}}(0) = 1$ and $\hat{\phi}'_{B_{N,1}}(0) = 0$. Moreover, by application of the dominated convergence theorem, we have that $\left(\hat{\phi}_{B_{N,1}}\right)''(s) = -\mathbb{E}_{k-1}[B_{N,1}^2 e^{isB_{N,1}}]$. We can therefore apply Lemma 2 to $s \mapsto \hat{\phi}_{B_{N,1}}(s)$ which yields the identity:

$$\begin{aligned} \hat{\phi}_{B_{N,1}}(s) &= 1 + s^2 \int_0^1 \int_0^1 u \left(\hat{\phi}_{B_{N,1}}\right)''(suv) \, du \, dv \\ &= 1 - s^2 \int_0^1 \int_0^1 u \mathbb{E}_{k-1}[B_{N,1}^2 e^{isuvB_{N,1}}] \, du \, dv. \end{aligned}$$

Choosing $s = \frac{t}{\sqrt{N}}$ for a fixed $t \in \mathbb{R}$, we have:

$$\hat{\phi}_{B_{N,1}}(s) = 1 - \frac{t^2}{N} \int_0^1 \int_0^1 u \mathbb{E}_{k-1}[B_{N,1}^2 e^{i\frac{t}{\sqrt{N}}uvB_{N,1}}] \, du \, dv.$$

By Lemma 3 and for any $s \in \mathbb{R}$, we know that $\mathbb{E}_{k-1}[B_{N,1}^2 e^{isB_{N,1}}]$ converges in probability towards $A_{\infty}(s)$ defined as:

$$A_{\infty}(s) := \gamma_k^2 [1] \pi_k \left[(f - \pi_k[f])^2 e^{is\gamma_k[1](f - \pi_k[f])} \right].$$

Moreover, Lemma 3 also ensures this convergence to be uniform in s over the interval $[-|t|, |t|]$. Hence, we can write:

$$\hat{\phi}_{B_{N,1}}(s) = 1 - \frac{t^2}{N} \int_0^1 \int_0^1 u A_{\infty}\left(\frac{t}{\sqrt{N}}uv\right) \, du \, dv + \frac{t^2}{N} o_{\mathbb{P}}(1). \quad (25)$$

By the dominated convergence theorem we know that $A_{\infty}\left(\frac{t}{\sqrt{N}}uv\right) \xrightarrow{N} A_{\infty}(0)$ for any fixed t, u and v . Moreover, since A_{∞} is bounded, we can apply the dominated convergence theorem a second time to conclude that:

$$\int_0^1 \int_0^1 u A_{\infty}\left(\frac{t}{\sqrt{N}}uv\right) \, du \, dv \xrightarrow{N} \int_0^1 \int_0^1 u A_{\infty}(0) \, du \, dv = \frac{1}{2} \gamma_k^2 [1] \pi_k \left[(f^2 - \pi_k[f]^2) \right]. \quad (26)$$

Using (26) in (25), we have shown so far that:

$$\hat{\phi}_{B_{N,1}}(s) = 1 - \frac{t^2}{2N} \gamma_k^2 [1] \text{Var}_{\pi_k}[f] + \frac{t^2}{N} o_{\mathbb{P}}(1).$$

Recalling that $(1 + \frac{x}{N} + \frac{x}{N}o(1))^N \rightarrow e^x$ as $N \rightarrow \infty$, we can therefore conclude that:

$$\hat{\phi}_{R_N}(t) = \left(\hat{\phi}_{B_{N,1}}\left(\frac{t}{\sqrt{N}}\right)\right)^N \xrightarrow{P} \exp\left(-\frac{t^2}{2} \gamma_k [1]^2 \text{Var}_{\pi_k}[f]\right).$$

which is the desired result. \square

Lemma 2. Let $f : \mathbb{R} \rightarrow \mathbb{C}$ be a function that is twice differentiable and that $f(0) = 0$ and $f'(0) = 0$. Then the following identity holds:

$$f(s) = 1 + s^2 \int_0^1 \int_0^1 u f''(suv) \, du \, dv.$$

Proof. The identity follows by direct integration. □

Lemma 3. Let f be a function in \mathcal{C}_2 and define:

$$B_{N,i} = \gamma_k^N [1] (U_{N,i} - \mathbb{E}_{k-1}[U_{N,i}]),$$

where, conditionally on \mathcal{F}_{k-1}^N , the variables $(U_{N,i})_{1 \leq i \leq N}$ are independent and identically distributed as $f(X_k)$ with:

$$P(X_k \in dx | \mathcal{F}_{k-1}^N) = \sum_{i=1}^N \frac{G_{k,T_k}(X_{k-1}^i)}{\sum_{j=1}^N G_{k,T_k}(X_{k-1}^j)} K_k(T_k(X_{k-1}^i), dx).$$

Define the limiting function

$$A_\infty(s) := \gamma_k^2 [1] \pi_k \left[(f - \pi_k[f])^2 e^{is\gamma_k[1](f - \pi_k[f])} \right]$$

then

$$\mathbb{E}_{k-1} [B_{N,1}^2 e^{isB_{N,1}}] \xrightarrow{P} A_\infty(s),$$

where convergence is in probability and is uniform in s over any compact interval.

Proof. By definition of $U_{N,1}$, we have

$$\mathbb{E}_{k-1}[U_{N,1}] = \frac{\pi_{k-1}^N [Q_{k,T_k}[f]]}{\pi_{k-1}^N [G_{k,T_k}]} := \bar{f}^N,$$

where we introduced \bar{f}^N as a shorthand notation. Since f belongs to \mathcal{C}_2 we can apply the weak law of large numbers in Theorem 1 which implies that

$$\bar{f}^N \xrightarrow{P} \frac{\pi_{k-1} [Q_{k,T_k}[f]]}{\pi_{k-1} [G_{k,T_k}]} = \pi_k[f],$$

where we used Proposition 6 to get the last equality. We also have $\gamma_k^N [1] \xrightarrow{P} \gamma_k [1]$ by Theorem 1. Furthermore, by definition of $B_{N,1}$ and $U_{N,1}$, we can write

$$\mathbb{E}_{k-1} [B_{N,1}^2 e^{isB_{N,1}}] = (\gamma_k^N [1])^2 \frac{\pi_{k-1}^N \left[Q_{k,T_k} \left[(f - \bar{f}^N)^2 e^{is\gamma_k^N [1](f - \bar{f}^N)} \right] \right]}{\pi_{k-1}^N [G_{k,T_k}]}.$$

Recalling that f^2 belongs to \mathcal{C}_4 , we can again apply the weak law of large numbers in Theorem 1 along with the continuous mapping theorem to conclude that

$$\begin{aligned} \mathbb{E}_{k-1} [B_{N,1}^2 e^{isB_{N,1}}] &\xrightarrow{P} \gamma_k [1]^2 \frac{\pi_{k-1} \left[Q_{k,T_k} \left[(f - \pi_k[f])^2 e^{is\gamma_k [1](f - \pi_k[f])} \right] \right]}{\pi_{k-1} [G_{k,T_k}]} \\ &= \gamma_k [1]^2 \pi_k \left[(f - \pi_k[f])^2 e^{is\gamma_k [1](f - \pi_k[f])} \right] = A_\infty(s), \end{aligned}$$

where the second line is obtained by application of Proposition 6. Moreover, using Proposition 9 (as in the proof of Proposition 5), we can conclude that convergence is uniform over $s \in [-|t|, |t|]$. □

C.5. Convergence of the flow transport

Proposition 14. *Under Assumptions (A) to (D) and (F) to (I), it holds that:*

$$\theta_k^N \xrightarrow{P} \theta_k^*$$

Proof. To simplify notations, we write $d(\theta, \Theta^*) = \inf_{\theta' \in \Theta^*} \|\theta - \theta'\|$. We will first show that θ_k^N approaches the set of critical points Θ^* , i.e. $d(\theta_k^N, \Theta^*) \xrightarrow{P} 0$. We assume by contradiction that $d(\theta_k^N, \Theta^*)$ does not converge to 0 in probability. Hence, there exist $\epsilon > 0$ and $\eta > 0$ as well as a subsequence of $\theta_k^{\phi_N}$ with $\phi_N \rightarrow +\infty$ such that:

$$\mathbb{P}\left[d(\theta_k^{\phi_N}, \Theta^*) \geq \epsilon\right] \geq \eta, \forall N \in \mathbb{N} \quad (27)$$

However, we also know that the sequence θ^{ϕ_N} is tight as it is supported on Θ which is compact by Assumption (C). Hence, it admits a subsequence that converges in distribution towards a r.v. θ_k^∞ . Without loss of generality, we assume θ^{ϕ_N} to be such convergent subsequence. Since $\theta \mapsto d(\theta, \Theta^*)$ is continuous we have by the continuous mapping theorem that

$$d(\theta_k^{\phi_N}, \Theta^*) \xrightarrow{d} d(\theta_k^\infty, \Theta^*) \quad (28)$$

We will now show that θ_k^∞ must be supported on Θ^* , the set of local minima of $\theta \mapsto \mathcal{L}_\theta$. This would imply $d(\theta_k^\infty, \Theta^*) = 0$.

We know by Assumptions (F) and (G) that $\theta \mapsto \nabla \mathcal{L}(\theta)$ and $\theta \mapsto H\mathcal{L}(\theta)$ are continuous functions, hence using the continuous mapping theorem, it holds that

$$\begin{aligned} \nabla \mathcal{L}(\theta_k^{\phi_N}) &\xrightarrow{d} \nabla \mathcal{L}(\theta_k^\infty), \\ H\mathcal{L}(\theta_k^{\phi_N}) &\xrightarrow{d} H\mathcal{L}(\theta_k^\infty). \end{aligned} \quad (29)$$

Moreover, by Lemma 4 we can express the approximate local optimality assumption (H) in terms of the population loss $\theta \mapsto \mathcal{L}_k(\theta)$ instead of the empirical loss $\theta \mapsto \mathcal{L}_k^N(\theta)$:

$$\begin{aligned} \nabla \mathcal{L}(\theta^{\phi_N}) &= o_{\mathbb{P}}(1), \\ H\mathcal{L}(\theta^{\phi_N}) &\geq o_{\mathbb{P}}(1). \end{aligned} \quad (30)$$

Combining (29) and (30) it follows that $\nabla \mathcal{L}(\theta_k^\infty) = 0$ and $H\mathcal{L}(\theta_k^\infty) \geq 0$. This precisely means that θ_k^∞ is supported on the set of local minimizers Θ^* so that $d(\theta_k^\infty, \Theta^*) = 0$. Hence, (28) implies that $d(\theta_k^{\phi_N}, \Theta^*)$ converges in distribution to a deterministic value 0. This, in turn, means convergence in probability

$$d(\theta_k^{\phi_N}, \Theta^*) \xrightarrow{P} 0. \quad (31)$$

We have extracted a subsequence that satisfies both (27) and (31), which is contradictory. We can therefore conclude that $d(\theta_k^N, \Theta^*) \xrightarrow{P} 0$. We introduce now the decomposition

$$\|\theta_k^N - \theta_k^*\| = (\|\theta_k^N - \theta_k^*\| - d(\theta_k^N, \Theta^*)) + d(\theta_k^N, \Theta^*). \quad (32)$$

We already know that the second term in (32) converges to 0 in probability. Moreover, we know by Assumption (I) that θ_k^* is asymptotically the closest point in Θ^* to θ_k^N , hence, the first term also converges to 0 in probability, concluding the proof. □

Lemma 4. *Under Assumptions (A) to (D) and (F) to (H) It holds that:*

$$\begin{aligned} \nabla \mathcal{L}(\theta^N) &= o_{\mathbb{P}}(1), \\ H\mathcal{L}(\theta^N) &\geq o_{\mathbb{P}}(1). \end{aligned}$$

Proof. For simplicity, we introduce the function:

$$l_\theta(x) = V_k(\tau_\theta(x)) - V_{k-1}(x) - \log |\nabla_x \tau_\theta(x)| + \log\left(\frac{Z_k}{Z_{k-1}}\right).$$

Hence, by definition of $\mathcal{L}(\theta)$ and $\mathcal{L}^N(\theta)$, we have $\mathcal{L}(\theta) = \pi_{k-1}[l_\theta]$ and $\mathcal{L}^N(\theta) = \pi_{k-1}^N[l_\theta]$. Under Assumptions **(F)** and **(G)** the gradient $\nabla_\theta l_\theta(x)$ and Hessian $H_\theta l_\theta(x)$ are well defined and admit a growth of at most 2 in x uniformly in θ . Since, π_{k-1} admit a finite second order moment by Assumption **(B)**, we can apply the dominated convergence theorem to write

$$\begin{aligned} (\pi_{k-1}^N - \pi_{k-1})[\nabla_\theta l_\theta] &= \nabla_\theta \mathcal{L}^N(\theta) - \nabla_\theta \mathcal{L}(\theta), \\ (\pi_{k-1}^N - \pi_{k-1})[H_\theta l_\theta] &= H_\theta \mathcal{L}^N(\theta) - H_\theta \mathcal{L}(\theta). \end{aligned} \quad (33)$$

Moreover, recalling that, under Assumptions **(A)** to **(D)**, the particle estimator π_k^N is consistent by Theorem 1 and the families of functions $(\nabla_\theta l_\theta(x))_{\theta \in \Theta}$ and $(H_\theta l_\theta(x))_{\theta \in \Theta}$ are indexed by a compact set Θ by Assumption **(C)** and admit a quadratic growth in x , we can apply Proposition 9 to both families of functions to get a uniform convergence in probability

$$\begin{aligned} \sup_\theta |(\pi_{k-1}^N - \pi_{k-1})[\nabla_\theta l_\theta]| &\xrightarrow{P} 0, \\ \sup_\theta |(\pi_{k-1}^N - \pi_{k-1})[H_\theta l_\theta]| &\xrightarrow{P} 0. \end{aligned} \quad (34)$$

Therefore, combining (33) and (34) it follows that:

$$\begin{aligned} \sup_\theta \nabla_\theta \mathcal{L}^N(\theta) - \nabla_\theta \mathcal{L}(\theta) &\xrightarrow{P} 0, \\ \sup_\theta H_\theta \mathcal{L}^N(\theta) - H_\theta \mathcal{L}(\theta) &\xrightarrow{P} 0. \end{aligned}$$

We can rely on Assumption **(H)** to directly write

$$\begin{aligned} \nabla \mathcal{L}(\theta_k^N) &= (\nabla \mathcal{L}(\theta_k^N) - \nabla \mathcal{L}^N(\theta_k^N)) + o_{\mathbb{P}}(1) = o_{\mathbb{P}}(1), \\ H\mathcal{L}(\theta_k^N) &\geq (H\mathcal{L}(\theta_k^N) - H\mathcal{L}^N(\theta_k^N)) + o_{\mathbb{P}}(1) = o_{\mathbb{P}}(1). \end{aligned}$$

□

D. Asymptotic stochastic equi-continuity

In this section, we establish asymptotic stochastic equi-continuity (ASEC) of a process defined by the fluctuations of the particle approximation when applied to a suitable class of functions \mathcal{G} . More precisely, we would like to establish ASEC for the empirical process E^N indexed by a class of functions \mathcal{G} and defined as follows:

$$E_k^N : \mathcal{G} \rightarrow \mathbb{R} \quad f \mapsto E_k^N[f] = \sqrt{N}(\pi_k^N[f] - \pi_k[f]).$$

This property will be useful for proving the CLT result in Theorem 2. We start by introducing some notions used in this section.

D.1. Class of functions with finite locally uniform entropy

For some positive constant C , we consider \mathcal{G} a subset of $\mathcal{L}\mathcal{C}_p(C)$ or $\mathcal{C}_p(C)$ with a measurable envelope function $F(x) := \sup_{f \in \mathcal{G}} |f(x)|$. For $\epsilon > 0$ and a probability distribution \mathbb{P} , we denote by $\mathcal{N}(\epsilon, \mathcal{G}, L_2(P))$ the covering number of \mathcal{G} w.r.t. $L_2(P)$, meaning the smallest number p of $L_2(P)$ -balls centered on functions f_1, \dots, f_p and of radius smaller than ϵ needed to cover \mathcal{G} . The uniform covering number defined by taking the supremum of $\mathcal{N}(\epsilon, \mathcal{G}, L_2(P))$ over all probability distributions P . Unlike in (Del Moral, 2004, Lemma 9.6.1.), we will not use the uniform covering number as this quantity will be infinite in our setting. Instead we will consider *local* version of it, which we define next. We first consider the functions $H_2(x) = 1 + \|x\|^{p+2}$ and the set $S_R := \{Q \in \mathcal{P} : Q[H_2^2] \leq R^2\}$ and define the *locally uniform covering number* $\mathcal{N}^{LU}(\epsilon, \mathcal{G}, R)$ to be

$$\mathcal{N}^{LU}(\epsilon, \mathcal{G}, R) := \begin{cases} \sup_{Q \in S_R} \mathcal{N}(\epsilon, \mathcal{G}, L_2(Q)), & S_R \neq \emptyset, \\ 1, & S_R = \emptyset. \end{cases} \quad (35)$$

We define the *locally uniform entropy* of the class \mathcal{G} to be

$$\mathcal{J}(\delta, \mathcal{G}, R) := \int_0^\delta \sqrt{\log(\mathcal{N}^{LU}(\epsilon, \mathcal{G}, R))} d\epsilon. \quad (36)$$

We have the following basic properties of the *locally uniform entropy*.

Proposition 15. *Assume the envelope function $F(x)$ of \mathcal{G} satisfies $F(x) \leq C(1 + \|x\|^p)$ and $J(\delta, \mathcal{G}, r) < +\infty$ for any $r > 1$ and $\delta > 0$. Then for any $\delta > 0$, $0 \leq r \leq r'$, it holds that*

$$\begin{aligned} \mathcal{J}(\delta, \mathcal{G}, r) &\leq \mathcal{J}(\delta, \mathcal{G}, r'), \\ \sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{G}, r) &< +\infty. \end{aligned}$$

Proof. For the first statement, recall that the constraint set S_r defining $\mathcal{N}^{LU}(\epsilon, \mathcal{G}, r)$ in (35) trivially satisfies $S_r \subset S_{r'}$, hence $\mathcal{N}^{LU}(\epsilon, \mathcal{G}, r) \leq \mathcal{N}^{LU}(\epsilon, \mathcal{G}, r')$. The result follows directly by definition of the entropy in (36).

For the second statement, we first note that $\mathcal{N}^{LU}(c, \mathcal{G}, r) = 1$ for any $c \geq 4Cr$. Indeed, either $S_{4Cr} = \emptyset$ in which case $\mathcal{N}^{LU}(c, \mathcal{G}, r) = 1$ holds by definition, or $S_{4Cr} \neq \emptyset$. In the later case, for any element $P \in S_{4Cr}$ any $f, g \in \mathcal{G}$ we have $P[(f - g)^2]^{\frac{1}{2}} \leq 2P[F^2]^{\frac{1}{2}}$. Moreover, since $F(x) \leq C(1 + \|x\|^p) \leq 2C(1 + \|x\|^{p+2})$ we can conclude that $P[(f - g)^2]^{\frac{1}{2}} \leq 2P[F^2]^{\frac{1}{2}} \leq 4Cr \leq c \leq 2Cr \leq c$. Hence, any $L_2(P)$ -ball centered in an element f of \mathcal{G} and of radius c covers \mathcal{G} . This directly implies that $\mathcal{N}^{LU}(c, \mathcal{G}, r) = 1$. The result follows directly by definition of the entropy. \square

We will be in particular interested in classes of functions \mathcal{G} for which the *locally uniform entropy* $\mathcal{J}(\delta, \mathcal{G}, r)$ is finite for any $r > 1$ and $\delta > 0$ and satisfies a growth condition $\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{G}, r) \lesssim r$. Note that this condition always holds when the uniform entropy is finite and is therefore a weaker requirement. Next, we show the stability of this condition when applying the operator $Q_{k,T}$ to functions in f with T varying over \mathcal{T} . More precisely, we control the entropy of the set \mathcal{QG} of the form

$$\mathcal{QG} := \{Q_{k, \tau_\theta}[f] \mid \theta \in \Theta, f \in \mathcal{G}\}. \quad (37)$$

Proposition 16. *Let \mathcal{G} be a subset of $\mathcal{LC}_p(C)$ for some positive constant. Assume \mathcal{G} has a finite bracketing number $\mathcal{N}_{[]}(\epsilon, \mathcal{G}, L_2(\pi_k))$, that the locally uniform entropy $J(\delta, \mathcal{G}, r)$ defined in (36) is finite for any $r > 1$ and $\delta > 0$ and satisfies the linear growth condition $\sup_{\delta > 0} J(\delta, \mathcal{G}, r) \lesssim r$. Under Assumptions (A) and (C) to (G), the class \mathcal{QG} defined in (37) is a subset of $\mathcal{LC}_p(C')$ for some $C' > 0$, has a finite bracketing number $\mathcal{N}_{[]}(\delta, \mathcal{QG}, L_2(\pi_{k-1}))$ and a finite locally uniform entropy $J(\delta, \mathcal{QG}, r)$ for any $r > 1$ and $\delta > 0$ satisfying the linear growth condition:*

$$\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{QG}, r) \lesssim r.$$

In particular, the result holds if \mathcal{G} is a single element in $\mathcal{LC}_p(C)$.

Proof. Bounding the locally uniform entropy. Fix $\epsilon > 0$ and $r > 1$. Let P be a probability distribution in S_r , ie. such that $P[H_2^2] \leq r^2$. We consider an ϵr -net f_1, \dots, f_p of \mathcal{G} with respect to $L_2(PK_k)$. We choose a covering of minimal size, i.e. such that $p = \mathcal{N}(\epsilon r, \mathcal{G}, L_2(PK_k))$. We also consider $\theta_1, \dots, \theta_{p'}$, an ϵ -covering of Θ with with minimum cardinality, i.e. $p' = \mathcal{N}(\epsilon, \Theta, \|\cdot\|)$ where $\mathcal{N}(\epsilon, \Theta, \|\cdot\|)$ is the covering number of Θ with ϵ -balls under the Euclidean distance. We also denote by $\mathcal{J}(\delta, \Theta, \|\cdot\|)$ the entropy of set Θ defined as

$$\mathcal{J}(\delta, \Theta, \|\cdot\|) := \int_0^\delta \sqrt{\log(\mathcal{N}(\epsilon, \Theta, \|\cdot\|))} d\epsilon.$$

Since Θ is bounded and finite dimensional, $\sup_{\delta > 0} \mathcal{J}(\delta, \Theta, \|\cdot\|)$ is finite.

Let $g \in \mathcal{G}$, hence by definition of \mathcal{G} , there exists $\theta \in \Theta$ and $f \in \mathcal{G}$ such that $g = Q_{k, \tau_{\theta_i}}[f]$.

By definition of the ϵ -covering of Θ , there exists $j' \in \{1, \dots, p'\}$ such that $\|\theta - \theta_{j'}\| \leq \epsilon$. We can then use Proposition 17 which holds under Assumptions (A) and (C) to (G) to write:

$$Q_{k, \tau_{\theta_{j'}}}[f](x) - \epsilon c_2 H_2(x) \leq g(x) \leq Q_{k, \tau_{\theta_{j'}}}[f](x) + \epsilon c_2 H_2(x), \quad \forall x \in \mathcal{X}. \quad (38)$$

Moreover, by definition of the ϵr -net, there exists $j \in \{1, \dots, q\}$ such that: $PK_k[(f - f_j)^2] \leq \epsilon^2 r^2$. Subtracting $Q_{k, \tau_{\theta_{j'}}}[f_j](x)$ from (38) yields:

$$Q_{k, \tau_{\theta_{j'}}}[f - f_j](x) - \epsilon c_2 H_2(x) \leq g(x) - Q_{k, \tau_{\theta_{j'}}}[f_j] \leq Q_{k, \tau_{\theta_{j'}}}[f - f_j](x) + \epsilon c_2 H_2(x), \quad \forall x \in \mathcal{X}. \quad (39)$$

We need to quantify $P\left[\left(g - Q_{k, \tau_{\theta_{j'}}}[f_j]\right)^2\right]$. Using (39) it is easy to see that

$$\begin{aligned} P\left[\left(g - Q_{k, \tau_{\theta_{j'}}}[f_j]\right)^2\right]^{\frac{1}{2}} &\leq MP\left[\left(K_k[(f - f_j)]\right)^2\right]^{\frac{1}{2}} + \epsilon c_2 P[H_2^2]^{\frac{1}{2}} \\ &\leq MPK_k[(f - f_j)^2]^{\frac{1}{2}} + \epsilon c_2 P[H_2^2]^{\frac{1}{2}} \\ &\leq \epsilon\left(Mr + c_2 P[H_2^2]^{\frac{1}{2}}\right) \leq \epsilon r(M + c_2) := \epsilon r c_3. \end{aligned}$$

Hence, $Q_{k, \tau_{\theta_{j'}}}[f_j]$ form an $\epsilon r c_3$ net of \mathcal{QG} . This allows us to write:

$$\mathcal{N}(\epsilon r c_3, \mathcal{QG}, L_2(P)) \leq \mathcal{N}(\epsilon r, \mathcal{G}, L_2(PK_k)) \mathcal{N}(\epsilon, \Theta, \|\cdot\|).$$

We will now upper-bound $\mathcal{N}(\epsilon r, \mathcal{G}, L_2(PK_k))$ by a locally uniform covering number. To achieve this we need to find a scalar r' such that $PK_k \in S_{r'}$ whenever $P \in S_r$. By Assumption (A), we know that K_k preserves the classes \mathcal{C}_p and \mathcal{C}_{p+1} . Hence, there exists a constant c_4 large enough so that $K_k[H_2^2] \leq c_4 H_2^2(x)$, implying that $PK_k[H_2^2] \leq c_4 P[H_2^2]$. By choosing $r' = c_4 r$, we are guaranteed that $PK_k \in S_{r'}$. Hence, it follows that

$$\mathcal{N}(\epsilon r c_3, \mathcal{QG}, L_2(P)) \leq \mathcal{N}^{LU}(\epsilon r, \mathcal{G}, c_4 r) \mathcal{N}(\epsilon, \Theta, \|\cdot\|). \quad (40)$$

Taking the supremum over S_r in the l.h.h. of (40), we get

$$\mathcal{N}^{LU}(\epsilon r c_3, \mathcal{QG}, r) \leq \mathcal{N}^{LU}(\epsilon r, \mathcal{G}, c_4 r) \mathcal{N}(\epsilon, \Theta, \|\cdot\|). \quad (41)$$

Taking the logarithm of (41) and using the inequality $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ for any non-negative numbers a and b , we obtain a bound on the entropy after a simple change of variables

$$\mathcal{J}(\delta, \mathcal{QG}, r) \leq c_3 \mathcal{J}(c_3^{-1} \delta, \mathcal{G}, c_4 r) + c_3 r \mathcal{J}((rc_3)^{-1} \delta, \Theta, \|\cdot\|)$$

By assumption, we have that $\sup_{\delta > 0} \mathcal{J}(c_3^{-1} \delta, \mathcal{G}, c_4 r) \lesssim r$. Moreover, since Θ is bounded and finite dimensional, it must hold that $\sup_{\delta > 0, r \geq 1} \mathcal{J}((rc_3)^{-1} \delta, \Theta, \|\cdot\|) < +\infty$. Thus we have shown that $\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{QG}, r) \lesssim r$.

Bounding the bracketing number. Similarly, fix $\epsilon > 0$ and let $\{[l_i, u_i]\}_{1:p}$ be ϵ -brackets covering \mathcal{G} w.r.t. $L_2(\pi_k)$, i.e. $\pi_k[(l_i - u_i)^2] \leq \epsilon^2$ and for any $f \in \mathcal{G}$ there exists $i \in \{1, p\}$ such that

$$l_i(x) \leq f(x) \leq u_i(x), \forall x \in \mathcal{X}.$$

Moreover, we assume the ϵ -brackets to be of minimal size, i.e. $p = \mathcal{N}(\epsilon, \mathcal{G}, L_{\pi_k})$. From (38) we directly have

$$Q_{k, \tau_{\theta_{j'}}}[l_j](x) - \epsilon c_2 H_2(x) \leq g(x) \leq Q_{k, \tau_{\theta_{j'}}}[u_j](x) + \epsilon c_2 H_2(x), \quad \forall x \in \mathcal{X}.$$

Using Proposition 18 we deduce that

$$\pi_{k-1}\left[\left(Q_{k, \tau_{\theta_{j'}}}[u_j - l_j](x) - 2\epsilon c_2 H_2(x)\right)^2\right]^{\frac{1}{2}} \lesssim \epsilon\left(C + \pi_{k-1}[H_2^2]^{\frac{1}{2}}\right).$$

Thus we have shown that the bracketing number $\mathcal{N}_{[]}(\epsilon, \mathcal{QG}, L_{\pi_{k-1}})$ must be finite. \square

Proposition 17. Let \mathcal{G} be a subset $\mathcal{LC}_p(C)$ for some positive constant C . Under Assumptions **(A)** and **(C)** to **(G)**, the class \mathcal{QG} belongs to $\mathcal{LC}_p(C')$ for some positive C' . Moreover, the following holds for any function $f \in \mathcal{G}$

$$|Q_{k,\tau_\theta}[f](x) - Q_{k,\tau_{\theta'}}[f](x)| \lesssim \|\theta - \theta'\| H_2(x)$$

where $H_2(x) = 1 + \|x\|^{p+2}$.

Proof. Fix $1 \leq k \leq K$. We first start by controlling G_{k,τ_θ} . For conciseness, we write $G_{k,\tau_\theta} = e^{h_\theta(x)}$ with $h_\theta(x)$ given by

$$h_\theta(x) = \nabla V_k(\tau_\theta(x)) - \log |\nabla_x \tau_\theta(x)|.$$

The function $\theta \mapsto h_\theta(x)$ is differentiable by Assumptions **(F)** and **(G)** and its gradient is given by:

$$\nabla_\theta h_\theta(x) = \nabla_x V_k(\tau_\theta(x)) \nabla_\theta \tau_\theta(x) - \text{Tr} \left((\nabla_x \tau_\theta(x))^{-1} \nabla_\theta \nabla_x \tau_\theta(x) \right).$$

Moreover, by Assumption **(F)**, we also know that the singular values of $\nabla_x \tau_\theta(x)$ are lower-bounded by a positive constant c uniformly in x and θ . Hence $\left| \text{Tr} \left((\nabla_x \tau_\theta(x))^{-1} \nabla_\theta \nabla_x \tau_\theta(x) \right) \right| \leq c^{-1} \|\nabla_\theta \nabla_x \tau_\theta(x)\|$. Moreover, again by Assumptions **(C)**, **(F)** and **(G)** $\nabla_x V_k(x)$, $\tau_\theta(x)$, $\nabla_\theta \tau_\theta(x)$, $\nabla_\theta \nabla_x \tau_\theta(x)$ all have a linear growth in x uniformly in θ . Hence, we deduce that $\nabla_\theta h_\theta(x)$ has at most a quadratic growth in x . In addition, G_{k,τ_θ} is uniformly bounded in θ and x by Assumption **(D)**, therefore, we deduce that

$$\|\nabla_\theta G_{k,\tau_\theta}(x)\| \leq G_{k,\tau_\theta}(x) \|\nabla_\theta h_\theta(x)\| \leq C(1 + \|x\|^2).$$

For $\theta, \theta' \in \Theta$ and setting $\theta_t = t\theta + (1-t)\theta'$, we use the fundamental theorem of calculus to write:

$$|G_{k,\tau_\theta}(x) - G_{k,\tau_{\theta'}}(x)| \leq \|\theta - \theta'\| \int_0^1 \|\nabla_\theta G_{k,\tau_{\theta_t}}(x)\| dt \lesssim \|\theta - \theta'\| (1 + \|x\|^2). \quad (42)$$

For the second part of the proof, we simply use the following decomposition:

$$Q_{k,\tau_\theta}[f](x) - Q_{k,\tau_{\theta'}}[f](x) = (G_{k,\tau_\theta}(x) - G_{k,\tau_{\theta'}}(x)) K_k[f](x) + G_{k,\tau_\theta}(K_k[f](\tau_\theta) - K_k[f](\tau_{\theta'})). \quad (43)$$

For the first term in the r.h.s. of (43), we use (42) and that, by Assumption **(A)**, $K_k[f](x)$ belongs to $\mathcal{C}_p(C')$ for some constant C' independent of f in $\mathcal{LC}_p(C)$. Hence, we deduce that:

$$|(G_{k,\tau_\theta}(x) - G_{k,\tau_{\theta'}}(x)) K_k[f](x)| \lesssim \|\theta - \theta'\| (1 + \|x\|^{p+2}). \quad (44)$$

For the second term in the r.h.s. of (43), we know by Assumption **(E)** that $K_k[f](x)$ belongs to $\mathcal{LC}_p(C')$ for some constant C' independent of f in $\mathcal{LC}_p(C)$. Since $G_{k,\tau_\theta}(x)$ is bounded uniformly in x and θ , we get

$$|G_{k,\tau_\theta}(K_k[f](\tau_\theta) - K_k[f](\tau_{\theta'}))| \lesssim (1 + \|\tau_\theta(x)\|^{p+1} + \|\tau_{\theta'}(x)\|^{p+1}) \|\tau_\theta(x) - \tau_{\theta'}(x)\|.$$

Moreover, we know that $\tau_\theta(x)$ has at most a linear growth in x uniformly in θ and that

$$\|\tau_\theta - \tau_{\theta'}\| \leq \|\theta' - \theta\| \int_0^1 \|\nabla_\theta \tau_{\theta_t}(x)\| dt \lesssim \|\theta' - \theta\| \|1 + \|x\|\|.$$

We can therefore deduce that:

$$|G_{k,\tau_\theta}(K_k[f](\tau_\theta) - K_k[f](\tau_{\theta'}))| \lesssim \|\theta - \theta'\| (1 + \|x\|^{p+2}). \quad (45)$$

Combining (44) and (45) yields the desired result. \square

Proposition 18. Under Assumption **(D)** for any $f \in L_2(\pi_k)$ we have

$$\|Q_{k,T}[f]\|_{L_2(\pi_{k-1})} \leq C \|f\|_{L_2(\pi_k)}$$

with C being independent of f .

Proof. Let f be a function in $L_2(\pi_k)$, we can directly write:

$$\begin{aligned} \pi_{k-1} \left[(Q_{k,T}[f])^2 \right] &= \int \pi_{k-1}(x) G_{k,T}^2(x) (K_k[f](T(x)))^2 dx \\ &= \frac{Z_k}{Z_{k-1}} \int \pi_k(y) G_{k,T}(T^{-1}(y)) (K_k[f](y))^2 dy \\ &\leq M \frac{Z_k}{Z_{k-1}} \int \pi_k(y) (K_k[f](y))^2 dy \\ &\leq M \frac{Z_k}{Z_{k-1}} \int \pi_k(y) K_k[f^2](y) dy \\ &= M \frac{Z_k}{Z_{k-1}} \pi_k[f^2]. \end{aligned}$$

The second line is using the change of variables $y = T(x)$ and the expression of the importance weight $G_{k,T}$. The third line follows by Assumption **(D)** stating the weights $G_{k,T}(x)$ are bounded by a positive constant $M > 0$ independently from $x \in \mathcal{X}$ and $T \in \mathcal{T}$. The fourth line follows by application of Cauchy–Schwarz while the last line is a consequence of the Markov kernel K_k being invariant under π_k . Choosing $C = M \frac{Z_k}{Z_{k-1}}$ gives the desired result. \square

D.2. Asymptotic stochastic equi-continuity

Definition of asymptotic stochastic equicontinuity For a positive scalar $\delta > 0$ we denote by $\mathcal{G}_k(\delta)$ the intersection of \mathcal{G} with the $L_2(\pi_k)$ -ball of radius δ

$$\mathcal{G}_k(\delta) := \left\{ f \in \mathcal{G} \mid \|f\|_{L_2(\pi_k)} \leq \delta \right\}.$$

Consider a stochastic processes X indexed by \mathcal{G} , we define the uniform semi-norm over the set $\mathcal{G}_k(\delta)$ to be

$$\|X\|_{\mathcal{G}_k(\delta)} := \sup_{f \in \mathcal{G}_k(\delta)} |X(f)|.$$

We will always be in the setting where $\|X\|_{\mathcal{G}_k(\delta)}$ is a random variable (i.e. measurable). Let now X^N be a sequence of stochastic processes that are linear in their index (i.e. $f \mapsto X^N[f]$ is a linear map). In this case the sequence X^N is said to be asymptotically stochastically equicontinuous if for any sequence δ^N of positive real numbers converging to 0, it holds that $\|X^N\|_{\mathcal{G}_k(\delta^N)}$ converges to 0 in probability.

We are now ready to state the main result of this section establishing asymptotic stochastic equicontinuity of the sequence of processes E_k^N .

Proposition 19 (Asymptotic stochastic equicontinuity). *Set $p = 2$. Under Assumptions **(A)** to **(G)**, and for any positive $C > 0$ and any subclass \mathcal{G} of $\mathcal{L}\mathcal{C}_p(C)$ with finite bracketing number $\mathcal{N}_{[]}(\delta, \mathcal{G}, L_2(\pi_k))$ and finite locally uniform entropy **(36)** $\mathcal{J}(\delta, \mathcal{G}, r)$ satisfying $\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{G}, r) \lesssim r$, the process $f \mapsto (E_k^N[f])$ is asymptotically stochastically equicontinuous. In other words, for any sequence δ^N of positive real numbers converging to 0, $\|E_k^N\|_{\mathcal{G}_k(\delta^N)}$ converges to 0 in probability.*

Proof. We proceed by induction over k . The result holds by Proposition 20 for $k = 0$. By induction, we assume the property holds for the process E_{k-1}^N . We then consider a function class \mathcal{G} satisfying the conditions of the result. We will show that $(E_k^N[f])_{\mathcal{G}}$ is asymptotically stochastically equicontinuous.

To achieve this, we fix a sequence δ^N converging to 0 and consider a function $f \in \mathcal{G}_k(\delta^N)$. We then use the following decomposition of $E_k^N[f]$

$$E_k^N[f] = \underbrace{E_k^N[f] - \mathbb{E}_{k-1}[E_k^N[f]]}_{R_k^N[f]} + \underbrace{\mathbb{E}_{k-1}[E_k^N[f]]}_{P_k^N[f]}.$$

Taking the supremum over $\mathcal{G}_k(\delta^N)$ on both sides of the above inequality yields

$$\|E_k^N\|_{\mathcal{G}_k(\delta^N)} \leq \|R_k^N\|_{\mathcal{G}_k(\delta^N)} + \|P_k^N\|_{\mathcal{G}_k(\delta^N)}.$$

By Proposition 20 we have that $\|R_k^N\|_{\mathcal{G}_k(\delta^N)} \xrightarrow[N \rightarrow +\infty]{P} 0$. Thus, we only need to prove that $\|P_k^N\|_{\mathcal{G}_k(\delta^N)} \xrightarrow[N \rightarrow +\infty]{P} 0$. By direct computation, we have:

$$P_k^N[f] = \frac{\gamma_{k-1}^N[1]}{\gamma_k^N[1]} E_{k-1}^N[Q_{k,T_k}[f - \pi_k[f]]].$$

Now let us call \mathcal{QG} be the set of functions of the form $Q_{k,T}[f]$ for some $f \in \mathcal{G}$ and $T \in \mathcal{T}$. Proposition 18 ensures also that $\pi_{k-1}[Q_{k,T_k}[f]] \leq C\pi_k[f] \leq C\delta^N$, hence $Q_{k,T_k}[f] \in \mathcal{QG}_{k-1}(C\delta^N)$. Therefore, we have shown that

$$\|P_k^N\|_{\mathcal{QG}_k(\delta^N)} \leq \frac{\gamma_{k-1}^N[1]}{\gamma_k^N[1]} \|E_{k-1}^N\|_{\mathcal{QG}_{k-1}(C\delta^N)}.$$

Finally, by Proposition 16, we know that \mathcal{QG} belongs to $\mathcal{LC}_p(C')$ for some positive C' , has a finite bracketing number $\mathcal{N}(\epsilon, \mathcal{QG}, L_2(\pi_{k-1}))$ and a finite locally covering entropy $\mathcal{J}(\delta, \mathcal{QG}, r)$ for any $\delta > 0$ and $r > 1$ satisfying $\mathcal{J}(\delta, \mathcal{QG}, r) \lesssim r$. Therefore, by the recursion assumption, we know that $\|E_{k-1}^N\|_{\mathcal{QG}_{k-1}(C\delta^N)} \xrightarrow[N \rightarrow +\infty]{P} 0$. Moreover, recalling that $\gamma_k^N[1] \xrightarrow[N \rightarrow +\infty]{P} Z_{k-1}$ for any $0 \leq k \leq K$ by Theorem 1, we can directly conclude that $\|P_k^N\|_{\mathcal{G}_k(\delta^N)} \xrightarrow[N \rightarrow +\infty]{P} 0$. \square

Proposition 20. *Let \mathcal{G} be the class defined in Proposition 19 and denote by R_k^N to be the process index by \mathcal{G} and defined by:*

$$R_k[f] = \sqrt{N}(\pi_k^N[f] - \mathbb{E}_{k-1}[\pi_k^N[f]])$$

with the convention that for $k = 0$, \mathbb{E}_{-1} is the expectation over the samples $\{X_0^{(i)}\}_{1:N}$ from π_0 . Under Assumptions (A) to (D) and for any sequence δ^N converging to 0 and any $0 \leq k \leq K$, it holds that

$$\|R_k^N\|_{\mathcal{G}_k(\delta^N)} \xrightarrow[N \rightarrow +\infty]{P} 0.$$

Proof. Let $0 \leq k \leq K$. We use the convention that \mathcal{F}_{-1} is the empty sigma algebra and that \mathbb{E}_{-1} is the expectation over the samples $X_0^{(i)}_{1:N}$ from π_0 . This will allow the same treatment for both $k = 0$ and $k > 0$. Note that, $\mathbb{E}_{k-1}[R_k^N[f]] = 0$ for any $f \in \mathcal{G}$. Hence, the process is centered conditionally on \mathcal{F}_{k-1} . Moreover, the particles $\{X_k^{(i)}\}_{1:N}$ defining R_k^N are i.i.d. conditionally on \mathcal{F}_{k-1} . Therefore, we can use a symmetrization approach as in (Del Moral, 2004, Lemma 9.6.1.). Let $\{\epsilon^i\}_{1:N}$ be i.i.d. Bernoulli variables with $\mathbb{P}(\epsilon^i = 1) = \mathbb{P}(\epsilon^i = -1) = \frac{1}{2}$ independent from \mathcal{F}_k . Conditionally on \mathcal{F}_{k-1} the samples $\{X_k^i\}_{1:N}$ are i.i.d., hence, the symmetrization inequality in (Sen, 2018, Theorem 3.14) holds conditionally on \mathcal{F}_{k-1} :

$$\mathbb{E}_{k-1} \left[\|R_k^N\|_{\mathcal{G}_k(\delta^N)} \right] \leq 2\sqrt{N} \mathbb{E}_{k-1} \left[\|m_\epsilon^N\|_{\mathcal{G}_k(\delta^N)} \right] \quad (46)$$

where $m_\epsilon^N = \frac{1}{N} \sum_{i=1}^N \epsilon_i \delta_{X_k^i}$. Now by conditioning on \mathcal{F}_k , we can apply Hoeffding inequality (Sen, 2018, Lemma 3.11) which implies that

$$\mathbb{P}_k \left[\sqrt{N} (m_\epsilon^N[f] - m_\epsilon^N[h]) > s \right] \leq 2e^{-\frac{s^2}{2\pi_k^N[(f-h)^2]}}.$$

Following the proof of (Sen, 2018, Theorem 4.8), we can use the Maximal inequality for sub-Gaussian processes (Sen, 2018, Theorem 4.5), which implies that:

$$\sqrt{N} \mathbb{E}_k \left[\|m_\epsilon^N\|_{\mathcal{G}_k(\delta^N)} \right] \lesssim \int_0^{\psi^N} \sqrt{\log(\mathcal{N}(r, \mathcal{G}_k(\delta^N), L_2(\pi_k^N)))} dr, \quad (47)$$

where $\psi^N := \sup_{f \in \mathcal{G}(\delta^N)} \pi_k^N[f^2]$. By definition of the locally uniform covering number in (35), we have that

$$\mathcal{N}(r, \mathcal{G}_k(\delta^N), L_2(\pi_k^N)) \leq \mathcal{N}^{LU}(r, \mathcal{G}, r^N),$$

with $(r^N)^2 = \pi_k^N[H_2^2]$. Hence, we can upper-bound $\sqrt{N}\mathbb{E}_k[\|m_\epsilon^N\|_{\mathcal{G}(\delta^N)}]$ using the locally uniform entropy so that (47) becomes

$$\sqrt{N}\mathbb{E}_k[\|m_\epsilon^N\|_{\mathcal{G}_k(\delta^N)}] \lesssim \mathcal{J}(\psi^N, \mathcal{G}, r^N). \quad (48)$$

We can now take the expectation of (48) conditionally on \mathcal{F}_{k-1} and combine the resulting inequality with (46) yielding

$$\mathbb{E}_{k-1}[\|R_k^N\|_{\mathcal{G}_k(\delta^N)}] \lesssim \mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, r^N)].$$

We only need to prove that $\mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, r^N)] \xrightarrow{N \rightarrow +\infty} 0$. To this end, we define the event

$$S_\epsilon^N = \{|\pi_k^N[H_2^2] - \pi_k[H_2^2]| \geq \epsilon\}$$

for some $\epsilon > 0$ and its complement \bar{S}_ϵ^N . We then decompose the expected entropy into to terms

$$\mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, r^N)] = \mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, r^N)\bar{S}_\epsilon^N] + \mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, r^N)S_\epsilon^N].$$

Conditionally on \bar{S}_ϵ^N we have that $r^N \leq \sqrt{\epsilon + \pi_k[H_2^2]} := R$. Hence, by monotonicity of the entropy w.r.t. to the last argument (Proposition 15), it holds that $\mathcal{J}(\psi^N, \mathcal{G}, r^N)\bar{S}_\epsilon^N \leq \mathcal{J}(\psi^N, \mathcal{G}, R)$. Conditionally on S_ϵ^N we will rely on the assumption that $\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{G}, r) \lesssim r$ which implies that $\mathcal{J}(\psi^N, \mathcal{G}, r^N)S_\epsilon^N \lesssim r^N S_\epsilon^N$. Hence, we can write

$$\begin{aligned} \mathbb{E}_{k-1}[\|R_k^N\|_{\mathcal{G}_k(\delta^N)}] &\lesssim \mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, r^N)] \lesssim \mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, R)] + \mathbb{E}_{k-1}[\pi_k^N[H_2^2]^{\frac{1}{2}} S_\epsilon^N] \\ &\lesssim \mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, R)] + \mathbb{E}_{k-1}[\pi_k^N[H_2^2]^{\frac{1}{2}}] \mathbb{P}_{k-1}[S_\epsilon^N]^{\frac{1}{2}}, \end{aligned} \quad (49)$$

where we used Cauchy–Schwarz inequality to get the second line. To control the first term, we first note that $\psi^N \xrightarrow{P} 0$. Indeed, by construction, $\pi_k[f^2] \leq (\delta^N)^2$ for any $f \in \mathcal{G}_k(\delta^N)$. Moreover, by assumption, $\mathcal{G}_k(\delta^N)$ has a finite bracketing number $\mathcal{N}_{[]}(\epsilon, \mathcal{G}_k(\delta^N), L_2(\pi_k))$ for any $\epsilon > 0$ and the functions f^2 have a growth of at most $2p$. We can then apply Proposition 8, which holds under Assumptions (A) to (D) to get the following uniform convergence in probability of $\pi_k^N[f^2]$ over $\mathcal{G}_k(\delta^N)$, i.e. $\sup_{f \in \mathcal{G}_k(\delta^N)} (\pi_k^N - \pi_k)[f^2] \xrightarrow{P} 0$. This allows to conclude that

$$\psi^N \leq (\delta^N)^2 + \sup_{f \in \mathcal{G}_k(\delta^N)} (\pi_k^N - \pi_k)[f^2] \xrightarrow{P} 0.$$

Second, we know by Proposition 15 that $\mathcal{J}(\psi^N, \mathcal{G}, R)$ is bounded by a finite quantity, i.e. $\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{G}, R) < +\infty$. Hence, by the dominated convergence theorem, it follows that $\mathbb{E}_{k-1}[\mathcal{J}(\psi^N, \mathcal{G}, R)] \xrightarrow{P} 0$. For the second term we have that $\mathbb{E}_{k-1}[\pi_k^N[H_2^2]] \xrightarrow{P} \pi_k[H_2^2]$ by Theorem 1 with $\pi_k[H_2^2]$ being finite under Assumption (B). Moreover, since we also have $\pi_k^N[H_2^2] \xrightarrow{P} \pi_k[H_2^2]$, this necessarily implies that $\mathbb{E}[\mathbb{P}_{k-1}[S_\epsilon^N]] = \mathbb{P}[S_\epsilon^N] \xrightarrow{P} 0$ and henceforth that $\mathbb{P}_{k-1}[S_\epsilon^N] \xrightarrow{P} 0$.

We have shown that the r.h.s. of (49) converges to 0 in probability. This directly implies $\mathbb{E}_{k-1}[\|R_k^N\|_{\mathcal{G}_k(\delta^N)}] \xrightarrow{P} 0$ and henceforth that $a_\epsilon^N := \mathbb{P}_{k-1}[\|R_k^N\|_{\mathcal{G}_k(\delta^N)} \geq \epsilon] \rightarrow 0$ for any choice of $\epsilon > 0$ by Markov inequality. Noting that a_ϵ^N is bounded for a given $\epsilon > 0$, we get that $\mathbb{E}[a_\epsilon^N] \rightarrow 0$ which exactly means $\|R_k^N\|_{\mathcal{G}_k(\delta^N)} \xrightarrow{P} 0$. \square

We then have the following result which is a consequence of Proposition 19

Proposition 21. *Let f be in $\mathcal{L}\mathcal{C}_p$ and consider the family of function $\mathcal{Q}\mathcal{G}$ of the form $S_\theta(x) = Q_{k, \tau_\theta}[f](x) - Q_{k, \tau_{\theta^*}}[f](x)$ indexed by the parameter $\theta \in \Theta$. Under Assumptions (A) to (G) and for any random sequence g^N in $\mathcal{Q}\mathcal{G}$ such that $\pi_{k-1}[(g^N)^2] \xrightarrow{P} 0$ it holds that $E_{k-1}^N[g^N] \xrightarrow{P} 0$.*

Proof. By Proposition 16 we know that \mathcal{QG} belongs to $\mathcal{LC}_p(C)$ for some positive C , has a finite bracketing number $\mathcal{N}(\delta, \mathcal{QG}, L_2(\pi_{k-1}))$, a finite *locally uniform entropy* $\mathcal{J}(\delta, \mathcal{QG}, r)$ for any $r \geq 1$ and $\delta > 0$ and satisfying a linear growth condition $\sup_{\delta > 0} \mathcal{J}(\delta, \mathcal{QG}, r) \lesssim r$. Hence, by direct application of Proposition 19, we know that the process $g \mapsto \sqrt{N}(\pi_{k-1}^N - \pi_{k-1})[g] := E_{k-1}^N[g]$ indexed by functions $g \in \mathcal{QG}$ is asymptotically stochastically equicontinuous, meaning that $\|E_{k-1}^N\|_{\mathcal{QG}(\delta^N)} \xrightarrow{P} 0$ for any sequence of positive numbers δ^N converging to 0 and where by definition:

$$\|E_{k-1}^N\|_{\mathcal{QG}(\delta^N)} := \sup_{\substack{g \in \mathcal{QG} \\ \pi_{k-1}[g^2] \leq (\delta^N)^2}} |E_{k-1}^N[g]|.$$

Denote by $b^N := \pi_k[(g^N)^2]$ and fix $\epsilon > 0$. We know that for any sequence R^N , the following inclusion of events hold:

$$\begin{aligned} \{E_k^N[g^N] > \epsilon\} \cap \{b^N > R^N\} &\subset \{b^N > R^N\} \\ \{E_k^N[g^N] > \epsilon\} \cap \{b^N \leq R^N\} &\subset \left\{ \|E_{k-1}^N\|_{\mathcal{QG}(R^N)} \geq \epsilon \right\}. \end{aligned}$$

Hence, the following inequality holds

$$\begin{aligned} \mathbb{P}[E_k^N[g^N] > \epsilon] &= \mathbb{P}[\{E_k^N[g^N] > \epsilon\} \cap \{b^N > R^N\}] + \mathbb{P}[\{E_k^N[g^N] > \epsilon\} \cap \{b^N \leq R^N\}] \\ &\leq \mathbb{P}[b^N > R^N] + \mathbb{P}\left[\|E_{k-1}^N\|_{\mathcal{QG}(R^N)} \geq \epsilon\right]. \end{aligned}$$

Moreover, by assumption, we have $b^N := \pi_k[(g^N)^2] \xrightarrow{P} 0$. Hence, by Lemma 5, there exists a deterministic sequence R^N converging to 0 and such that

$$\mathbb{P}[b^N > R^N] \leq R^N.$$

We therefore have

$$\mathbb{P}[E_k^N[g^N] > \epsilon] \leq R^N + \mathbb{P}\left[\|E_{k-1}^N\|_{\mathcal{QG}(R^N)} \geq \epsilon\right].$$

The first term converges to 0 by definition of R^N so does the second term by asymptotic stochastic equicontinuity of E_{k-1}^N and since $R^N \rightarrow 0$. Therefore, we have shown that $\mathbb{P}[E_k^N[g^N] > \epsilon] \rightarrow 0$. □

Lemma 5. *If b^N is non-negative r.v. such that $b^N \xrightarrow{P} 0$, then there exists a deterministic sequence R^N converging to 0 such that:*

$$\mathbb{P}[b^N > R^N] \leq R^N.$$

Proof. Using (Dudley, 2018, Theorem 9.2.2.), $b^N \xrightarrow{P} 0$ is equivalent to having $r^N \rightarrow 0$ where

$$r^N := \inf \{ \epsilon \geq 0 \mid \mathbb{P}[b^N > \epsilon] \leq \epsilon \}.$$

Moreover, by definition of r^N , for any N there exists $a^N \leq \frac{1}{N}$ such that $R^N := r^N + a^N$ satisfies

$$\mathbb{P}[b^N > R^N] \leq R^N.$$

Since $R^N \rightarrow 0$, we have constructed a sequence R^N converging slowly enough to 0 so that $\mathbb{P}[b^N > R^N] \rightarrow 0$. □

D.3. Proof of Proposition 10

Proof of Proposition 10. We first decompose P_N as $P_N = A_N + B_N$ with

$$\begin{aligned} A_N &= \sqrt{N} \gamma_{k-1}^N [1] (\pi_{k-1}^N - \pi_{k-1}) [Q_{k, T_k} [f] - Q_{k, T_k^*} [f]], \\ B_N &= \sqrt{N} (\gamma_{k-1}^N [Q_{k, T_k^*} [f]] - \gamma_{k-1} [Q_{k, T_k^*} [f]]), \end{aligned}$$

where we used the identity $\pi_{k-1}[Q_{k,T_k}[f]] = \pi_{k-1}[Q_{k,T_k^*}[f]]$ from Proposition 6 to express A_N as a difference in expectations under π_{k-1}^N and π_{k-1} . By induction, B_N converges in distribution to a centered Gaussian with variance $\mathbb{V}_{k-1}^N[Q_{k,T_k^*}[f]]$.

We only need to show that $A_N \xrightarrow{P} 0$. To this end, for the given function $f \in \mathcal{L}\mathcal{C}_p$, we consider the family of functions $\mathcal{Q}\mathcal{G}$ of the form $S_\theta(x) = Q_{k,\tau_\theta}[f](x) - Q_{k,\tau_{\theta^*}}[f](x)$ indexed by the parameter $\theta \in \Theta$. We will first prove that $\pi_{k-1}[S_{\theta^N}^2] \xrightarrow{P} 0$ then apply Proposition 21 to show that $\sqrt{N}(\pi_{k-1}^N - \pi_{k-1})[S_{\theta^N}] \xrightarrow{P} 0$. By Proposition 17, we have that

$$|S_{\theta^N}(x)| = |Q_{k,T_k}[f](x) - Q_{k,T_k^*}[f](x)| \lesssim \|\theta_k^N - \theta_k^*\| H_2(x) \quad (50)$$

where $H_2(x) = 1 + \|x\|^{p+2}$. Moreover, we know by Assumption (B) that $\pi_{k-1}[H_2^2]$ is finite. We can square (50) and take the expectation under π_{k-1} to get

$$\pi_{k-1}[S_{\theta^N}^2] \lesssim \|\theta_k^N - \theta_k^*\|^2 \pi_{k-1}[H_2^2] < +\infty.$$

Since $\|\theta_k^N - \theta_k^*\| \xrightarrow{P} 0$ by Proposition 14 we conclude that $\pi_{k-1}[S_{\theta^N}^2] \xrightarrow{P} 0$. We can then apply Proposition 21 to the sequence $g^N := S_{\theta^N}$ to get that $\sqrt{N}(\pi_{k-1}^N - \pi_{k-1})[S_{\theta^N}] \xrightarrow{P} 0$. We also know by Theorem 1 that $\gamma_{k-1}^N[1] \xrightarrow{P} \gamma_{k-1}[1]$. Therefore, it follows directly that $A_N \xrightarrow{P} 0$ concluding the proof. \square

E. Continuous-time limit

E.1. Notation and General setting

We start by introducing some notations.

Continuous-time path of probability measures. Let $(\Pi_t)_{[0,1]}$ be a continuous-time path of probabilities connecting π_0 to π . We write $V_t(x)$ and Z_t to denote the known potential and unknown normalizing constant of Π_t . We assume that the discrete auxiliary targets π_k are of the form $\pi_k = \Pi_{t_k}$ where $t_k = k\lambda$ and $\lambda = \frac{1}{K}$. In this case, increasing K amounts to decreasing the *step-size* λ .

Markov Kernels. We consider the ULA kernels K_k of the form:

$$K_k(x, dy) = \mathcal{N}(y; x - \lambda V_k(x), 2\lambda I) dy$$

where $\mathcal{N}(x; m, \Sigma)$ is the Gaussian density with mean m and covariance Σ .

Normalizing Flows. For a given fixed step-size $\lambda > 0$, we will assume that the normalizing flows T are of the form:

$$T(x) = x + \lambda A_\theta(x)$$

where $(\theta, x) \mapsto A_\theta(x)$ is a function defined from $\Theta \times \mathcal{X}$ to \mathcal{X} , with Θ being a compact subset of \mathbb{R}^p . We then call \mathcal{T}_λ the set of normalizing flows defined by varying the parameter θ :

$$\mathcal{T}_\lambda = \{x \mapsto x + \lambda A_\theta(x) \mid \theta \in \Theta\}. \quad (51)$$

To ensure that all functions in \mathcal{T}_λ are diffeomorphisms, we will require $A_\theta(x)$ to satisfy Assumption (b) and for λ small enough, the following lemma ensures that all functions in \mathcal{T}_λ are indeed diffeomorphisms and thus define a valid family of normalizing flows:

Lemma 6. *Under Assumption (b) and if $\lambda \leq \frac{1}{2L}$, then any element in \mathcal{T}_λ is a diffeomorphism.*

Proof. Let θ be in Θ . The map $T(x) = x + \lambda A_\theta(x)$ satisfies $\|T(x)\| \rightarrow +\infty$ whenever $\|x\| \rightarrow +\infty$ since $x \mapsto \lambda A_\theta(x)$ has a growth in x of at most $\frac{1}{2}\|x\|$. Therefore, T is proper. Moreover, the Jacobian $\nabla T(x)$ is invertible for any $x \in \mathcal{X}$. Thus by the Hadamard–Caccioppoli Theorem, one can conclude that T is a diffeomorphism. \square

Time-dependent control. To describe the continuous-time limit, we consider time dependent controls of the form $\alpha_t(x) = A_{\theta_t}(x)$, where $t \mapsto \theta_t$ is a trajectory in parameter space Θ . More precisely, we consider the set $\text{Lip}([0, 1], \Theta)$ of 1-Lipschitz functions $t \mapsto \theta_t$ from $[0, 1]$ to Θ and we define the set \mathcal{A} of admissible controls as

$$\mathcal{A} := \{\alpha : (t, x) \mapsto A_{\theta_t}(x) \mid t \mapsto \theta_t \in \text{Lip}([0, 1], \Theta)\}.$$

The following smoothness properties are a direct consequence of Assumption **(b)**.

Lemma 7. *Under Assumption **(b)**, any control in \mathcal{A} satisfies:*

$$\begin{aligned} \|\alpha_t(x) - \alpha_{t'}(x')\| &\leq L(\|x - x'\| + |t - t'|), \\ \|\nabla \alpha_t(x) - \nabla \alpha_{t'}(x')\| &\leq L(1 + \|x\| + \|x'\|)(\|x - x'\| + |t - t'|), \\ \|\alpha_0(0)\| &\leq C, \end{aligned} \tag{52}$$

with $C = \sup_{\theta \in \Theta} |A_\theta(0)| < +\infty$.

Optimal controls and Normalizing Flows. For a given step-size $\frac{1}{2L} \geq \lambda > 0$, we denote by T_k^λ an NF that minimizes \mathcal{L}_k over the set \mathcal{T}_λ defined in (51) and we denote by θ_k^λ its corresponding parameter in Θ ,

$$T_k^\lambda(x) = x + \lambda A_{\theta_k^\lambda}(x).$$

The NF T_k^λ and parameter θ_k^λ exist since the set Θ is compact and the maps $\theta \mapsto \mathcal{L}_k(I + \lambda A_\theta)$ is continuous. For any $\lambda = \frac{1}{K}$ with K being a positive integer, we will later assume there exists at least a solution $\alpha^\lambda \in \mathcal{A}$ that interpolates between the optimal NFs T_k^λ for all $0 \leq k \leq K$, i.e.:

$$T_k^\lambda(x) = x + \lambda \alpha_{t_k}^\lambda(x), \quad \forall k \in [1 : K].$$

It is easy to see that the loss $\mathcal{L}_\lambda^{\text{tot}}(\alpha)$ defined in (53) admits α^λ as a minimizer over the set of admissible controls \mathcal{A} :

$$\mathcal{L}_\lambda^{\text{tot}}(\alpha) := \frac{1}{\lambda} \sum_{k=1}^K \mathcal{L}_k(I + \lambda \alpha_{t_k}). \tag{53}$$

Feynman–Kac measures. We will introduce a family of Feynman–Kac measures indexed by \mathcal{A} . For any element $\alpha \in \mathcal{A}$, we denote by $b_t^\alpha(x)$ the controlled drift:

$$b_t^\alpha(x) = \alpha_t(x) - \nabla_x V_t(x).$$

We first start by introducing a family of non-homogeneous stochastic differential equations:

$$dX_t^\alpha = b_t^\alpha(X_t^\alpha) dt + \sqrt{2} dB_t, \quad X_0^\alpha \sim \pi_0, \tag{54}$$

where $(B_t)_{t \geq 0}$ is a standard multivariate Brownian motion. Here, we emphasize the dependence of the controlled drift b_t^α and the process $(X_t^\alpha)_{t \geq 0}$ on the control α . We will sometimes simply write b_t and X_t when clear from context. We denote by $\bar{\Lambda}_t^\alpha$ the joint distribution of such process up to time t and by Λ_t^α its marginal at time t which satisfies the following continuity equation:

$$\partial_t \Lambda_t^\alpha = \nabla \cdot b_t^\alpha + \Delta \Lambda_t^\alpha.$$

We can then consider the following continuous-time importance weights over the process X_t^α :

$$w_t^\alpha = \exp\left(\int_0^t g_s^\alpha(X_s^\alpha) ds\right), \quad g_t^\alpha = \nabla \cdot \alpha_t - \nabla_x V_t^\top \alpha_t - \partial_t V_t.$$

As shown later in Proposition 22 in Appendix E.3, the normalizing constant $\bar{\Lambda}_t^\alpha[w_t^\alpha]$ of the weights w_t^α is given by $\exp\left(-\int_0^t \Pi_s[\partial_s V_s] ds\right)$. This allows to define the normalized weights as:

$$\bar{w}_t^\alpha = \exp\left(\int_0^t \bar{g}_s^\alpha(X_s^\alpha) ds\right), \quad \bar{g}_t^\alpha = g_t^\alpha + \Pi_t[\partial_t V_t].$$

The function $g_t^\alpha(x)$ is called the *instantaneous work* and measures how much Λ_t is different from Π_t . We can now introduce the family of Feynman–Kac measures defined by the importance weights over the sample path $(X_t^\alpha)_{[0,1]}$ from the process defined in (54):

$$\bar{\Pi}_t^\alpha[f] = \frac{\bar{\Lambda}_t^\alpha[w_t^\alpha f]}{\bar{\Lambda}_t^\alpha[w_t^\alpha]} = \bar{\Lambda}_t^\alpha[\bar{w}_t^\alpha f], \quad (55)$$

where f is any bounded continuous functional defined over the space of admissible sample paths $(X_t^\alpha)_{[0,1]}$. We will see in Proposition 22 that the marginal of $\bar{\Pi}_t^\alpha$ at time t is exactly Π_t .

We consider now the *expected instantaneous variance* of $g_{t,\alpha}$ defined as:

$$\mathcal{M}(\alpha) := \frac{1}{2} \int_0^1 \left(\Pi_t \left[(g_t^\alpha)^2 \right] - (\Pi_t[g_t^\alpha])^2 \right) dt. \quad (56)$$

We will show that $\mathcal{M}(\alpha)$ is the limit function of (53). We will assume that the infimum of $\mathcal{M}(\alpha)$ over \mathcal{A} is achieved for an element $\alpha^* \in \mathcal{A}$:

$$\mathcal{M}(\alpha^*) = \inf_{\alpha \in \mathcal{A}} \mathcal{M}(\alpha).$$

Interpolating measures. For a given $\lambda = \frac{1}{K}$ and a control $\alpha \in \mathcal{A}$, we define the following functions:

$$\begin{aligned} \beta_t^{\alpha,\lambda}(x) &= \alpha_t(x) - \nabla V_t(x + \lambda \alpha_t(x)), \\ \delta_t^{\alpha,\lambda}(x) &= \frac{1}{\lambda} (V_{t-\lambda}(x) - V_t(x + \lambda \alpha_t(x)) + \log |I + \lambda \nabla \alpha_t(x)|) \end{aligned} \quad (57)$$

with $\delta_t^{\alpha,\lambda}(x)$ being defined for $t \geq \lambda$. The function $\beta_t^{\alpha,\lambda}$ allows us to introduce the non-anticipative drift function $b_t^{\alpha,\lambda}$ which depends on the path of a process $X_{[0,1]}$ up to time t :

$$b_t^{\alpha,\lambda}(X) = \sum_{k=1}^K \beta_{t_k}^{\alpha,\lambda}(X_{t_{k-1}}) \mathbb{1}_{[t_{k-1}, t_k)}(t).$$

We can then consider the continuous-time process $(X_t^{\alpha,\lambda})_{[0,1]}$ defined as:

$$dX_t^{\alpha,\lambda} = b_t^{\alpha,\lambda}(X^{\alpha,\lambda}) dt + \sqrt{2} dB_t, \quad X_0^{\alpha,\lambda} \sim \pi_0. \quad (58)$$

We denote by $\bar{\Lambda}_t^{\alpha,\lambda}$ the joint distribution of $X_{[0,t]}^{\alpha,\lambda}$. To introduce the interpolating measure, we start by defining the instantaneous work:

$$g_s^{\alpha,\lambda}(X) = \delta_{t_k}^{\alpha,\lambda}(X_{t_{k-1}}), \quad t_{k-1} \leq s < t_k. \quad (59)$$

It is easy to see that $\delta_{t_k}^{\alpha,\lambda}(x) = \frac{1}{\lambda} \log G_{k, T_k^{\alpha,\lambda}}(x)$ where $T_k^{\alpha,\lambda}$ are defined from the control α using $T_k^{\alpha,\lambda} = x + \lambda \alpha_{t_k}(x)$. We define the truncated time $\tau_\lambda(t) = \lambda \lfloor \frac{t}{\lambda} \rfloor$ and the corresponding index $k_\lambda(t) = \lfloor \frac{t}{\lambda} \rfloor$. We then introduce the following functions:

$$w_t^{\alpha,\lambda}(X_{[0,t]}) := \exp \left(\int_0^{\tau_\lambda(t)} g_s^{\alpha,\lambda}(X) ds \right), \quad r_t^\lambda(X_{[0,t]}) = \exp \left(- \sum_{k=1}^{k_\lambda(t)} h_{t_k}^\lambda(X_{t_k}) \right) \quad (60)$$

where h_t^λ is defined as

$$h_t^\lambda(x) = \log \left(\int \exp(V_t(x) - V_t(y) - \frac{1}{4\lambda} \|x - y - \lambda \nabla V_t(y)\|^2) dy \right) - \frac{d}{2} \log(4\pi\lambda).$$

The function $w_t^{\alpha,\lambda}$ represents the correction due to the use of the control α , while r_t^λ is the correction due to the transition kernel K_k being only approximately invariant w.r.t. π_k . We can now introduce the interpolating measures $\bar{\Pi}_t^\lambda$ defined as:

$$\bar{\Pi}_t^{\alpha,\lambda}[f] = \frac{\bar{\Lambda}_t^{\alpha,\lambda}[w_t^{\alpha,\lambda} r_t^\lambda f]}{\bar{\Lambda}_t^\lambda[w_t^{\alpha,\lambda} r_t^\lambda]}. \quad (61)$$

Additional notations We introduce discrepancy measures between controls α and α' :

$$\begin{aligned} \mathcal{S}(\alpha, \alpha') &= \int_0^1 \Pi_t \left[\|\alpha_t - \alpha'_t\|^2 + \|\nabla \alpha_t - \nabla \alpha'_t\|^2 \right] dt, \\ \mathcal{S}_t^\lambda(\alpha, \alpha') &= \lambda \sum_{k=1}^{k_\lambda(t)} \Pi_{t_{k-1}} \left[\|\alpha_{t_k} - \alpha'_{t_k}\|^2 + \|\nabla \alpha_{t_k} - \nabla \alpha'_{t_k}\|^2 \right]. \end{aligned} \quad (62)$$

For a function $c : [0, 1] \times \mathcal{X} \rightarrow \mathbb{R}$, we define the scalar:

$$\mathbb{M}_t^\lambda[c] := \lambda \sum_{k=1}^{k_\lambda(t)} \Pi_{t_{k-1}}[c_{t_k}]. \quad (63)$$

In all what follows, for two real numbers A and B , the relation $A \lesssim B$ mean that there exists a positive constant $C > 0$ that is uniform over $t \in [0, 1]$ and over the set of admissible controls \mathcal{A} and a value λ_0 such that $A \leq CB$ for all $\lambda \leq \lambda_0$.

In the rest of the paper, when we write \mathbb{E} , the expectation is w.r.t. $\bar{\Lambda}^\alpha$ unless stated otherwise. To simplify notation, we also write X_s in place of X_s^α for the process satisfying the SDE in (54) with control α .

E.2. Assumptions

- (a) For some fixed $L > 0$, $V_t(x)$ is continuously differentiable in (t, x) with $\nabla V_t(x)$ being L -Lipschitz jointly in (t, x) . Moreover, $\partial_t V_t(x)$ satisfies:

$$|\partial_t V_t(x) - \partial_t V_{t'}(x')| \leq L(1 + \|x\| + \|x'\|)(\|x - x'\| + |t - t'|).$$

- (b) $A_\theta(x)$ is continuously differentiable in (θ, x) , L -Lipschitz jointly in (θ, x) and satisfies:

$$\|\nabla A_\theta(x) - \nabla A_{\theta'}(x')\| \leq L(1 + \|x\| + \|x'\|)(\|x - x'\| + |\theta - \theta'|),$$

for all $x, x' \in \mathcal{X}$ and $\theta, \theta' \in \Theta$.

- (c) Π_t admits finite 4-th order moments uniformly bounded for $t \in [0, 1]$.

To provide the main convergence result, we need to strengthen the assumption on the moments of Π_t :

- (d) There exists $c > 0$ such that for any $0 \leq \lambda < c$ the expectation $\Pi_t \left[\exp \left(\lambda \|\nabla V_t\|^2 \right) \right]$ is finite and uniformly bounded for $t \in [0, 1]$.

We will then assume the existence of admissible solutions α^* minimizing \mathcal{M} and α^λ minimizing $\mathcal{L}_\lambda^{\text{tot}}(\alpha)$ for any $\lambda = \frac{1}{K}$. We will also need an assumption on the local behavior of the loss $\mathcal{M}(\alpha)$ near the optimum α^* . This local behavior will be controlled in terms of the following discrepancy:

$$\mathcal{S}(\alpha, \alpha') = \int_0^1 \Pi_t \left[\|\alpha_t - \alpha'_t\|^2 + \|\nabla \alpha_t - \nabla \alpha'_t\|^2 \right] dt.$$

- (e) The minimizer α^* of $\mathcal{M}(\alpha)$ exists in \mathcal{A} and is unique. Moreover, for any $\delta > 0$ it holds that

$$\mathcal{M}(\alpha^*) < \inf_{\mathcal{S}(\alpha, \alpha^*) > \delta} \mathcal{M}(\alpha).$$

Finally, there exists $\delta_0 > 0$ such that for all $\alpha \in \mathcal{A}$ satisfying $\mathcal{S}(\alpha^*, \alpha) \leq \delta_0$, it holds that

$$\mathcal{S}(\alpha^*, \alpha) \lesssim \mathcal{M}(\alpha) - \mathcal{M}(\alpha^*).$$

- (f) For any $\lambda = \frac{1}{K}$ with K a positive integer, there exists at least a solution $\alpha^\lambda \in \mathcal{A}$, such that α^λ interpolates between the optimal NF T_k^λ that minimizes $\mathcal{L}_k(T)$, i.e.: $T_k^\lambda(x) = x + \lambda \alpha_{t_k}^\lambda(x)$.

Finally, a control $\alpha \in \mathcal{A}$ is said to induce bounded weight if the following assumption hold:

- (g) The functions $g_t^{\alpha^*}$ are bounded from above by a constant C at all time $t \in [0, 1]$.

E.3. Continuous-time importance sampling

Proposition 22. *Under Assumptions (a) to (c) and if the control α satisfies Assumption (g), then the marginal at time t of $\bar{\Pi}_t^\alpha$ defined in (55) is equal to Π_t . Moreover, we have $\bar{\Lambda}_t^\alpha[w_t^\alpha]$ of the weights w_t^α is given by $\exp\left(-\int_0^t \Pi_s[\partial_s V_s] ds\right)$.*

Proof. We know that α satisfies (52) of Lemma 7 thanks to Assumption (b). This, in addition to Assumptions (a) and (c), ensures the SDE defined in (54) is well defined and admits finite 4th order moments by Lemma 13. Denote by ρ_t the marginal of $\bar{\Pi}_t$ at time t . Let f be a bounded smooth function f of X_t at time t . Then, by definition of ρ_t , we have:

$$\rho_t[f] = \frac{\bar{\Lambda}_t^\alpha[w_t^\alpha f]}{\bar{\Lambda}_t^\alpha[w_t^\alpha]} = \frac{\mathbb{E}[w_t^\alpha f(X_t)]}{\mathbb{E}[w_t^\alpha]}.$$

Such quantity is finite since the importance weights w_t^α are bounded by Assumption (g). It is sufficient to show that both ρ_t and Π_t satisfy the same partial differential equation with the same initial condition. Let us first express the time derivative of the process $w_s^\alpha f(X_s)$, which is obtained using Itô's formula:

$$\begin{aligned} \frac{d}{ds}[w_s^\alpha f(X_s)] &= \frac{d}{ds} \left[\exp\left(\int_0^s g_u^\alpha(X_u) du\right) f(X_s) \right] \\ &= w_s^\alpha (g_s^\alpha(X_s) f(X_s) + \nabla f(X_s)^\top (\alpha_s(X_s) - \nabla V_s(X_s)) + \Delta f(X_s)) \\ &\quad + \sqrt{2} w_s^\alpha \nabla f(X_s)^\top dB_s. \end{aligned}$$

Integrating the above expression in time on the interval $[t, t+h]$ for some $h > 0$ and taking the expectation w.r.t. the process X_t yields:

$$\begin{aligned} \mathbb{E}[w_{t+h}^\alpha f(X_{t+h})] - \mathbb{E}[w_t^\alpha f(X_t)] &= \int_t^{t+h} w_s^\alpha (\nabla f(X_s)^\top (\alpha_s(X_s) - \nabla V_s(X_s)) + \Delta f(X_s)) ds \\ &\quad + \int_t^{t+h} w_s^\alpha g_s^\alpha(X_s) f(X_s) ds, \end{aligned}$$

where we used $\mathbb{E}[\int_t^{t+h} \sqrt{2} w_s^\alpha \nabla f(X_s)^\top dB_s] = 0$. A similar expression can be obtained for $h < 0$. By continuity of the integrands of the r.h.s., we can divide by h and take the limit $h \rightarrow 0$ which yields:

$$\frac{d}{dt} \mathbb{E}[w_t^\alpha f(X_t)] = \mathbb{E}[w_t^\alpha (g_t^\alpha(X_t) f(X_t) + \nabla f(X_t)^\top (\alpha_t(X_t) - \nabla V_t(X_t)) + \Delta f(X_t))]. \quad (64)$$

Now we can compute the time derivative of $\rho_t[f]$:

$$\begin{aligned} \partial_t \rho_t[f] &= \frac{1}{\mathbb{E}[w_t^\alpha]} \frac{d}{dt} \mathbb{E}[w_t^\alpha f(X_t)] - \frac{\mathbb{E}[w_t^\alpha f(X_t)]}{\mathbb{E}[w_t^\alpha]^2} \frac{d}{dt} \mathbb{E}[w_t^\alpha] \\ &= \rho_t [g_t^\alpha f + \nabla f^\top (\alpha_t - \nabla V_t) + \Delta f] - \rho_t[f] \rho_t[g_t^\alpha], \end{aligned}$$

where we used (64) and the definition of ρ_t to obtain the final expression. Hence, we have the following partial differential equation:

$$\partial_t \rho_t[f] = -\rho_t [\nabla f^\top (\nabla V_t - \alpha_t)] + \rho_t [\Delta f] + \rho_t [(g_t^\alpha - \rho_t[g_t^\alpha]) f], \quad \rho_0[f] = \Pi_0[f].$$

This implies the following partial differential equation on ρ_t using integration by parts:

$$\partial_t \rho_t = \nabla \cdot (\rho_t (\nabla V_t - \alpha_t)) + \Delta \rho_t + (g_t^\alpha - \rho_t[g_t^\alpha]) \rho_t, \quad \rho_0 = \Pi_0.$$

It is easy to check that Π_t satisfies the same partial differential equation. One then concludes by uniqueness of the solution.

To get the expression of the normalizing constant $\mathcal{Z}_t = \bar{\Lambda}_t^\alpha[w_t^\alpha]$, we take the time derivative of \mathcal{Z}_t and using Itô's formula, we get:

$$\frac{d\mathcal{Z}_t}{dt} = \bar{\Lambda}_t^\alpha [g_t^\alpha w_t^\alpha] = \mathcal{Z}_t \bar{\Pi}_t^\alpha [g_t^\alpha].$$

Since g_t^α depends only on the process at time t , and since the marginal of $\bar{\Pi}_t$ at time t is equal to Π_t , we obtain:

$$\frac{dZ_t}{dt} = Z_t \Pi_t [g_t^\alpha],$$

where we have used integration by parts to conclude that $\Pi_t [\nabla \cdot \alpha_t - \nabla_x V_t^\top \alpha_t] = 0$. Using the fact that $Z_0 = 1$ and solving the above differential equation, we get $Z_t = \exp\left(-\int_0^t g_s^\alpha ds\right) = \exp\left(-\int_0^t \Pi_s [\partial_s V_s] ds\right)$. \square

Lemma 8. *Under Assumptions (a) and (c) and if the control α satisfies Assumption (g), then there exists a positive constant C such that for all $t \in [0, 1]$ and $x \in \mathcal{X}$:*

$$\bar{g}_t^\alpha(x) \leq C.$$

In particular, this implies that the normalized weights \bar{w}_t^α are uniformly bounded over $t \in [0, 1]$.

Proof. We have by definition $\bar{g}_t^\alpha = g_t^\alpha + \Pi_t [\partial_t V_t]$. Since α satisfies Assumption (g), this implies that g_t^α is bounded from above by some constant number C . We only need to control $\Pi_t [\partial_t V_t]$ in time. By Assumption (a), we have the estimate:

$$|\partial_t V_t(x)| \leq |\partial_t V_0(0)| + L(1 + 2\|x\|)^2. \quad (65)$$

Moreover, Assumption (c) states that the 4-th, order moments of Π_t (hence the lower order moments) are finite and uniformly bounded for $t \in [0, 1]$. Therefore, using (65) implies $|\Pi_t [\partial_t V_t(x)]|$ is also finite and uniformly bounded on $[0, 1]$. This concludes the proof. \square

The next proposition shows that the interpolating measures $\bar{\Pi}_t^{\alpha, \lambda}$ and $\bar{\Lambda}_t^{\alpha, \lambda}$ admit $\bar{\eta}_k$ and $\bar{\pi}_k$ as marginals. It is a direct consequence of the definition of (58) and importance weights (60) and is thus provided without a proof.

Proposition 23. *Let α be a continuous-time control in \mathcal{A} . Let $\bar{\Lambda}_t^{\alpha, \lambda}$ and $\bar{\Pi}_t^{\alpha, \lambda}$ defined by (58) and (61). Under Assumptions (a) to (c) and provided α satisfies Assumption (g), then the joint distribution of the vector $(X_{t_0}, \dots, X_{t_k})$ is equal to $\bar{\eta}_k$, that is the joint distribution of the Markov chain defined by $X_{0:k}$, with NFs $T_k^{\alpha, \lambda}$ given by:*

$$T_k^{\alpha, \lambda}(x) = x + \lambda \alpha_{t_k}(x). \quad (66)$$

Moreover, consider the joint distribution $\bar{\pi}_k$ defined by

$$\bar{\pi}_k[f] = \frac{\bar{\eta}_k[w_k f]}{\bar{\eta}_k[w_k]},$$

where the IS weights w_k are given by (11) using the same choice of NFs T_k as in (66). Then, for any bounded smooth function f of $(x_{t_0}, \dots, x_{t_k})$ it holds that:

$$\bar{\Pi}_{t_k}^{\alpha, \lambda}[f] = \bar{\pi}_k[f].$$

E.4. Relative entropy computation

In this section, we are interested in computing the relative entropy between the limit distribution $\bar{\Pi}_t^\alpha$ with control α and the interpolating measure $\bar{\Pi}_t^{\alpha', \lambda}$ using the control α' . For clarity, we introduce the following notation:

$$\begin{aligned} \mathcal{D}_t^\lambda(\alpha, \alpha') &:= \text{KL}(\bar{\Pi}_t^\alpha \parallel \bar{\Pi}_t^{\alpha', \lambda}), \\ \Delta_t^\lambda(\alpha, \alpha') &:= \mathcal{D}_t^\lambda(\alpha, \alpha) - \mathcal{D}_t^\lambda(\alpha, \alpha'). \end{aligned}$$

$\mathcal{D}_t^\lambda(\alpha, \alpha')$ represents the relative entropy between the limit distribution $\bar{\Pi}_t^\alpha$ with control α and the interpolating measure $\bar{\Pi}_t^{\alpha', \lambda}$ using the control α' . The error term $\Delta_t^\lambda(\alpha, \alpha')$ represents how much additional error is introduced by using a different control α' for the interpolating measure instead of the reference control α . In Proposition 24, we provide an expression for $\mathcal{D}_t^\lambda(\alpha, \alpha')$, then in Proposition 26 we control the difference in relative entropies $\Delta_t^\lambda(\alpha, \alpha')$ when the control α' of the interpolating measure is replaced by α .

Proposition 24. Let α and α' be two continuous time controls in \mathcal{A} . Under Assumptions (a) to (c) and provided α satisfies Assumption (g), then the relative entropy between $\bar{\Lambda}^\alpha$ and $\bar{\Lambda}^{\alpha',\lambda}$ is well defined and given by:

$$\text{KL}(\bar{\Lambda}^\alpha \|\bar{\Lambda}^{\alpha',\lambda}) = \frac{1}{4} \mathbb{E} \left[\int_0^t \|b_s^\alpha(X_s) - b_s^{\alpha',\lambda}(X)\|^2 ds \right].$$

Moreover, the relative entropy between $\bar{\Pi}_t^\alpha$ and $\bar{\Pi}_t^{\alpha',\lambda}$ is well defined and given by:

$$\begin{aligned} \mathcal{D}_t^\lambda(\alpha, \alpha') = & \mathbb{E} \left[\bar{w}_t^\alpha \left(\int_0^{\tau_\lambda(t)} (g_s^\alpha(X_s) - g_s^{\alpha',\lambda}(X)) ds \right) \right] + \mathbb{E} \left[\bar{w}_t^\alpha \int_{\tau_\lambda(t)}^t \bar{g}_s^\alpha(X_s) ds \right] \\ & + \frac{1}{4} \mathbb{E} \left[\bar{w}_t^\alpha \int_0^t \|b_s^\alpha(X_s) - b_s^{\alpha',\lambda}(X)\|^2 ds \right] + \frac{1}{\sqrt{2}} \mathbb{E} \left[\bar{w}_t^\alpha \int_0^t (b_s^\alpha(X_s) - b_s^{\alpha',\lambda}(X))^\top dB_s \right] \\ & + \mathbb{E} \left[\bar{w}_t^\alpha \left(\sum_{k=1}^{k_\lambda(t)} h_{t_k}^\lambda(X_{t_k}) \right) \right]. \end{aligned}$$

Proof. Let us express the relative entropy between $\bar{\Pi}_t$ and $\bar{\Pi}_t^\lambda$ defined as the expectation under $\bar{\Pi}_t$ of the logarithm of the Radon-Nykodim of $\bar{\Pi}_t$ w.r.t. $\bar{\Pi}_t^\lambda$. By a simple chaining argument, we have that:

$$\log \left(\frac{d\bar{\Pi}_t^\alpha}{d\bar{\Pi}_t^{\alpha',\lambda}} \right) = \log \left(\frac{d\bar{\Pi}_t^\alpha}{d\bar{\Lambda}_t^\alpha} \right) + \log \left(\frac{d\bar{\Lambda}_t^\alpha}{d\bar{\Lambda}_t^{\alpha',\lambda}} \right) + \log \left(\frac{d\bar{\Lambda}_t^{\alpha',\lambda}}{d\bar{\Pi}_t^{\alpha',\lambda}} \right). \quad (67)$$

Hence, $\mathcal{D}_t^\lambda(\alpha, \alpha')$ is obtained by taking the expectation of (67) w.r.t $\bar{\Pi}_t^\alpha$, that is the expectation w.r.t $\bar{\Lambda}_t^\alpha$ weighted by \bar{w}_t^α :

$$\mathcal{D}_t^\lambda(\alpha, \alpha') = \mathbb{E} \left[\bar{w}_t^\alpha \log \left(\frac{d\bar{\Pi}_t^\alpha}{d\bar{\Lambda}_t^\alpha} \right) \right] + \mathbb{E} \left[\bar{w}_t^\alpha \log \left(\frac{d\bar{\Lambda}_t^\alpha}{d\bar{\Lambda}_t^{\alpha',\lambda}} \right) \right] + \mathbb{E} \left[\bar{w}_t^\alpha \log \left(\frac{d\bar{\Lambda}_t^{\alpha',\lambda}}{d\bar{\Pi}_t^{\alpha',\lambda}} \right) \right]. \quad (68)$$

We only need to express each term in (67). The first and last terms in (67) are given by definition of $\bar{\Pi}_t^\alpha$ and $\bar{\Pi}_t^{\alpha',\lambda}$:

$$\begin{aligned} \log \left(\frac{d\bar{\Pi}_t^\alpha}{d\bar{\Lambda}_t^\alpha} \right) &= \int_0^t g_s^\alpha(X_s) ds - \log \left(\frac{\mathcal{Z}_0}{\mathcal{Z}_t} \right), \\ \log \left(\frac{d\bar{\Lambda}_t^{\alpha',\lambda}}{d\bar{\Pi}_t^{\alpha',\lambda}} \right) &= \log \left(\frac{\mathcal{Z}_0}{\mathcal{Z}_{\tau_\lambda(t)}} \right) - \int_0^{\tau_\lambda(t)} g_s^{\alpha',\lambda}(X) ds + \sum_{k=1}^{k_\lambda(t)} h_{t_k}^\lambda(X_{t_k}). \end{aligned} \quad (69)$$

The second term is obtained by application of Girsanov's formula, since $\bar{\Lambda}_t^\alpha$ and $\bar{\Lambda}_t^{\alpha',\lambda}$ are mutually absolutely continuous as they share the same volatility term by construction:

$$\log \left(\frac{d\bar{\Lambda}_t^\alpha}{d\bar{\Lambda}_t^{\alpha',\lambda}} \right) := \frac{1}{\sqrt{2}} \int_0^t (b_s^\alpha(X_s) - b_s^{\alpha',\lambda}(X))^\top dB_s + \frac{1}{4} \int_0^t \|b_s^\alpha(X_s) - b_s^{\alpha',\lambda}(X)\|^2 ds. \quad (70)$$

We obtain the desired expression for $\mathcal{D}_t^\lambda(\alpha, \alpha')$ by plugging (69) and (70) in (68). Finally, the relative entropy between $\bar{\Lambda}_t^\alpha$ and $\bar{\Lambda}_t^{\alpha',\lambda}$ is obtained directly by taking the expectation of (70) under $\bar{\Lambda}_t^\alpha$. \square

E.5. Relative entropy bounds

In this section, we provide bounds on the relative entropy $\mathcal{D}_t^\lambda(\alpha, \alpha)$ in Proposition 27 and difference in relative entropies $\Delta_t^\lambda(\alpha, \alpha')$ in Proposition 26. We start by Proposition 25 which will be crucial in the proofs of Propositions 26 and 27. Proposition 25 provides estimates of the expectations under $\bar{\Lambda}_t^\alpha$ of product of functions, where one of the functions depend only on the value of the process at an earlier time s . We defer the proof of Proposition 25 to Appendix E.8.1 which crucially relies on a coupling argument later provided in Lemma 14.

Proposition 25. Let $0 \leq s \leq s' \leq t$. Consider a function f of the process X_s at time s such that $\mathbb{E}[\|f(X_s)\|^2] < \infty$ and let $h_u(X_u, X_s)$ be a function of the sample path $(X_t)_{[s, s']}$ satisfying $\mathbb{E}[\int_s^{s'} \|h_u(X_u, X_s)\|^2 du] < \infty$. Under Assumptions (a) to (c) and provided α satisfies Assumption (g), the following holds:

$$\|\mathbb{E}[\bar{w}_t^\alpha f(X_s)]\| \lesssim \Pi_s[\|f(X_s)\|] \quad (71)$$

$$\left| \mathbb{E} \left[\bar{w}_t^\alpha f(X_s)^\top \int_s^{s'} dB_u \right] \right| \lesssim (s' - s) \Pi_s[\|f\|^2]^{\frac{1}{2}} \quad (72)$$

$$\left| \mathbb{E} \left[\bar{w}_t^\alpha f(X_s)^\top \int_s^{s'} h_u(X_u, X_s) du \right] \right| \lesssim \sqrt{s' - s} \Pi_s[\|f\|^2]^{\frac{1}{2}} \mathbb{E} \left[\int_s^{s'} \|h_u(X_u, X_s)\|^2 du \right]^{\frac{1}{2}} \quad (73)$$

In particular, if $t_{k-1} \leq s \leq s' \leq t_k$ for some $k \leq k_\lambda(t) + 1$ and $h_u(x, y) = b_u^\alpha(x) - \beta_{u'}^{\alpha, \lambda}(y)$ with $t_{k-1} \leq u' \leq t_k$, then

$$\left| \mathbb{E} \left[\bar{w}_t^\alpha f(X_s)^\top \int_s^{s'} h(X_u, X_s) du + dB_u \right] \right| \lesssim \lambda \Pi_s[\|f\|^2]^{\frac{1}{2}}.$$

We can now provide an upper-bound on the difference in relative entropies $\Delta_t^\lambda(\alpha, \alpha')$ in terms of the time discretization step-size λ and the time-discrete discrepancy $\mathcal{S}_t^\lambda(\alpha, \alpha')$ between the controls α and α' , defined in (62).

Proposition 26. Let α and α' be two continuous time controls in \mathcal{A} . Under Assumptions (a) to (c) and provided α satisfies Assumption (g), then the following upper-bound holds:

$$|\Delta_t^\lambda(\alpha, \alpha')| \lesssim \mathcal{S}_t^\lambda(\alpha, \alpha') + \mathcal{S}_t^\lambda(\alpha, \alpha')^{\frac{1}{2}} + \lambda,$$

where \mathcal{S}_t^λ is defined in (62).

Proof. By direct computation using the expression of the relative entropy in Proposition 24, we have that:

$$\begin{aligned} -\Delta_t^\lambda(\alpha, \alpha') &= \mathcal{D}_t^\lambda(\alpha, \alpha') - \mathcal{D}_t^\lambda(\alpha, \alpha) \\ &= \mathbb{E} \left[\bar{w}_t^\alpha \left(\int_0^{\tau_\lambda(t)} (g_s^{\alpha, \lambda}(X) - g_s^{\alpha', \lambda}(X)) ds \right) \right] \\ &\quad + \frac{1}{4} \mathbb{E} \left[\bar{w}_t^\alpha \int_0^t \left(\|b_s^\alpha(X_s) - b_s^{\alpha', \lambda}(X)\|^2 - \|b_s^\alpha(X_s) - b_s^{\alpha, \lambda}(X)\|^2 \right) ds \right] \end{aligned}$$

We further introduce:

$$\Delta \delta_s(x) := \delta_s^{\alpha, \lambda}(x) - \delta_s^{\alpha', \lambda}(x), \quad \Delta \beta_s(x) := \beta_s^{\alpha, \lambda}(x) - \beta_s^{\alpha', \lambda}(x),$$

$$h_{u,s}(x, y) = b_u^\alpha(x) - \beta_{u'}^{\alpha, \lambda}(y).$$

Recalling the definition of $g_s^{\alpha, \lambda}(X)$ from (59), we can express $\Delta_t^\lambda(\alpha, \alpha')$ in terms of $\Delta \delta$, $\Delta \beta$ and $h_{u,s}(x, y)$:

$$\begin{aligned} -\Delta_t^\lambda(\alpha, \alpha') &= \lambda \sum_{k=1}^{k_\lambda(t)} \mathbb{E} \left[\bar{w}_t^\alpha \left(\Delta \delta_{t_k}(X_{t_{k-1}}) + \frac{1}{4} \|\Delta \beta_{t_k}(X_{t_{k-1}})\|^2 \right) \right] \\ &\quad + \frac{1}{2} \sum_{k=1}^{k_\lambda(t)} \mathbb{E} \left[\bar{w}_t^\alpha (\Delta \beta_{t_k}(X_{t_{k-1}}))^\top \int_{t_{k-1}}^{t_k} (h_{u, t_k}(X_u, X_{t_{k-1}}) du + \sqrt{2} dB_u) \right] \\ &\quad + \frac{1}{2} \mathbb{E} \left[\bar{w}_t^\alpha (\Delta \beta_{\tau_\lambda(t)+\lambda}(X_{\tau_\lambda(t)}))^\top \int_{\tau_\lambda(t)}^t (h_{u, \tau_\lambda(t)+\lambda}(X_u, X_{\tau_\lambda(t)}) du + \sqrt{2} dB_u) \right]. \end{aligned} \quad (74)$$

By Proposition 25, we know that the remainder term in the last line of (74) is of order λ . Moreover, Proposition 25 allows us to control the first two terms in (74) so that the following bound holds:

$$\begin{aligned}
 |\Delta_t^\lambda(\alpha, \alpha')| &\lesssim \lambda \sum_{k=1}^{k_\lambda(t)} \Pi_{t_{k-1}} \left[|\Delta\delta_{t_k}| + \|\Delta\beta_{t_k}\|^2 \right] + \lambda \sum_{k=1}^{k_\lambda(t)} \Pi_{t_{k-1}} \left[\|\Delta\beta_{t_k}\|^2 \right]^{\frac{1}{2}} + \lambda \\
 &\lesssim \lambda \sum_{k=1}^{k_\lambda(t)} \Pi_{t_{k-1}} \left[|\Delta\delta_{t_k}| + \|\Delta\beta_{t_k}\|^2 \right] + \left(\lambda \sum_{k=1}^{k_\lambda(t)} \Pi_{t_{k-1}} \left[\|\Delta\beta_{t_k}\|^2 \right] \right)^{\frac{1}{2}} + \lambda \\
 &\lesssim \mathbb{M}_t^\lambda[|\Delta\delta|] + \mathbb{M}_t^\lambda[\|\Delta\beta\|^2] + \mathbb{M}_t^\lambda[\|\Delta\beta\|^2]^{\frac{1}{2}} + \lambda,
 \end{aligned} \tag{75}$$

where we used Cauchy–Schwarz inequality in the second line and introduced the notation \mathbb{M}_t^λ from (63) in the last line, $|\Delta\delta|$ and $\|\Delta\beta\|$ being viewed as functions from $[0, 1] \times \mathcal{X}$ to \mathbb{R} . Thus, we only need to control $\mathbb{M}_t^\lambda[|\Delta\delta|]$ and $\mathbb{M}_t^\lambda[\|\Delta\beta\|^2]$. To control $\mathbb{M}_t^\lambda[|\Delta\delta|]$, we can rely on the following pointwise estimate from Lemma 17:

$$\left| \delta_{t_k}^{\alpha, \lambda}(x) - \delta_{t_k}^{\alpha', \lambda}(x) \right| \lesssim ((1 + \|x\|) \|\alpha_{t_k}(x) - \alpha'_{t_k}(x)\| + \|\nabla\alpha_{t_k}(x) - \nabla\alpha'_{t_k}(x)\|).$$

Further defining $\bar{L} : [0, 1] \times \mathbb{X} \rightarrow \mathbb{R}$ to be $(s, x) \mapsto \bar{L}_s(x) = 1 + \|x\|$, this allows us to write

$$\begin{aligned}
 \mathbb{M}_t^\lambda[\|\Delta\beta\|^2] &\lesssim \mathbb{M}_t^\lambda[\|\alpha - \alpha'\|^2] \\
 \mathbb{M}_t^\lambda[|\Delta\delta|] &\lesssim \mathbb{M}_t^\lambda[\bar{L}\|\alpha - \alpha'\|] + \mathbb{M}_t^\lambda[\|\nabla\alpha - \nabla\alpha'\|] \\
 &\lesssim \mathbb{M}_t^\lambda[\bar{L}^2]^{\frac{1}{2}} \mathbb{M}_t^\lambda[\|\alpha - \alpha'\|^2]^{\frac{1}{2}} + \mathbb{M}_t^\lambda[\|\nabla\alpha - \nabla\alpha'\|]
 \end{aligned}$$

where we used Cauchy–Schwarz inequality to get the last inequality. The factor $\mathbb{M}_t^\lambda[\bar{L}^2]$ is also bounded as it is a Riemann sum and converges towards:

$$\mathbb{M}_t^\lambda[\bar{L}^2] \xrightarrow{\lambda \rightarrow 0} \int_0^t \Pi_t \left[(1 + \|X\|)^2 \right] dt,$$

which is finite by Assumption (c). Again using the pointwise estimate from Lemma 17:

$$\|\beta_t^{\alpha, \lambda}(x) - \beta_t^{\alpha', \lambda}(x)\| \lesssim \|\alpha_t(x) - \alpha'_t(x)\|,$$

it follows directly that $\mathbb{M}_t^\lambda[\|\Delta\beta\|^2] \lesssim \mathbb{M}_t^\lambda[\|\alpha - \alpha'\|^2]$. Therefore, we have shown:

$$\begin{aligned}
 \mathbb{M}_t^\lambda[\|\Delta\beta\|^2] &\lesssim \mathbb{M}_t^\lambda[\|\alpha - \alpha'\|^2] \\
 \mathbb{M}_t^\lambda[|\Delta\delta|] &\lesssim \mathbb{M}_t^\lambda[\|\alpha - \alpha'\|^2]^{\frac{1}{2}} + \mathbb{M}_t^\lambda[\|\nabla\alpha - \nabla\alpha'\|].
 \end{aligned} \tag{76}$$

The desired upper-bound follows using (76) in (75). \square

Next we control the relative entropy $\mathcal{D}_t^\lambda(\alpha, \alpha)$ between the Feynman-Kac measure $\bar{\Pi}_t^\alpha$ and the interpolating measure $\bar{\Pi}_t^{\alpha, \lambda}$ using the same control α .

Proposition 27. *Under Assumptions (a) to (d) and if α satisfies Assumption (g), the following upper bound holds:*

$$\mathcal{D}_t^\lambda(\alpha, \alpha) := \text{KL}(\bar{\Pi}_t^\alpha \|\bar{\Pi}_t^{\alpha, \lambda}) \lesssim \sqrt{\lambda}.$$

Proof. **of Proposition 27.** We use the expression of the relative entropy $\mathcal{D}_t^\lambda(\alpha, \alpha)$ provided in Proposition 24, then recalling that the normalized weights \bar{w}_t^α are bounded by Lemma 8 we have

$$\begin{aligned} \mathcal{D}_t^\lambda(\alpha, \alpha') &\lesssim \mathbb{E} \left[\left| \int_0^{\tau_\lambda(t)} (g_s^\alpha(X_s) - g_s^{\alpha', \lambda}(X)) \, ds \right| \right] + \mathbb{E} \left[\left| \int_{\tau_\lambda(t)}^t \bar{g}_s^\alpha(X_s) \, ds \right| \right] \\ &\quad + \frac{1}{4} \mathbb{E} \left[\int_0^t \|b_s^\alpha(X_s) - b_s^{\alpha', \lambda}(X)\|^2 \, ds \right] + \frac{1}{\sqrt{2}} \mathbb{E} \left[\left| \int_0^t (b_s^\alpha(X_s) - b_s^{\alpha', \lambda}(X))^\top \, dB_s \right| \right] \\ &\quad + \mathbb{E} \left[\bar{w}_t^\alpha \left(\sum_{k=1}^{k_\lambda(t)} |h_{t_k}^\lambda(X_{t_k})| \right) \right]. \end{aligned}$$

Now by a direct application of Cauchy–Schwarz inequality and triangular inequalities, we get the following upper-bound on $\mathcal{D}_t^\lambda(\alpha, \alpha)$

$$D_t^\lambda(\alpha, \alpha) \lesssim \left(\mathcal{E}_W + \mathcal{E}_R + \mathcal{E}_H + \frac{1}{4} \mathcal{E}_{KL} + \frac{1}{\sqrt{2}} \mathcal{E}_{KL}^{\frac{1}{2}} \right)$$

where \mathcal{E}_W , \mathcal{E}_R , \mathcal{E}_H and \mathcal{E}_{KL} are given by

$$\begin{aligned} \mathcal{E}_W &:= \mathbb{E} \left[\left| \int_0^{\tau_\lambda(t)} g_s^\alpha(X_s) - g_s^{\alpha, \lambda}(X) \, ds \right| \right], & \mathcal{E}_R &:= \mathbb{E} \left[\int_{\tau_\lambda(t)}^t \bar{g}_s^\alpha(X_s) \, ds \right], \\ \mathcal{E}_{KL} &:= \mathbb{E} \left[\int_0^t \|b_s^\alpha(X_s) - b_s^{\alpha, \lambda}(X)\|^2 \, ds \right], & \mathcal{E}_H &:= \mathbb{E} \left[w_t^\alpha \sum_{k=1}^{k_\lambda(t)} |h_{t_k}^\lambda| \right]. \end{aligned}$$

Bound on \mathcal{E}_R . $\mathbb{E}[|\bar{g}_s^\alpha(X_s)|]$ has a quadratic growth by Lemma 17 and the process X_s has a bounded second moment by Lemma 13, therefore $\mathcal{E}_R \lesssim |t - \tau_\lambda(t)| \leq \lambda$.

Bound on \mathcal{E}_W . We first start by applying a triangular inequality:

$$\mathcal{E}_W \leq \sum_{k=1}^{k_\lambda(t)} \mathbb{E} \left[\underbrace{\int_{t_{k-1}}^{t_k} |g_s^\alpha(X_s) - \delta_s^{\alpha, \lambda}(X_{t_{k-1}})| \, ds}_{\mathcal{E}_{W,k}} \right]$$

where by definition of $g_s^{\alpha, \lambda}(X) = \delta_s^{\alpha, \lambda}(X_{t_{k-1}})$ for $t_{k-1} \leq s \leq t_k$. We then use the point-wise upper-bound on the difference $g_s^\alpha(x') - \delta_t^{\alpha, \lambda}(x)$ provided in Lemma 15:

$$\left| g_s^\alpha(x') - \delta_t^{\alpha, \lambda}(x) \right| \lesssim (1 + \|x\| + \|x'\|) \|x - x'\| + \lambda(1 + \|x\| + \|x'\|)^2.$$

This allows to control each integral $\mathcal{E}_{W,k}$ after a simple application of Cauchy–Schwarz inequality:

$$\begin{aligned} \mathcal{E}_{W,k} &\lesssim \lambda \int_{t_{k-1}}^{t_k} \mathbb{E} \left[(1 + \|X_s\| + \|X_{t_{k-1}}\|)^2 \right] \, ds \\ &\quad + \left(\int_{t_{k-1}}^{t_k} \mathbb{E} \left[\|X_s - X_{t_{k-1}}\|^2 \right] \, ds \right)^{\frac{1}{2}} \left(\int_{t_{k-1}}^{t_k} \mathbb{E} \left[(1 + \|X_s\| + \|X_{t_{k-1}}\|)^2 \right] \, ds \right)^{\frac{1}{2}} \end{aligned}$$

where we also used Fubini’s theorem to exchange the order of the expectation and time integral. By Lemma 13, we know that the second moments of the process X_t are bounded over the time interval $[0, 1]$ and that $\mathbb{E}[\|X_s - X_{t_{k-1}}\|^2] \lesssim |s - t_{k-1}|$. Therefore, we get the upper-bound $\mathbb{E}[\mathcal{E}_{W,k}] \lesssim \lambda^2 + \lambda\sqrt{\lambda}$. Finally, summing over k ranging from 1 to $k_\lambda(t)$ yields:

$$\mathcal{E}_W \leq \sum_{k=1}^{k_\lambda(t)} \mathbb{E}[\mathcal{E}_{W,k}] \lesssim \sqrt{\lambda}.$$

Bound on \mathcal{E}_H . By a direct application of Proposition 25, we know that:

$$\mathcal{E}_H \lesssim \sum_{k=1}^{k_\lambda(t)} \Pi_{t_k} [|h_{t_k}^\lambda|].$$

We just need to control each term $\Pi_{t_k} [|h_{t_k}^\lambda|]$ as a function of λ . The technical Lemma 18 provides an upper-bound on $|h_t^\lambda(x)|$ that is point-wise in x and of the form:

$$|h_t^\lambda(x)| \leq \lambda \sqrt{\lambda} \frac{Q_1(x)}{1 + Q_2(x)} \exp(C\lambda \|\nabla V_t(x)\|^2)$$

where Q_1 and Q_2 are non-negative functions of $\|x\|$ of polynomial growth and independent from λ and t and C is a non-negative constant. Taking the expectation under Π_t and using the integrability condition of Assumption (d), it follows that $\Pi_{t_k} [|h_{t_k}^\lambda|] \lesssim \lambda \sqrt{\lambda}$. Thus summing over k directly yields $\mathcal{E}_H \lesssim \sqrt{\lambda}$.

Bound on \mathcal{E}_{KL} . We start by decomposing \mathcal{E}_{KL} as a sum of $k_\lambda(t) + 1$ integral terms over intervals of size less or equal to λ :

$$\mathcal{E}_{KL} = \int_{\tau_\lambda(t)}^t \mathbb{E} \left[\|b_s^\alpha(X_s) - \beta_{\tau_\lambda(t)+1}^{\alpha,\lambda}(X_{\tau_\lambda(t)})\|^2 \right] ds + \sum_{k=1}^{k_\lambda(t)} \int_{t_{k-1}}^{t_k} \mathbb{E} \left[\|b_s^\alpha(X_s) - \beta_{t_k}^{\alpha,\lambda}(X_{t_{k-1}})\|^2 \right] ds$$

where we used that $b_s^{\alpha,\lambda}(X) = \beta_{t_k}^{\alpha,\lambda}(X_{t_{k-1}})$ for $t_{k-1} \leq s \leq t_k$ by definition of $b_s^{\alpha,\lambda}(X)$. We can then use Lemma 9 which allows to control each term by λ^2 . Hence, after summing, we directly get that $\mathcal{E}_{KL} \lesssim \lambda |t - \tau_\lambda(t)| + \sum_{k=1}^{k_\lambda(t)} \lambda^2 \lesssim \lambda$.

We finally get the desired result by combining upper-bounds on each quantity $\mathcal{E}_W, \mathcal{E}_R, \mathcal{E}_H$ and \mathcal{E}_{KL} which are all of order $\sqrt{\lambda}$ at least. \square

Lemma 9. Let s and s' be such that $t_{k-1} \leq s \leq s' \leq t_k$ for $0 \leq k \leq K$ and $t_{k-1} \leq u' \leq t_k$. Under Assumptions (a) to (c) it holds that:

$$\mathbb{E} \left[\int_s^{s'} \|b_u^\alpha(X_u) - \beta_{u'}^{\alpha,\lambda}(X_s)\|^2 du \right] \lesssim \lambda |s' - s|.$$

Proof. By a direct application of Proposition 31, we have

$$\mathbb{E} \left[\int_s^{s'} \|b_u^\alpha(X_u) - \beta_{u'}^{\alpha,\lambda}(X_s)\|^2 du \right] \lesssim \int_s^{s'} \mathbb{E} \left[\|X_u - X_s\|^2 + \lambda^2 (1 + \|X_s\|^2) \right] du.$$

Moreover, by Lemma 13, we directly know that the second moments of the process are bounded at any time and that $\mathbb{E}[\|X_{u'} - X_u\|^2] \leq |u - u'|$, hence:

$$\mathbb{E} \left[\int_s^{s'} \|b_u^\alpha(X_u) - \beta_{u'}^{\alpha,\lambda}(X_s)\|^2 du \right] \lesssim \int_s^{s'} (|u - s| + \lambda^2) du \lesssim \lambda |s' - s|.$$

\square

E.6. Convergence towards an optimal Feynman-Kac measure

In this section, we show that the optimal interpolating controls α^λ converge towards the continuous-time limit optimal control α^* . We will first need proposition Proposition 28 below showing that the discrete-time loss $\mathcal{L}_\lambda^{tot}(\alpha)$ converges towards $\mathcal{M}(\alpha)$ uniformly over the class of admissible controls with a rate of λ . Proposition 28 is a restatement of Proposition 1, we defer its proof to Appendix E.7.

Proposition 28. Under Assumptions (a) to (c) and for $\lambda \leq \frac{1}{2L}$, it holds that:

$$|\mathcal{L}_\lambda^{tot}(\alpha) - \mathcal{M}(\alpha)| \lesssim \lambda,$$

where $\mathcal{M}(\alpha)$ and $\mathcal{L}_\lambda^{tot}(\alpha)$ are defined in (13) and (14).

Using Proposition 28 we can show that the interpolating control α^λ , which exists by Assumption (f), converges towards the continuous-time limit optimal control α^* . This essentially relies on the local behavior of the objective \mathcal{M} near the optimum α^* as described by Assumption (e). Proposition 29 makes this idea more precise.

Proposition 29 (Convergence of the controls). *Under Assumptions (a) to (c), (e) and (f), it holds that:*

$$\mathcal{S}_t^\lambda(\alpha^*, \alpha^\lambda) \lesssim \lambda, \quad \mathcal{S}(\alpha^*, \alpha^\lambda) \lesssim \lambda,$$

where \mathcal{S} and \mathcal{S}_t^λ are defined in (62).

Proof. Recall that α^λ is a minimizer of $\mathcal{L}_\lambda^{\text{tot}}(\alpha)$ defined in (53) over \mathcal{A} , while α^* is a minimizer of $\mathcal{M}(\alpha)$ (defined in (56)) over \mathcal{A} . By Assumptions (c), (e) and (f) both exist and belong to \mathcal{A} . We will first show that $|\mathcal{M}(\alpha^*) - \mathcal{M}(\alpha^\lambda)| \lesssim \lambda$, then, we will use the coercivity property of the loss \mathcal{M} (Assumption (e)) to obtain convergence rates in terms of the divergence \mathcal{S} between α^* and α^λ . Finally, we establish the convergence of the discrete sums \mathcal{S}_t^λ , by interpreting it as a Riemann sum.

Bounds on $|\mathcal{M}(\alpha^*) - \mathcal{M}(\alpha^\lambda)|$ and $|\mathcal{L}_\lambda^{\text{tot}}(\alpha^*) - \mathcal{L}_\lambda^{\text{tot}}(\alpha^\lambda)|$. By definition of the minimizers, the following lower bound holds $0 \leq \mathcal{M}(\alpha^\lambda) - \mathcal{M}(\alpha^*)$ holds. It remains to find an upper-bound. For this purpose, we use the following decomposition:

$$\begin{aligned} \mathcal{M}(\alpha^\lambda) - \mathcal{M}(\alpha^*) &= (\mathcal{M}(\alpha^\lambda) - \mathcal{L}_\lambda^{\text{tot}}(\alpha^\lambda)) + (\mathcal{L}_\lambda^{\text{tot}}(\alpha^\lambda) - \mathcal{L}_\lambda^{\text{tot}}(\alpha^*)) + (\mathcal{L}_\lambda^{\text{tot}}(\alpha^*) - \mathcal{M}(\alpha^*)) \\ &\leq (\mathcal{M}(\alpha^\lambda) - \mathcal{L}_\lambda^{\text{tot}}(\alpha^\lambda)) + (\mathcal{L}_\lambda^{\text{tot}}(\alpha^*) - \mathcal{M}(\alpha^*)), \end{aligned}$$

where we used that $\mathcal{L}_\lambda^{\text{tot}}(\alpha^\lambda) - \mathcal{L}_\lambda^{\text{tot}}(\alpha^*) \leq 0$ to get the second line. Moreover, Proposition 28 provides an error bound of the form $|\mathcal{M}(\alpha) - \mathcal{L}_\lambda^{\text{tot}}(\alpha)| \lesssim \lambda$ that holds uniformly over the set \mathcal{A} . This directly implies $\mathcal{M}(\alpha^\lambda) - \mathcal{M}(\alpha^*) \lesssim \lambda$. Thus, we have shown:

$$|\mathcal{M}(\alpha^\lambda) - \mathcal{M}(\alpha^*)| \lesssim \lambda. \quad (77)$$

Convergence in Sobolev norm $\mathcal{S}(\alpha^*, \alpha^\lambda)$. By Assumption (e), we know that the minimizer α^* is well-separated, i.e. for any $\delta > 0$ it holds that

$$\mathcal{M}(\alpha^*) < \inf_{\mathcal{S}(\alpha, \alpha^*) > \delta} \mathcal{M}(\alpha) \quad (78)$$

Along with (77), this necessarily implies that $\mathcal{S}(\alpha^\lambda, \alpha^*) \xrightarrow{\lambda \rightarrow 0} 0$. Indeed, if by contradiction, $\mathcal{S}(\alpha^\lambda, \alpha^*)$ does not converge to 0, then there exists a positive number δ_0 and a sequence of elements $(\lambda_l)_{l \geq 0}$ such that $\lambda_l \rightarrow 0$ and $\mathcal{S}(\alpha^{\lambda_l}, \alpha^*) > \delta_0$. If we set $M_{\delta_0} = \inf_{\mathcal{S}(\alpha, \alpha^*) > \delta_0} \mathcal{M}(\alpha)$, then (78) implies that $M_{\delta_0} > \mathcal{M}(\alpha^*)$. Hence, $\mathcal{M}(\alpha^{\lambda_l}) > M_{\delta_0} > \mathcal{M}(\alpha^*)$. This contradicts the fact that $\mathcal{M}(\alpha^{\lambda_l}) \xrightarrow{\lambda \rightarrow 0} \mathcal{M}(\alpha^*)$.

Convergence rate in Sobolev norm. Now, we can get a convergence rate for $\mathcal{S}(\alpha^\lambda, \alpha^*)$. By Assumption (e), we know that for δ small enough, the following local coercivity property holds:

$$\mathcal{S}(\alpha, \alpha^*) \lesssim \mathcal{M}(\alpha) - \mathcal{M}(\alpha^*), \quad \forall \alpha \in \mathcal{A}: \quad \mathcal{S}(\alpha, \alpha^*) \leq \delta$$

Since $\mathcal{S}(\alpha^\lambda, \alpha^*) \xrightarrow{\lambda \rightarrow 0} 0$, there exists λ_0 small enough such that $\mathcal{S}(\alpha^\lambda, \alpha^*) \leq \delta$ for all $\lambda \leq \lambda_0$. The convergence rate follows directly using (77): $\mathcal{S}(\alpha^\lambda, \alpha^*) \lesssim \lambda$ for all $\lambda \leq \lambda_0$.

Convergence of discrete sums. We will first show $\mathcal{S}_t^\lambda(\alpha^*, \alpha^\lambda) \lesssim \lambda$ for $t = 1$. The result will follow for any $t \leq 1$ since by definition we have $\mathcal{S}_t^\lambda(\alpha, \alpha') \leq \mathcal{S}_1^\lambda(\alpha, \alpha')$ for any α and α' in \mathcal{A} . By the triangular inequality, we have:

$$\begin{aligned} \mathcal{S}_1^\lambda(\alpha^*, \alpha^\lambda) &\leq \mathcal{S}(\alpha^*, \alpha^\lambda) + |\mathcal{S}_1^\lambda(\alpha^*, \alpha^\lambda) - \mathcal{S}(\alpha^*, \alpha^\lambda)| \\ &\lesssim \lambda + |\mathcal{S}_1^\lambda(\alpha^*, \alpha^\lambda) - \mathcal{S}(\alpha^*, \alpha^\lambda)| \end{aligned}$$

where we used that $\mathcal{S}(\alpha^*, \alpha^\lambda) \lesssim \lambda$ to get the second line. Therefore, it suffices to show that:

$$|\mathcal{S}_1^\lambda(\alpha, \alpha') - \mathcal{S}(\alpha, \alpha')| \lesssim \lambda.$$

uniformly over α and α' in \mathcal{A} . The term $\mathcal{S}_1^\lambda(\alpha, \alpha')$ is simply a Riemann sum converging towards $\mathcal{S}(\alpha, \alpha')$ and whose convergence rate depends on the smoothness in time of the integrand. The finiteness of the moments of Π_t , along with the controls α^* and α^λ and their gradient being locally Lipschitz in time and the variations of the potential \tilde{V}_t having a linear growth allows to get the desired rate. \square

We are now ready to prove our main result of Theorem 3 which we restate as Theorem 6 below for convenience and to keep a consistent notation. The proof simply combines the estimates in Propositions 26 and 27 along with the convergence results of the controls in Proposition 29 and loss (Proposition 28).

Theorem 6. *Under Assumptions (a) to (f) and if α^* satisfies Assumption (g), then it holds that:*

$$\text{KL}\left(\bar{\Pi}_t^{\alpha^*} \parallel \bar{\Pi}_t^{\alpha^\lambda, \lambda}\right) \lesssim \sqrt{\lambda}$$

Moreover, the discrete-time objective $\mathcal{L}_\lambda^{\text{tot}}(\alpha)$ converges towards $\mathcal{M}(\alpha)$ uniformly over the class \mathcal{A} with at rate λ .

Proof. By Assumptions (e) and (f), we know that both α^λ and α^* exist and belong to the class of admissible controls \mathcal{A} . We can then use the following decomposition of the relative entropy $\mathcal{D}_t^\lambda(\alpha^*, \alpha^\lambda)$:

$$|\mathcal{D}_t^\lambda(\alpha^*, \alpha^\lambda)| \leq |\mathcal{D}_t^\lambda(\alpha^*, \alpha^*)| + |\mathcal{D}_t^\lambda(\alpha^*, \alpha^\lambda) - \mathcal{D}_t^\lambda(\alpha^*, \alpha^*)|.$$

Proposition 27 ensures that $|\mathcal{D}_t^\lambda(\alpha^*, \alpha^*)| \lesssim \sqrt{\lambda}$, while Proposition 26 shows that:

$$|\mathcal{D}_t^\lambda(\alpha^*, \alpha^\lambda) - \mathcal{D}_t^\lambda(\alpha^*, \alpha^*)| \lesssim \lambda + \mathcal{S}_t^\lambda(\alpha^*, \alpha^\lambda) + \mathcal{S}_t^\lambda(\alpha^*, \alpha^\lambda)^{\frac{1}{2}}.$$

We can then use Proposition 29 to ensure the discrepancy $\mathcal{S}_t^\lambda(\alpha^*, \alpha^\lambda) \lesssim \lambda$ for λ small enough. This directly yields:

$$|\mathcal{D}_t^\lambda(\alpha^*, \alpha^\lambda)| \lesssim \sqrt{\lambda}.$$

Finally, Proposition 28 shows that the discrete time losses $\mathcal{L}_\lambda^{\text{tot}}$ converge towards \mathcal{M} uniformly over the class \mathcal{A} at rate λ . This concludes the proof. \square

E.7. Uniform convergence of the objective: Proof of Proposition 1 (Proposition 28)

We will first introduce some notations that are used only in this section. In all what follows, T_k are normalizing flows of the form $T_k(x) = x + \lambda\alpha_{t_k}(x)$ for some fixed $\alpha \in \mathcal{A}$. We further denote by $T_{t,s}(x) = x + (s-t)\alpha_s(x)$ and introduce $R_{s,t}$:

$$R_{t,s} := \log\left(\frac{Z_t}{Z_s}\right) + \Pi_t[(V_s \circ T_{t,s} - V_t - \log|\nabla T_{t,s}|)]$$

With the above notations we clearly have $T_k = T_{t_{k-1}, t_k}$ and $R_{t_k, t_{k+1}} = \text{KL}((T_k)_\# \Pi_{t_k} \parallel \Pi_{t_{k+1}})$. The discrete time objective can then be expressed in terms of R :

$$\mathcal{F}_\lambda^{\text{tot}}(\alpha) = \frac{1}{\lambda} \sum_{k=0}^{K-1} R_{t_k, t_{k+1}}. \quad (79)$$

We will show that such sum can be written as a double integral of the form:

$$\mathcal{F}_\lambda^{\text{tot}}(\alpha) = -\frac{1}{\lambda} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \int_{t_k}^s \partial_t \partial_s R_{t,s} dt ds.$$

This allows to view (79) as a Riemann's sum whose limit is determined by the local behavior of the function $\partial_t \partial_s R_{t,s}$ when $|t-s|$ is small. We first show in Proposition 30 that $-\partial_t \partial_s R_{t,s}$ approaches $\Pi_t[(\bar{g}_t^\alpha)^2]$ when t and s are close. The proof of Proposition 30 is deferred to Appendix E.7.1 and relies on Lemmas 10 and 11. Lemma 10 provides closed form expressions for the derivatives $\partial_s R_{t,s}$ and $\partial_t \partial_s R_{t,s}$ as expectations of some integrable functions under Π_t . Then Lemma 11 shows that the integrand in $-\partial_t \partial_s R_{t,s}$ approaches $(\bar{g}_t^\alpha)^2(x)$ with an error that is polynomial in x and proportional to the distance $|t-s|$.

Proposition 30. Under Assumptions (a) to (c) and for t and s such that $|t - s| \leq \frac{1}{2L}$, it holds that:

$$\left| \partial_t \partial_s R_{t,s} + \Pi_t \left[(\bar{g}_t^\alpha)^2 \right] \right| \lesssim |t - s| \quad (80)$$

$$\left| \Pi_t \left[(\bar{g}_t^\alpha)^2 \right] - \Pi_s \left[(\bar{g}_s^\alpha)^2 \right] \right| \lesssim |t - s| \quad (81)$$

Moreover, when $s = t$, we have: $\partial_s R_{s,s} = 0$ and $\partial_t \partial_s R_{t,t} = \Pi_t \left[(\bar{g}_t^\alpha)^2 \right]$.

Now that we have described the behavior of $\partial_t \partial_s R_{t,s}$ when $|t - s|$ is small, we can proceed to the proof of Proposition 28.

Proof. of Proposition 28 By the fundamental theorem of calculus, we can directly write:

$$R_{t_k, t_{k+1}} = \int_{t_k}^{t_{k+1}} \partial_s R_{t_k, s} ds = - \int_{t_k}^{t_{k+1}} \int_{t_k}^s \partial_t \partial_s R_{t, s} dt ds$$

where we used that $R_{t_k, t_k} = 0$ and that $\partial_s R_{s,s} = 0$ by Proposition 30. Moreover, define \mathcal{M}_λ such that:

$$\mathcal{M}_\lambda(\alpha) := \frac{1}{2} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \Pi_{t_k} \left[(\bar{g}_{t_k}^\alpha)^2 \right] dt = \frac{1}{\lambda} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \int_{t_k}^s \Pi_{t_k} \left[(\bar{g}_{t_k}^\alpha)^2 \right] dt ds,$$

where the second expression of $\mathcal{M}_\lambda(\alpha)$ is obtained by direct calculation. We will control both errors $|\mathcal{F}_\lambda^{tot}(\alpha) - \mathcal{M}_\lambda(\alpha)|$ and $|\mathcal{M}_\lambda(\alpha) - \mathcal{M}(\alpha)|$. We get an upper-bound on $|\mathcal{M}_\lambda(\alpha) - \mathcal{M}(\alpha)|$ directly using the first expression of $\mathcal{M}_\lambda(\alpha)$ and using the Lipschitz smoothness of $\Pi_t \left[(\bar{g}_t^\alpha)^2 \right]$ as shown in (81) of Proposition 30:

$$\begin{aligned} |\mathcal{M}_\lambda(\alpha) - \mathcal{M}(\alpha)| &= \left| \frac{1}{2} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \Pi_{t_k} \left[(\bar{g}_{t_k}^\alpha)^2 \right] - \Pi_t \left[(\bar{g}_t^\alpha)^2 \right] dt \right| \\ &\lesssim \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} |t_k - t| dt \lesssim \lambda \end{aligned}$$

To control $|\mathcal{F}_\lambda^{tot}(\alpha) - \mathcal{M}_\lambda(\alpha)|$ we use the second expression of $\mathcal{M}_\lambda(\alpha)$ and rely on the following decomposition:

$$\begin{aligned} |\mathcal{F}_\lambda^{tot}(\alpha) - \mathcal{M}_\lambda(\alpha)| &\leq \frac{1}{\lambda} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \int_{t_k}^s \left| \partial_t \partial_s R_{t, s} + \Pi_{t_k} \left[(\bar{g}_{t_k}^\alpha)^2 \right] \right| dt ds \\ &\leq \frac{1}{\lambda} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \int_{t_k}^s \left(\left| \partial_t \partial_s R_{t, s} + \Pi_t \left[(\bar{g}_t^\alpha)^2 \right] \right| + \left| \Pi_t \left[(\bar{g}_t^\alpha)^2 \right] - \Pi_{t_k} \left[(\bar{g}_{t_k}^\alpha)^2 \right] \right| \right) dt ds \\ &\lesssim \frac{1}{\lambda} \sum_{k=0}^{K-1} \int_{t_k}^{t_{k+1}} \int_{t_k}^s |t - s| + |t - t_k| dt ds \lesssim \lambda \end{aligned}$$

where we used the estimates (80) and (81) in Proposition 30 to get the last line. The result follows by direct application of triangular inequality. \square

E.7.1. PROOF OF THE TECHNICAL LEMMAS

We first provide expressions for $\partial_s R_{t,s}$ and $\partial_t \partial_s R_{t,s}$

Lemma 10. Under Assumptions (a) to (c), we have the following expressions for $\partial_s R_{t,s}$ and $\partial_t \partial_s R_{t,s}$:

$$\begin{aligned} \partial_s R_{t,s} &= \int c_{t,s}(y) q_{t,s}(y) dy, \\ \partial_t \partial_s R_{t,s} &= E_{t,s}^{(1)} + E_{t,s}^{(2)}, \end{aligned}$$

where we define $c_{t,s}$, $q_{t,s}$, $E_{t,s}^{(1)}$ and $E_{t,s}^{(2)}$ as:

$$\begin{aligned} q_{t,s}(y) &:= \Pi_t(T_{t,s}^{-1}(y)) |\nabla T_{t,s}(T_{t,s}^{-1}(y))|^{-1} \\ c_{t,s}(y) &:= \dot{V}_s(y) - \Pi_s[\dot{V}_s] + \nabla V_s(y)^\top \alpha_t(T_{t,s}^{-1}(y)) - \text{Tr}(\nabla T_{t,s}^{-1}(y) \nabla \alpha_t(T_{t,s}^{-1}(y))) \\ E_{t,s}^{(1)} &= \Pi_t \left[\left(\partial_t \alpha_t + (\nabla T_{t,s})^{-1} \nabla \alpha_t \right)^\top \nabla \log \left(\frac{q_{t,s}}{\Pi_s} \right) \circ T_{t,s} \right] \\ E_{t,s}^{(2)} &= \Pi_t [(c_{t,s} \circ T_{t,s})(\partial_t \log(q_{t,s}) \circ T_{t,s})] \end{aligned}$$

In particular, when $s = t$, we have: $\partial_s R_{t,t} = 0$ and $\partial_t \partial_s R_{t,t} = \Pi_t [(\bar{g}_t^\alpha)^2]$.

Proof. • Computing $\partial_s R_{t,s}$.

First, recall that the control α satisfies (52) of Lemma 7 under Assumption (b). In addition and under Assumptions (a) and (c), it is possible to apply the dominated convergence theorem, thus yielding:

$$\partial_s R_{t,s} = -\Pi_s [\dot{V}_s] + \Pi_t [\dot{V}_s \circ T_{t,s} + (\nabla V_s \circ T_{t,s})^\top \alpha_t - \text{Tr}(\nabla T_{t,s}^{-1} \nabla \alpha_t)]$$

We can then perform a change of variables $y = T_{t,s}(x)$ in the above expression to get:

$$\begin{aligned} \partial_s R_{t,s} &= \int \left(\dot{V}_s(y) - \Pi_s[\dot{V}_s] + \nabla V_s(y)^\top \alpha_t(T_{t,s}^{-1}(y)) - \text{Tr}(\nabla T_{t,s}^{-1}(y) \nabla \alpha_t(T_{t,s}^{-1}(y))) \right) q_{t,s}(y) \, dy \\ &= \int c_{t,s}(y) q_{t,s}(y) \, dy. \end{aligned}$$

When $s = t$, we get $c_{t,t}(x) = \bar{g}_t^\alpha(x)$ and $q_{t,t}(x) = \Pi_t(x)$, thus $\partial_s R_{t,t} = 0$.

• Computing $\partial_t \partial_s R_{t,s}$.

To compute $\partial_t \partial_s R_{t,s}$ we will introduce an auxiliary function $r_{t,t',s}$:

$$\begin{aligned} r_{t,t',s} &:= \int \left(\dot{V}_s(y) - \Pi_s[\dot{V}_s] + \nabla V_s(y)^\top \alpha_t(T_{t,s}^{-1}(y)) \right) q_{t',s}(y) \, dy \\ &\quad - \int \left(\text{Tr}(\nabla T_{t,s}^{-1}(y) \nabla \alpha_t(T_{t,s}^{-1}(y))) \right) q_{t',s}(y) \, dy. \\ &= \int c_{t,s}(y) q_{t,s}(y) \, dy. \end{aligned} \tag{82}$$

Hence, from the expression of $\partial_s R_{t,s}$, it always holds that $\partial_s R_{t,s} = r_{t,t,s}$. Moreover, provided the partial derivatives $\partial_t r_{t,t',s}$ and $\partial_{t'} r_{t,t',s}$ are continuous, the partial derivative $\partial_t \partial_s R_{t,s}$ is simply given by:

$$\partial_t \partial_s R_{t,s} = \partial_t r_{t,t',s}|_{t'=t} + \partial_{t'} r_{t,t',s}|_{t'=t}.$$

Thus, we only need to compute each term $E_{t,s}^{(1)} = \partial_t r_{t,t',s}|_{t'=t}$ and $E_{t,s}^{(2)} = \partial_{t'} r_{t,t',s}|_{t'=t}$ separately.

Computing $E_{t,s}^{(1)} = \partial_t r_{t,t',s}|_{t'=t}$. First, a simple computation shows that

$$\nabla \cdot [\alpha_t(T_{t,s}^{-1}(y))] = \text{Tr}(\nabla \alpha_t(T_{t,s}^{-1}(y)) \nabla T_{t,s}^{-1}(y))$$

hence, we can perform integration by parts on the last term of (82) to obtain a second expression for $r_{t,t',s}$

$$r_{t,t',s} = \int \left(\dot{V}_s(y) - \Pi_s[\dot{V}_s] \right) q_{t',s}(y) \, dy + \int \alpha_t(T_{t,s}^{-1}(y))^\top \nabla \log \left(\frac{q_{t',s}(y)}{\Pi_s(y)} \right) q_{t',s}(y) \, dy. \tag{83}$$

After differentiating (83) w.r.t. t , only the second integral in (83) contributes to the expression of the derivative:

$$\partial_t r_{t,t',s} = \int (\partial_t \alpha_t(T_{t,s}^{-1}(y)) + \nabla T_{t,s}^{-1}(y) \nabla \alpha_t(T_{t,s}^{-1}(y)))^\top \nabla \log \left(\frac{q_{t',s}(y)}{\Pi_s(y)} \right) q_{t',s}(y) dy.$$

Taking $t' = t$ and performing again a change of variables $y = T_{t,s}(x)$, it follows that:

$$\begin{aligned} E_{t,s}^{(1)} &= \int \left(\partial_t \alpha_t(x) + (\nabla T_{t,s}(x))^{-1} \nabla \alpha_t(x) \right)^\top \nabla \log \left(\frac{q_{t,s}}{\Pi_s} \right) (T_{t,s}(x)) \Pi_t(x) dx. \\ &= \Pi_t \left[\left(\partial_t \alpha_t + (\nabla T_{t,s})^{-1} \nabla \alpha_t \right)^\top \nabla \log \left(\frac{q_{t,s}}{\Pi_s} \right) \circ T_{t,s} \right] \end{aligned}$$

Computing $E_{t,s}^{(2)} = \partial_{t'} r_{t,t',s}|_{t'=t}$. From (82) we have that $r_{t,t',s} = \int c_{t,s}(y) q_{t',s}(y) dy$. Therefore, by the dominated convergence theorem, we can write:

$$\partial_{t'} r_{t,t',s} = \int c_{t,s}(y) \partial_t \log(q_{t',s}(y)) q_{t',s}(y) dy.$$

Taking $t' = t$ and applying a change of variables $y = T_{t,s}(x)$, we get:

$$\begin{aligned} E_{t,s}^{(2)} &= \int c_{t,s}(T_{t,s}(x)) \partial_t \log(q_{t,s})(T_{t,s}(x)) \Pi_t(x) dx \\ &= \Pi_t [c_{t,s} \circ T_{t,s} (\partial_t \log(q_{t,s}) \circ T_{t,s})]. \end{aligned}$$

□

Next we show that the functions $\partial_t \log(q_{t,s}(T_{t,s}(x)))$ and $-c_{t,s}(T_{t,s}(x))$ appearing in the expression of $\partial_t \partial_s R_{t,s}$ provided in Lemma 10, approach $\bar{g}_t^\alpha(x)$ point-wise with an error proportional to $|t - s|$.

Lemma 11. *Under Assumptions (a) to (c) and for t and s such that $|t - s| \leq \frac{1}{2L}$, it holds:*

$$|\partial_t \log(q_{t,s}(T_{t,s}(x))) - \bar{g}_t^\alpha(x)| \lesssim |t - s| (1 + \|x\| + \|x\|^2) \quad (84)$$

$$|c_{t,s}(T_{t,s}(x)) + \bar{g}_t^\alpha(x)| \lesssim |t - s| (1 + \|x\| + \|x\|^2) \quad (85)$$

$$\left\| \nabla \log \left(\frac{q_{t',s}}{\Pi_s} \right) (T_{t,s}(x)) \right\| \lesssim |t - s| (1 + \|x\| + \|x\|^2) \quad (86)$$

The proof of Lemma 11 relies on the expressions of the time derivatives $\partial_t T_{t,s}^{-1}(y)$, $\partial_t \nabla T_{t,s}^{-1}(y)$ and $\partial_t \log |\nabla T_{t,s}^{-1}(y)|$ which we provide in Lemma 12 without proof as they follow by direct calculations.

Lemma 12. *Let y be in \mathcal{X} and denote $x_{t,s} := T_{t,s}^{-1}(y)$ and $C_{t,s} = \nabla T_{t,s}^{-1}(y)$. Under Assumptions (a) to (c) and using (52) of Lemma 7, it holds that:*

$$\begin{aligned} \partial_t T_{t,s}^{-1}(y) &= C_{t,s}(\alpha_t(x) - (s - t)\partial_t \alpha_t(x)) \\ \partial_t \nabla T_{t,s}^{-1}(y) &= C_{t,s} \nabla \alpha_t(x_{t,s}) C_{t,s} - (s - t) C_{t,s} \mathcal{G}_{t,s}(y) C_{t,s} \\ \partial_t \log |\nabla T_{t,s}^{-1}(y)| &= \text{tr}(C_{t,s} \nabla \alpha_t(x)) - (s - t) \text{Tr}(C_{t,s} \mathcal{G}_{t,s}(y)) \end{aligned}$$

with

$$\mathcal{G}_{t,s}(y) := \partial_t \nabla \alpha_t(x_{t,s}) + H \alpha_t(x_{t,s}) \partial_t T_{t,s}^{-1}(y)$$

We can now prove Lemma 11.

Proof. of Lemma 11 Let $y \in \mathcal{X}$ and define $x_{t,s} = T_{t,s}^{-1}(y)$ and $C_{t,s} = \nabla T_{t,s}^{-1}(y)$.

Equation (84). We first express the partial derivative $\partial_t \log(q_{t,s})$:

$$\begin{aligned} \partial_t \log(q_{t,s}(y)) &= - \left(\dot{V}_t(x_{t,s}) - \Pi_t \left[\dot{V}_t \right] + \nabla V_t(x_{t,s}) \partial_t T_{t,s}^{-1}(y) - \partial_t \log |\nabla T_{t,s}^{-1}(y)| \right) \\ &= - \left(\dot{V}_t(x_{t,s}) - \Pi_t \left[\dot{V}_t \right] + \nabla V_t(x_{t,s}) C_{t,s} \alpha_t(x_{t,s}) - \text{Tr}(C_{t,s} \nabla \alpha_t(x_{t,s})) \right) \\ &\quad + (s-t) \nabla V_t(x_{t,s}) \nabla C_{t,s} \partial_t \alpha_t(x_{t,s}) \\ &\quad - (s-t) \text{Tr}(C_{t,s} (\partial_t \nabla \alpha_t(x_{t,s}) + H \alpha_t(x_{t,s}) C_{t,s} (\alpha_t(x_{t,s}) - (s-t) \partial_t \alpha_t(x_{t,s})))) \end{aligned}$$

where the expressions of the time derivatives $\partial_t T_{t,s}^{-1}(y)$ and $\partial_t \log |\nabla T_{t,s}^{-1}(y)|$ are given by Lemma 12. We can then evaluate $\partial_t \log(q_{t,s}(y))$ at $y = T_{t,s}(x)$ for some $x \in \mathcal{X}$:

$$\begin{aligned} \partial_t \log(q_{t,s}(T_{t,s}(x))) &:= - \left(\dot{V}_t(x) - \Pi_t \left[\dot{V}_t \right] + \nabla V_t(x) C_{t,s} \alpha_t(x) - \text{Tr}(C_{t,s} \nabla \alpha_t(x)) \right) \\ &\quad + (s-t) \nabla V_t(x) C_{t,s} \partial_t \alpha_t(x) \\ &\quad - (s-t) \text{Tr}(C_{t,s} (\partial_t \nabla \alpha_t(x) + H \alpha_t(x) C_{t,s} (\alpha_t(x) - (s-t) \partial_t \alpha_t(x)))) \end{aligned}$$

Hence, by definition of \bar{g}_t^α , we have:

$$\begin{aligned} \partial_t \log(q_{t,s}(T_{t,s}(x))) - \bar{g}_t^\alpha(x) &= - \nabla V_t(x)^\top (C_{t,s} - I) \alpha_t(x) + \text{Tr}((C_{t,s} - I) \nabla \alpha_t(x)) \\ &\quad + (s-t) \nabla V_t(x) C_{t,s} \partial_t \alpha_t(x) \\ &\quad - (s-t) \text{Tr}(C_{t,s} (\partial_t \nabla \alpha_t(x) + H \alpha_t(x) C_{t,s} (\alpha_t(x) - (s-t) \partial_t \alpha_t(x)))) \end{aligned}$$

We first note that:

$$\begin{aligned} C_{t,s} - I &= \nabla T_{t,s}^{-1}(y) - I = (\nabla T_{t,s}(x))^{-1} - I \\ &= (I + (s-t) \nabla \alpha_t(x)) - I \\ &= (t-s) C_{t,s} \nabla \alpha_t(x). \end{aligned}$$

This implies:

$$\begin{aligned} \partial_t \log(q_{t,s}(T_{t,s}(x))) - \bar{g}_t^\alpha(x) &= (s-t) (\nabla V_t(x)^\top C_{t,s} \nabla \alpha_t(x) \alpha_t(x) - \text{Tr}(C_{t,s} \nabla \alpha_t(x)^2)) \\ &\quad + (s-t) \nabla V_t(x) C_{t,s} \partial_t \alpha_t(x) \\ &\quad - (s-t) \text{Tr}(C_{t,s} (\partial_t \nabla \alpha_t(x) + H \alpha_t(x) C_{t,s} (\alpha_t(x) - (s-t) \partial_t \alpha_t(x)))). \end{aligned}$$

Finally, we know by Lemma 16 that ∇V_t and α_t have at most a linear growth in x and by Lemma 7 we also have that $\nabla \alpha_t$ and $\partial_t \alpha_t$ are bounded, that $H \alpha_t(x)$ and $\partial_t \nabla \alpha_t$ have at most a linear growth in x . This directly yields the desired upper-bound:

$$|\partial_t \log(q_{t,s}(T_{t,s}(x))) - \bar{g}_t^\alpha(x)| \lesssim |t-s| (1 + \|x\| + \|x\|^2).$$

Equation (85) We will first control the term $c_{t,s}(T_{t,s}(x)) + \bar{g}_s^\alpha(x)$, then we will show that $\bar{g}_t^\alpha(x) - \bar{g}_s^\alpha(x)$ is of the same order. By definition of $c_{t,s}(x)$ and \bar{g}_s^α , we can write:

$$\begin{aligned} c_{t,s}(T_{t,s}(x)) + \bar{g}_s^\alpha(x) &= \left(\dot{V}_s(T_{t,s}(x)) - \dot{V}_s(x) \right) + (\nabla V_s(T_{t,s}(x)) - \nabla V_s(x))^\top \alpha_t(x) \\ &\quad + \nabla V_s(x)^\top (\alpha_t(x) - \alpha_s(x)) + \text{Tr}(\nabla \alpha_s(x) - \nabla \alpha_t(x)) \\ &\quad + (s-t) \text{Tr}(C_{t,s} \nabla \alpha_t(x)^2). \end{aligned}$$

Hence, using Assumption (a) and Lemma 7 we have:

$$\|c_{t,s}(T_{t,s}(x)) + \bar{g}_s^\alpha(x)\| \lesssim |t-s| (1 + \|x\| + \|x\|^2).$$

Similarly, we can show that $\bar{g}_t^\alpha(x) - \bar{g}_s^\alpha(x)$ satisfies a similar bound. This allows to get (85) using triangular inequality.

Equation (86) By direct calculation, we have:

$$\nabla \log \left(\frac{q_{t,s}}{\Pi_s} \right) (y) = \nabla V_s(y) - C_{t,s} \nabla V_t(x_{t,s}) - (s-t) \text{Tr}(C_{t,s} H \alpha_t(x_{t,s}) C_{t,s}).$$

In particular, choosing $y = T_{t,s}(x)$, we get:

$$\begin{aligned} \nabla \log \left(\frac{q_{t,s}}{\Pi_s} \right) (T_{t,s}(x)) &= \nabla V_s(T_{t,s}(x)) - C_{t,s} \nabla V_t(x) - (s-t) \text{Tr}(C_{t,s} H \alpha_t(x) C_{t,s}) \\ &= \nabla V_s(T_{t,s}(x)) - V_t(x) - (C_{t,s} - I) \nabla V_t(x) - (s-t) \text{Tr}(C_{t,s} H \alpha_t(x) C_{t,s}). \end{aligned}$$

The result follows directly by a similar argument as done previously. \square

Proof. of Proposition 30

Equation (80) For the first inequality, we use the expression of $\partial_t \partial_s R_{t,s}$ from Lemma 10:

$$\partial_t \partial_s R_{t,s} = E_{t,s}^{(1)} + E_{t,s}^{(2)}$$

where

$$\begin{aligned} E_{t,s}^{(1)} &= \Pi_t \left[\left(\partial_t \alpha_t + (\nabla T_{t,s})^{-1} \nabla \alpha_t \right)^\top \nabla \log \left(\frac{q_{t,s}}{\Pi_s} \right) \circ T_{t,s} \right], \\ E_{t,s}^{(2)} &= \Pi_t [(c_{t,s} \circ T_{t,s}) (\partial_t \log(q_{t,s}) \circ T_{t,s})]. \end{aligned}$$

Recall that $\nabla \alpha_t$ and $\partial_t \alpha_t$ are bounded by Lemma 7 and that since $\lambda \leq \frac{1}{2L}$, we also have that $\|\nabla T_{t,s}\|^{-1}$ is bounded. Moreover, (86) allows to write:

$$\left| E_{t,s}^{(1)} \right| \lesssim |t-s| \Pi_t [1 + \|X\| + \|X\|^2]. \quad (87)$$

For $E_{t,s}^{(2)}$, we use the following decomposition:

$$\begin{aligned} E_{t,s}^{(2)} + \Pi_t [(\bar{g}_t^\alpha)^2] &= \Pi_t [(c_{t,s} \circ T_{t,s} + \bar{g}_t^\alpha) (\partial_t \log(q_{t,s}) \circ T_{t,s})] \\ &\quad - \Pi_t [\bar{g}_t^\alpha (\partial_t \log(q_{t,s}) \circ T_{t,s} - \bar{g}_t^\alpha)]. \end{aligned}$$

The functions $c_{t,s} \circ T_{t,s}$ and \bar{g}_t^α have at most a quadratic growth and by (84) and (85) we get:

$$\left| E_{t,s}^{(2)} + \Pi_t [(\bar{g}_t^\alpha)^2] \right| \lesssim |t-s| \Pi_t [1 + \|X\| + \|X\|^2 + \|X\|^3 + \|X\|^4]. \quad (88)$$

Both bound (87) and (88) are finite due to Assumption (c), which implies the result.

Equation (81). To get the last inequality it suffice to differentiate in time and use the growth assumptions and integrability conditions. \square

E.8. Diffusion estimates

Lemma 13. Under Assumptions (a) to (c), the following inequalities hold:

$$\mathbb{E} \left[\|X_t\|^2 \right] \lesssim 1 \quad \mathbb{E} \left[\|X_t\|^4 \right] \lesssim 1, \quad \mathbb{E} \left[\|X_t - X_s\|^2 \right] \lesssim |t-s|$$

Proof. This is a direct consequence of the drift $b_t(x)$ being jointly Lipschitz in t and x and the initial distribution Π_0 having finite moments of order 4. \square

E.8.1. PROOF OF PROPOSITION 25

Proof. of Proposition 25

Bound on Equation (71) We use the decomposition of w_t^α as a product of w_s^α and a bounded term:

$$\mathbb{E}[\bar{w}_t^\alpha f(X_s)] = \mathbb{E}\left[\bar{w}_s^\alpha f(X_s) \exp\left(\int_s^t \bar{g}_u^\alpha du\right)\right].$$

By Lemma 8, we know that the normalized weight \bar{w}_t^α is upper-bounded by a term of the form $\exp(C(t-s))$, hence we get

$$\|\mathbb{E}[\bar{w}_t^\alpha f(X_s)]\| \lesssim \mathbb{E}[\bar{w}_s^\alpha \|f(X_s)\|] = \Pi_s[\|f\|]$$

where we used Proposition 22 to get the final result.

Bound on Equation (72) We first define the random variable:

$$\Psi_{s,s'}^t := \mathbb{E}\left[\exp\left(\int_s^t \bar{g}_u^\alpha(X_u) du\right) \int_s^{s'} dB_u \middle| \mathcal{F}_s\right],$$

where \mathcal{F}_s is the filtration associated to the Brownian motion B_t defining the process X_t . Using the tower property for conditional expectations and Cauchy-Schwartz, we have

$$\left|\mathbb{E}\left[\bar{w}_t^\alpha f(X_s)^\top \int_s^{s'} dB_u\right]\right| = |\mathbb{E}[\bar{w}_t^\alpha f(X_s)^\top \Psi_{s,s'}^t]| \leq \mathbb{E}[\bar{w}_t^\alpha \|f\|^2]^{\frac{1}{2}} \mathbb{E}[\bar{w}_t^\alpha \|\Psi_{s,s'}^t\|^2]^{\frac{1}{2}}.$$

Using (71) for both terms on the r.h.s., we obtain

$$\left|\mathbb{E}\left[\bar{w}_t^\alpha f(X_s)^\top \int_s^{s'} dB_u\right]\right| \lesssim \Pi_s[\|f\|^2]^{\frac{1}{2}} \mathbb{E}[\|\Psi_{s,s'}^t\|^2]^{\frac{1}{2}}.$$

It remains to show that $\mathbb{E}[\|\Psi_{s,s'}^t\|^2] \lesssim (s'-s)^2$. To achieve this, we first introduce the notation \mathbb{E}_s for the conditional expectation knowing the process up to time s . We also write $w_{s,s'}(X) = \exp(\int_s^{s'} \bar{g}_u^\alpha(X_u) du)$ and $r_{s,t}(X_s) = \mathbb{E}_s[w_{s,t}(X)]$. We consider now a second process $(X'_t)_{t \geq 0}$ that is coupled to $(X_t)_{t \geq 0}$ as follows. $X_u = X'_u$ up to time s , then on the interval $[s, s']$, X'_t is driven by a Brownian motion B'_t that is independent from B_t . Finally, starting from time s' up to time t , the process X'_t is driven again by the same Brownian motion B_t as X_t . Hence, (X_s, X'_s) defines a coupling of two SDEs. We denote by $\mathbb{E}_{s',s}$ the conditional expectation knowing the process (X_t) up to time s' and the process (X'_t) up to time s . Such coupling is shown in Lemma 14 to satisfy $\mathbb{E}[\int_s^t \|X_u - X'_u\|^2 du] \lesssim |s' - s|$. To use this property, we will express $\Psi_{s,s'}^t$ in terms of the two coupled processes. We start by noting that $\mathbb{E}_s[\mathbb{E}_{s'}[w_{s,t}(X')](B_{s'} - B_s)] = 0$ by independence of the increments of Brownian motions. Hence, we can write

$$\begin{aligned} \|\Psi_{s,s'}^t\| &= \|\mathbb{E}_s[(w_{s,t}(X) - \mathbb{E}_s[w_{s,t}(X')])(B_{s'} - B_s)]\| \\ &\leq \mathbb{E}_s[\|\mathbb{E}_{s',s}[w_{s,t}(X) - w_{s,t}(X')]\| \|B_{s'} - B_s\|] \\ &\leq e^{C(t-s)} \mathbb{E}_{s,s} \left[\|B_{s'} - B_s\| \int_s^t |g_u^\alpha(X_u) - g_u^\alpha(X'_u)| du \right] \end{aligned}$$

For the last line, we rely on local Lipschitzness of the exponential along with the fact $\bar{g}_u^\alpha(X_u) \leq C$ by Assumption (g). Moreover, Lemma 17 shows that

$$|g_t^\alpha(x) - g_t^\alpha(x')| \lesssim (1 + \|x\| + \|x'\|)\|x - x'\|,$$

hence we have:

$$\begin{aligned}
 \|\Psi_{s,s'}^t\| &\lesssim \mathbb{E}_{s,s} \left[\|B_{s'} - B_s\| \int_s^t (1 + \|X_u\| + \|X'_u\|) \|X_u - X'_u\| du \right] \\
 &\lesssim \mathbb{E}_{s,s} \left[\|B_{s'} - B_s\|^2 \int_s^t (1 + \|X_u\| + \|X'_u\|)^2 du \right]^{\frac{1}{2}} \mathbb{E}_{s,s} \left[\int_s^t \|X_u - X'_u\|^2 du \right]^{\frac{1}{2}} \\
 &\lesssim |s' - s| \mathbb{E}_{s,s} \left[\int_s^t \|X_u - X'_u\|^2 du \right]^{\frac{1}{2}}
 \end{aligned}$$

where for the second line we used Cauchy-Schwarz inequality. For the last line, we used the fact that the processes X_s and X'_s have finite moments of order 4 by Lemma 13 and that the increment $B_{s'} - B_s$ follows a zero-mean multivariate normal of covariance given by $(s' - s)I_d$. Finally, by application of Lemma 14, we get the desired bound and conclude the proof

$$\mathbb{E} \left[\|\Psi_{s,s'}^t\|^2 \right] \lesssim |s' - s| \mathbb{E} \left[\int_s^t \|X_u - X'_u\|^2 du \right] \lesssim |s' - s|^2.$$

Bound on Equation (73) We first apply Cauchy-Schwarz inequality:

$$\begin{aligned}
 \left| \mathbb{E} \left[\bar{w}_t^\alpha f(X_s)^\top \int_s^{s'} h_u(X_u, X_s) du \right] \right| &\leq \mathbb{E} \left[\bar{w}_t^\alpha \|f(X_s)\|^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\bar{w}_t^\alpha \left\| \int_s^{s'} h_u(X_u, X_s) du \right\|^2 \right]^{\frac{1}{2}} \\
 &\lesssim \Pi_s \left[\|f\|^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\left\| \int_s^{s'} h_u(X_u, X_s) du \right\|^2 \right]^{\frac{1}{2}} \\
 &\lesssim \sqrt{s' - s} \Pi_s \left[\|f\|^2 \right]^{\frac{1}{2}} \mathbb{E} \left[\int_s^{s'} \|h_u(X_u, X_s)\|^2 du \right]^{\frac{1}{2}}
 \end{aligned}$$

To get the second line in the above inequality, we use (71) for the first expectation in the r.h.s. and use that \bar{w}_t^α is bounded by Lemma 8 to bound the second expectation. The last line follows by direct application of Cauchy-Schwarz inequality in time and yields the desired result. In the particular case when $h_u(x, y) = b_u^\alpha(x) - \beta_{u'}^{\alpha,\lambda}(y)$ we further use Lemma 9 which yields:

$$\mathbb{E} \left[\int_s^{s'} \|h_u(X_u, X_s)\|^2 \right] = \mathbb{E} \left[\int_s^{s'} \left\| b_u^\alpha(X_u) - \beta_{u'}^{\alpha,\lambda}(X_s) \right\|^2 du \right] \lesssim \lambda |s' - s|$$

The result then follows directly by applying this bound to (73) and recalling that $|s' - s| \leq \lambda$ by assumption. \square

Lemma 14. *Let s and s' be two numbers in $[0, t]$ with $s \leq s'$. Let $(X_t)_{t \geq 0}$ and $(X'_t)_{t \geq 0}$ be two coupled processes following the SDE (54) and such that $X_t = X'_t$ up to time s , then on the interval $[s, s']$ the two processes are driven with two independent motions B_t and B'_t and finally, starting from time s' , the two processes are again diffused with the same Brownian motion. Then, under Assumption (a) and Lemma 7, we have the following:*

$$\mathbb{E} \left[\int_s^t \|X_u - X'_u\|^2 du \right] \lesssim |s - s'|$$

Proof. By definition of the X_t and X'_t , since they share the same Brownian motion starting from time s' , we have

$$\begin{aligned}
 X_u &= X_{s'} + \int_{s'}^u b_l(X_l) dl + \sqrt{2} \int_{s'}^u dB_l, \\
 X'_u &= X'_{s'} + \int_{s'}^u b_l(X'_l) dl + \sqrt{2} \int_{s'}^u dB_l.
 \end{aligned}$$

Hence,

$$\|X_u - X'_u\| \leq \|X_{s'} - X'_s\| + \int_{s'}^u \|b_l(X_l) - b_l(X'_l)\| dl.$$

By Assumption **(a)** and Lemma 7, we have that $\|b_l(x) - b_l(x')\| \lesssim \|x - x'\|$, hence we have:

$$\|X_u - X'_u\| \lesssim \|X_{s'} - X'_s\| + \int_{s'}^u \|X_l - X'_l\| dl.$$

By application of Gronwall's lemma, we have for any $s' \leq u \leq t$:

$$\|X_u - X'_u\| \lesssim \|X_{s'} - X'_s\|. \quad (89)$$

Now, for $s \leq u \leq s'$, we have

$$\begin{aligned} X_u &= X_s + \int_s^u b_l(X_l) dl + \sqrt{2} \int_s^u dB_l, \\ X'_u &= X_s + \int_s^u b_l(X'_l) dl + \sqrt{2} \int_s^u dB'_l. \end{aligned}$$

Hence, we have the following bound:

$$\|X_u - X'_u\| \lesssim \int_s^u \|X_l - X'_l\| dl + \sqrt{2}(\|B_u - B_s\| + \|B'_u - B'_s\|).$$

This allows to upper-bound the expectation $\mathbb{E}[\|X_u - X'_u\|^2]^{\frac{1}{2}}$:

$$\mathbb{E}[\|X_u - X'_u\|^2]^{\frac{1}{2}} \lesssim \int_s^u \mathbb{E}[\|X_l - X'_l\|^2]^{\frac{1}{2}} + 2\sqrt{2}\sqrt{(u-s)d}$$

Using Gronwall's lemma a second time, we get for any $s \leq u \leq s'$:

$$\mathbb{E}[\|X_u - X'_u\|^2]^{\frac{1}{2}} \lesssim \sqrt{u-s}. \quad (90)$$

For $u > s'$, we can then use (89) along with (90) to write

$$\mathbb{E}[\|X_u - X'_u\|^2] \lesssim s' - s. \quad (91)$$

Finally, using (90) and (91) and integrating over u on the interval $[s, t]$, we get the desired result after applying Fubini's theorem to exchange the order of the expectation and time integral:

$$\mathbb{E}\left[\int_s^t \|X_u - X'_u\|^2 du\right] \lesssim |s' - s|.$$

□

E.9. Pointwise estimates

Lemma 15. *Under Assumptions **(a)** to **(c)**, and for any $t_{k-1} \leq s \leq t_k$ and $0 \leq u \leq 1$, it holds that*

$$\left|g_s^\alpha(x') - \delta_t^{\alpha,\lambda}(x)\right| \lesssim (1 + \|x\| + \|x'\|)\|x - x'\| + \lambda(1 + \|x\| + \|x'\|)^2.$$

Proof. By definition of $\delta_t^{\alpha,\lambda}(x)$ in (57) and using the fundamental theorem of calculus we have:

$$\begin{aligned} \delta_t^{\alpha,\lambda}(x) &= \int_0^1 \text{Tr}\left(\left((I + \lambda u \nabla \alpha_t(x))^{-1} - I\right) \nabla \alpha_t(x)\right) + \text{Tr}(\nabla \alpha_t(x)) du \\ &\quad - \int_0^1 \left(\dot{V}_{t-\lambda u}(y_u) + \nabla V_{t-\lambda u}(y_u)^\top \alpha_t(x)\right) du, \end{aligned}$$

where we introduced $y_u = x + \lambda(1 - u)\alpha_t(x)$. This directly yields:

$$g_s^\alpha(x) - \delta_t^{\alpha, \lambda}(x') = \int_0^1 (B_u + C_u + D_u + E_u) du \quad (92)$$

where:

$$\begin{aligned} B_u &= \text{Tr}(\nabla \alpha_s(x')) - \text{Tr}(\nabla \alpha_t(x)), \\ C_u &= -\text{Tr}\left(\left((I + \lambda u \nabla \alpha_t(x))^{-1} - I\right) \nabla \alpha_t(x)\right), \\ D_u &= \nabla V_{t-\lambda u}(y_u)^\top \alpha_t(x) - \nabla V_s(x')^\top \alpha_s(x'), \\ E_u &= \dot{V}_{t-\lambda u}(y_u) - \dot{V}_s(x'). \end{aligned}$$

Bound on B_u . Under Assumption **(b)**, Lemma 7 applies and we directly have:

$$\begin{aligned} |B_u| &\lesssim (1 + \|x\| + \|x'\|)(\|x - x'\| + |t - s|) \\ &\lesssim (1 + \|x\| + \|x'\|)(\|x - x'\| + \lambda) \end{aligned} \quad (93)$$

where we used that $t - \lambda \leq s \leq t$ to get the last line.

Bound on $|C_u|$. By direct calculation we get:

$$\begin{aligned} C_u &= \text{Tr}\left(\left((I + \lambda u \nabla \alpha_t(x))^{-1} - I\right) \nabla \alpha_t(x)\right) \\ &= \lambda u \text{Tr}\left(\left((I + \lambda u \nabla \alpha_t(x))^{-1} \nabla \alpha_t(x)\right)^2\right) \end{aligned}$$

By Lemma 7, we know that $\nabla \alpha_t(x)$ is bounded by a constant term L . Hence, when $\lambda < \frac{1}{2L}$, it follows that:

$$\|(I + u \lambda \nabla \alpha_t(x))^{-1}\| \leq \frac{1}{1 - u \lambda L} \leq 2.$$

This directly implies that:

$$|C_u| \leq 2L^2 \lambda. \quad (94)$$

Bound on D_u .

$$\begin{aligned} |D_u| &= \left| \left(\nabla V_{t-\lambda u}(y_u)^\top - \nabla V_s(x')^\top\right) \alpha_t(x) + \nabla V_s(x')^\top (\alpha_t(x) - \alpha_s(x')) \right| \\ &\leq \|\nabla V_{t-\lambda u}(y_u)^\top - \nabla V_s(x')^\top\| \|\alpha_t(x)\| + \|\nabla V_s(x')^\top\| \|\alpha_t(x) - \alpha_s(x')\| \\ &\lesssim (|t - s - \lambda u| + \|y_u - x'\|) \|\alpha_t(x)\| + \|\nabla V_s(x')\| (\|x - x'\| + |t - s|) \\ &\lesssim (\lambda + \|y_u - x'\|) \|\alpha_t(x)\| + \|\nabla V_s(x')\| (\|x - x'\| + \lambda) \end{aligned}$$

where we used Assumption **(a)** and (52) of Lemma 7 for the third line and that $|t - s - \lambda u| \leq \lambda$ and $|t - s| \leq \lambda$ to get the last line. Moreover, by Lemma 16, we have that:

$$\|\alpha_t(x)\| \lesssim (1 + \|x\|), \quad \|\nabla V_t(x)\| \lesssim (1 + \|x\|). \quad (95)$$

This allows to further write:

$$|D_u| \lesssim (\lambda(1 + \|x\|) + \|x - x'\|)(1 + \|x\|). \quad (96)$$

Bound on E_u . Using Assumption **(a)**, we directly have:

$$\begin{aligned} |E_u| &\lesssim (1 + \|y_u\| + \|x'\|)(\|y_u - x'\| + |t - s - \lambda u|) \\ &\lesssim (1 + \|x\| + \|x'\| + \lambda u \|\alpha_t(x)\|)(\|x - x'\| + \lambda u \|\alpha_t(x)\| + |t - s - \lambda u|) \\ &\lesssim (1 + \|x\| + \|x'\| + \lambda u \|\alpha_t(x)\|)(\|x - x'\| + \lambda u \|\alpha_t(x)\| + \lambda), \end{aligned}$$

where we used that $|t - s - \lambda u| \leq \lambda$ since $t - \lambda \leq s \leq t$ and $0 \leq u \leq 1$. Moreover, using that α has a linear growth in x (by (95)) we have:

$$|E_u| \lesssim (1 + \|x\| + \|x'\|)(\|x - x'\| + \lambda(1 + \|x\|)). \quad (97)$$

Final bound. Using (93), (94), (96) and (97) in (92), it follows:

$$\left| g_s^\alpha(x) - \delta_t^{\alpha, \lambda}(x') \right| \lesssim (1 + \|x\| + \|x'\|)(\|x - x'\| + \lambda(1 + \|x\|)).$$

□

Proposition 31. Let t and s be in $[0, 1]$ such that $|t - s| \leq \lambda$. Under Assumptions (a) and (b), it holds that:

$$\left\| b_s^\alpha(x') - \beta_t^{\alpha, \lambda}(x) \right\| \lesssim \|x - x'\| + \lambda(1 + \|x\|).$$

Proof. By definition of $b_s^\alpha(x')$ and $\beta_t^{\alpha, \lambda}(x)$, the following identity holds:

$$\left\| b_s^\alpha(x') - \beta_t^{\alpha, \lambda}(x) \right\| = \left\| \alpha_s(x') - \alpha_t(x) - \nabla V_s(x') + \nabla V_t(x + \lambda \alpha_t(x)) \right\|.$$

Therefore, since $\alpha_t(x)$ and $\nabla V_t(x)$ are Lipschitz by Assumption (a) and Lemma 7, we have:

$$\begin{aligned} \left\| b_s^\alpha(x') - \beta_t^{\alpha, \lambda}(x) \right\| &\leq L(2\|x - x'\| + 2|t - s| + \lambda\|\alpha_t(x)\|) \\ &\leq L(2\|x - x'\| + 2\lambda + \lambda\|\alpha_t(x)\|). \end{aligned}$$

Moreover, Lemma 16 below ensures that $\|\alpha_t(x)\| \lesssim (1 + \|x\|)$, therefore:

$$\left\| b_s^\alpha(x') - \beta_t^{\alpha, \lambda}(x) \right\| \lesssim \|x - x'\| + \lambda(1 + \|x\|).$$

□

Lemma 16. Under Assumption (a) the gradient of the potential has a linear growth in x , i.e:

$$\|\nabla_x V_t(x)\| \lesssim (1 + \|x\|), \quad \forall x \in \mathcal{X}, \forall t \in [0, 1]. \quad (98)$$

Under Assumption (b), there exists a positive constant $D > 0$, such that for all $\alpha \in \mathcal{A}$.

$$\|\alpha_t(x)\| \leq D(1 + \|x\|), \quad \forall x \in \mathcal{X}, \forall t \in [0, 1]. \quad (99)$$

Proof. The equation (98) is a simple consequence of the Lipschitz assumption on ∇V_t . For (99), we recall that under Assumption (b), Lemma 7 holds. Therefore we have

$$\begin{aligned} \|\alpha_t(x)\| &\leq \|\alpha_0(0)\| + \|\alpha_t(x) - \alpha_0(0)\| \\ &\leq C + L(\|x\| + |t|) \\ &\leq C + L(1 + \|x\|). \end{aligned}$$

The result follows since the constants C and L are independent of the choice of the control α in \mathcal{A} . □

Lemma 17. Under Assumptions (a) and (b), the following bounds hold:

$$\begin{aligned} \left\| \beta_t^{\alpha, \lambda}(x) - \beta_t^{\alpha', \lambda}(x) \right\| &\lesssim \|\alpha_t(x) - \alpha'_t(x)\|, \\ \left| \delta_t^{\alpha, \lambda}(x) - \delta_t^{\alpha', \lambda}(x) \right| &\lesssim ((1 + \|x\|)\|\alpha_t(x) - \alpha'_t(x)\| + \|\nabla \alpha_t(x) - \nabla \alpha'_t(x)\|), \\ |g_t^\alpha(x) - g_t^{\alpha'}(x)| &\lesssim (1 + \|x\| + \|x'\|)\|x - x'\|. \end{aligned}$$

Proof. **Bound on $|\beta_t^{\alpha, \lambda} - \beta_t^{\alpha', \lambda}|$.** By Assumption (a) and Lemma 7 we have that:

$$\left\| \beta_t^{\alpha, \lambda}(x) - \beta_t^{\alpha', \lambda}(x) \right\| \leq (1 + L\lambda)\|\alpha_t(x) - \alpha'_t(x)\|.$$

Bound on $|\delta_t^{\alpha,\lambda} - \delta_t^{\alpha',\lambda}|$. For the second inequality, we use that:

$$\begin{aligned} \lambda \left(\delta_t^{\alpha,\lambda}(x) - \delta_t^{\alpha',\lambda}(x) \right) &= V_t(x + \lambda \alpha'_t(x)) - V_t(x + \lambda \alpha_t(x)) \\ &\quad + \log |I + \lambda \nabla \alpha_t(x)| - \log |I + \lambda \nabla \alpha'_t(x)|. \end{aligned}$$

Hence, for $\lambda \leq \frac{1}{2L}$, we get by Assumption (a) and Lemma 7 that

$$\begin{aligned} \lambda \left| \delta_t^{\alpha,\lambda}(x) - \delta_t^{\alpha',\lambda}(x) \right| &\lesssim \lambda [(1 + \|x\|) + \lambda (\|\alpha_t(x)\| + \|\alpha'_t(x)\|)] \|\alpha_t(x) - \alpha'_t(x)\| \\ &\quad + \lambda \|\nabla \alpha_t(x) - \nabla \alpha'_t(x)\|. \end{aligned}$$

Moreover, by Lemma 16, we know that α and α' have at most a linear growth. This allows to further write:

$$\left| \delta_t^{\alpha,\lambda}(x) - \delta_t^{\alpha',\lambda}(x) \right| \lesssim ((1 + \|x\|) \|\alpha_t(x) - \alpha'_t(x)\| + \|\nabla \alpha_t(x) - \nabla \alpha'_t(x)\|).$$

Bound on $|g_t^\alpha(x) - g_t^\alpha(x')$. It follows by direct computation. □

Lemma 18. *Under Assumptions (a) and (c) there exists a constant λ_0 independent on t such that for $\lambda \leq \lambda_0$, there exists $M > 0$ such that for all $0 \leq t \leq 1$ it holds that:*

$$|h_t^\lambda(x)| \leq \lambda \sqrt{\lambda} \frac{Q_1(x)}{1 + Q_2(x)} \exp(C\lambda \|\nabla V_t(x)\|^2)$$

where Q_1 and Q_2 are non-negative functions of $\|x\|$ of polynomial growth and independent from λ and t and C is a non-negative constants.

Proof. Let us perform a change of variables $u = \frac{x-y}{\sqrt{2\lambda}}$ in the integral appearing in the definition of $h_t^\lambda(x)$:

$$h_t^\lambda(x) := \log \int h(x, u, \lambda) \, du,$$

where we introduce the function $h(x, u, \lambda)$:

$$h(x, u, \lambda) = \frac{1}{(\sqrt{2\pi})^d} \exp \left(V_t(x) - V_t(y_\lambda) - \frac{1}{2} \left\| u + \sqrt{\frac{\lambda}{2}} \nabla V_t(y_\lambda) \right\|^2 \right),$$

and where $y_\lambda := x - \sqrt{2\lambda}u$ for conciseness. By the fundamental theorem of calculus and using the dominated convergence theorem we can write:

$$h_t^\lambda(x) = \lambda \int_0^1 \frac{\int \partial_\lambda h(x, u, s\lambda) \, du}{\int h(x, u, s\lambda) \, du} \, ds. \quad (100)$$

Upper-bound on $|\int \partial_\lambda h(x, u, s\lambda) \, du|$. We now find an expression for $\int \partial_\lambda h(x, u, s\lambda) \, du$ which exhibits an explicit factor $\sqrt{\lambda}$. By simple computation, we have that $\partial_\lambda h(x, u, s\lambda)$ is given by:

$$\partial_\lambda h(x, u, \lambda) = \frac{1}{2} \left(\frac{1}{\sqrt{2\lambda}} \nabla V_t(y_\lambda)^\top u + u^\top H V_t(y_\lambda) u - \frac{1}{2} \|\nabla V_t(y_\lambda)\|^2 \right) h(x, u, \lambda),$$

where $H V_t$ is the Hessian of V . Integrating w.r.t u , we get:

$$\int \partial_\lambda h(x, u, \lambda) \, du = \frac{1}{2} \frac{1}{\sqrt{2\lambda}} \phi_\lambda(1) + \frac{1}{2} \int \left(u^\top H V_t(y_\lambda) u - \frac{1}{2} \|\nabla V_t(y_\lambda)\|^2 \right) h(x, u, \lambda) \, du,$$

where the function $s \mapsto \phi_\lambda(s)$ is defined for $0 \leq s \leq 1$ as

$$\phi_\lambda(s) = \int \nabla V_t(y_{s\lambda})^\top u h(x, u, s\lambda) \, du,$$

with $\phi_\lambda(0) = 0$. By the fundamental theorem of calculus and using the dominated convergence theorem, we have that:

$$\begin{aligned} \phi_\lambda(1) &= \int_0^1 \partial_s \phi_\lambda(s) \, ds \\ &= -\sqrt{2\lambda} \int_0^1 \frac{1}{2\sqrt{s}} \int u^\top HV_t(y_{s\lambda}) u h(x, u, s\lambda) \, du \, ds \\ &\quad + \frac{\lambda}{2} \int_0^1 \frac{1}{\sqrt{2s\lambda}} \int (\nabla V_t(y_{s\lambda})^\top u)^2 h(x, u, s\lambda) \, du \, ds \\ &\quad + \frac{\lambda}{2} \int_0^1 \int \nabla V_t(y_{s\lambda})^\top u \left(u^\top HV_t(y_{s\lambda}) u - \frac{1}{2} \|\nabla V_t(y_{s\lambda})\|^2 \right) h(x, u, s\lambda) \, du \, ds. \end{aligned}$$

We can further use that $h(x, u, s\lambda) = h(x, u, \lambda) + \lambda \int_s^1 \partial_\lambda h(x, u, s'\lambda) \, ds'$ and express $\int \partial_\lambda h(x, u, \lambda) \, du$ as:

$$\int \partial_\lambda h(x, u, \lambda) \, du = A_\lambda(x) + B_\lambda(x) + C_\lambda(x) + D_\lambda(x),$$

where $A_\lambda, B_\lambda, C_\lambda$ and D_λ are given by:

$$\begin{aligned} A_\lambda(x) &= \frac{1}{2} \int_0^1 \int \frac{1}{2\sqrt{s}} u^\top (HV_t(y_\lambda) - HV_t(y_{s\lambda})) u h(x, u, \lambda) \, du \, ds, \\ B_\lambda(x) &= \frac{1}{4} \int_0^1 \frac{1}{2\sqrt{s}} \int ((\nabla V_t(y_{s\lambda})^\top u)^2 - \|\nabla V_t(y_\lambda)\|^2) h(x, u, \lambda) \, du \, ds, \\ C_\lambda(x) &= \frac{1}{4\sqrt{2}} \sqrt{\lambda} \int_0^1 \int \nabla V_t(y_{s\lambda})^\top u \left(u^\top HV_t(y_{s\lambda}) u - \frac{1}{2} \|\nabla V_t(y_{s\lambda})\|^2 \right) h(x, u, s\lambda) \, du \, ds, \\ D_\lambda(x) &= \frac{\lambda}{4} \int_0^1 \int \int_s^1 \partial_\lambda h(x, u, s'\lambda) ((\nabla V_t(y_{s\lambda})^\top u)^2 - 2HV_t(y_{s\lambda})) \, du \, ds' \, ds. \end{aligned}$$

We can further decompose $B_\lambda(x)$ in two terms $B_\lambda^{(1)}(x)$ and $B_\lambda^{(2)}(x)$ by adding and subtracting $\nabla V_t(x)$ inside each square:

$$B_\lambda(x) := B_\lambda^{(1)}(x) + B_\lambda^{(2)}(x),$$

with:

$$\begin{aligned} B_\lambda^{(1)}(x) &:= \frac{1}{4} \int_0^1 \frac{1}{2\sqrt{s}} \int u^\top ((\nabla V_t(y_{s\lambda}) - \nabla V_t(x))(\nabla V_t(y_{s\lambda}) - \nabla V_t(x))^\top) u \, du \, ds \\ &\quad - \frac{1}{4} \int_0^1 \frac{1}{2\sqrt{s}} \int (\nabla V_t(y_\lambda) - \nabla V_t(x))^\top (\nabla V_t(y_\lambda) + \nabla V_t(x)) \, du \, ds, \\ B_\lambda^{(2)}(x) &:= \frac{1}{4} \int (\nabla V_t(x)^\top u)^2 (h(x, u, \lambda) - h(x, u, 0)) \, du \\ &= \frac{\lambda}{4} \int_0^1 \frac{1}{2\sqrt{s}} \int (\nabla V_t(x)^\top u)^2 \partial_\lambda h(x, u, s\lambda) \, du \, ds. \end{aligned}$$

For $B_\lambda^{(2)}(x)$ we used that $\int (\nabla V_t(x)^\top u)^2 h(x, u, 0) \, du = \|\nabla V_t(x)\|^2$ and then applied the fundamental theorem of calculus for the difference $(h(x, u, \lambda) - h(x, u, 0))$ to exhibit a factor λ . Using that $\nabla V_t(x)$ and $HV_t(x)$ are Lipschitz in x and the growth assumptions, it can be shown that each of the terms $A_\lambda, B_\lambda^{(1)}, B_\lambda^{(2)}, C_\lambda$ and D_λ :

$$\begin{aligned} |B_\lambda^{(2)}(x)| + |C_\lambda(x)| &\leq \sqrt{\lambda} \int E_\lambda^2(x, u) \, du, & |A_\lambda(x)| + |B_\lambda^{(1)}(x)| &\leq \sqrt{\lambda} \int E_\lambda^1(x, u) \, du, \\ |D_\lambda(x)| &\leq \sqrt{\lambda} \int E_\lambda^3(x, u) \, du, \end{aligned}$$

with $E_\lambda^1(x, u)$, $E_\lambda^2(x, u)$ and $E_\lambda^3(x, u)$ given by:

$$\begin{aligned} E_\lambda^1(x, u) &= Q_1(x, u) \left(\int_0^1 \int_s^1 h(x, u, s'\lambda) ds' ds \right), \\ E_\lambda^2(x, u) &= Q_2(x, u) \left(\int_0^1 h(x, u, s\lambda) ds \right), \\ E_\lambda^3(x, u) &= Q_3(x, u) h(x, u, s\lambda). \end{aligned}$$

Here $Q_1(x, u)$, $Q_2(x, u)$ and $Q_3(x, u)$ are polynomials in $\|x\|$ and $\|u\|$ with non-negative coefficients independent from λ and t . Using again the Lipschitz regularity of $\nabla V_t(x)$, we can further find an upper-bound on $h(x, u, \lambda)$ of the form:

$$\begin{aligned} h(x, u, \lambda) &\leq C \exp(M\lambda \|\nabla V_t(x)\|^2 - \alpha(\lambda)\|u\|^2 + \beta(\lambda)\|u\| \|\nabla V_t(x)\|), \\ h(x, u, \lambda) &\geq C' \exp(M'\lambda \|\nabla V_t(x)\|^2 - \alpha'(\lambda)\|u\|^2 - \beta(\lambda)\|u\| \|\nabla V_t(x)\|). \end{aligned}$$

Here, C, C' and M, M' are independent of λ and t . The functions $\alpha(\lambda)$ and $\alpha'(\lambda)$ are positive for λ small enough and converge to $\frac{1}{2}$ with rate λ when $\lambda \rightarrow 0$. Finally $\beta(\lambda)$ converges to 0 with rate $\sqrt{\lambda}$. Using those bounds for $E_\lambda^1(x, u)$, $E_\lambda^2(x, u)$ and $E_\lambda^3(x, u)$ and $h(x, y, \lambda)$ and integrating over u , it follows that:

$$\begin{cases} \left| \int \partial_\lambda h(x, u, \lambda) du \right| \leq \sqrt{\lambda} Q_1(\|x\|) \exp(L\lambda \|\nabla V_t(x)\|^2), \\ \int h(x, u, \lambda) du \geq (1 + Q_2(\|x\|)) \exp(L'\lambda \|\nabla V_t(x)\|^2), \end{cases} \quad (101)$$

where Q_1 and Q_2 are non-negative functions of $\|x\|$ of polynomial growth and independent from λ and t and L, L' are some positive constants.

Upper-bound on $|h_t^\lambda(x)|$. Recalling (100), and using (101) we directly have:

$$|h_t^\lambda(x)| \leq \lambda \sqrt{\lambda} \frac{Q_1(\|x\|)}{1 + Q_2(\|x\|)} \exp(|L - L'|\lambda \|\nabla V_t(x)\|^2).$$

□

F. Variant of the algorithm for unbiased estimates

Algorithm 2 and Algorithm 3 give details of the full practical implementation described in Section 6 of the main paper.

G. Additional experimental details and discussion

The algorithm is implemented in JAX (Bradbury et al., 2018), Haiku (Hennigan et al., 2020) and Optax (Hessel et al., 2020). In all cases we use the Adam optimizer (Kingma and Ba, 2014) for learning the flow with the Optax defaults for all parameters except for the learning rates which are chosen for each example and are given below. For the MCMC kernels we used the TensorFlow Probability (Dillon et al., 2017) JAX substrate.

Each experimental configuration was performed using an NVIDIA v100 GPU and 4 CPUs. All experimental configurations took under 30 minutes and most were much shorter.

In all cases $N_{\text{test}} = 2000$ and this was the batch size for SMC and for VI ensuring fair comparison of trained estimators. Note VI is not sensitive to the training batch size because samples are replenished at each training step.

G.1. Two dimensional example

For each transition, we used 10 iterations of Hamiltonian Monte Carlo with 10 leapfrog steps per iteration. The rational quadratic splines closely match the implementation described by (Durkan et al., 2019). We used ten bins and padded with the identity outside of the range $[-4, 4]$. The bins height and widths were parameterized in terms of unconstrained real values. We then took the soft-max of these values and scaled and shifted them so that they had a minimal value of 10^{-4} and fitted with the range $[-4, 4]$. The derivatives were again parameterized using unconstrained real values and then made positive using the transformation $\nu + \log(1 + \exp(x))$ where $\nu = 10^{-4}$.

Algorithm 2 Annealed Flow Transport: Detailed Version

- 1: **Input:** Number of training, test and validation particles $N_{\text{train}}, N_{\text{test}}, N_{\text{val}}$, unnormalized annealed targets $\{\gamma_k\}_{k=0}^K$ such that $\gamma_0 = \pi_0$ and $\gamma_K = \gamma$, resampling thresholds $A_a \in [1/N_a, 1)$ for $a \in \{\text{train, test, val}\}$, number of training iterations J .
 - 2: **Output:** Approximations $\pi_K^{N_{\text{test}}}$ and $Z_K^{N_{\text{test}}, \text{test}}$ of π and Z .
 - 3: **for** $a \in \{\text{train, test, val}\}$ **do**
 - 4: Sample $X_0^{i,a} \sim \pi_0$ and set $W_0^{i,a} \leftarrow \frac{1}{N_a}$ and $Z_0^{N,a} \leftarrow 1$.
 - 5: **end for**
 - 6: **for** $k = 1, \dots, K$ **do**
 - 7: Learn the flow $T_k \leftarrow \text{LearnFlow}\left(J, \left\{X_{k-1}^{i,\text{train}}, W_{k-1}^{i,\text{train}}\right\}_{i=1}^{N_{\text{train}}}, \left\{X_{k-1}^{i,\text{val}}, W_{k-1}^{i,\text{val}}\right\}_{i=1}^{N_{\text{val}}}\right)$
 - 8: **for** $a \in \{\text{train, test, val}\}$ **do**
 - 9: Transport particles: $\tilde{X}_k^{i,a} \leftarrow T_k(X_{k-1}^{i,a})$.
 - 10: Estimate normalizing constant Z_k :

$$Z_k^{N_a,a} \leftarrow Z_{k-1}^{N_a,a} \left(\sum_{i=1}^{N_a} W_{k-1}^{i,a} G_{k,T_k}(X_{k-1}^{i,a}) \right).$$
 - 11: Compute IS weights:

$$w_k^{i,a} \leftarrow W_{k-1}^{i,a} G_{k,T_k}(X_{k-1}^{i,a}) \text{ // unnormalized}$$

$$W_k^{i,a} \leftarrow \frac{w_k^{i,a}}{\sum_{j=1}^{N_a} w_k^{j,a}} \text{ // normalized}$$
 - 12: Compute effective sample size $\text{ESS}_k^{N_a}$

$$\text{ESS}_k^{N_a} \leftarrow \left(\sum_{i=1}^{N_a} \left(W_k^{i,a} \right)^2 \right)^{-1}.$$
 - 13: **if** $\text{ESS}_k^{N_a} / N_a \leq A_a$ **then**
 - 14: Resample N_a particles from split a denoted abusively also $\tilde{X}_k^{i,a}$ according to the weights $W_k^{i,a}$,
 - 15: Set $W_k^{i,a} \leftarrow \frac{1}{N_a}$.
 - 16: **end if**
 - 17: Sample $X_k^{i,a} \sim K_k(\tilde{X}_k^{i,a}, \cdot)$. // MCMC
 - 18: **end for**
 - 19: **end for**
-

Algorithm 3 LearnFlow

- 1: **Input:** Number of training iterations J , training and validation particles and weights $\left\{X_{k-1}^{i,\text{train}}, W_{k-1}^{i,\text{train}}\right\}_{i=1}^{N_{\text{train}}}$ and $\left\{X_{k-1}^{i,\text{val}}, W_{k-1}^{i,\text{val}}\right\}_{i=1}^{N_{\text{val}}}$.
 - 2: **Output:** Estimated flow T_k
 - 3: Initialize flow to identity $T_k = Id$.
 - 4: Initialize list of flows $\mathcal{T}_{\text{opt}} \leftarrow \{T_k\}$.
 - 5: Initialize list of validation losses

$$\mathcal{E} \leftarrow \left\{ \sum_{i=1}^{N_{\text{val}}} W_{k-1}^{i,\text{val}} h_{T_k}(X_{k-1}^{i,\text{val}}) \right\}$$
 - 6: **for** $j = 1, \dots, J$ **do**
 - 7: Compute training loss using (8)

$$\mathcal{L}_k^{N_{\text{train}}}(T_k) \leftarrow \sum_{i=1}^{N_{\text{train}}} W_{k-1}^{i,\text{train}} h_{T_k}(X_{k-1}^{i,\text{train}}).$$
 - 8: Update T_k using SGD to minimize $\mathcal{L}_k^{N_{\text{train}}}(T_k)$.
 - 9: Update list of flows $\mathcal{T}_{\text{opt}} \leftarrow \mathcal{T}_{\text{opt}} \cup \{T_k\}$
 - 10: Initialize list of validation losses \mathcal{E}

$$\mathcal{E} \leftarrow \mathcal{E} \cup \left\{ \sum_{i=1}^{N_{\text{val}}} W_{k-1}^{i,\text{val}} h_{T_k}(X_{k-1}^{i,\text{val}}) \right\}$$
 - 11: **end for**
 - 12: Return flow with smallest validation error from the list of flows \mathcal{T}_{opt} .
-




Figure 3: Larger version of Figure 1. Weighted samples for a 2-D target density with AFT. The colours show the normalized weights which are clipped at the 95th percentile for clarity.

The inverse autoregressive flow used the same autoregressive network detailed in the Funnel example below but with a changed input dimensionality and outputting the spline parameters instead of the parameters of an affine transformation.

The Adam learning rate was 10^{-3} the training period J was 1000. AFT flow estimation particle numbers were $N_{\text{train}} = N_{\text{val}} = 2000$.

With the total time number of time steps scaled to be in the interval $[0, 1]$ with 0 being the initial distribution and 1 the final distribution, the HMC step sizes were interpolated between the times $[0., 0.25, 0.5, 1.]$ using the step sizes $[0.5, 0.5, 0.5, 0.3]$.

G.2. Additional details of Funnel example

We used 1000 steps of slice sampling per temperature with a maximum of 5 step size doublings as defined in the TensorFlow probability interface. For the affine inverse regressive flow we used an autoregressive neural network (Germain et al., 2015) where the correct autoregressive structure is achieved by masking network weights. We used a Leaky Relu non-linearity. The unmasked network would have 30 hidden units per input dimension and we used 3 hidden layers. To achieve identity initialisation of the flow we initialised the final weights and biases of the network to zero. The weights of the rest of network were initialised using a truncated normal distribution scaled by the fan-in and biases were initialised to zero. The final output of the MLP was taken and mapped directly into the mean and also the scale of the network after adding one to give the identity transformation overall.

The Adam learning rate was 10^{-3} the training period J was 4000. AFT flow estimation particle numbers were $N_{\text{train}} = N_{\text{val}} = 6000$.

With the total time number of time steps scaled to be in the interval $[0, 1]$ with 0 being the initial distribution and 1 the final distribution, the slice sampling step sizes were interpolated between the times $[0., 0.25, 0.5, 0.75, 1.]$ using the step sizes $[0.9, 0.7, 0.6, 0.5, 0.4]$.

G.3. Additional details of Variational Autoencoder example

G.3.1. AUTOENCODER TRAINING AND ARCHITECTURE DETAILS

The VAE encoder architecture was as follows. The encoder, which parameterized the amortized variational distribution started with two convolutional layers each followed by a Rectified Linear non-linearity. The first convolution had kernel shape 4×4 , stride 2×2 , 16 output channels and ‘valid’ padding. The second convolution had 4×4 , stride 2×2 , 32 output channels and ‘valid’ padding. In all cases layers were initialized using the Haiku defaults. The output of the second convolution was fed into linear layers which parameterized the variational mean, and the value of the diagonal variational standard deviation. After each of these linear layers we used layer normalization (Lei Ba et al., 2016) with the standard additional scale and translation parameters to ensure no loss in expressivity. Positivity of the standard deviation parameter was ensured by transforming the real valued vector through a softplus non-linearity.

The VAE decoder architecture was as follows. We affinely projected the 30 dimensional latents into a $7 \times 7 \times 32 = 1568$ dimensional space, which we then followed again by layer normalization. We then reshaped this vector to $7 \times 7 \times 32$ ready for feeding into the deconvolutional layers. There were three such deconvolutional layers separated by two rectified linear non-linearity layers. The first deconvolutional layer had kernel shape 3×3 , a stride of 2×2 and 64 channels. The second




Figure 4: Larger version of results from the four different examples. Cyan lines denote gold standard values of the log normalizing constant. In (c) and (d) green horizontal lines denote the median value for an importance sampling estimate based on variational inference. Note that in (d) the small AFT error bars can make it difficult to see- it can be found next to the gold standard value in each case.

deconvolutional layer had kernel shape 3×3 , a stride of 2×2 and 32 channels. The third and final deconvolutional layer had kernel shape 3×3 , a stride of 1×1 and one output channel to match the target image shape.

Next we describe the training algorithm for learning the VAE, which is distinct from the experiments we ran evaluating the latent space. We used the ADAM optimizer with a small learning rate of 5×10^{-5} . We used gradient clipping with value of 10^5 and trained for 5×10^5 iterations. We used the reparameterization trick (Kingma and Welling, 2014; Rezende et al., 2014) for estimating the expectation of the log likelihood and the analytic expression for the KL divergence between the normal approximating distribution and the standard normal prior.

G.3.2. ADDITIONAL DETAILS AND DISCUSSION OF VAE EXPERIMENTAL RESULTS

Figure 5 shows reconstructions and samples for the trained autoencoder. Note that we adopt the common practice of plotting the pixel probabilities for the reconstructions and the samples. In the latter case, this has the effect of making the samples look smoother than if we sampled pixel values.

As discussed in the main text, the experiments on the test set were divided into two types. One group of experiments (Group A) was run over the whole test set. These experiments were used to verify the quality of the autoencoder and identify interesting/challenging examples for the more detailed study. The detailed study (Group B) mirrored the pattern of the other examples.

We now give more detail of the Group A experiments. For each image in the binarized MNIST test set we ran SMC once with 1000 temperatures. The variance in the log likelihood estimate from these runs was small relative to our the variational inference experiments. For each image in the binarized MNIST test set we did four repeats of the unamortized normal variational approximation with diagonal covariance and associated importance estimates. Figure 8 shows log likelihood estimates for SMC and the VI based importance sampler. We have averaged the VI estimates over the four repeats. The plot shows strong correlation- the variation in log likelihood from different digits is the chief source of variance as reported by (Wu et al., 2017). The four VI repeats were then used to estimate the mean absolute error in the variational approximation relative to SMC as shown in Figure 7. These errors were used to identify challenging inference digits. In particular we chose digits on the 99.8-th percentile of error, which implies about 1 in 500 digits is more challenging than the ones we chose. The three digits selected- one used for the main paper and two for repeats in this Appendix are shown in Figure 9. The experiment repeats are shown in Figures 9 and 10.

As discussed in the main text we found that the variational inference performed well considering the simple form of the approximate posterior. This is likely a consequence of the fact that the training objective favours posteriors that are well matched by the variational approximation. Whilst there can be bias from this effect relative to using an exact marginal likelihood (Hoffman, 2017) the ability of a flexible generative network to adapt to this constraint can make the effect less extreme than in simpler statistical models (Turner and Sahani, 2011).

G.4. Additional details of Cox process example

We experimented with transforming the problem to a whitened representation \tilde{x} where $x = L\tilde{x} + \mu$ and L is the Cholesky decomposition of K (Beskos et al., 2011; Neal, 2011). This can also be viewed as choosing a non-identity pre-conditioner or mass matrix. We found that this rendered the problem easier to the extent of no longer being a challenging benchmark, and that in this circumstance there was little that a NF could contribute. To maintain the difficulty of the baseline we therefore focused our comparison on the unwhitened space of the latent function.

The covariance K is taken to be $K(u, v) = \sigma^2 \exp\left(-\frac{\|u-v\|_2}{M\beta}\right)$, where $\sigma^2 = 1.91$ and the mean vector has a constant value of $\log(126) - \sigma^2$, and $a = 1/M^2$, matching (Møller et al., 1998).

The Adam learning rate was 10^{-2} the training period J was 500. AFT flow estimation particle numbers were $N_{\text{train}} = N_{\text{val}} = 2000$.

For each transition, we used ten iterations of Hamiltonian Monte Carlo with 10 leapfrog steps per iteration. With the total time number of time steps scaled to be in the interval $[0, 1]$ with 0 being the initial distribution and 1 the final distribution, the slice sampling step sizes were interpolated between the times $[0., 0.25, 0.5, 1.]$ using the step sizes $[0.3, 0.3, 0.2, 0.2]$. The gold standard value for the normalizing constant shown in Figure 2 was found using 1000 repeats of SMC with 1000 temperatures and using the Cholesky whitening.




Figure 5: Variational Autoencoder samples and reconstructions. Top line: Images from the binarized MNIST dataset. Middle line: Reconstructions of the same images. Bottom line. Free samples from the model- these have no particular relationship to the images above them.




Figure 6: Estimated log likelihood from long run and SMC and mean estimated log likelihood from VI proposal for the binarized MNIST test set. The VI means are estimated based on four repeats of the VI training and sampling. The dotted line shows equality for reference.




Figure 7: Histogram showing the mean absolute error of the variational proposal importance sampler for the binarized MNIST test set relative to a long run of SMC. The means are estimated based on four repeats of the VI training and sampling. The cyan line shows the discrepancy of the challenging digit chosen for detailed investigation in the main text



Figure 8: Binarized MNIST digits used in detailed experiments. These digits were selected as challenging for the variational approximation. Left: The digit used in the main paper. Centre and Right: Respectively the digits used for the repeat experiments in Figure 9 and 10.




Figure 9: Results for a second challenging digit VAE likelihood, selected from the tail of plot 7 similar to the result in the main text.




Figure 10: Results for a third challenging digit VAE likelihood, selected from the tail of plot 7 similar to the result in the main text.




Figure 11: Left: data points for the spatial point process example. Right: Inferred posterior rate for the Cox process with a 40×40 discretization using AFT. The plot area is defined as having unit area. The samples were taken from a single run of AFT with 30 temperatures and other parameters as discussed in the text.

We found that to obtain best performance for AFT in this example, it was important to have sufficient HMC updates per transition.

An example of the samples produced using AFT can be seen in Figure 11.