



HAL
open science

Full error analysis of policy gradient learning algorithms for exploratory linear quadratic mean-field control problem in continuous time with common noise

Noufel Frikha, Huy en Pham, Xuanye Song

► To cite this version:

Noufel Frikha, Huy en Pham, Xuanye Song. Full error analysis of policy gradient learning algorithms for exploratory linear quadratic mean-field control problem in continuous time with common noise. 2024. hal-04665587

HAL Id: hal-04665587

<https://hal.science/hal-04665587v1>

Preprint submitted on 31 Jul 2024

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 destin e au d p t et   la diffusion de documents scientifiques de niveau recherche, publi s ou non,  manant des  tablissements d'enseignement et de recherche fran ais ou  trangers, des laboratoires publics ou priv s.

Full error analysis of policy gradient learning algorithms for exploratory linear quadratic mean-field control problem in continuous time with common noise

Noufel FRIKHA* Huy en PHAM † Xuanye SONG‡

Abstract

We consider reinforcement learning (RL) methods for finding optimal policies in linear quadratic (LQ) mean field control (MFC) problems over an infinite horizon in continuous time, with common noise and entropy regularization. We study policy gradient (PG) learning and first demonstrate convergence in a model-based setting by establishing a suitable gradient domination condition. Next, our main contribution is a comprehensive error analysis, where we prove the global linear convergence and sample complexity of the PG algorithm with two-point gradient estimates in a model-free setting with unknown parameters. In this setting, the parameterized optimal policies are learned from samples of the states and population distribution. Finally, we provide numerical evidence supporting the convergence of our implemented algorithms.

Key words: Mean-field control; reinforcement learning; linear-quadratic; two-point gradient estimation; Polyak-Lojasiewicz inequality; gradient descent; sample complexity.

Contents

1	Introduction	3
2	Problem formulation	5
2.1	Setup and preliminaries	5
2.2	Model free perspective and reparametrization	7

*Universit e Paris 1 Panth on-Sorbonne, Centre d’Economie de la Sorbonne (CES), 106 Boulevard de l’H pital, 75642 Paris Cedex 13, noufel.frikha at univ-paris1.fr. The work of this author has benefited from the support of the Institut Europlace de Finance.

†LPSM, Universit e Paris Cit e and Sorbonne University, pham at lpsm.paris. The work of this author is partially supported by the BNP-PAR Chair “Futures of Quantitative Finance”, and the Chair Finance & Sustainable Development / the FiME Lab (Institut Europlace de Finance)

‡LPSM, Universit e Paris Cit e, xsong at lpsm.paris; Research assistant at Hong Kong Polytechnic University.

3	Model-based PG algorithm	10
3.1	Gradient domination condition	10
3.2	Global convergence of model-based PG algorithm	11
4	Model-free PG algorithm	14
4.1	Notations	14
4.2	Model-free PG algorithm with population simulator	15
5	Numerical example	19
5.1	Model-based algorithm	19
5.2	Model-free algorithm	20
A	Explicit solution to LQ MFC	22
B	Proof of Propositions 2.1 and 2.2	27
B.1	Proof of Proposition 2.1	27
B.2	Proof of the Proposition 2.2	29
C	Proofs of results for model-based algorithms	30
C.1	Proof of Proposition 3.1	30
C.2	Proof of Theorem 3.2	33
C.3	Proof of Proposition 3.3	34
C.3.1	Auxiliary results	35
C.3.2	Lipschitz regularity of Σ_θ , $\hat{\Sigma}_\zeta$, K_θ and Λ_ζ	38
C.3.3	Proof of Proposition 3.3	40
D	Proofs of the results for the model-free algorithm	40
D.1	Strategy of proof of Theorem 4.1	41
D.2	Proof of Theorem 4.1	42
D.3	Analysis of the single step GD algorithm	45
D.4	Error analysis of J and its gradient	47
D.4.1	Proof of Proposition D.2	48
D.4.2	Proof of Proposition D.3	49
D.4.3	Proof of Proposition D.4	52
D.4.4	Proof of Proposition D.5	55
D.4.5	Proof of Proposition D.6.	62

1 Introduction

The last decade has seen significant advances in solving optimal control of dynamical systems in unknown environments using reinforcement learning (RL) methods. The essence of RL is to learn optimal decisions through trial and error, which involves repeatedly trying a policy, observing the state, receiving and evaluating the reward, and subsequently improving the policy. There are two main approaches in RL: (i) Q -learning, which is based on dynamic programming, and (ii) policy gradient (PG), which is based on the parametrization of policies. A key feature of RL is the *exploration* of the unknown environment to broaden the search space, achievable through randomized policies. RL is a very active branch of machine learning. For an overview of this field in the discrete-time setting, we refer to the second edition of the monograph [20], and for recent advances in the continuous-time setting, see [14] and [15].

Mean-field control (MFC), also known as the *McKean-Vlasov* (MKV) control problem, is a class of stochastic control problems that focuses on the study of large population models of interacting agents who cooperate and act for collective welfare according to a central decision-maker (or social planner). This field has attracted growing interest over the last decade, resulting in a substantial body of literature on both its theory and its various applications in economics, finance, population dynamics, social sciences, and herd behavior. For a detailed treatment of the topic, we refer to the seminal two-volume monograph [4, 5].

RL for MFC has recently attracted attention in the research community, see e.g. [6], [11], [1], [9], [19]. The challenge lies in accurately learning optimal policies and value functions defined on the infinite-dimensional space of probability measures. The mathematical understanding and convergence analysis of these RL algorithms are still in their infancy.

In this paper, we aim to address questions surrounding convergence and sample complexity, focusing on policy gradient methods in RL within the context of infinite horizon linear quadratic (LQ) MFC with common noise for continuous time systems. The LQ problem is indeed the cornerstone of optimal control theory due to its tractability and can be viewed as an approximation of more general nonlinear control problems. To encourage exploration in unknown environments, we employ randomized policies and add entropy regularization, following the approach of recent papers [14], [12], [22], [10], [21].

Our main contributions. Our paper proposes and analyzes convergent PG algorithms to solve infinite horizon exploratory LQ MFC problems in a continuous time setting, with common noise and entropy regularization.

- Our first contribution is to derive the explicit form of the optimal solution using coupled algebraic Riccati equations, thereby generalizing the results in [2] to include entropy regularization for randomized policies (Theorem A.1). Motivated by the explicit form of the optimal randomized policy, we reformulate the LQ MFC problem into a minimization

problem over Gaussian policies. The mean of each Gaussian policy being linear in the state and conditional mean with respect to the common noise is parameterised using two matrix-valued coefficients $\Theta = (\theta, \zeta)$. The parameterized cost function is shown to be smooth and satisfy a gradient domination condition, also known as the Polyak-Lojasiewicz inequality (Propositions 3.1 and 3.2) following the approach of [8] and [23]. Such inequality is known to be crucial to ensure the convergence of PG algorithms in non-convex landscape.

- We then propose and study PG methods in both exact and model-free settings. Our work provides theoretical guarantees of convergence for the gradient descent (GD) algorithms with suitable step sizes (Theorems 3.1 and 4.1). In the model-free case, where the exact gradient is unavailable, we adapt the two-point gradient estimation method of [16] to our mean-field setting by relying on samples of discrete-time trajectories and population distributions. For the first time to the best of our knowledge, we provide a comprehensive error analysis accounting respectively for the error of perturbation with respect to the exact expected functional cost, for the horizon truncation, for the time and particle discretizations, for the statistical error and finally for the optimization error from gradient iterations, demonstrating global linear convergence with polynomial computational sample complexities.

Related works. The closest papers related to our work are [6] and [23]. In [6], the authors consider an LQ MFC problem with common noise in a discrete time setting and prove the convergence of PG algorithms for deterministic policies in both model-based and model-free settings. The paper [23] addresses an infinite-horizon time average LQ MFC control problem in a continuous time setting without an entropy regularizer and demonstrates the convergence of PG with deterministic policies in the exact model-based setting, using a varying step size at each iteration. Finally, it is worth mentioning that our proofs extend the arguments presented in [8] and [13] for the discrete time setting, covering both finite and infinite horizons, as well as those in [16] and [3] for the continuous time setting. In these references, the authors demonstrated convergence results for standard LQ problems.

Outline. The paper is organized as follows. In Section 2, we formulate the exploratory LQ MFC problem in continuous time with common noise, provide the theoretical optimal policy, and discuss parameterization in the model-free case. In Section 3, we demonstrate the convergence of the model-based gradient descent algorithm using the gradient domination condition. Section 4 presents the gradient estimation algorithm employing the population simulator and develops the convergence analysis of the model-free gradient descent algorithm. In Section 5, we provide numerical experiments that illustrate our convergence results for both the model-based and model-free algorithms. The proofs of all results are included in the Appendix.

Notations.

- We denote by $x \cdot y$ the scalar product between the two vectors x, y , and by $M : N = \text{tr}(MN^\top)$ the inner product of the two matrices M, N with compatible dimensions, where N^\top is the transpose matrix of N . The Frobenius norm of a matrix A is defined by $\|A\|_F := \sqrt{A : A}$.
- \mathbb{S}^d is the set of symmetric $d \times d$ matrices, and \mathbb{S}_+^d (resp. $\mathbb{S}_{>+}^d$) is the set of nonnegative (resp. positive definite) matrices in \mathbb{S}^d . The partial order \geq on \mathbb{S}^d is defined as: $M \geq N$ if $M - N \in \mathbb{S}_+^d$. We also write $M > 0$ to mean that $M \in \mathbb{S}_{>+}^d$.

2 Problem formulation

2.1 Setup and preliminaries

The linear dynamics of the mean-field state equation with randomized controls and common noise is described by

$$dX_t = \left[BX_t + \bar{B}\mathbb{E}_0[X_t] + D \int a\pi_t(da) \right] dt + \gamma dW_t + \gamma_0 dW_t^0, \quad (2.1)$$

on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ supporting two independent Brownian motions W (the *idiosyncratic noise*), and W^0 (the *common noise*), of dimension d and d_0 . For convenience, we choose the probability space in the product form $(\Omega^0 \times \Omega^1, \mathcal{F}^0 \otimes \mathcal{F}^1, \mathbb{P}^0 \otimes \mathbb{P}^1)$, and denote by $\mathbb{F}^1 = (\mathcal{F}_t^1)_{t \geq 0}$ the right-continuous \mathbb{P}^1 -completion of the canonical filtration generated by W , and by $\mathbb{F}^0 = (\mathcal{F}_t^0)_{t \geq 0}$ the right-continuous \mathbb{P}^0 -completion of the canonical filtration generated by W^0 . The initial condition X_0 is an \mathbb{R}^d -random variable, which is \mathcal{G} -measurable, where \mathcal{G} is a σ -algebra independent of (W, W^0) . We denote by $\mathbb{F} = \mathbb{F}^0 \vee \mathbb{F}^1 \vee \mathcal{G}$.

Here, $\mathbb{E}_0[\cdot]$ stands for the conditional expectation given \mathbb{F}^0 , B, \bar{B} are constant matrices in $\mathbb{R}^{d \times d}$, D is a constant matrix in $\mathbb{R}^{d \times m}$, γ is a constant matrix in $\mathbb{R}^{d \times d}$, γ^0 is a constant matrix in $\mathbb{R}^{d \times d_0}$, and in the sequel, we shall denote by $\hat{B} := B + \bar{B}$.

The randomized control $\pi = (\pi_t)_t$ is an \mathbb{F} -progressively measurable process in $\mathcal{P}_2(\mathbb{R}^m)$, the set of probability measures on the action space $A = \mathbb{R}^m$ with a finite second order moment. We shall consider randomized controls π with densities $a \mapsto \mathbf{p}_t(a)$, $t \geq 0$.

The infinite horizon LQ MFC problem consists in minimizing over such randomized controls π the quadratic cost functional with entropy regularizer of parameter $\lambda > 0$:

$$\begin{aligned} J(\pi; \lambda) = & \mathbb{E} \left[\int_0^\infty e^{-\beta t} \left(X_t^\top Q X_t + \mathbb{E}_0[X_t]^\top \bar{Q} \mathbb{E}_0[X_t] \right. \right. \\ & \left. \left. + \int a^\top R a \pi_t(da) + \lambda \int \log \mathbf{p}_t(a) \pi_t(da) \right) dt \right]. \end{aligned}$$

Here Q, \bar{Q} are constant matrices in \mathbb{S}^d such that $Q > 0, \hat{Q} := Q + \bar{Q} > 0$ and $R \in \mathbb{S}_{>+}^m$. Notice that the cost functional is written equivalently as

$$J(\boldsymbol{\pi}; \lambda) = \mathbb{E} \left[\int_0^\infty e^{-\beta t} \left((X_t - \mathbb{E}_0[X_t])^\top Q (X_t - \mathbb{E}_0[X_t]) + \mathbb{E}_0[X_t]^\top \hat{Q} \mathbb{E}_0[X_t] \right. \right. \quad (2.2) \\ \left. \left. + \int a^\top R a \boldsymbol{\pi}_t(\mathrm{d}a) + \lambda \int \log \boldsymbol{p}_t(a) \boldsymbol{\pi}_t(\mathrm{d}a) \right) \mathrm{d}t \right].$$

Assumption 2.1. *The two following Algebraic Riccati Equations (ARE) for $K \in \mathbb{S}^d$ and $\Lambda \in \mathbb{S}^d$ respectively admit a unique positive definite solutions:*

$$-\beta K + KB + B^\top K + Q - KDR^{-1}D^\top K = 0, \quad (2.3)$$

$$-\beta \Lambda + \Lambda \hat{B} + \hat{B}^\top \Lambda + \hat{Q} - \Lambda DR^{-1}D^\top \Lambda = 0. \quad (2.4)$$

Remark 2.1. *According to Section 6 in [2], the condition: $Q > 0, \hat{Q} = Q + \bar{Q} > 0, R > 0$, guarantees the existence of positive definite solutions to (2.3)-(2.4) and Assumption 2.1 is specifically for ensuring uniqueness.*

Under Assumption 2.1, the optimal randomized control, solution to (2.2) is given in feedback policy form as $\boldsymbol{\pi}_t^* = \pi^*(\cdot | X_t^* - \mathbb{E}_0[X_t^*], \mathbb{E}_0[X_t^*])$, where $\pi^*(\cdot | y, z)$ is the normal distribution

$$\pi^*(\cdot | y, z) = \mathcal{N} \left(-R^{-1}D^\top K y - R^{-1}D^\top \Lambda z; \frac{\lambda}{2} R^{-1} \right), \quad y, z \in \mathbb{R}^d, \quad (2.5)$$

and where $(X_t^*)_{t \in [0, T]}$ is the state process with randomized control $\boldsymbol{\pi}^*$ and (K, Λ) is the unique positive definite solution to (2.3)-(2.4). Moreover, the optimal cost is given by

$$J(\boldsymbol{\pi}^*; \lambda) = K : M + \Lambda : \hat{M} + v(\lambda) \quad (2.6)$$

where

$$M := \text{Var}(X_0) + \frac{1}{\beta} \gamma \gamma^\top, \quad \hat{M} := \mathbb{E}[X_0] \mathbb{E}[X_0]^\top + \frac{1}{\beta} \gamma_0 \gamma_0^\top, \quad (2.7) \\ v(\lambda) := \frac{1}{\beta} \left(-\frac{\lambda m}{2} \log(\pi \lambda) + \frac{\lambda}{2} \log |\det(R)| \right).$$

Here, $\text{Var}(X_0)$ denotes the covariance matrix of X_0 . In Appendix A, we state and prove this result in a more general case.

Throughout the remainder of this paper, we shall also assume that the following assumption is in force.

Assumption 2.2. *M, \hat{M} belongs to $\mathbb{S}_{>+}^d$, i.e. $\sigma_{\min}(M), \sigma_{\min}(\hat{M}) > 0$, where $\sigma_{\min}(\cdot)$ denotes the smallest eigenvalue of a square matrix.*

Remark 2.2. *The above assumption for M is satisfied when $\text{Var}(X_0) > 0$. As for \hat{M} , denoting by $\bar{\gamma}_0 = [\mathbb{E}[X_0], \frac{1}{\sqrt{\beta}} \gamma_0] \in \mathbb{R}^{d \times (d_0+1)}$, the augmented matrix formed by $\mathbb{E}[X_0]$ and γ_0 , we have $\hat{M} = \bar{\gamma}_0 \bar{\gamma}_0^\top$. Thus, $\hat{M} > 0$ if and only if $\bar{\gamma}_0$ is of rank d , which requires that that $d_0 \geq d - 1$.*

2.2 Model free perspective and reparametrization

In this section, we are interested in the model-free setting for the linear mean-field dynamics of state process, i.e., when the parameters $B, \bar{B}, D, \gamma, \gamma_0$ in (2.1) are unknown, and so the optimal policy in (2.5) cannot be implemented from the unique solution to the Riccati system (2.3)-(2.4).

Motivated by the Gaussian distribution of the optimal randomized policy, whose mean is a linear combination of $Y_t = X_t - \mathbb{E}_0[X_t]$ and $Z_t = \mathbb{E}_0[X_t]$, we propose the following parameterization of the randomized policy:

$$\pi^\Theta(\cdot|y, z) = \mathcal{N}(\theta y + \zeta z; \frac{\lambda}{2} R^{-1}), \quad y, z \in \mathbb{R}^d,$$

where $\Theta = (\theta, \zeta) \in (\mathbb{R}^{m \times d})^2$ are the two-parameter matrices to be optimized. The density of the parametrized randomized policy is explicitly given by

$$a \in \mathbb{R}^m \mapsto p^\Theta(y, z, a) = \sqrt{\frac{\det(R)}{(\pi\lambda)^m}} \exp\left(-\frac{1}{\lambda}(a - (\theta y + \zeta z))^\top R(a - (\theta y + \zeta z))\right).$$

The associated dynamics of the parametrized process (X_t^Θ) , starting from $X_0^\Theta = X_0$, is given by

$$dX_t^\Theta = \left[BX_t^\Theta + \bar{B}\mathbb{E}_0[X_t^\Theta] + D \int a \pi_t^\Theta(da) \right] dt + \gamma dW_t + \gamma_0 dW_t^0, \quad (2.8)$$

where $\pi_t^\Theta = \pi^\Theta(\cdot|X_t^\Theta - \mathbb{E}_0[X_t^\Theta], \mathbb{E}_0[X_t^\Theta])$, with density $\mathbf{p}_t^\Theta(a) = p^\Theta(X_t^\Theta - \mathbb{E}_0[X_t^\Theta], \mathbb{E}_0[X_t^\Theta], a)$. The corresponding cost function, now defined as a function on $(\mathbb{R}^{m \times d})^2$, is expressed (with a slight abuse of notation) as follows:

$$\begin{aligned} J(\theta, \zeta; \lambda) &:= J(\pi^\Theta; \lambda) \\ &= \mathbb{E} \left[\int_0^\infty e^{-\beta t} \left((X_t^\Theta - \mathbb{E}_0[X_t^\Theta])^\top Q (X_t^\Theta - \mathbb{E}_0[X_t^\Theta]) + \mathbb{E}_0[X_t^\Theta]^\top \hat{Q} \mathbb{E}_0[X_t^\Theta] \right. \right. \\ &\quad \left. \left. + \int a^\top R a \pi_t^\Theta(da) + \lambda \int \log \mathbf{p}_t^\Theta(a) \pi_t^\Theta(da) \right) dt \right] \end{aligned} \quad (2.9)$$

where $\hat{Q} := Q + \bar{Q}$, with the objective of minimizing the parameterized cost function $J(\theta, \zeta; \lambda)$ over $\Theta = (\theta, \zeta) \in (\mathbb{R}^{m \times d})^2$. It is clear from (2.5) that $\inf_{\pi} J(\pi; \lambda) = \inf_{\Theta} J(\theta, \zeta; \lambda)$, and the solution to this minimization problem is given by

$$\theta^* = -R^{-1} D^\top K, \quad \zeta^* = -R^{-1} D^\top \Lambda, \quad (2.10)$$

where $(K, \Lambda) \in (\mathbb{S}_{>+}^d)^2$ is the solution to the ARE (2.3) and (2.4), ensuring that $J(\theta^*, \zeta^*; \lambda) = J(\pi^*; \lambda)$.

Let us introduce the parametrized auxiliary processes

$$Y_t^\Theta = X_t^\Theta - \mathbb{E}_0[X_t^\Theta], \quad Z_t^\Theta = \mathbb{E}_0[X_t^\Theta], \quad t \geq 0, \quad \Theta = (\theta, \zeta) \in (\mathbb{R}^{m \times d})^2.$$

By observing from the definition of $\pi_t^\Theta = \pi^\Theta(\cdot | Y_t^\Theta, Z_t^\Theta)$ that

$$\int a \pi_t^\Theta(da) = \theta Y_t^\Theta + \zeta Z_t^\Theta, \quad \mathbb{E}_0 \left[\int a \pi_t^\Theta(da) \right] = \zeta Z_t^\Theta,$$

we see from (2.8) that the dynamics of (Y^Θ, Z^Θ) is decoupled and governed by

$$\begin{cases} dY_t^\Theta &= (B + D\theta)Y_t^\Theta dt + \gamma dW_t, \\ dZ_t^\Theta &= (\hat{B} + D\zeta)Z_t^\Theta dt + \gamma^0 dW_t^0, \end{cases} \quad t \geq 0 \quad (2.11)$$

with $Y_0^\Theta := Y_0 = X_0 - \mathbb{E}[X_0]$ and $Z_0^\Theta := Z_0 = \mathbb{E}[X_0]$. Notice that Y^Θ (resp. Z^Θ) depends on Θ only via the first (resp. second) component parameter θ (resp. ζ), and we shall then write $Y^\theta = Y^\Theta$, $Z^\zeta = Z^\Theta$.

Moreover, by noting that

$$\begin{aligned} \int a^\top R a \pi_t^\Theta(da) &= \theta Y_t^\theta + \zeta (Z_t^\zeta)^\top R \theta Y_t^\theta + \zeta Z_t^\zeta + \frac{\lambda m}{2} \\ \int \log \mathbf{p}_t^\Theta(a) \pi_t^\Theta(da) &= -\frac{m}{2}(1 + \log(2\pi)) - \frac{m}{2} \log \left| \frac{\lambda}{2 \det(R)} \right| \end{aligned}$$

the parametrized cost in (2.9) can be written as a quadratic function of (Y_t^θ, Z_t^ζ) , namely:

$$J(\theta, \zeta; \lambda) = \mathbb{E} \left[\int_0^\infty e^{-\beta t} \left((Y_t^\theta)^\top (Q + \theta^\top R \theta) Y_t^\theta + (Z_t^\zeta)^\top (\hat{Q} + \zeta^\top R \zeta) Z_t^\zeta \right) dt \right] + v(\lambda),$$

recalling that $v(\lambda)$ is given by (2.7). An important observation is that it can be decomposed as follows

$$J(\theta, \zeta; \lambda) = J_1(\theta) + J_2(\zeta) + v(\lambda),$$

where

$$J_1(\theta) := (Q + \theta^\top R \theta) : \Sigma_\theta, \quad J_2(\zeta) := (\hat{Q} + \zeta^\top R \zeta) : \hat{\Sigma}_\zeta, \quad (2.12)$$

with

$$\Sigma_\theta := \int_0^\infty e^{-\beta t} \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] dt, \quad \hat{\Sigma}_\zeta := \int_0^\infty e^{-\beta t} \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top] dt, \quad (2.13)$$

so that the minimization over $\Theta = (\theta, \zeta)$ amounts to separate minimization problems:

$$\inf_{\theta \in \mathbb{R}^{m \times d}} J_1(\theta), \quad \text{and} \quad \inf_{\zeta \in \mathbb{R}^{m \times d}} J_2(\zeta). \quad (2.14)$$

In fact, we now demonstrate that the two aforementioned minimization problems over θ and ζ can be reduced to suitable smaller sets defined by

$$\mathcal{S} = \left\{ \theta \in \mathbb{R}^{m \times d} : B - \frac{\beta}{2} I_d + D\theta \text{ is stable} \right\}, \quad \hat{\mathcal{S}} = \left\{ \zeta \in \mathbb{R}^{m \times d} : \hat{B} - \frac{\beta}{2} I_d + D\zeta \text{ is stable} \right\},$$

recalling that $\hat{B} = B + \bar{B}$, on which the problem is well-posed.

Lemma 2.1. *The minimizers in (2.10) satisfy $\theta^* \in \mathcal{S}$, $\zeta^* \in \hat{\mathcal{S}}$, and thus*

$$\inf_{\theta \in \mathbb{R}^{m \times d}} J_1(\theta) = \inf_{\theta \in \mathcal{S}} J_1(\theta), \quad \text{and} \quad \inf_{\zeta \in \mathbb{R}^{m \times d}} J_2(\theta) = \inf_{\zeta \in \hat{\mathcal{S}}} J_2(\theta).$$

Proof. The ARE (2.3) for K rewrites as

$$(B - \frac{\beta}{2} I_d - DR^{-1}D^\top K)^\top K + K(B - \frac{\beta}{2} I_d - DR^{-1}D^\top K) = -(Q + KDR^{-1}D^\top K). \quad (2.15)$$

Since $Q > 0$ then $Q + KDR^{-1}D^\top K > 0$, hence all the eigenvalues of the matrix on the r.h.s. of (2.15) are strictly negative. Moreover, since $K > 0$, if the matrix $(B - \frac{\beta}{2} I_d - DR^{-1}D^\top K)$ has a non-negative real part eigenvalue, then the largest eigenvalue of the matrix on the l.h.s. of (2.15) is non-negative and therefore cannot be equal to the matrix on the r.h.s. It follows that $B - \frac{\beta}{2} I_d - DR^{-1}D^\top K$ is stable, which means tht $\theta^* \in \mathcal{S}$.

Similarly, by using the ARE (2.4) for Λ , we show that $\hat{B} - \frac{\beta}{2} I_d - DR^{-1}D^\top \Lambda$ is also stable, i.e., $\zeta^* \in \hat{\mathcal{S}}$ \square

We now state a characterization of the two matrices Σ_θ and $\hat{\Sigma}_\zeta$ defined in (2.13).

Proposition 2.1. *For all $\theta \in \mathcal{S}, \zeta \in \hat{\mathcal{S}}$, the matrices Σ_θ and $\hat{\Sigma}_\zeta$ are well-defined in \mathbb{S}^d and are the unique positive definite solution to the following Algebraic Lyapunov Equations (ALE):*

$$\begin{aligned} -\beta \Sigma_\theta + (B + D\theta) \Sigma_\theta + \Sigma_\theta (B + D\theta)^\top + M &= 0, \\ -\beta \hat{\Sigma}_\zeta + (\hat{B} + D\zeta) \hat{\Sigma}_\zeta + \hat{\Sigma}_\zeta (\hat{B} + D\zeta)^\top + \hat{M} &= 0. \end{aligned} \quad (2.16)$$

Proof. Cf Appendix B.1 \square

We conclude this section by providing an alternate useful expression of the cost functions J_1 and J_2 defined by (2.12) on which the minimization over $\theta \in \mathcal{S}$ and $\zeta \in \hat{\mathcal{S}}$ will be performed.

Proposition 2.2. *For all $\theta \in \mathcal{S}, \zeta \in \hat{\mathcal{S}}$, we have*

$$J_1(\theta) = K_\theta : M, \quad J_2(\zeta) = \Lambda_\zeta : \hat{M},$$

where K_θ and Λ_ζ are the unique elements in $\mathbb{S}_{>+}^d$ solutions to the ALE

$$-\beta K_\theta + (B + D\theta)^\top K_\theta + K_\theta (B + D\theta) + Q + \theta^\top R \theta = 0 \quad (2.17)$$

$$-\beta \Lambda_\zeta + (\hat{B} + D\zeta)^\top \Lambda_\zeta + \Lambda_\zeta (\hat{B} + D\zeta) + \hat{Q} + \zeta^\top R \zeta = 0. \quad (2.18)$$

Proof. Cf Appendix B.2 \square

Remark 2.3. *Note that, according to (2.6), we have $K_{\theta^*} = K$ and $\Lambda_{\zeta^*} = \Lambda$.*

3 Model-based PG algorithm

In this section, we establish the convergence of the gradient descent algorithm towards the optimal parameters (θ^*, ζ^*) assuming full knowledge of the model parameters, thus allowing for the exact computation of the gradient. This foundation will facilitate the learning of optimal parameters in the model-free setting in the subsequent section, despite the nonconvex optimization framework. Furthermore, the gradient descent method in the model-based scenario offers an alternative to solving the Riccati system (2.3)-(2.4), which is computationally intensive, while the convergence result (Theorem 3.1) remains dimension-free.

Given $(\theta_{[0]}, \zeta_{[0]}) \in \mathcal{S} \times \hat{\mathcal{S}}$, the updating rule at step $k \geq 0$ for the exact gradient descent of the minimization problem (2.14) is given by

$$\begin{cases} \theta_{[k+1]} &= \theta_{[k]} - \rho \nabla J_1(\theta_{[k]}), \\ \zeta_{[k+1]} &= \zeta_{[k]} - \rho \nabla J_2(\zeta_{[k]}), \end{cases} \quad (3.1)$$

where $\rho > 0$ is the constant step size (learning rate), and $\nabla J_1(\theta)$ and $\nabla J_2(\zeta)$ are the gradients of J_1 and J_2 with respect to their respective parameters.

3.1 Gradient domination condition

We first provide an explicit formula for the gradient that will be used for the implementation of the (exact) gradient descent rule.

Proposition 3.1 (Expression of the gradients). *For $\theta \in \mathcal{S}$, $\zeta \in \hat{\mathcal{S}}$, it holds*

$$\nabla J_1(\theta) = 2E_\theta \Sigma_\theta, \quad \text{and} \quad \nabla J_2(\zeta) = 2\hat{E}_\zeta \hat{\Sigma}_\zeta,$$

where $E_\theta = R\theta + D^\top K_\theta$ and $\hat{E}_\zeta = R\zeta + D^\top \Lambda_\zeta$.

Proof. Cf. Appendix C.1 □

Next, to assess the convergence of the aforementioned gradient descent algorithm in this nonconvex landscape, we will demonstrate that the cost functions J_1 and J_2 satisfy a Polyak-Lojasiewicz (PL) inequality on \mathcal{S} and $\hat{\mathcal{S}}$, respectively, also known as the gradient domination condition.

Proposition 3.2 (Gradient domination). *There exist two positive constants $\kappa_1, \kappa_2 > 0$ such that for any $\theta \in \mathcal{S}, \zeta \in \hat{\mathcal{S}}$,*

$$J_1(\theta) - J_1(\theta^*) \leq \kappa_1 \|\nabla J_1(\theta)\|_F^2, \quad \text{and} \quad J_2(\zeta) - J_2(\zeta^*) \leq \kappa_2 \|\nabla J_2(\zeta)\|_F^2.$$

More precisely, the constants κ_1 and κ_2 are given by

$$\kappa_1 = \frac{\|\Sigma_{\theta^*}\|_F}{4\sigma_{\min}(R)\sigma_{\min}^2(M)}, \quad \kappa_2 = \frac{\|\hat{\Sigma}_{\zeta^*}\|_F}{4\sigma_{\min}(R)\sigma_{\min}^2(\hat{M})}.$$

Proof. Cf. Appendix C.2 □

3.2 Global convergence of model-based PG algorithm

For fixed $\ell, \hat{\ell} \in \mathbb{R}_+$, we define the level subsets of \mathcal{S} and $\hat{\mathcal{S}}$ as follows:

$$\mathcal{S}(\ell) = \left\{ \theta \in \mathcal{S} : J_1(\theta) \leq \ell \right\} \quad \text{and} \quad \hat{\mathcal{S}}(\hat{\ell}) = \left\{ \zeta \in \mathcal{S} : J_2(\zeta) \leq \hat{\ell} \right\}.$$

In the sequel, we shall naturally restrict to $\ell > J_1(\theta^*)$ and $\hat{\ell} > J_2(\zeta^*)$ to avoid empty level sets.

We now show that both gradient maps $\theta \mapsto \nabla J_1(\theta)$ and $\zeta \mapsto \nabla J_2(\zeta)$ are Lipschitz-continuous on $\mathcal{S}(a)$ and $\hat{\mathcal{S}}(\hat{a})$ respectively.

Proposition 3.3. *There exist explicit positive constants*

$$\begin{aligned} L(\ell) &= L(\ell; \sigma_{\min}(Q), \sigma_{\min}(M), \|R\|_F, \|D\|_F, J_1(\theta^*)), \\ \hat{L}(\hat{\ell}) &= \hat{L}(\hat{\ell}; \sigma_{\min}(\hat{Q}), \sigma_{\min}(\hat{M}), \|R\|_F, \|D\|_F, J_2(\zeta^*)), \end{aligned} \quad (3.2)$$

such that $\theta \mapsto \nabla J_1(\theta)$ is $L(\ell)$ -Lipschitz continuous on $\mathcal{S}(\ell)$ and $\zeta \mapsto \nabla J_2(\zeta)$ is $\hat{L}(\hat{\ell})$ -Lipschitz continuous on $\hat{\mathcal{S}}(\hat{\ell})$, that is, for all $\theta, \theta' \in \mathcal{S}(\ell)$ and all $\zeta, \zeta' \in \hat{\mathcal{S}}(\hat{\ell})$

$$\|\nabla J_1(\theta') - \nabla J_1(\theta)\|_F \leq L(\ell) \|\theta' - \theta\|_F, \quad \|\nabla J_2(\zeta') - \nabla J_2(\zeta)\|_F \leq \hat{L}(\hat{\ell}) \|\zeta' - \zeta\|_F.$$

Proof. Cf. Appendix C.3 □

The following theorem is the main result of this section and states the linear convergence rate of the gradient descent method. A key point is to prove that the sequence $(\theta_{[k]}, \zeta_{[k]})_{k \geq 0}$ always lies in $\mathcal{S}(\ell) \times \hat{\mathcal{S}}(\hat{\ell})$ provided that one starts from an initial point $(\theta_{[0]}, \zeta_{[0]}) \in \mathcal{S}(\ell) \times \hat{\mathcal{S}}(\hat{\ell})$ with a suitably chosen common step size $\rho_\theta = \rho_\zeta := \rho > 0$.

Theorem 3.1 (Global convergence of the exact gradient descent method). *Let $(\theta_{[0]}, \zeta_{[0]}) \in \mathcal{S}(\ell) \times \hat{\mathcal{S}}(\hat{\ell})$ and select a constant step size $\rho \in (0, \frac{2}{\tilde{L}(\ell, \hat{\ell})})$ where*

$$\tilde{L}(\ell, \hat{\ell}) = \max(L(\ell), \hat{L}(\hat{\ell})).$$

Then, the sequence $(\theta_{[k]}, \zeta_{[k]})_{k \geq 0}$ generated by the exact GD algorithm (3.1) stays in $\mathcal{S}(\ell) \times \hat{\mathcal{S}}(\hat{\ell})$. Moreover, for any fixed accuracy $\varepsilon > 0$, we achieve

$$J(\theta_{[k]}, \zeta_{[k]}; \lambda) - J(\theta^*, \zeta^*; \lambda) \leq \varepsilon$$

with a number of iterations satisfying

$$k \geq \frac{\log\left(\frac{\varepsilon}{J(\theta_{[0]}, \zeta_{[0]}) - J(\theta^*, \zeta^*)}\right)}{\log\left(1 - \frac{\rho(2 - \rho\tilde{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})}\right)},$$

where

$$\kappa(\ell, \hat{\ell}) = \max\left(\kappa_1, \kappa_2, \frac{2}{\tilde{L}(\ell, \hat{\ell})}\right) + \frac{1}{2}.$$

Remark 3.1. In the model-based case, since we can compute J_1 and J_2 explicitly, one can choose $a = J_1(\theta_{[0]})$ and $\hat{a} = J_2(\zeta_{[0]})$ for any $(\theta_{[0]}, \zeta_{[0]}) \in \mathcal{S} \times \hat{\mathcal{S}}$.

The proof of the above convergence result is based on the following stability result which shows that both \mathcal{S} and $\hat{\mathcal{S}}$ are stable by the one step transition of the GD algorithm.

Lemma 3.1. If $\theta \in \mathcal{S}(\ell), \theta \neq \theta^*$ and $\zeta \in \hat{\mathcal{S}}(\hat{\ell}), \zeta \neq \zeta^*$, then for all $\rho \in (0, \frac{2}{L(\ell, \hat{\ell})}) = (0, \min(\frac{2}{L(\ell)}, \frac{2}{L(\hat{\ell})}))$, it holds

$$\theta_\rho := \theta - \rho \nabla J_1(\theta) \in \mathcal{S}(\ell), \quad \zeta_\rho := \zeta - \rho \nabla J_2(\zeta) \in \hat{\mathcal{S}}(\hat{\ell}).$$

Proof. For a fixed $\theta \in \mathcal{S}(\ell)$, we let

$$\rho_{\max} = \sup \{ \rho' \geq 0 : \theta - \rho \nabla J_1(\theta) \in \mathcal{S}(\ell), \forall \rho \in [0, \rho'] \}.$$

From the first order Taylor expansion and the continuity of ∇J_1 , one has

$$J_1(\theta_\rho) = J_1(\theta) - \rho \|\nabla J_1(\theta)\|_F^2 + o(\rho).$$

Since $\theta \neq \theta^*$, we have $\|\nabla J_1(\theta)\|_F^2 > 0$ and then for all ρ small enough, we have $J_1(\theta_\rho) \leq J_1(\theta) \leq \ell$ which implies that $\theta_\rho \in \mathcal{S}(\ell)$ and thus $\rho_{\max} > 0$. Moreover, from Proposition C.1 where it is shown that $\mathcal{S}(\ell)$ is bounded, we deduce that $\rho_{\max} < +\infty$.

The definition of ρ_{\max} implies that for any $\rho \in (0, \rho_{\max}]$ and any $t \in [0, 1]$, $t\theta_\rho + (1-t)\theta = \theta - (t\rho)\nabla J_1(\theta) = \theta_{t\rho} \in \mathcal{S}(\ell)$ since $t\rho \in (0, \rho_{\max}]$.

Now the second-order Taylor's expansion for $J_1(\theta)$ combined with the local Lipschitz continuity of ∇J_1 stated in Proposition 3.3 guarantees that

$$J_1(\theta_\rho) \leq J_1(\theta) + \langle \nabla J_1(\theta), \theta_\rho - \theta \rangle + \frac{L(a)}{2} \|\theta_\rho - \theta\|_F^2. \quad (3.3)$$

Suppose that $\rho_{\max} < \frac{2}{L(\ell)}$. From the continuity of $\rho \mapsto J_1(\theta - \rho \nabla J_1(\theta))$, and the definition of $\mathcal{S}(\ell)$, we have $J_1(\theta - \rho_{\max} \nabla J_1(\theta)) = \ell$. Then, by (3.3), we get

$$J_1(\theta - \rho_{\max} \nabla J_1(\theta)) \leq J_1(\theta) - \rho_{\max} \|\nabla J_1(\theta)\|_F^2 + \frac{\rho_{\max}^2 L(a)}{2} \|\nabla J_1(\theta)\|_F^2, \quad (3.4)$$

which implies (as $2 - \rho_{\max} L(\ell) > 0$) that

$$J_1(\theta - \rho_{\max} \nabla J_1(\theta)) < J_1(\theta) \leq \ell.$$

This contradicts the fact that $J_1(\theta - \rho_{\max} \nabla J_1(\theta)) = \ell$. Therefore, we conclude that $\rho_{\max} \geq \frac{2}{L(\ell)}$ and for all $\rho \in (0, \frac{2}{L(\ell)})$, $\theta_\rho = \theta - \rho \nabla J_1(\theta) \in \mathcal{S}(\ell)$. Similar arguments show the result for

ζ_ρ . In particular, one shows that for all $\rho \in (0, \frac{2}{\bar{L}(\hat{\ell})})$, $\zeta_\rho \in \hat{\mathcal{S}}(\hat{\ell})$. In conclusion, for all $\rho \in (0, \min(\frac{2}{\bar{L}(\hat{\ell})}, \frac{2}{\bar{L}(\hat{a})}))$, it holds at the same time

$$\theta_\rho := \theta - \rho \nabla J_1(\theta) \in \mathcal{S}(\ell) \quad \text{and} \quad \zeta_\rho := \zeta - \rho \nabla J_2(\zeta) \in \hat{\mathcal{S}}(\hat{\ell}).$$

□

Proof of Theorem 3.1. Start from some fixed stabilizing initial parameter $(\theta_{[0]}, \zeta_{[0]}) \in \mathcal{S}(\ell) \times \hat{\mathcal{S}}(\hat{\ell})$. Since $\check{L}(\ell, \hat{\ell}) = \max(L(\ell), \hat{L}(\hat{\ell}))$, by choosing a step size $\rho \in (0, \frac{2}{\bar{L}(\ell, \hat{\ell})})$, Lemma 3.1 guarantees that

$$\theta_{[1]} = \theta_{[0]} - \rho \nabla J_1(\theta_{[0]}) \in \mathcal{S}(\ell), \quad \text{and} \quad \zeta_{[1]} = \zeta_{[0]} - \rho \nabla J_2(\zeta_{[0]}) \in \hat{\mathcal{S}}(\hat{\ell}).$$

Moreover, by (3.4), the following inequalities hold

$$\begin{cases} J_1(\theta_{[1]}) - J_1(\theta_{[0]}) & \leq \frac{-\rho(2-\rho\check{L}(\ell, \hat{\ell}))}{2} \|\nabla J_1(\theta_{[0]})\|_F^2 \\ J_2(\zeta_{[1]}) - J_2(\zeta_{[0]}) & \leq \frac{-\rho(2-\rho\check{L}(\ell, \hat{\ell}))}{2} \|\nabla J_2(\zeta_{[0]})\|_F^2. \end{cases} \quad (3.5)$$

The definition of $\kappa(\ell, \hat{\ell})$ together with the gradient domination condition of Proposition 3.2, ensures that for all $\theta \in \mathcal{S}(\ell)$, $\zeta \in \hat{\mathcal{S}}(\hat{\ell})$

$$J_1(\theta) - J_1(\theta^*) \leq \kappa(\ell, \hat{\ell}) \|\nabla J_1(\theta)\|_F^2, \quad J_2(\zeta) - J_2(\zeta^*) \leq \kappa(\ell, \hat{\ell}) \|\nabla J_2(\zeta)\|_F^2,$$

and that for all $\rho \in (0, \frac{2}{\bar{L}(\ell, \hat{\ell})})$

$$\frac{\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})} \in (0, 1).$$

Hence, coming back to (3.5), we deduce

$$\begin{aligned} J_1(\theta_{[1]}) - J_1(\theta_{[0]}) & \leq \frac{-\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})} (J_1(\theta_{[0]}) - J_1(\theta^*)), \\ J_2(\zeta_{[1]}) - J_2(\zeta_{[0]}) & \leq \frac{-\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})} (J_2(\zeta_{[0]}) - J_2(\zeta^*)). \end{aligned}$$

Summing the two previous inequalities, we obtain

$$\begin{aligned} J(\theta_{[1]}, \zeta_{[1]}; \lambda) - J(\theta_{[0]}, \zeta_{[0]}; \lambda) & = (J_1(\theta_{[1]}) - J_1(\theta_{[0]})) + (J_2(\zeta_{[1]}) - J_2(\zeta_{[0]})) \\ & \leq \frac{-\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})} ((J_1(\theta_{[0]}) - J_1(\theta^*)) + (J_2(\zeta_{[0]}) - J_2(\zeta^*))) \\ & = \frac{-\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})} (J(\theta_{[0]}, \zeta_{[0]}; \lambda) - J(\theta^*, \zeta^*; \lambda)). \end{aligned}$$

By a direct induction argument, at each step $k \geq 0$, it holds

$$J(\theta_{[k+1]}, \zeta_{[k+1]}; \lambda) - J(\theta_{[k]}, \zeta_{[k]}; \lambda) \leq \frac{-\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})} (J(\theta_{[k]}, \zeta_{[k]}; \lambda) - J(\theta^*, \zeta^*; \lambda)),$$

so that

$$J(\theta_{[k+1]}, \zeta_{[k+1]}; \lambda) - J(\theta^*, \zeta^*; \lambda) \leq \left(1 - \frac{\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})}\right) (J(\theta_{[k]}, \zeta_{[k]}; \lambda) - J(\theta^*, \zeta^*; \lambda)).$$

We thus conclude

$$J(\theta_{[k]}, \zeta_{[k]}; \lambda) - J(\theta^*, \zeta^*; \lambda) \leq \left(1 - \frac{\rho(2 - \rho\check{L}(\ell, \hat{\ell}))}{2\kappa(\ell, \hat{\ell})}\right)^k (J(\theta_{[0]}, \zeta_{[0]}; \lambda) - J(\theta^*, \zeta^*; \lambda)).$$

The above inequality readily leads to the rest of the assertions in the theorem. \square

Remark 3.2. *The problem of the solvability of the Lyapunov equations (2.17)-(2.18) have been investigated in [23]. The main difference with our result is that the authors in [23] have shown that if $\theta \in \mathcal{S}$ and $\|\theta' - \theta\| \leq c(\theta)$ with the perturbation $c(\theta)$ depending on θ then $\theta' \in \mathcal{S}$. Therefore, at each step $\theta' = \theta - \rho\nabla J(\theta)$, they choose the step size $\rho > 0$ according to $\nabla J(\theta)$ in order to make $\|\rho\nabla J(\theta)\| = \|\theta' - \theta\|$ the perturbation less than $c(\theta)$.*

Here, we adopt a different approach: we truncate \mathcal{S} and define a compact subset $\mathcal{S}(\ell)$. We demonstrate that ∇J is Lipschitz continuous on $\mathcal{S}(\ell)$ with a Lipschitz constant $L(\ell)$ and directly choose the step size $\rho > 0$ according to $L(\ell)$, ensuring that the sequence $(\theta_k)_{k \geq 0}$ remains in $\mathcal{S}(\ell)$. Thus, $\rho > 0$ remains constant throughout all iterations. We refer to [16] for a detailed description of the method we have extended to solve an LQ MFC problem. It is important to note that in [16], the only source of randomness is the initial value.

4 Model-free PG algorithm

4.1 Notations

In the model-free setting, we do not have access to the values of the functions J_1 and J_2 , or their gradients, as the model coefficients are unknown. Hence, we combine the two parameters $\Theta = (\theta, \zeta) \in H = \mathbb{R}^{m \times d} \times \mathbb{R}^{m \times d}$. The value function $J_1 + J_2$ is then regarded as a function defined on H .

For all $\Theta = (\theta, \zeta) \in H$, we define the norm of Θ by

$$\|\Theta\|_H = \sqrt{\|\theta\|_F^2 + \|\zeta\|_F^2},$$

and for $\Theta_1 = (\theta_1, \zeta_1), \Theta_2 = (\theta_2, \zeta_2) \in H$, their inner product is defined by

$$\langle \Theta_1, \Theta_2 \rangle_H := \theta_1 : \theta_2 + \zeta_1 : \zeta_2.$$

We let

$$\mathcal{R} := \left\{ \Theta = (\theta, \zeta) \in H : (\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}} \right\} \subset H,$$

and, for $b > 0$, we define the level subset of \mathcal{R} by

$$\mathcal{R}(b) := \left\{ \Theta = (\theta, \zeta) \in \mathcal{R} : \check{J}(\Theta) \leq b \right\}, \quad (4.1)$$

where

$$\check{J}(\Theta) := J(\Theta; \lambda) - v(\lambda) = J_1(\theta) + J_2(\zeta). \quad (4.2)$$

Note that \check{J} is obtained by removing the constant part that depends only on λ from J .

Since J_1 and J_2 are non-negative functions, one has $\mathcal{R}(b) \subset \mathcal{S}(b) \times \hat{\mathcal{S}}(b)$. Having shown that both $\mathcal{S}(b)$ and $\hat{\mathcal{S}}(b)$ are compact subsets of $\mathbb{R}^{m \times d}$, we get that $\mathcal{R}(b)$ is bounded. The continuity of \check{J} on H also implies that $\mathcal{R}(b)$ is closed. Since H is of finite dimension, we deduce that $\mathcal{R}(b)$ is compact.

Finally, one can define the total gradient of J and \check{J} as

$$\nabla \check{J}(\Theta) = \nabla_{\Theta} J(\Theta; \lambda) := (\nabla_{\theta} J(\theta, \zeta; \lambda), \nabla_{\zeta} J(\theta, \zeta; \lambda)) = (\nabla J_1(\theta), \nabla J_2(\zeta)).$$

According to Theorem 3.2, the gradient domination inequality holds:

$$J(\Theta; \lambda) - J(\Theta^*; \lambda) = J_1(\theta) - J_1(\theta^*) + J_2(\zeta) - J_2(\zeta^*) \leq \kappa \|\nabla J(\Theta)\|_H^2, \quad (4.3)$$

where $\kappa = \max\{\kappa_1, \kappa_2\}$. The same inequality also holds for \check{J} .

4.2 Model-free PG algorithm with population simulator

In the current model-free setting, we do not explicitly know the gradient $\nabla \check{J}(\theta, \zeta)$. Therefore, the GD algorithm from the previous section cannot be directly applied. Instead, we rely on a stochastic PG algorithm based on a stochastic population simulator, which provides an approximation of the controlled MKV dynamics (2.8) along with the associated cost J (2.9).

For a finite terminal horizon $T > 0$ and a positive integer n , we consider the uniform time grid of the interval $[0, T]$ given by $\Delta := \Delta(T, n) = \{0 = t_0 < t_1 < \dots < t_n = T\}$ where $t_l = lh$, $l = 0, 1, \dots, n$, and $h := \frac{T}{n}$. For a given positive integer N , the N interacting agents system with states $(X_{t_l}^{\Theta, \Delta, (j)})_{1 \leq j \leq N; 0 \leq l \leq n}$ evolves according to the dynamics

$$\begin{aligned} X_{t_{l+1}}^{\Theta, \Delta, (j)} &= X_{t_l}^{\Theta, \Delta, (j)} + (BX_{t_l}^{\Theta, \Delta, (j)} + \bar{B}\hat{\mu}_{t_l}^{\Theta, \Delta, N} + D\alpha_{t_l}^{\Theta, \Delta, (j)})h \\ &\quad + \gamma(W_{t_{l+1}}^{(j)} - W_{t_l}^{(j)}) + \gamma_0(W_{t_{l+1}}^0 - W_{t_l}^0), \quad j = 1, \dots, N, l = 0, \dots, n-1, \end{aligned} \quad (4.4)$$

where $X_0^{\Theta, \Delta, (j)} = X_0^{(j)}$, $\hat{\mu}_{t_l}^{\Theta, \Delta, N} := \frac{1}{N} \sum_{j=1}^N X_{t_l}^{\Theta, \Delta, (j)}$, and the action that the j -th agent takes at time $t = t_l$ is drawn as $\alpha_{t_l}^{\Theta, \Delta, (j)} \sim \pi^{\Theta}(\cdot | X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}, \hat{\mu}_{t_l}^{\Theta, \Delta, N})$ independently

of $((W^{(j)})_{j=1,\dots,N}, W^0)$. Here, $(W^{(j)})_{j=1,\dots,N}$ are *i.i.d.* d -dimensional Brownian motions on $(\Omega^1, \mathcal{F}^1, \mathbb{P}^1)$ and W^0 is the common noise on $(\Omega^0, \mathcal{F}^0, \mathbb{P}^0)$, and $(X_0^{(j)})_{1 \leq j \leq N}$ is independent of $((W^{(j)})_{j=1,\dots,N}, W^0)$.

In the spirit of [21], the population simulator samples the randomized actions $\alpha_{t_l}^{\Theta, \Delta, (j)}$ as follows

$$\begin{aligned} \alpha_{t_l}^{\Theta, \Delta, (j)} &= \theta(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}) \\ &\quad + \zeta \hat{\mu}_{t_l}^{\Theta, \Delta, N} + \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}^{(j)}, \quad j = 1, \dots, N, l = 0, \dots, n-1, \end{aligned} \quad (4.5)$$

where the *action simulation noises* $(\boldsymbol{\xi}_{t_l}^{(j)})_{1 \leq j \leq N; 0 \leq l \leq n-1}$ are *i.i.d.* random variables independent of $((X_0^{(j)}, W^{(j)})_{j=1,\dots,N}, W^0)$ with law $\mathcal{N}(0, \mathbf{I}_m)$. In particular, one may assume that \mathcal{F}_0 is rich enough to support not only $(X_0^{(j)})_{1 \leq j \leq N}$ but also the sequence $(\boldsymbol{\xi}_{t_l}^{(j)})_{1 \leq j \leq N; 0 \leq l \leq n}$.

Introducing the notations $\mathbf{w}_l^{(j)} = \frac{1}{\sqrt{h}}(W_{t_{l+1}}^{(j)} - W_{t_l}^{(j)})$, $\mathbf{w}_l^0 = \frac{1}{\sqrt{h}}(W_{t_{l+1}}^0 - W_{t_l}^0)$, the dynamics (4.4) writes

$$\begin{aligned} X_{t_{l+1}}^{\Theta, \Delta, (j)} &= X_{t_l}^{\Theta, \Delta, (j)} + \left(B X_{t_l}^{\Theta, \Delta, (j)} + \bar{B} \hat{\mu}_{t_l}^{\Theta, \Delta, N} + D(\theta(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}) + \zeta \hat{\mu}_{t_l}^{\Theta, \Delta, N}) \right. \\ &\quad \left. + \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}^{(j)} \right) h + \sqrt{h} \gamma \mathbf{w}_l^{(j)} + \sqrt{h} \gamma_0 \mathbf{w}_l^0, \quad j = 1, \dots, N, l = 0, \dots, n-1, \end{aligned} \quad (4.6)$$

and we notice that $(X^{\Theta, \Delta, (j)})_{j=1,\dots,N}$ are exchangeable in law.

The discounted running cost at time t_l of the j -th agent is defined by

$$\begin{aligned} \text{Runc}_{t_l}^{\Theta, \Delta, (j)} &= e^{-\beta t_l} \left((X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top Q (X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}) + (\hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top \hat{Q} \hat{\mu}_{t_l}^{\Theta, \Delta, N} \right. \\ &\quad \left. + (\alpha_{t_l}^{\Theta, \Delta, (j)})^\top R \alpha_{t_l}^{\Theta, \Delta, (j)} \right), \end{aligned}$$

and the average cost estimation of the population over the interval $[0, T]$ for a given $\Theta \in H$ is given by

$$\begin{aligned} \mathcal{J}_{pop}^{\Delta, N}(\Theta) &:= \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} \left(\text{Runc}_{t_l}^{\Theta, \Delta, (j)} \right. \\ &\quad \left. + \lambda e^{-\beta t_l} \log p^\Theta(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}, \hat{\mu}_{t_l}^{\Theta, \Delta, N}, \alpha_{t_l}^{\Theta, \Delta, (j)}) \right) \\ &= \frac{h}{N} \sum_{j=1}^N \left(\sum_{l=1}^n e^{-\beta t_l} \left((X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top Q (X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}) \right. \right. \\ &\quad \left. \left. + (\hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top \hat{Q} \hat{\mu}_{t_l}^{\Theta, \Delta, N} + (\alpha_{t_l}^{\Theta, \Delta, (j)})^\top R \alpha_{t_l}^{\Theta, \Delta, (j)} \right. \right. \\ &\quad \left. \left. + \lambda \log p^\Theta(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}, \hat{\mu}_{t_l}^{\Theta, \Delta, N}, \alpha_{t_l}^{\Theta, \Delta, (j)}) \right) \right). \end{aligned} \quad (4.7)$$

Note that, since $\beta, \lambda > 0$, R and the density p^Θ are known, the last term in the above expression can be computed explicitly.

In the spirit of [8], see also [7] and [6], we compute in Algorithm 1 a biased estimator of the true gradient $\nabla \check{J} = (\nabla_\theta \check{J}, \nabla_\zeta \check{J})$ based solely on the average cost estimation (4.7). We let $\mathbb{B}_r \subset \mathbb{R}^{m \times d}$ be the ball of radius r (with respect to the Frobenius norm) centered at the origin, and $\mathbb{S}_r = \partial \mathbb{B}_r$ be its boundary. The uniform distribution on \mathbb{S}_r is denoted by $\mu_{\mathbb{S}_r}$.

Algorithm 1: Model Free Population-Based Gradient Estimation

Input data: Feedback parameters $(\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$, number of agents N , number of perturbations \tilde{N} , finite horizon $T > 0$, number of period of discrete grid n , radius r .

for $i = 1, \dots, \tilde{N}$ **do**

• Sample U_i and V_i i.i.d. $\sim \mu_{\mathbb{S}_r}$

• Define perturbed feedback parameters

$$\theta_i = \theta + U_i ; \zeta_i = \zeta + V_i.$$

• Sample $\mathcal{J}_{pop}^{\Delta, N, i}$ via the population simulator (4.7) for $\Theta = \Theta_i = (\theta_i, \zeta_i)$.

end

Return: The estimations of the gradient of the functional cost with respect to $\Theta = (\theta, \zeta)$:

$$\begin{aligned} \tilde{\nabla}_\theta^{\Delta, N, pop} J(\Theta) &= \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathcal{J}_{pop}^{\Delta, N, i} U_i, \\ \tilde{\nabla}_\zeta^{\Delta, N, pop} J(\Theta) &= \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathcal{J}_{pop}^{\Delta, N, i} V_i. \end{aligned}$$

We are now in position to define our (stochastic) model-free PG algorithm. Given $(\theta_{[0]}, \zeta_{[0]}) \in \mathcal{R}(b)$, for some $b > 0$, recalling (4.1), we update the parameters as follows

$$\theta_{[k+1]} = \theta_{[k]} - \rho \tilde{\nabla}_\theta^{\Delta, N, pop} J(\theta_{[k]}, \zeta_{[k]}), \quad \zeta_{[k+1]} = \zeta_{[k]} - \rho \tilde{\nabla}_\zeta^{\Delta, N, pop} J(\theta_{[k]}, \zeta_{[k]}), \quad k \geq 0, \quad (4.8)$$

for some well-chosen constant positive learning rate ρ .

Our main contribution is the following convergence result of the model-free PG sequence $(\Theta_{[k]} := (\theta_{[k]}, \zeta_{[k]}))_{k \geq 0}$ as given by (4.8).

Theorem 4.1. *Let $\varepsilon > 0$ be a prescribed accuracy and $\rho \in (0, \frac{4}{9\bar{L}(b)})$ be the constant learning rate. Then, choosing r small enough such that*

$$r \leq \check{r}(\varepsilon) := \min \left(\check{r}(b), \frac{1}{h_r(b, 1/(\frac{1}{10\sqrt{2}}\sqrt{\frac{\varepsilon}{\kappa}}))}, \frac{1}{h'_r(b, 1/(\frac{1}{10\sqrt{2}}\sqrt{\frac{\varepsilon}{\kappa}}))} \right),$$

recalling that κ is the coefficient in the gradient domination condition appearing in (4.3), choosing T large enough such that

$$T \geq \check{T}(\varepsilon) := \frac{1}{c_2(2b)} \log \left(\frac{10\sqrt{2}dc_1(2b)}{\check{r}(\varepsilon)} \sqrt{\frac{\kappa}{\varepsilon}} \right),$$

choosing the number of periods in the grid n large enough such that

$$n \geq \check{n}(\varepsilon) := \frac{10\sqrt{2}d\check{T}(\varepsilon)c_3(2b)}{\check{r}(\varepsilon)} \sqrt{\frac{\kappa}{\varepsilon}},$$

taking the number of particles N large enough such that

$$N \geq \check{N}(\varepsilon) := \frac{10\sqrt{2}dc_4(2b, \check{T}(\varepsilon), \check{n}(\varepsilon))}{\check{r}(\varepsilon)} \sqrt{\frac{\kappa}{\varepsilon}},$$

and finally taking the number of samples \tilde{N} in Algorithm 1 large enough such that

$$\tilde{N} \geq \check{\tilde{N}}(\varepsilon) := \max \left\{ h_{\tilde{N}}(b, \frac{1}{\check{r}(\varepsilon)}, 1/(\frac{1}{10\sqrt{2}}\sqrt{\frac{\varepsilon}{\kappa}})), h'_{\tilde{N}}(b, \frac{1}{\check{r}(\varepsilon)}, \check{T}(\varepsilon), \check{n}(\varepsilon), \check{N}(\varepsilon), 1/(\frac{1}{10\sqrt{2}}\sqrt{\frac{\varepsilon}{\kappa}})) \right\},$$

it holds

$$J(\Theta_{[k]}) - J(\Theta^*) \leq \varepsilon$$

with probability at least $1 - 4k(\frac{10\sqrt{2}d\sqrt{\kappa}}{\sqrt{\varepsilon}})^{-d}$ in at most

$$k \geq \log \left(\frac{\varepsilon}{J(\Theta_{[0]}) - J(\Theta^*)} \right) / \log \left(1 - \frac{\rho(4 - 9\rho\bar{L}(b))}{8\bar{\kappa}(b)} \right)$$

iterations.

The proof is deferred to Section D of the appendix. In the above convergence result, $\check{r}(b)$, h_r , h'_r , $c_1(2b)$, $c_2(2b)$, $c_3(2b)$, $c_4(2b)$, $h_{\tilde{N}}$ and $h'_{\tilde{N}}$ are explicit functions whose definition is provided in the proof. In particular, $\check{r}(b)$ is defined in Lemma D.1, h_r, h'_r are defined in Propositions D.2, D.5, $c_1(2b)$, $c_2(2b)$ are defined in Proposition D.3, $c_3(2b)$ is defined in Proposition D.4, $c_4(2b)$ is defined in Proposition D.6 and $h_{\tilde{N}}, h'_{\tilde{N}}$ are defined in Propositions D.2, D.5.

Remark 4.1. *Regarding the choice of b , we recall the definition of the level set, in (4.1)-(4.2)- and $b > 0$ is the superior limit of \check{J} , where \check{J} is the expected function cost J in which we removed the known constant $v(\lambda)$ (since $\beta > 0, \lambda > 0, R$ are known). Hence, a natural choice is to take $b \geq \check{J}(\Theta_{[0]}) = J(\Theta_{[0]}) - v(\lambda)$ and one can replace $J(\Theta_{[0]})$ by its sample average approximation using the average cost estimation of the population (4.7).*

5 Numerical example

We demonstrate the convergence analysis of our algorithms using a one-dimensional example with the following model parameters:

Q	\hat{Q}	B	\hat{B}	β	γ	γ_0	D	R	X_0
0.1	0.2	0.1	0.2	20	0.05	0.05	0.05	0.2	$\mathcal{N}(1, 1)$

Table 1: The parameters of the model

The optimal parameters (2.10) are $(\theta^*, \zeta^*) = (-0.00126, -0.00255)$. Thus, the optimal cost is given by $J_1(\theta^*) = 0.005051$, $J_2(\zeta^*) = 0.010205$. Both model-based and model-free algorithms are initialized with $(\theta_{[0]}, \zeta_{[0]}) = (-2, -2)$.

5.1 Model-based algorithm

In the model-based case, we can derive explicit formulas for $\theta \mapsto J_1(\theta)$ and $\zeta \mapsto J_2(\zeta)$, along with their derivatives $J_1'(\theta)$ and $J_2'(\zeta)$, using the results obtained in Sections 2 and 3.

We evaluate the PG algorithm (Algorithm 3.1) using three different step sizes: $\rho = 0.5, 0.9, 1.2$. Figure 1 illustrates the convergence of the sequence $(\theta_{[k]}, \zeta_{[k]})_{n \geq 0}$ towards (θ^*, ζ^*) in terms of the number of iterations.




Figure 1: Convergence of the sequence $(\theta_{[k]}, \zeta_{[k]})_{k \geq 0}$ towards (θ^*, ζ^*) (dashed lines).

Then, Figure 2 illustrates the convergence of the cost functions J_1 and J_2 by plotting the relative errors $k \mapsto (J_1(\theta_{[k]}) - J_1(\theta^*)) / J_1(\theta^*)$ and $k \mapsto (J_2(\zeta_{[k]}) - J_2(\zeta^*)) / J_2(\zeta^*)$.




Figure 2: Convergence of the error rates $(J_1(\theta_{[k]}) - J_1(\theta^*)) / J_1(\theta^*)$ and $(J_2(\zeta_{[k]}) - J_2(\zeta^*)) / J_2(\theta^*)$.

The PG algorithm 3.1 demonstrates very good performance. The error and the error rate of the value function converge to zero after roughly 200 iterations for different values of the learning rate.

5.2 Model-free algorithm

We use the following parameter values for the model-free PG algorithm:

T	n	N	\tilde{N}	r	λ
1	100	100	100	0.05	0.001

Table 2: Parameters of the model-free PG algorithm.

The optimal cost corresponding to the parameters given in Table 1 is $J(\theta^*, \zeta^*) = 0.015360$. First, we test whether the parameters in Table 2 allow for accurate gradient estimations. In Figure 3, we evaluate this for $(\theta_{[0]}, \zeta_{[0]}) = (-2, -2)$, by generating 100 gradient estimations using the gradient estimation algorithm (Algorithm 1).




Figure 3: 100 gradient estimations vs exact values of $\nabla_{\theta}J(\theta, \zeta)$ and $\nabla_{\zeta}J(\theta, \zeta)$ (dashed lines).

In the figures above, the vertical coordinates of the points represent the estimated gradients, with the dashed line indicating the exact gradients. We observe that most points cluster around the dashed line, remaining within 10^{-2} above and below it. This suggests that, with high probability, the error in gradient estimation is controlled within 10^{-2} , indicating that the parameters in Table 2 are well-chosen.

Our main result (Theorem 4.1) is established for a constant step size $\rho > 0$. However, from a numerical standpoint, one might consider exploring adaptive selection of ρ by adjusting it based on the observed behavior of the cost function J . Specifically, if $\rho > 0$ is in an acceptable interval (an open interval), the cost function J should decrease. According to the findings of the model-based algorithm, increasing ρ (while ensuring it remains within the acceptable interval) may accelerate convergence.

Therefore, we can initialize with a $\rho > 0$ and run the algorithm for several iterations (e.g., 100) initially. If we observe a clear downward trend in the (estimated) cost function J , it indicates that our chosen $\rho > 0$ falls within the acceptable interval. We can then consider increasing ρ slightly to potentially accelerate convergence. Conversely, if we do not observe a decreasing trend, we should consider reducing $\rho > 0$.

The above results are based on the following choice of the step size as a function of the iteration index k :

$$\rho(k) = \begin{cases} 0.5 & \text{if } k \leq 100, \\ 0.9 & \text{if } 100 < k \leq 200, \\ 1.2 & \text{if } 200 < k \leq 350 \end{cases}$$

In Figure 4, we observe the convergence of both sequences $(\theta_{[k]})_{k \geq 0}$ and $(\zeta_{[k]})_{k \geq 0}$ towards θ^* and ζ^* respectively.




Figure 4: Convergence of the sequences $(\theta_{[k]})_{k \geq 0}$ and $(\zeta_{[k]})_{k \geq 0}$ towards θ^* and ζ^* (dashed lines).

In Figure 5, we plot the error $(J(\theta_{[k]}, \zeta_{[k]}) - J(\theta^*, \zeta^*))_{k \geq 0}$ (on the left side) and the error rate $(J(\theta_{[k]}, \zeta_{[k]}) - J(\theta^*, \zeta^*)) / J(\theta^*, \zeta^*)_{k \geq 0}$ (on the right side) as a function of the number of iterations. As in the model-based case, we observe excellent performance of the policy gradient algorithm. Both the error and the error rate on the value function vanish after approximately 200 iterations.




Figure 5: Convergence of the cost error $J(\theta_{[k]}, \zeta_{[k]}) - J(\theta^*, \zeta^*)$ and error rate $(J(\theta_{[k]}, \zeta_{[k]}) - J(\theta^*, \zeta^*)) / J(\theta^*, \zeta^*)$.

A Explicit solution to LQ MFC

Let us consider the more general linear mean-field dynamics

$$\begin{cases} dX_t &= (BX_t + \bar{B}\mathbb{E}_0[X_t] + D \int a \pi_t(da)) dt \\ &+ (\gamma + FX_t + \bar{F}\mathbb{E}_0[X_t]) dW_t + (\gamma_0 + F_0X_t + \bar{F}_0\mathbb{E}_0[X_t]) dW_t^0, \\ X_0 &\sim \mu, \end{cases} \quad (\text{A.1})$$

and a quadratic cost with entropy regularizer

$$J(\boldsymbol{\pi}) = \mathbb{E} \left[\int_0^\infty e^{-\beta t} \left(X_t^\top Q X_t + \mathbb{E}_0[X_t]^\top \bar{Q} \mathbb{E}_0[X_t] \right. \right. \\ \left. \left. + \int a^\top R a \boldsymbol{\pi}_t(da) + \lambda \int \log \boldsymbol{p}_t(a) \boldsymbol{\pi}_t(da) \right) dt \right]$$

to be minimized over randomized controls $\boldsymbol{\pi} = (\boldsymbol{\pi}_t)_t$ with density $a \in \mathbb{R}^m \mapsto (\boldsymbol{p}_t(a))_t$ of feedback form: $\boldsymbol{\pi}_t = \pi(\cdot|t, X_t, \mathbb{E}_0[X_t])$ for some randomized policy $(t, x, \mu) \mapsto \pi(\cdot|t, x, \mu) \in \mathcal{P}_2(\mathbb{R}^m)$. For simplicity of presentation, we assume here that W and W^0 are one-dimensional Brownian motions, see Remark A.2 for the extension to the multi-dimensional case.

Theorem A.1. *Assume that the two following coupled Riccati equations for $K \in \mathbb{S}^d$ and $\Lambda \in \mathbb{S}^d$ have a unique positive definite solutions*

$$\begin{cases} -\beta K + Q + KB + B^\top K + F^\top KF + F_0^\top KF_0 - K^\top DR^{-1}D^\top K = 0 \\ -\beta \Lambda + (Q + \bar{Q}) + \Lambda(B + \bar{B}) + (B + \bar{B})^\top \Lambda + (F + \bar{F})^\top K(F + \bar{F}) + (F_0 + \bar{F}_0)^\top \Lambda(F_0 + \bar{F}_0) \\ -\Lambda^\top DR^{-1}D^\top \Lambda = 0, \end{cases} \quad (\text{A.2})$$

Let $(K, \Lambda) \in (\mathcal{S}_{>+}^d)^2$ be the unique positive definite solutions to (A.2) and $Y \in \mathbb{R}^d$ be the solution to

$$-\beta Y + (B + \bar{B})^\top Y + (F + \bar{F})^\top K \gamma^\top + (F_0 + \bar{F}_0)^\top \Lambda \gamma_0^\top - \Lambda^\top DR^{-1}D^\top Y = 0 \quad (\text{A.3})$$

Then the optimal randomised control is of feedback form with Gaussian distribution, namely for all $t \geq 0$, $\pi_t^*(\cdot) = \pi^*(\cdot|X_t^*, \mathbb{E}_0[X_t^*])$ with

$$\pi^*(\cdot|x, \bar{x}) = \mathcal{N} \left(-R^{-1}D^\top K(x - \bar{x}) - R^{-1}D^\top \Lambda \bar{x} - R^{-1}D^\top Y; \frac{\lambda}{2} R^{-1} \right)$$

where X^* is the unique solution to (A.1) with $\boldsymbol{\pi} = \boldsymbol{\pi}^*$. Moreover, the optimal functional cost satisfies

$$J(\boldsymbol{\pi}^*) = \mathbb{E}[(X_0 - \mathbb{E}[X_0])^\top K(X_0 - \mathbb{E}[X_0]) + \mathbb{E}[X_0]^\top \Lambda \mathbb{E}[X_0] + 2Y^\top X_0] \\ + \frac{1}{\beta} \left(\gamma^\top K \gamma + \gamma_0^\top \Lambda \gamma_0 - Y^\top DR^{-1}D^\top Y - \frac{\lambda m}{2} \log(2\pi) - \frac{\lambda}{2} \log \left| \frac{\lambda}{2 \det(R)} \right| \right).$$

Proof. We employ the same Martingale approach as in [9] and [2] to prove that $\boldsymbol{\pi}^*$ is the optimal policy, noting that the value function is time-independent in the infinite horizon case. In what follows, we omit the dependence of the process w.r.t the control.

Step 1. Let us consider the function defined on $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$ by $w(x, \mu) = \bar{w}(x, \bar{\mu})$, recalling that $\bar{\mu}$ stands for the mean of μ , where \bar{w} is defined by

$$\bar{w}(x, \bar{x}) = (x - \bar{x})^\top K(x - \bar{x}) + \bar{x}^\top \Lambda \bar{x} + 2Y^\top x + r, \quad \mathbb{R}^d \times \mathbb{R}^d$$

for some matrices/vectors $K, \Lambda \in \mathbb{S}_+^d$, $Y \in \mathbb{R}^d$ and $r \in \mathbb{R}$. Then, given $\boldsymbol{\pi} \in \Pi$ with density \boldsymbol{p} and $X := X^\boldsymbol{\pi}$ the unique solution to (A.1), we introduce the following process

$$\begin{aligned} \mathcal{S}_t^\boldsymbol{\pi} &:= e^{-\beta t} \bar{w}(X_t, \mathbb{E}_0[X_t]) + \int_0^t e^{-\beta s} \left((X_s)^\top Q X_s + \mathbb{E}_0[X_s]^\top \bar{Q} \mathbb{E}_0[X_s] \right. \\ &\quad \left. + \int a^\top R a \boldsymbol{\pi}_s(da) + \lambda \int \log \boldsymbol{p}_s(a) \boldsymbol{\pi}_s(da) \right) ds, \quad t \geq 0, \end{aligned}$$

where we set $\boldsymbol{p}_t(a) = \boldsymbol{\pi}_t(a)$ the density function of $\boldsymbol{\pi}$ for $t \geq 0$. Note that $\mathbb{E}_0[X_t]$ satisfies

$$d\mathbb{E}_0[X_t] = ((B + \bar{B})\mathbb{E}_0[X_t] + D\mathbb{E}_0[\int a \boldsymbol{\pi}_t(da)])dt + (\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])dW_t^0.$$

Step 2. The derivatives of w are given by

$$\begin{aligned} \partial_x \bar{w}(x, \bar{x}) &= 2K(x - \bar{x}) + 2Y, \quad \partial_{\bar{x}} \bar{w}(x, \bar{x}) = -2K(x - \bar{x}) + 2\Lambda \bar{x}, \\ \partial_{xx}^2 \bar{w}(x, \bar{x}) &= 2K, \quad \partial_{\bar{x}\bar{x}}^2 \bar{w}(x, \bar{x}) = 2(K + \Lambda), \quad \partial_{x\bar{x}}^2 \bar{w}(x, \bar{x}) = \partial_{\bar{x}x}^2 \bar{w}(x, \bar{x}) = -2K. \end{aligned}$$

Moreover, using the dynamics of X_t and $\mathbb{E}_0[X_t]$, we obtain

$$\begin{aligned} d\langle X, X \rangle_t &= ((\gamma + F X_t + \bar{F} \mathbb{E}_0[X_t])(\gamma + F X_t + \bar{F} \mathbb{E}_0[X_t])^\top \\ &\quad + (\gamma_0 + F_0 X_t + \bar{F}_0 \mathbb{E}_0[X_t])(\gamma_0 + F_0 X_t + \bar{F}_0 \mathbb{E}_0[X_t])^\top) dt, \\ d\langle \mathbb{E}_0[X], \mathbb{E}_0[X] \rangle_t &= (\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])(\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])^\top dt, \\ d\langle X, \mathbb{E}_0[X] \rangle_t &= (\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])(\gamma_0 + F_0 X_t + \bar{F}_0 \mathbb{E}_0[X_t])^\top dt. \end{aligned}$$

Applying Ito's rule to $(w(X_t, \mathbb{E}_0[X_t]))_{t \geq 0}$, we obtain

$$\begin{aligned} d\bar{w}(X_t, \mathbb{E}_0[X_t]) &= \partial_x \bar{w}(X_t, \mathbb{E}_0[X_t]) \cdot dX_t + \partial_{\bar{x}} \bar{w}(X_t, \mathbb{E}_0[X_t]) \cdot d\mathbb{E}_0[X_t] + \frac{1}{2} \partial_{xx}^2 \bar{w}(X_t, \mathbb{E}_0[X_t]) : d\langle X, X \rangle_t \\ &\quad + \frac{1}{2} \partial_{\bar{x}\bar{x}}^2 \bar{w}(X_t, \mathbb{E}_0[X_t]) : d\langle \mathbb{E}_0[X], \mathbb{E}_0[X] \rangle_t + \partial_{x\bar{x}}^2 \bar{w}(X_t, \mathbb{E}_0[X_t]) : d\langle X, \mathbb{E}_0[X] \rangle_t \\ &= \left(2((X_t - \mathbb{E}_0[X_t])^\top K + Y^\top)(B X_t + \bar{B} \mathbb{E}_0[X_t] + D \int a \boldsymbol{\pi}_t(da)) \right. \\ &\quad + 2(- (X_t - \mathbb{E}_0[X_t])^\top K + \mathbb{E}_0[X_t]^\top \Lambda) ((B + \bar{B})\mathbb{E}_0[X_t] + D\mathbb{E}_0[\int a \boldsymbol{\pi}_t(da)]) \\ &\quad + \langle K, ((\gamma + F X_t + \bar{F} \mathbb{E}_0[X_t])(\gamma + F X_t + \bar{F} \mathbb{E}_0[X_t])^\top + (\gamma_0 + F_0 X_t + \bar{F}_0 \mathbb{E}_0[X_t])(\gamma_0 + F_0 X_t + \bar{F}_0 \mathbb{E}_0[X_t])^\top) \rangle \\ &\quad + \langle K + \Lambda, (\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])(\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])^\top \rangle \\ &\quad \left. + \langle -2K, (\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])(\gamma_0 + F_0 X_t + \bar{F}_0 \mathbb{E}_0[X_t])^\top \rangle \right) dt + \left(\dots \right) dW_t + \left(\dots \right) dW_t^0 \end{aligned} \tag{A.4}$$

and taking conditional expectation \mathbb{E}_0 on both sides of the preceding equality

$$\begin{aligned}
d\mathbb{E}_0[\bar{w}(X_t, \mathbb{E}_0[X_t])] &= \left(\mathbb{E}_0 \left[2((X_t - \mathbb{E}_0[X_t])^\top K + Y^\top)(B(X_t - \mathbb{E}_0[X_t]) + (B + \bar{B})\mathbb{E}_0[X_t] + D \int a\pi_t(da)) \right. \right. \\
&+ 2(- (X_t - \mathbb{E}_0[X_t])^\top K + \mathbb{E}_0[X_t]^\top \Lambda)((B + \bar{B})\mathbb{E}_0[X_t] + D\mathbb{E}_0[\int a\pi_t(da)]) \\
&+ \langle K, ((\gamma + FX_t + \bar{F}\mathbb{E}_0[X_t])(\gamma + FX_t + \bar{F}\mathbb{E}_0[X_t])^\top + (\gamma_0 + F_0X_t + \bar{F}_0\mathbb{E}_0[X_t])(\gamma_0 + F_0X_t + \bar{F}_0\mathbb{E}_0[X_t])^\top) \rangle \\
&\left. \left. + \langle -K + \Lambda, (\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])(\gamma_0 + (F_0 + \bar{F}_0)\mathbb{E}_0[X_t])^\top \rangle \right] \right) dt + \left(\dots \right) dW_t^0.
\end{aligned} \tag{A.5}$$

Step 3. Referring to Lemma 6.3 in [2], we try to show that $t \in \mathbb{R}_+ \mapsto \mathbb{E}[\mathcal{S}_t^\pi]$ is non-decreasing for all $\pi \in \Pi$ and $t \in \mathbb{R}_+ \mapsto \mathbb{E}[\mathcal{S}_t^{\pi^*}]$ is constant for π^* . Applying Ito's formula to \mathcal{S}_t^π and taking expectation we get

$$d\mathbb{E}[\mathcal{S}_t^\pi] = e^{-\beta t} \mathbb{E}[\mathcal{D}_t^\pi] dt.$$

with

$$\begin{aligned}
\mathcal{D}_t^\pi &= -\beta \bar{w}(X_t, \mathbb{E}_0[X_t]) + \frac{d}{dt} \mathbb{E}[\bar{w}(X_t, \mathbb{E}_0[X_t])] + X_t^\top Q X_t + \mathbb{E}_0[X_t]^\top \bar{Q} \mathbb{E}_0[X_t] \\
&+ \int a^\top R a \pi_t(da) + \lambda \int (\log \mathbf{p}_t(a)) \pi_t(da).
\end{aligned}$$

Then, using (A.4) and (A.5), we obtain

$$\begin{aligned}
\frac{d}{dt} \mathbb{E}[\bar{w}(X_t, \mathbb{E}_0[X_t])] &= \frac{d}{dt} \mathbb{E}[\mathbb{E}_0[\bar{w}(X_t, \mathbb{E}_0[X_t])]] \\
&= \mathbb{E} \left[(X_t - \mathbb{E}_0[X_t])^\top (KB + B^\top K + F^\top KF + F_0^\top KF_0)(X_t - \mathbb{E}_0[X_t]) \right. \\
&+ \mathbb{E}_0[X_t]^\top (\Lambda(B + \bar{B}) + (B + \bar{B})^\top \Lambda + (F + \bar{F})^\top K(F + \bar{F}) + (F_0 + \bar{F}_0)^\top \Lambda(F_0 + \bar{F}_0)) \mathbb{E}_0[X_t] \\
&+ 2(Y^\top(B + \bar{B}) + \gamma^\top K(F + \bar{F}) + \gamma_0^\top \Lambda(F_0 + \bar{F}_0)) X_t \\
&\left. + 2((X_t - \mathbb{E}_0[X_t])^\top KD + \mathbb{E}_0[X_t]^\top \Lambda D + Y^\top D) \int_A a \pi_t(da) \right] + \gamma^\top K \gamma + \gamma_0^\top \Lambda \gamma_0
\end{aligned}$$

so that

$$\begin{aligned}
\mathbb{E}[\mathcal{D}_t^\pi] &= \mathbb{E} \left[(X_t - \mathbb{E}_0[X_t])^\top (-\beta K + KB + B^\top K + F^\top KF + F_0^\top KF_0 + Q)(X_t - \mathbb{E}_0[X_t]) \right. \\
&+ \mathbb{E}_0[X_t]^\top (-\beta \Lambda + \Lambda(B + \bar{B}) + (B + \bar{B})^\top \Lambda + (F + \bar{F})^\top K(F + \bar{F}) + (F_0 + \bar{F}_0)^\top \Lambda(F_0 + \bar{F}_0) + Q + \bar{Q}) \mathbb{E}_0[X_t] \\
&+ 2(-\beta Y^\top + Y^\top(B + \bar{B}) + \gamma^\top K(F + \bar{F}) + \gamma_0^\top \Lambda(F_0 + \bar{F}_0)) X_t \\
&+ \int a^\top R a \pi_t(da) + 2((X_t - \mathbb{E}_0[X_t])^\top KD + \mathbb{E}_0[X_t]^\top \Lambda D + Y^\top D) \int a \pi_t(da) + \lambda \int (\log p_t(a)) \pi_t(da) \\
&\left. + \gamma^\top K \gamma + \gamma_0^\top \Lambda \gamma_0 - \beta r \right].
\end{aligned}$$

Step 4. We introduce the map

$$\mathbb{I}(\pi_t) = \mathbb{E}_0 \left[\int a^\top R a \pi_t(da) + 2((X_t - \mathbb{E}_0[X_t])^\top KD + \mathbb{E}_0[X_t]^\top \Lambda D + Y^\top D) \int a \pi_t(da) + \lambda \int (\log p_t(u)) \pi_t(du) \right].$$

Following the approach outlined in Step 3 of the proof of Theorem B.1 in [9], the optimal policy that minimizes \mathbb{I} is given by:

$$\pi_t^*(\cdot | X_t, \mathbb{E}_0[X_t]) = \mathcal{N}(-R^{-1}D^\top K(X_t - \mathbb{E}_0[X_t]) - R^{-1}D^\top \Lambda \mathbb{E}_0[X_t] - R^{-1}D^\top Y; \frac{\lambda}{2}R^{-1})$$

and the infimum writes

$$\begin{aligned}
\mathbb{I}(\pi_t^*) &= \mathbb{E}_0 \left[- (X_t - \mathbb{E}_0[X_t])^\top K D R^{-1} D^\top K (X_t - \mathbb{E}_0[X_t]) - \mathbb{E}_0[X_t]^\top \Lambda D R^{-1} D^\top \Lambda \mathbb{E}_0[X_t] \right. \\
&\quad \left. - 2Y^\top D R^{-1} D^\top \Lambda X_t - Y^\top D R^{-1} D^\top Y \right] - \frac{\lambda m}{2} \log(2\pi) - \frac{\lambda}{2} \log \left| \frac{\lambda}{2 \det(R)} \right|.
\end{aligned}$$

Step 5. Noting that R is strictly positive-definite, for all $\pi \in \Pi$, we have

$$\begin{aligned}
\mathbb{E}[\mathcal{D}_t^\pi] &= \mathbb{E} \left[(X_t - \mathbb{E}_0[X_t])^\top (-\beta K + KB + B^\top K + F^\top KF + F_0^\top KF_0 + Q)(X_t - \mathbb{E}_0[X_t]) \right. \\
&+ \mathbb{E}_0[X_t]^\top (-\beta \Lambda + \Lambda(B + \bar{B}) + (B + \bar{B})^\top \Lambda + (F + \bar{F})^\top K(F + \bar{F}) + (F_0 + \bar{F}_0)^\top \Lambda(F_0 + \bar{F}_0) + Q + \bar{Q}) \mathbb{E}_0[X_t] \\
&+ 2(-\beta Y^\top + Y^\top(B + \bar{B}) + \gamma^\top K(F + \bar{F}) + \gamma_0^\top \Lambda(F_0 + \bar{F}_0)) X_t + \langle K, \gamma \gamma^\top \rangle + \langle \Lambda, \gamma_0 \gamma_0^\top \rangle - \beta r + \mathbb{I}(\pi_t) \left. \right] \\
&\geq \mathbb{E} \left[(X_t - \mathbb{E}_0[X_t])^\top (-\beta K + KB + B^\top K + F^\top KF + F_0^\top KF_0 + Q - K D R^{-1} D^\top K)(X_t - \mathbb{E}_0[X_t]) \right. \\
&+ \mathbb{E}_0[X_t]^\top (-\beta \Lambda + \Lambda(B + \bar{B}) + (B + \bar{B})^\top \Lambda + (F + \bar{F})^\top K(F + \bar{F}) + (F_0 + \bar{F}_0)^\top \Lambda(F_0 + \bar{F}_0) \\
&+ Q + \bar{Q} - \Lambda D R^{-1} D^\top \Lambda) \mathbb{E}_0[X_t] + 2(-\beta Y^\top + Y^\top(B + \bar{B}) + \gamma^\top K(F + \bar{F}) + \gamma_0^\top \Lambda(F_0 + \bar{F}_0) \\
&\quad \left. - Y^\top D R^{-1} D^\top \Lambda) X_t + \gamma^\top K \gamma + \gamma_0^\top \Lambda \gamma_0 - \beta r - Y^\top D R^{-1} D^\top Y - \frac{\lambda m}{2} \log(2\pi) - \frac{\lambda}{2} \log \left| \frac{\lambda}{2 \det(R)} \right|. \right.
\end{aligned} \tag{A.6}$$

Now, taking (K, Λ, Y) as the solutions to (A.2)-(A.3) and letting

$$r = \frac{1}{\beta} \left(\gamma^\top K \gamma + \gamma_0^\top \Lambda \gamma_0 - Y^\top D R^{-1} D^\top Y - \frac{\lambda m}{2} \log(2\pi) - \frac{\lambda}{2} \log \left| \frac{\lambda}{2 \det(R)} \right| \right),$$

we observe that the right hand side of (A.6) vanishes, which means that for all $\pi \in \Pi$, $\mathbb{E}[\mathcal{D}_t^\pi] \geq 0$. Moreover, equality in (A.6) holds true for $\pi = \pi^*$ so that

$$\inf_{\pi \in \Pi} \mathbb{E}[\mathcal{D}_t^\pi] = \mathbb{E}[\mathcal{D}_t^{\pi^*}] = 0.$$

We thus conclude that for any $\pi \in \Pi$, $t \mapsto \mathbb{E}[\mathcal{S}_t^\pi]$ is non-decreasing on $[0, +\infty)$ and $t \mapsto \mathbb{E}[\mathcal{S}_t^{\pi^*}]$ is constant on $[0, +\infty)$ which eventually implies

$$J(\pi^*) = \mathbb{E}[\mathcal{S}_\infty^{\pi^*}] = \mathbb{E}[\mathcal{S}_0^{\pi^*}] = \mathbb{E}[\bar{w}(X_0, \mathbb{E}_0[X_0])].$$

The proof is now complete. \square

Remark A.1. According to Section 6 [2], if $R > 0$, $Q > 0$ and $\hat{Q} = Q + \bar{Q} > 0$, the existence of a positive definite solution to the coupled Riccati equations (A.2) is guaranteed and the assumption is specifically for the uniqueness part.

Remark A.2. As noted in Remark 5.2 in [2], one may extend Theorem A.1 to the case where W and W^0 are multi-dimensional Brownian motions. One may consider $W = (W_t^{\{1\}}, \dots, W_t^{\{q\}})_{t \geq 0}$ (resp. $W^0 = (W_t^{0,\{1\}}, \dots, W_t^{0,\{q_0\}})_{t \geq 0}$) is a q -dimensional (resp. q_0 -dimensional) standard Brownian motion and the dynamics of the controlled state process writes as

$$\begin{cases} dX_t^\pi = (BX_t^\pi + \bar{B}\mathbb{E}_0[X_t^\pi] + D \int_A a\pi_t(da))dt + \sum_{i=1}^q (\gamma^{\{i\}} + F^{\{i\}} X_t^\pi + \bar{F}^{\{i\}} \mathbb{E}_0[X_t^\pi]) dW_t^{\{i\}} \\ + \sum_{i=1}^{q_0} (\gamma_0^{\{i\}} + F_0^{\{i\}} X_t^\pi + \bar{F}_0^{\{i\}} \mathbb{E}_0[X_t^\pi]) dW_t^{0,\{i\}}, \\ X_0^\pi \sim \mu. \end{cases}$$

B Proof of Propositions 2.1 and 2.2

B.1 Proof of Proposition 2.1

We first provide two auxiliary technical results.

Lemma B.1. (Solution to the ODE for symmetric matrix-valued function)

Let $\Gamma : [0, T] \mapsto \mathbb{S}^d$ be the solution the following ODE

$$\frac{d\Gamma_t}{dt} = \Gamma_t A^\top + A \Gamma_t + O \tag{B.1}$$

for some $A \in \mathbb{R}^{d \times d}$ and $O \in \mathbb{S}^d$. Then, it holds

$$\Gamma_t = \exp(tA) \Gamma_0 \exp(tA^\top) + \int_0^t \exp(-(s-t)A) O \exp(-(s-t)A^\top) ds,$$

where $\exp(A) := \sum_{n=0}^{+\infty} \frac{A^n}{n!}$.

Lemma B.2. *If $(\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$, then it holds*

$$\lim_{t \rightarrow +\infty} e^{-\beta t} \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] = 0, \quad \lim_{t \rightarrow +\infty} e^{-\beta t} \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top] = 0.$$

Proof. From the dynamics (2.11) of $(Y_t^\theta, Z_t^\zeta)_{t \geq 0}$ and Itô's rule for $t \mapsto \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top]$ and $t \mapsto \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top]$, we get

$$\begin{aligned} d(\mathbb{E}[Y_t^\theta (Y_t^\theta)^\top]) &= (\mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] (B + D\theta)^\top + (B + D\theta) \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] + \gamma \gamma^\top) dt \\ d(\mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top]) &= (\mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top] (\hat{B} + D\zeta)^\top + (\hat{B} + D\zeta) \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top] + \gamma_0 \gamma_0^\top) dt. \end{aligned} \quad (\text{B.2})$$

Noting that both $t \mapsto \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top]$ and $t \mapsto \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top]$ are solution of an ODE of the form (B.1), Lemma B.1 guarantees that

$$\begin{aligned} \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] &= \exp(t(B + D\theta)) \mathbb{E}[Y_0 (Y_0)^\top] \exp(t(B + D\theta)^\top) \\ &\quad + \int_0^t \exp(-(s-t)(B + D\theta)) \gamma \gamma^\top \exp(-(s-t)(B + D\theta)^\top) ds, \\ \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top] &= \exp(t(\hat{B} + D\zeta)) \mathbb{E}[Z_0 (Z_0)^\top] \exp(t(\hat{B} + D\zeta)^\top) \\ &\quad + \int_0^t \exp(-(s-t)(\hat{B} + D\zeta)) \gamma_0 \gamma_0^\top \exp(-(s-t)(\hat{B} + D\zeta)^\top) ds \end{aligned}$$

so that

$$\begin{aligned} e^{-\beta t} \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] &= \exp(t(B - \frac{\beta}{2} I_d + D\theta)) \mathbb{E}[Y_0 (Y_0)^\top] \exp(t(B - \frac{\beta}{2} I_d + D\theta)^\top) \\ &\quad + \int_0^t \exp(-(s-t)(B - \frac{\beta}{2} I_d + D\theta)) \gamma \gamma^\top \exp(-(s-t)(B - \frac{\beta}{2} I_d + D\theta)^\top) ds, \\ e^{-\beta t} \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top] &= \exp(t(\hat{B} - \frac{\beta}{2} I_d + D\zeta)) \mathbb{E}[Z_0 (Z_0)^\top] \exp(t(\hat{B} - \frac{\beta}{2} I_d + D\zeta)^\top) \\ &\quad + \int_0^t \exp(-(s-t)(\hat{B} - \frac{\beta}{2} I_d + D\zeta)) \gamma_0 \gamma_0^\top \exp(-(s-t)(\hat{B} - \frac{\beta}{2} I_d + D\zeta)^\top) ds. \end{aligned}$$

Now, if $(\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$, then both $B - \frac{\beta}{2} I_d + D\theta$ and $\hat{B} - \frac{\beta}{2} I_d + D\zeta$ are stable which allows to conclude. \square

Proof of Proposition 2.1: Using the integration by parts, we have

$$\int_0^\infty e^{-\beta t} d(\mathbb{E}[Y_t^\theta (Y_t^\theta)^\top]) + (-\beta) \int_0^\infty e^{-\beta t} \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] dt = e^{-\beta t} \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top] \Big|_{t=0}^{t=+\infty} = -\mathbb{E}[Y_0 (Y_0)^\top].$$

Using (B.2) together with the previous identity, we get

$$\begin{aligned} \int_0^\infty e^{-\beta t} d(\mathbb{E}[Y_t^\theta (Y_t^\theta)^\top]) &= \Sigma_\theta (B + D\theta)^\top + (B + D\theta) \Sigma_\theta + \frac{1}{\beta} \gamma \gamma^\top \\ &= -\mathbb{E}[Y_0 (Y_0)^\top] + \beta \Sigma_\theta \end{aligned}$$

so that Σ_θ defined by (2.13) satisfies

$$-\beta\Sigma_\theta + \Sigma_\theta(B + D\theta)^\top + (B + D\theta)\Sigma_\theta + \mathbb{E}[Y_0(Y_0)^\top] + \frac{1}{\beta}\gamma\gamma^\top = 0$$

which in turn, recalling that $M = \mathbb{E}[Y_0(Y_0)^\top] + \frac{1}{\beta}\gamma\gamma^\top$, allows to conclude. The proof for $\hat{\Sigma}_\zeta$ is similar. We thus omit the remaining technical details.

Expressing (2.16) in the standard form of the Lyapounov equation

$$(B - \frac{\beta}{2}I_d + D\theta)\Sigma_\theta + \Sigma_\theta(B - \frac{\beta}{2}I_d + D\theta)^\top + M = 0,$$

$$(\hat{B} - \frac{\beta}{2}I_d + D\zeta)\hat{\Sigma}_\zeta + \hat{\Sigma}_\zeta(\hat{B} - \frac{\beta}{2}I_d + D\zeta)^\top + \hat{M} = 0$$

and, recalling that $(\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$, we conclude that they admit a unique positive definite solution. \square

B.2 Proof of the Proposition 2.2

We write (2.17) and (2.18) as

$$(B - \frac{\beta}{2}I_d + D\theta)^\top K_\theta + K_\theta(B - \frac{\beta}{2}I_d + D\theta) + Q + \theta^\top R\theta = 0,$$

$$(\hat{B} - \frac{\beta}{2}I_d + D\zeta)^\top \Lambda_\zeta + \Lambda_\zeta(\hat{B} - \frac{\beta}{2}I_d + D\zeta) + \hat{Q} + \zeta^\top R\zeta = 0.$$

Since $(\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$, the above equations admit a unique positive definite solutions.

From the dynamics (2.11) of $(Y_t^\theta, Z_t^\zeta)_{t \geq 0}$ and Itô's rule, for all positive definite symmetric matrix Γ , one has

$$\begin{aligned} d(e^{-\beta t}(Y_t^\theta)^\top \Gamma Y_t^\theta) &= e^{-\beta t} \left((Y_t^\theta)^\top (-\beta\Gamma + (B + D\theta)^\top \Gamma + \Gamma(B + D\theta)) Y_t^\theta + \gamma^\top \Gamma \gamma \right) dt \\ &\quad + e^{-\beta t} ((Y_t^\theta)^\top \Gamma \gamma + \gamma^\top \Gamma Y_t^\theta) dW_t, \\ d(e^{-\beta t}((Z_t^\zeta)^\top \Gamma Z_t^\zeta)) &= e^{-\beta t} \left((Z_t^\zeta)^\top (-\beta\Gamma + (\hat{B} + D\zeta)^\top \Gamma + \Gamma(\hat{B} + D\zeta)) Z_t^\zeta + \gamma_0^\top \Gamma \gamma_0 \right) dt \\ &\quad + e^{-\beta t} ((Z_t^\zeta)^\top \Gamma \gamma_0 + \gamma_0^\top \Gamma Z_t^\zeta) dW_t^0. \end{aligned}$$

Then, from the very definition (2.12) of J_1 and J_2 , denoting by K_θ and Λ_ζ the respective

unique solutions to (2.17) and (2.18), we get

$$\begin{aligned}
J_1(\theta) &= (Q + \theta^\top R\theta) : \Sigma_\theta \\
&= \mathbb{E} \left[\int_0^\infty e^{-\beta t} ((Y_t^\theta)^\top (Q + \theta^\top R\theta)) Y_t^\theta dt \right] \\
&= \mathbb{E} \left[\int_0^\infty e^{-\beta t} ((Y_t^\theta)^\top (\beta K_\theta - (B + D\theta)^\top K_\theta - K_\theta (B + D\theta))) Y_t^\theta dt \right] \\
&= \mathbb{E} \left[\int_0^\infty e^{-\beta t} ((Y_t^\theta)^\top (\beta K_\theta - (B + D\theta)^\top K_\theta - K_\theta (B + D\theta))) Y_t^\theta dt + (\dots) dW_t \right] \\
&= -\mathbb{E} \left[\int_0^\infty d \left(e^{-\beta t} (Y_t^\theta)^\top K_\theta Y_t^\theta \right) \right] + \int_0^\infty e^{-\beta t} (\gamma^\top K_\theta \gamma) dt = \mathbb{E}[(Y_0)^\top K_\theta Y_0] + \frac{1}{\beta} \gamma^\top K_\theta \gamma,
\end{aligned}$$

and

$$\begin{aligned}
J_2(\zeta) &= (\hat{Q} + \zeta^\top R\zeta) : \hat{\Sigma}_\zeta \\
&= \mathbb{E} \left[\int_0^\infty e^{-\beta t} ((Z_t^\zeta)^\top (\hat{Q} + \zeta^\top R\zeta)) Z_t^\zeta dt \right] \\
&= \mathbb{E} \left[\int_0^\infty e^{-\beta t} ((Z_t^\zeta)^\top (\beta \Lambda_\zeta - (\hat{B} + D\zeta)^\top \Lambda_\zeta - \Lambda_\zeta (\hat{B} + D\zeta))) Z_t^\zeta dt \right] \\
&= \mathbb{E} \left[\int_0^\infty e^{-\beta t} ((Z_t^\zeta)^\top (\beta \Lambda_\zeta - (\hat{B} + D\zeta)^\top \Lambda_\zeta - \Lambda_\zeta (\hat{B} + D\zeta))) Z_t^\zeta dt + (\dots) dW_t^0 \right] \\
&= -\mathbb{E} \left[\int_0^\infty d \left(e^{-\beta t} (Z_t^\zeta)^\top \Lambda_\zeta Z_t^\zeta \right) \right] + \int_0^\infty e^{-\beta t} \gamma_0^\top \Lambda_\zeta \gamma_0 dt = \mathbb{E}[(Z_0)^\top \Lambda_\zeta Z_0] + \frac{1}{\beta} \gamma_0^\top \Lambda_\zeta \gamma_0.
\end{aligned}$$

The proof of Proposition 2.2 is now complete. \square

C Proofs of results for model-based algorithms

C.1 Proof of Proposition 3.1

We prove Proposition 3.1 only for J_1 , as the part concerning J_2 follows similar reasoning. First, we recall an auxiliary result regarding the exponential form of the solution to continuous Lyapunov equations.

Lemma C.1. *(Solution of continuous Lyapunov equation). Let W be a stable matrix and Q a symmetric matrix. The following continuous Lyapunov equation*

$$WY + YW^\top + Q = 0$$

admits a unique solution Y which satisfies

$$Y = \int_0^\infty e^{Wt} Q e^{W^\top t} dt.$$

Hence, recalling the algebraic Riccati equation (2.18) for K_θ , we deduce that

$$K_\theta = \int_0^\infty e^{(B - \frac{\beta}{2}I_d + D\theta)^\top t} (Q + \theta^\top R\theta) e^{(B - \frac{\beta}{2}I_d + D\theta)t} dt.$$

In order to differentiate J_1 , we will rely on the following definition of the derivatives of a matrix-valued map.

Definition C.1. (*Differentiability of the matrix applications*) The map $\mathbf{M} : \mathbb{R}^{m_1 \times n_1} \mapsto \mathbb{R}^{m_2 \times n_2}$ is differentiable at \mathbf{X} if there exists a matrix mapping: $\mathbf{Z} : \mathbb{R}^{m_1 \times n_1} \mapsto \mathbb{R}^{m_2 n_2 \times m_1 n_1}$ such that for any $\Delta \mathbf{X} \in \mathbb{R}^{m_1 \times n_1}$

$$\text{vec}(\mathbf{M}(\mathbf{X} + \Delta \mathbf{X}) - \mathbf{M}(\mathbf{X})) = \mathbf{Z}(\mathbf{X})\text{vec}(\Delta \mathbf{X}) + O(\text{tr}((\Delta \mathbf{X})^\top (\Delta \mathbf{X})))$$

where 'vec' stands for the vectorization operator defined for $Y = (y_{ij})_{1 \leq i \leq n, 1 \leq j \leq m}$ by

$$\text{vec}(Y) = (y_{11}, \dots, y_{n1}, y_{12}, \dots, y_{n2}, \dots, y_{1m}, \dots, y_{nm})^\top.$$

Lemma C.2. (*Differentiability of K_θ with respect to θ*) Denoting $E_\theta = R\theta + D^\top K_\theta$, for all $\theta, \theta' \in \mathcal{S}$, we have

$$K_{\theta'} - K_\theta = \int_0^\infty e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} [(E_{\theta'})^\top (\theta' - \theta) + (\theta' - \theta)^\top E_\theta + (\theta' - \theta)^\top R(\theta' - \theta)] e^{(B - \frac{\beta}{2}I_d + D\theta')t} dt$$

from which it readily follows that there exists a matrix $Z_{\theta'}$ depending only on θ' such that

$$\text{vec}(K_\theta - K_{\theta'}) = Z_{\theta'} \text{vec}(\theta - \theta') + O(\|\theta' - \theta\|_F^2).$$

Proof. Taking the difference between the two equations (2.17) solved by K_θ and $K_{\theta'}$, we get

$$\begin{aligned} 0 &= (B - \frac{\beta}{2}I_d + D\theta')^\top K_{\theta'} + K_{\theta'}(B - \frac{\beta}{2}I_d + D\theta') - ((B - \frac{\beta}{2}I_d + D\theta)^\top K_\theta + K_\theta(B - \frac{\beta}{2}I_d + D\theta)) \\ &\quad + (\theta')^\top R\theta' - \theta^\top R\theta \\ &= (B - \frac{\beta}{2}I_d + D\theta')^\top K_{\theta'} + K_{\theta'}(B - \frac{\beta}{2}I_d + D\theta') \\ &\quad - (((B - \frac{\beta}{2}I_d + D\theta') - D(\theta' - \theta))^\top K_\theta + K_\theta((B - \frac{\beta}{2}I_d + D\theta') - D(\theta' - \theta))) \\ &\quad + (\theta' - \theta + \theta)^\top R(\theta' - \theta + \theta) - \theta^\top R\theta \\ &= (B - \frac{\beta}{2}I_d + D\theta')^\top (K_{\theta'} - K_\theta) + (K_{\theta'} - K_\theta)(B - \frac{\beta}{2}I_d + D\theta') \\ &\quad + (\theta' - \theta)^\top D^\top K_\theta + K_\theta D(\theta' - \theta) + (\theta' - \theta + \theta)^\top R(\theta' - \theta + \theta) - \theta^\top R\theta \\ &= (B - \frac{\beta}{2}I_d + D\theta')^\top (K_{\theta'} - K_\theta) + (K_{\theta'} - K_\theta)(B - \frac{\beta}{2}I_d + D\theta') \\ &\quad + (\theta' - \theta)^\top (R\theta + D^\top K_\theta) + (R\theta + D^\top K_\theta)^\top (\theta' - \theta) + (\theta' - \theta)^\top R(\theta' - \theta) \\ &= (B - \frac{\beta}{2}I_d + D\theta')^\top (K_{\theta'} - K_\theta) + (K_{\theta'} - K_\theta)(B - \frac{\beta}{2}I_d + D\theta') \\ &\quad + (\theta' - \theta)^\top E_\theta + E_\theta^\top (\theta' - \theta) + (\theta' - \theta)^\top R(\theta' - \theta) \end{aligned}$$

so that $K_{\theta'} - K_{\theta}$ is the unique solution to the following Algebraic Lyapounov equation for Y

$$(B - \frac{\beta}{2}I_d + D\theta')^\top Y + Y(B - \frac{\beta}{2}I_d + D\theta') + (\theta' - \theta)^\top E_{\theta'} + E_{\theta'}^\top(\theta' - \theta) + (\theta' - \theta)^\top R(\theta' - \theta) = 0.$$

From Lemma C.1, we thus obtain

$$K_{\theta'} - K_{\theta} = \int_0^\infty e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} [(\theta' - \theta)^\top E_{\theta'} + E_{\theta'}^\top(\theta' - \theta) + (\theta' - \theta)^\top R(\theta' - \theta)] e^{(B - \frac{\beta}{2}I_d + D\theta')t} dt.$$

The previous identity directly gives

$$\begin{aligned} \text{vec}(K_{\theta'} - K_{\theta}) &= \int_0^\infty \text{vec}(e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} [(\theta' - \theta)^\top E_{\theta'} + E_{\theta'}^\top(\theta' - \theta) + (\theta' - \theta)^\top R(\theta' - \theta)] e^{(B - \frac{\beta}{2}I_d + D\theta')t}) dt \\ &= \left(\int_0^\infty e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} \otimes e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} dt \right) \text{vec}((\theta' - \theta)^\top E_{\theta'} + E_{\theta'}^\top(\theta' - \theta) + (\theta' - \theta)^\top R(\theta' - \theta)) \\ &= \left(\int_0^\infty e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} \otimes e^{(B - \frac{\beta}{2}I_d + D\theta')^\top t} dt \right) \text{vec}((\theta' - \theta)^\top E_{\theta'} + (E_{\theta'}^\top)^\top(\theta' - \theta) + U) \end{aligned}$$

where

$$\begin{aligned} U &= (\theta' - \theta)^\top R(\theta' - \theta) + (E_{\theta} - E_{\theta'})^\top(\theta' - \theta) + (\theta' - \theta)^\top (E_{\theta} - E_{\theta'}) \\ &= -(\theta' - \theta)^\top R(\theta' - \theta) + (K_{\theta'} - K_{\theta})D(\theta' - \theta) + (\theta' - \theta)^\top D^\top(K_{\theta'} - K_{\theta}) \\ &= O(\|\theta' - \theta\|_F^2). \end{aligned}$$

Note that in the last line we again used the expression of $K_{\theta'} - K_{\theta}$. Thus, there exists $Z_{\theta'}$ that depends on $B - \frac{\beta}{2}I_d + D\theta'$ and $E_{\theta'}$ such that

$$\text{vec}(K_{\theta'} - K_{\theta}) = Z_{\theta'}(\theta - \theta') + O(\|\theta' - \theta\|_F^2).$$

We conclude that K_{θ} is indeed differentiable with respect to θ . □

We recall the rules of the total differentiation of the matrices:

- (1) $d(\mathbf{X} \pm \mathbf{Y}) = d\mathbf{X} \pm d\mathbf{Y}$
- (2) $d(\mathbf{X}\mathbf{Y}) = (d\mathbf{X})\mathbf{Y} + \mathbf{X}(d\mathbf{Y})$
- (3) $d(\mathbf{X}^\top) = (d\mathbf{X})^\top$
- (4) $d(\mathbf{X}^{-1}) = -\mathbf{X}^{-1}(d\mathbf{X})\mathbf{X}^{-1}$

We are now in position to prove Proposition 3.1. We only prove the identity for the gradient of J_1 , as the gradient of J_2 can be treated similarly.

By the definition of the differentiation

$$dJ_1(\theta) = \text{tr}(\nabla J_1(\theta)^\top d\theta) = \nabla J_1(\theta) : d\theta.$$

Since $J_1(\theta) = K_\theta : M$, we have $dJ_1(\theta) = dK_\theta : M$ so that

$$\nabla J_1(\theta) : d\theta = dK_\theta : M.$$

Then, we differentiate totally the Lyapunov equation of K_θ and obtain

$$-\beta dK_\theta + dK_\theta(B + D\theta) + (B + D\theta)^\top dK_\theta + K_\theta D d\theta + d\theta^\top D^\top K_\theta + \theta^\top R d\theta + d\theta^\top R \theta = 0$$

which writes

$$-\beta dK_\theta + dK_\theta(B + D\theta) + (B + D\theta)^\top dK_\theta + E_\theta^\top d\theta + d\theta^\top E_\theta = 0.$$

We then multiply from the right the above equation by Σ_θ . We obtain

$$-\beta dK_\theta \Sigma_\theta + dK_\theta(B + D\theta) \Sigma_\theta + (B + D\theta)^\top dK_\theta \Sigma_\theta + (E_\theta^\top d\theta + d\theta^\top E_\theta) \Sigma_\theta = 0.$$

Using the fact that Σ_θ satisfies the equation

$$-\beta \Sigma_\theta + (B + D\theta) \Sigma_\theta + \Sigma_\theta (B + D\theta)^\top + M = 0$$

we get

$$\begin{aligned} \text{tr}((dK_\theta)M) &= \text{tr}((dK_\theta)(\beta \Sigma_\theta - (B + D\theta) \Sigma_\theta - \Sigma_\theta (B + D\theta)^\top)) \\ &= \text{tr}(\beta (dK_\theta) \Sigma_\theta - dK_\theta (B + D\theta) \Sigma_\theta - (B + D\theta)^\top dK_\theta \Sigma_\theta) \\ &= \text{tr}((E_\theta d\theta + d\theta^\top E_\theta) \Sigma_\theta) \\ &= 2 \text{tr}((E_\theta \Sigma_\theta)^\top d\theta) \end{aligned}$$

where in the second and the fourth lines we used the commutative property of the trace operator. Comparing the left and the right sides of the previous identity, we eventually get

$$\nabla J_1(\theta) = 2E_\theta \Sigma_\theta.$$

The proof of Proposition 3.1 is now complete. \square

C.2 Proof of Theorem 3.2

We here prove the Polyak-Lojasiewicz inequality stated in Theorem 3.2. We will only prove it for J_1 , as J_2 can be treated in a completely similar manner. We first need the following auxiliary result.

Lemma C.3. (*Perturbation analysis of J_1*) For all $\theta, \theta' \in \mathcal{S}$, it holds

$$J_1(\theta') - J_1(\theta) = \text{tr}[\Sigma_{\theta'}((\theta' - \theta + R^{-1}E_\theta)^\top R(\theta' - \theta + R^{-1}E_\theta) - E_\theta^\top R^{-1}E_\theta)].$$

Proof. We have

$$\begin{aligned}
J_1(\theta') - J_1(\theta) &= \text{tr}((K_{\theta'} - K_{\theta})M) \\
&= \text{tr}\left(\int_0^\infty e^{(B-\frac{\beta}{2}I_d+D\theta')^\top t} [(E_\theta)^\top(\theta' - \theta) + (\theta' - \theta)^\top E_\theta + (\theta' - \theta)^\top R(\theta' - \theta)] e^{(B-\frac{\beta}{2}I_d+D\theta')t} M dt\right) \\
&= \text{tr}\left(\left(\int_0^\infty e^{(B-\frac{\beta}{2}I_d+D\theta')t} M e^{(B-\frac{\beta}{2}I_d+D\theta')^\top t} dt\right) [(E_\theta)^\top(\theta' - \theta) + (\theta' - \theta)^\top E_\theta \right. \\
&\quad \left. + (\theta' - \theta)^\top R(\theta' - \theta)]\right) \\
&= \text{tr}\left(\Sigma_{\theta'}((\theta' - \theta + R^{-1}E_\theta)^\top R(\theta' - \theta + R^{-1}E_\theta) - E_\theta^\top R^{-1}E_\theta)\right),
\end{aligned}$$

where for the last equality we used the algebraic Riccati equation of Σ_θ . \square

Proof of Theorem 3.2. From Lemma C.3, we deduce

$$\begin{aligned}
J_1(\theta) - J_1(\theta^*) &= \text{tr}[\Sigma_{\theta^*}(E_\theta^\top R^{-1}E_\theta - (\theta^* - \theta + R^{-1}E_\theta)^\top R(\theta^* - \theta + R^{-1}E_\theta))] \\
&\leq \text{tr}[\Sigma_{\theta^*}(E_\theta^\top R^{-1}E_\theta)] \\
&\leq \frac{\|\Sigma_{\theta^*}\|_F}{\sigma_{\min}(R)} \text{tr}(E_\theta^\top E_\theta) \\
&\stackrel{(\diamond)}{\leq} \frac{\|\Sigma_{\theta^*}\|_F}{4\sigma_{\min}(R)\sigma_{\min}^2(\Sigma_\theta)} \text{tr}(\nabla J_1(\theta)^\top \nabla J_1(\theta)) \\
&\stackrel{(\nabla)}{\leq} \frac{\|\Sigma_{\theta^*}\|_F \text{tr}(\nabla J_1(\theta)^\top \nabla J_1(\theta))}{4\sigma_{\min}(R)\sigma_{\min}^2(M)}.
\end{aligned}$$

The inequality (\diamond) is a consequence of the inequality

$$\text{tr}(\nabla J_1(\theta)^\top \nabla J_1(\theta)) = 4\text{tr}(\Sigma_\theta E_\theta^\top E_\theta \Sigma_\theta) \geq 4\sigma_{\min}^2(\Sigma_\theta) \text{tr}(E_\theta E_\theta)$$

which stems from Proposition 3.1. The inequality (∇) follows from the fact that $B - \frac{\beta}{2}I_d + D\theta$ is stable which combined with Lemma C.1 guarantees that

$$\Sigma_\theta = \int_0^\infty e^{(B-\frac{\beta}{2}I_d+D\theta)t} M e^{(B-\frac{\beta}{2}I_d+D\theta)^\top t} dt \succeq M.$$

\square

C.3 Proof of Proposition 3.3

Let us recall that $\nabla J_1(\theta)$ and $\nabla J_2(\zeta)$ are given by

$$\begin{aligned}
\nabla J_1(\theta) &= 2(R\theta + D^\top K_\theta)\Sigma_\theta = 2E_\theta \Sigma_\theta, \\
\nabla J_2(\zeta) &= 2(R\zeta + D^\top \Lambda_\zeta)\hat{\Sigma}_\zeta = 2\hat{E}_\zeta \hat{\Sigma}_\zeta,
\end{aligned}$$

with $E_\theta = R\theta + D^\top K_\theta$, $\hat{E}_\zeta = R\zeta + D^\top \Lambda_\zeta$ and where K_θ, Λ_ζ are the unique solutions to

$$\begin{aligned} -\beta K_\theta + (B + D\theta)^\top K_\theta + K_\theta(B + D\theta) + Q + \theta^\top R\theta &= 0, \\ -\beta \Lambda_\zeta + (\hat{B} + D\zeta)^\top \Lambda_\zeta + \Lambda_\zeta(\hat{B} + D\zeta) + \hat{Q} + \zeta^\top R\zeta &= 0, \end{aligned}$$

and $\Sigma_\theta, \hat{\Sigma}_\zeta$ are the unique solutions to

$$\begin{aligned} -\beta \Sigma_\theta + (B + D\theta)\Sigma_\theta + \Sigma_\theta(B + D\theta)^\top + M &= 0, \\ -\beta \hat{\Sigma}_\zeta + (\hat{B} + D\zeta)\hat{\Sigma}_\zeta + \hat{\Sigma}_\zeta(\hat{B} + D\zeta)^\top + \hat{M} &= 0. \end{aligned}$$

Here again, we will only prove the Lipschitz continuity of ∇J_1 , as ∇J_2 can be treated in a similar manner.

C.3.1 Auxiliary results

We start by proving some useful bounds on $\|\theta\|_F, \|K_\theta\|_F, \|\Sigma_\theta\|_F, \|E_\theta\|_F$ when $\theta \in \mathcal{S}(\ell)$, and similar bounds for $\|\zeta\|_F, \|\Lambda_\zeta\|_F, \|\hat{\Sigma}_\zeta\|_F, \|\hat{E}_\zeta\|_F$ when $\zeta \in \hat{\mathcal{S}}(\hat{\ell})$,

Proposition C.1. *For all $\theta \in \mathcal{S}(\ell)$, it holds*

$$\begin{aligned} \|K_\theta\|_F &\leq \mathbf{B}d_K(\ell) = \frac{\ell}{\sigma_{\min}(M)}, \\ \|\Sigma_\theta\|_F &\leq \mathbf{B}d_\Sigma(\ell) = \frac{\ell}{\sigma_{\min}(Q)}, \\ \|E_\theta\|_F &\leq \mathbf{B}d_E(\ell) = \sqrt{\frac{\|R\|_F(\ell - J_1(\theta^*))}{\sigma_{\min}(M)}}, \\ \|\theta\|_F &\leq \mathbf{B}d_\theta(\ell) = \frac{1}{\sigma_{\min}(R)} \left(\sqrt{\frac{\|R\|_F(\ell - J_1(\theta^*))}{\sigma_{\min}(M)}} + \frac{\|D\|_F \ell}{\sigma_{\min}(M)} \right). \end{aligned}$$

Similarly, for all $\zeta \in \hat{\mathcal{S}}(\hat{\ell})$, it holds

$$\begin{aligned} \|\Lambda_\zeta\|_F &\leq \mathbf{B}d_\Lambda(\hat{\ell}) = \frac{\hat{\ell}}{\sigma_{\min}(\hat{M})}, \\ \|\hat{\Sigma}_\zeta\|_F &\leq \mathbf{B}d_{\hat{\Sigma}}(\hat{\ell}) = \frac{\hat{\ell}}{\sigma_{\min}(\hat{Q})}, \\ \|\hat{E}_\zeta\|_F &\leq \mathbf{B}d_{\hat{E}}(\hat{\ell}) = \sqrt{\frac{\|R\|_F(\hat{\ell} - J_2(\zeta^*))}{\sigma_{\min}(\hat{M})}}, \\ \|\zeta\|_F &\leq \mathbf{B}d_\zeta(\hat{\ell}) = \frac{1}{\sigma_{\min}(R)} \left(\sqrt{\frac{\|R\|_F(\hat{\ell} - J_2(\zeta^*))}{\sigma_{\min}(\hat{M})}} + \frac{\|D\|_F \hat{\ell}}{\sigma_{\min}(\hat{M})} \right). \end{aligned}$$

Before proving the above proposition, we will need the following technical result.

Lemma C.4. *For all $\theta \in \mathcal{S}$, the following lower bound holds*

$$J_1(\theta) - J_1(\theta^*) \geq \frac{\sigma_{\min}(M)}{\|R\|_F} \|E_\theta\|_F^2.$$

Proof. Applying Lemma C.3 for $\theta' = \theta - R^{-1}E_\theta = -R^{-1}D^\top K_\theta$, we get

$$J_1(\theta) - J_1(\theta') = \text{tr}(\Sigma_{\theta'}(E_\theta)^\top R^{-1}E_\theta)$$

so that

$$\begin{aligned} J_1(\theta) - J_1(\theta^*) &\geq J_1(\theta) - J_1(\theta') \\ &= \text{tr}(\Sigma_{\theta'} E_\theta^\top R^{-1} E_\theta) \\ &\geq \frac{\sigma_{\min}(M)}{\|R\|_F} \text{tr}(E_\theta^\top E_\theta) = \frac{\sigma_{\min}(M)}{\|R\|_F} \|E_\theta\|_F^2 \end{aligned}$$

where for the last inequality we used the fact that $\Sigma_\theta \succeq M$. □

Proof of Proposition C.1

Step 1 (Bound for K_θ). Recalling that $J_1(\theta) = \text{Tr}(K_\theta M)$, we directly deduce

$$J_1(\theta) \geq \text{Tr}(K_\theta) \sigma_{\min}(M) \geq \|K_\theta\|_F \sigma_{\min}(M)$$

so that

$$\|K_\theta\|_F \leq \frac{J_1(\theta)}{\sigma_{\min}(M)} \leq \frac{\ell}{\sigma_{\min}(M)}.$$

Step 2 (Bound for Σ_θ). We use the fact that $J_1(\theta) = \text{tr}(\Sigma_\theta(Q + \theta^\top R\theta))$ to deduce

$$J_1(\theta) \geq \text{Tr}(\Sigma_\theta) \sigma_{\min}(Q + \theta^\top R\theta) \geq \|\Sigma_\theta\|_F \sigma_{\min}(Q)$$

which in turn implies

$$\|\Sigma_\theta\|_F \leq \frac{J_1(\theta)}{\sigma_{\min}(Q)} \leq \frac{\ell}{\sigma_{\min}(Q)}.$$

Step 3 (Bound for E_θ). From Lemma C.4, we get

$$\|E_\theta\|_F \leq \sqrt{\frac{\|R\|_F (J_1(\theta) - J_1(\theta^*))}{\sigma_{\min}(M)}} \leq \sqrt{\frac{\|R\|_F (\ell - J_1(\theta^*))}{\sigma_{\min}(M)}}.$$

Step 4 (Bound for $\|\theta\|_F$). We first write

$$\begin{aligned} \|\theta\|_F &\leq \|R\theta\|_F \|R^{-1}\|_F \\ &\leq \frac{\|R\theta + D^\top K_\theta\|_F + \|-D^\top K_\theta\|_F}{\sigma_{\min}(R)} \\ &\leq \frac{\|E_\theta\|_F + \|D\|_F \|K_\theta\|_F}{\sigma_{\min}(R)} \end{aligned}$$

which combined with Lemma C.4 and the bound on $\|K_\theta\|_F$ clearly yields

$$\begin{aligned}\|\theta\|_F &\leq \frac{1}{\sigma_{\min}(R)} \left(\sqrt{\frac{\|R\|_F(J_1(\theta) - J_1(\theta^*))}{\sigma_{\min}(M)}} + \frac{\|D\|_F J_1(\theta)}{\sigma_{\min}(M)} \right) \\ &\leq \frac{1}{\sigma_{\min}(R)} \left(\sqrt{\frac{\|R\|_F(\ell - J_1(\theta^*))}{\sigma_{\min}(M)}} + \frac{\|D\|_F \ell}{\sigma_{\min}(M)} \right).\end{aligned}$$

□

New operators. Taking a given $d \times d$ stable matrix S and a symmetric matrix X , we define the two operators \mathcal{T} and \mathcal{F} by

$$\mathcal{T}^S(X) = \int_0^\infty e^{Sr} X e^{S^\top r} dr, \quad \mathcal{F}^S(X) = SX + XS^\top.$$

It is known that

$$\mathcal{F}^S \circ \mathcal{T}^S + I = 0$$

which means that $\mathcal{T}^S = -(\mathcal{F}^S)^{-1}$.

Lemma C.5. For all $H \in \mathbb{S}_{>+}^d$, it holds

$$\|\mathcal{T}^S\|_F \leq \frac{\|\mathcal{T}^S(H)\|_F}{\sigma_{\min}(H)}.$$

Proof. For any unit vector $v \in \mathbb{R}^d$ and any unit spectral norm matrix X , it holds

$$\begin{aligned}v^\top \mathcal{T}^S(X)v &= \int_0^\infty \text{tr}(X e^{S^\top r} v v^\top e^{Sr}) dr \\ &\leq \int_0^\infty \text{tr}(H e^{S^\top r} v v^\top e^{Sr}) dr \|H^{-1/2} X H^{-1/2}\|_F \\ &= v^\top \mathcal{T}^S(H)v \cdot \|H^{-1/2} X H^{-1/2}\|_F \\ &\leq \|\mathcal{T}^S(H)\|_F \sigma_{\min}^{-1}(H)\end{aligned}$$

which clearly yields the conclusion. □

Lemma C.6. It holds

$$\|\mathcal{T}^S\|_F = \|\mathcal{T}^{S^\top}\|_F.$$

Proof. From the commutative property of the trace operator, for any unit vector $v \in \mathbb{R}^d$ and any unit spectral norm matrix X , one has

$$\begin{aligned}v^\top \mathcal{T}^{S^\top}(X)v &= \int_0^\infty \text{tr}(X e^{Sr} v v^\top e^{S^\top r}) dr \\ &= \text{tr}(X \mathcal{T}^S(vv^\top)) \\ &\leq \|X\|_F \|\mathcal{T}^S(vv^\top)\|_F \\ &\leq \|X\|_F \|\mathcal{T}^{S^\top}\|_F \|vv^\top\|_F = \|\mathcal{T}^{S^\top}\|_F\end{aligned}$$

which yields

$$\|\mathcal{T}^{S^\top}\|_F \leq \|\mathcal{T}^S\|_F.$$

Noting that $S = (S^\top)^\top$, the reverse inequality directly follows. \square

In what follows, we introduce the notations

$$\Xi_\theta = B - \frac{\beta}{2}I_d + D\theta \quad \text{and} \quad \hat{\Xi}_\zeta = \hat{B} - \frac{\beta}{2}I_d + D\zeta.$$

C.3.2 Lipschitz regularity of Σ_θ , $\hat{\Sigma}_\zeta$, K_θ and Λ_ζ .

Perturbation analysis for Σ_θ and $\hat{\Sigma}_\zeta$

Recalling that Σ_θ and $\hat{\Sigma}_\zeta$ are the unique solution to the ALE (2.16), from Lemma C.1, it holds

$$\Sigma_\theta = \mathcal{T}^{\Xi_\theta}(M), \quad \hat{\Sigma}_\zeta = \mathcal{T}^{\hat{\Xi}_\zeta}(\hat{M}).$$

Hence, from Lemma C.5, we directly get

$$\|\mathcal{T}^{\Xi_\theta}\|_F \leq \frac{\|\Sigma_\theta\|_F}{\sigma_{\min}(M)}, \quad \|\mathcal{T}^{\hat{\Xi}_\zeta}\|_F \leq \frac{\|\hat{\Sigma}_\zeta\|_F}{\sigma_{\min}(\hat{M})}. \quad (\text{C.1})$$

Proposition C.2. *For all $\theta, \theta' \in \mathcal{S}(\ell)$, it holds*

$$\|\Sigma_{\theta'} - \Sigma_\theta\|_F \leq \text{Lip}_\Sigma(\ell) \|\theta' - \theta\|_F,$$

where $\text{Lip}_\Sigma(\ell) = \frac{2\|D\|_F(\mathbf{B}d_\Sigma(\ell))^2}{\sigma_{\min}(M)}$. Similarly, for all $\zeta, \zeta' \in \hat{\mathcal{S}}(\hat{\ell})$, it holds

$$\|\hat{\Sigma}_{\zeta'} - \hat{\Sigma}_\zeta\|_F \leq \text{Lip}_{\hat{\Sigma}}(\hat{\ell}) \|\zeta' - \zeta\|_F,$$

where $\text{Lip}_{\hat{\Sigma}}(\hat{\ell}) = \frac{2\|D\|_F(\mathbf{B}d_{\hat{\Sigma}}(\hat{\ell}))^2}{\sigma_{\min}(\hat{M})}$.

Proof. From

$$\begin{aligned} \|\mathcal{F}^{\Xi_{\theta'}}(X) - \mathcal{F}^{\Xi_\theta}(X)\|_F &= \|D(\theta' - \theta)X + X(\theta' - \theta)^\top D^\top\| \\ &\leq 2\|D\|_F \|X\|_F \|\theta' - \theta\|_F \end{aligned}$$

we get

$$\|\mathcal{F}^{\Xi_{\theta'}} - \mathcal{F}^{\Xi_\theta}\|_F \leq 2\|D\|_F \|\theta' - \theta\|_F.$$

Then, using Proposition C.1 and the inequality (C.1), we get

$$\|\mathcal{T}^{\Xi_\theta}\|_F \leq \frac{\|\Sigma_\theta\|_F}{\sigma_{\min}(M)} \leq \frac{\mathbf{B}d_\Sigma(\ell)}{\sigma_{\min}(M)}.$$

Since $\mathcal{F}^{\Xi_\theta}(\Sigma_\theta) = \mathcal{F}^{\Xi_{\theta'}}(\Sigma_{\theta'}) = -M$ we have

$$\begin{aligned}
\|\Sigma_{\theta'} - \Sigma_\theta\|_F &= \|\mathcal{T}^{\Xi_\theta}(\mathcal{F}^{\Xi_\theta}(\Sigma_{\theta'} - \Sigma_\theta))\|_F \\
&= \|\mathcal{T}^{\Xi_\theta}(\mathcal{F}^{\Xi_\theta}(\Sigma_{\theta'}) - \mathcal{F}^{\Xi_\theta}(\Sigma_\theta))\|_F \\
&= \|\mathcal{T}^{\Xi_\theta}(\mathcal{F}^{\Xi_\theta}(\Sigma_{\theta'}) - \mathcal{F}^{\Xi_{\theta'}}(\Sigma_{\theta'}))\|_F \\
&\leq \|\mathcal{T}^{\Xi_\theta}\|_F \|\mathcal{F}^{\Xi_{\theta'}} - \mathcal{F}^{\Xi_\theta}\|_F \|\Sigma_{\theta'}\|_F \\
&\leq \frac{2\|D\|_F(\mathbf{Bd}_\Sigma(\ell))^2}{\sigma_{\min}(M)} \|\theta' - \theta\|_F.
\end{aligned}$$

The Lipschitz regularity of $\hat{\Sigma}_\zeta$ is handled by similar arguments. The remaining technical details are omitted. \square

Perturbation analysis of K_θ and Λ_ζ .

Proposition C.3. *For all $\theta, \theta' \in \mathcal{S}(\ell)$, one has*

$$\|K_{\theta'} - K_\theta\|_F \leq Lip_K(\ell) \|\theta' - \theta\|_F$$

where

$$Lip_K(\ell) = 2(\mathbf{Bd}_E(\ell) + \|R\|_F \mathbf{Bd}_\theta(\ell)) \sup_{\theta \in \mathcal{S}(\ell)} \left\| \int_0^\infty e^{\Xi_\theta t} e^{(\Xi_\theta)^\top t} dt \right\|_F < +\infty.$$

Similarly, one can show that for all $\zeta, \zeta' \in \hat{\mathcal{S}}(\hat{a})$, it holds

$$\|\Lambda_{\zeta'} - \Lambda_\zeta\|_F \leq Lip_\Lambda(\hat{\ell}) \|\zeta' - \zeta\|_F$$

where

$$Lip_\Lambda(\hat{\ell}) = 2(\mathbf{Bd}_{\hat{E}}(\hat{\ell}) + \|R\|_F \mathbf{Bd}_\zeta(\hat{\ell})) \sup_{\zeta \in \hat{\mathcal{S}}(\hat{\ell})} \left\| \int_0^\infty e^{\hat{\Xi}_\zeta t} e^{(\hat{\Xi}_\zeta)^\top t} dt \right\|_F < +\infty.$$

Proof. From Lemma C.2, we get

$$K_{\theta'} - K_\theta = \int_0^\infty e^{\Xi_{\theta'}^\top t} [(E_\theta)^\top (\theta' - \theta) + (\theta' - \theta)^\top E_\theta + (\theta' - \theta)^\top R (\theta' - \theta)] e^{\Xi_{\theta'} t} dt$$

so that

$$\|K_{\theta'} - K_\theta\|_F \leq \|(E_\theta)^\top (\theta' - \theta) + (\theta' - \theta)^\top E_\theta + (\theta' - \theta)^\top R (\theta' - \theta)\|_F \left\| \int_0^\infty e^{\Xi_{\theta'} t} e^{\Xi_{\theta'}^\top t} dt \right\|_F.$$

The compactness of $\mathcal{S}(\ell)$ together with the continuity of $\theta \in \mathcal{S}(\ell) \mapsto \left\| \int_0^\infty e^{\Xi_\theta t} e^{\Xi_\theta^\top t} dt \right\|_F$ yields $\sup_{\theta \in \mathcal{S}(\ell)} \left\| \int_0^\infty e^{\Xi_\theta t} e^{\Xi_\theta^\top t} dt \right\|_F < +\infty$. The triangle inequality together with Proposition C.1 gives

$$\|\theta' - \theta\|_F \leq \|\theta'\|_F + \|\theta\|_F \leq 2\mathbf{Bd}_\theta(\ell)$$

so that

$$\|E_\theta^\top(\theta' - \theta) + (\theta' - \theta)^\top E_\theta + (\theta' - \theta)^\top R(\theta' - \theta)\|_F \leq 2(\mathbf{Bd}_E(\ell) + \|R\|_F \mathbf{Bd}_\theta(\ell))\|\theta' - \theta\|_F.$$

Gathering the previous bounds allows to conclude that $\theta \mapsto K_\theta$ is $Lip_K(\ell)$ -Lipschitz continuous on $\mathcal{S}(\ell)$. Similar arguments yield the Lipschitz regularity of Λ . We omit the remaining technical details. \square

C.3.3 Proof of Proposition 3.3

In order to establish the Lipschitz regularity of ∇J_1 and ∇J_2 on $\mathcal{S}(\ell)$, we start from the representation provided by Proposition 3.1. Namely, one has

$$\nabla J_1(\theta) = 2E_\theta \Sigma_\theta$$

where $E_\theta = R\theta + D^\top K_\theta$. Hence,

$$\nabla J_1(\theta') - \nabla J_1(\theta) = 2(E_{\theta'} - E_\theta)\Sigma_{\theta'} + 2E_\theta(\Sigma_{\theta'} - \Sigma_\theta).$$

Noting that

$$E_{\theta'} - E_\theta = R(\theta' - \theta) + D^\top(K_{\theta'} - K_\theta)$$

and using Proposition C.3, we obtain

$$\begin{aligned} \|E_{\theta'} - E_\theta\|_F &\leq \|R\|_F\|\theta' - \theta\|_F + \|D\|_F\|K_{\theta'} - K_\theta\|_F \\ &\leq (\|R\|_F + \|D\|_F Lip_K(\ell))\|\theta' - \theta\|_F. \end{aligned}$$

The previous bound, along again with Proposition C.3, provides

$$\begin{aligned} \|\nabla J_1(\theta') - \nabla J_1(\theta)\|_F &\leq 2\|E_{\theta'} - E_\theta\|_F\|\Sigma_{\theta'}\|_F + 2\|E_\theta\|_F\|\Sigma_{\theta'} - \Sigma_\theta\|_F \\ &\leq 2((\|R\|_F + \|D\|_F Lip_K(\ell))\mathbf{Bd}_\Sigma(\ell) + \mathbf{Bd}_E(\ell) Lip_\Sigma(\ell))\|\theta' - \theta\|_F. \end{aligned}$$

Using similar arguments, we find that for all $\zeta, \zeta' \in \hat{\mathcal{S}}(\hat{\ell})$,

$$\|\nabla J_2(\zeta') - \nabla J_2(\zeta)\|_F \leq \hat{L}(\hat{\ell})\|\zeta' - \zeta\|_F$$

where

$$\hat{L}(\hat{\ell}) = 2((\|R\|_F + \|D\|_F Lip_\Lambda(\hat{\ell}))\mathbf{Bd}_{\hat{\Sigma}}(\hat{\ell}) + \mathbf{Bd}_{\hat{E}}(\hat{\ell}) Lip_{\hat{\Sigma}}(\hat{\ell})).$$

D Proofs of the results for the model-free algorithm

In the first subsection, we describe the strategy of the proof of Theorem 4.1. In Subsection D.2, we provide its proof. Several auxiliary results are postponed to the next subsections.

D.1 Strategy of proof of Theorem 4.1

We first solve the well-posedness problem of the perturbation of the parameters. The following lemma is directly taken from Lemma 4 of [16].

Lemma D.1. *There exists $\check{r}(b)$ depending only upon $b \in \mathbb{R}_+$, m , d and the model parameters such that for all $0 < r \leq \check{r}(b)$, for any $\mathbf{U} = (U, V) \in H$, such that $\|U\|_F \leq r$, $\|V\|_F \leq r$ and for all $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$, we have*

$$\Theta + \mathbf{U} = (\theta + U, \zeta + V) \in \mathcal{R}(2b).$$

From now on, we assume that $r > 0$ is small enough ($r < \check{r}(b)$, $\check{r}(b)$ defined in Lemma D.1), so that for all $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$, for all U, V with $\|U\|_F = \|V\|_F = r$, $(\theta + U, \zeta + V) \in \mathcal{R}(2b)$.

The main ingredients of the proof can be described as follows. Denoting by $\tilde{\nabla}^{\Delta, N, pop} J(\Theta) = (\tilde{\nabla}_\theta^{\Delta, N, pop} J(\Theta), \tilde{\nabla}_\zeta^{\Delta, N, pop} J(\Theta)) \in H$ where $\tilde{\nabla}_\theta^{\Delta, N, pop} J(\Theta)$ and $\tilde{\nabla}_\zeta^{\Delta, N, pop} J(\Theta)$ are the results of Algorithm 1, the key idea is to show that the two following events

$$\begin{aligned} \mathbb{M}_1 &= \left\{ \langle \tilde{\nabla}^{\Delta, N, pop} J(\Theta), \nabla J(\Theta) \rangle_H \geq \nu_1 \|\nabla J(\Theta)\|_H^2 \right\} \\ \mathbb{M}_2 &= \left\{ \|\tilde{\nabla}^{\Delta, N, pop} J(\Theta)\|_H^2 \leq \nu_2 \|\nabla J(\Theta)\|_H^2 \right\} \end{aligned} \tag{D.1}$$

occur together with high probability, for some $\nu_1, \nu_2 > 0$. More precisely, by writing

$$\begin{aligned} \langle \tilde{\nabla}^{\Delta, N, pop} J(\Theta), \nabla J(\Theta) \rangle_H &= \|\nabla J(\Theta)\|_H^2 + \langle (\tilde{\nabla}^{\Delta, N, pop} - \nabla) J(\Theta), \nabla J(\Theta) \rangle_H \\ &\geq \|\nabla J(\Theta)\|_H^2 - \|(\tilde{\nabla}^{\Delta, N, pop} - \nabla) J(\Theta)\|_H \|\nabla J(\Theta)\|_H \end{aligned}$$

and

$$\|\tilde{\nabla}^{\Delta, N, pop} J(\Theta)\|_H \leq \|\nabla J(\Theta)\|_H + \|(\tilde{\nabla}^{\Delta, N, pop} - \nabla) J(\Theta)\|_H$$

we notice that if the event

$$\|(\tilde{\nabla}^{\Delta, N, pop} - \nabla) J(\Theta)\|_H \leq \frac{1}{2} \|\nabla J(\Theta)\|_H \tag{D.2}$$

occurs with high probability, then so does $\mathbb{M}_1 \cap \mathbb{M}_2$ with $\nu_1 = \frac{1}{2}$ and $\nu_2 = \frac{9}{4}$.

Next, for any fixed $\varepsilon > 0$, the P.-L. inequality (4.3) guarantees that if $\Theta \in \mathcal{R}$ is such that $J(\Theta) - J(\Theta^*) > \varepsilon$, then (D.2) is satisfied as soon as

$$\|(\tilde{\nabla}^{\Delta, N, pop} - \nabla) J(\Theta)\|_H \leq \varepsilon'$$

with $\varepsilon' = \frac{1}{2}\sqrt{\frac{\varepsilon}{\kappa}}$. Indeed, it suffices to notice that

$$\|(\tilde{\nabla}^{\Delta,N,pop} - \nabla)J(\Theta)\|_H \leq \varepsilon' = \frac{1}{2}\sqrt{\frac{\varepsilon}{\kappa}} \leq \frac{1}{2}\sqrt{\frac{J(\Theta) - J(\Theta^*)}{\kappa}} \leq \frac{1}{2}\|\nabla J(\Theta)\|_H.$$

Noting that

$$\|(\tilde{\nabla}^{\Delta,N,pop} - \nabla)J(\Theta)\|_H = \sqrt{\|(\tilde{\nabla}_\theta^{\Delta,N,pop} - \nabla_\theta)J(\Theta)\|_F^2 + \|(\tilde{\nabla}_\zeta^{\Delta,N,pop} - \nabla_\zeta)J(\Theta)\|_F^2}$$

we will show that for all $\varepsilon > 0$, the event

$$\|(\tilde{\nabla}_\theta^{\Delta,N,pop} - \nabla_\theta)J(\Theta)\|_F \leq \varepsilon \quad \text{and} \quad \|(\tilde{\nabla}_\zeta^{\Delta,N,pop} - \nabla_\zeta)J(\Theta)\|_F \leq \varepsilon$$

holds with high probability as soon as the parameters T, n, N, \tilde{N} are large enough and r is small enough.

D.2 Proof of Theorem 4.1

Step 1: The central idea is to use the following decomposition

$$\begin{aligned} (\tilde{\nabla}_\theta^{\Delta,N,pop} - \nabla_\theta)J(\Theta) &= (\tilde{\nabla}_\theta^{\Delta,N,pop} - \hat{\nabla}_\theta^{\Delta,N,pop})J(\Theta) + (\hat{\nabla}_\theta^{\Delta,N,pop} - \hat{\nabla}_\theta^{(T),\Delta})J(\Theta) \\ &\quad + (\hat{\nabla}_\theta^{(T),\Delta} - \hat{\nabla}_\theta^{(T)})J(\Theta) + (\hat{\nabla}_\theta^{(T)} - \hat{\nabla}_\theta)J(\Theta) \\ &\quad + (\hat{\nabla}_\theta - \nabla_\theta)J(\Theta) \\ &=: \mathbf{E}_\theta^{(1)} + \mathbf{E}_\theta^{(2)} + \mathbf{E}_\theta^{(3)} + \mathbf{E}_\theta^{(4)} + \mathbf{E}_\theta^{(5)}, \end{aligned}$$

and similarly

$$\begin{aligned} \|(\tilde{\nabla}_\zeta^{\Delta,N,pop} - \nabla_\zeta)J(\Theta)\|_F &= (\tilde{\nabla}_\zeta^{\Delta,N,pop} - \hat{\nabla}_\zeta^{\Delta,N,pop})J(\Theta) + (\hat{\nabla}_\zeta^{\Delta,N,pop} - \hat{\nabla}_\zeta^{(T),\Delta})J(\Theta) \\ &\quad + (\hat{\nabla}_\zeta^{(T),\Delta} - \hat{\nabla}_\zeta^{(T)})J(\Theta) + (\hat{\nabla}_\zeta^{(T)} - \hat{\nabla}_\zeta)J(\Theta) \\ &\quad + (\hat{\nabla}_\zeta - \nabla_\zeta)J(\Theta) \\ &=: \mathbf{E}_\zeta^{(1)} + \mathbf{E}_\zeta^{(2)} + \mathbf{E}_\zeta^{(3)} + \mathbf{E}_\zeta^{(4)} + \mathbf{E}_\zeta^{(5)}, \end{aligned}$$

where we introduced the following notations which will be useful for our convergence analysis:

- The perturbed policy gradients:

$$\hat{\nabla}_\theta J(\Theta) = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} J(\Theta_i) U_i \quad \text{and} \quad \hat{\nabla}_\zeta J(\Theta) = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} J(\Theta_i) V_i.$$

- The truncated policy gradients:

$$\hat{\nabla}_\theta^{(T)} J(\Theta) = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} J^{(T)}(\Theta_i) U_i \quad \text{and} \quad \hat{\nabla}_\zeta^{(T)} J(\Theta) = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} J^{(T)}(\Theta_i) V_i.$$

- The policy gradients with time discretization:

$$\hat{V}_\theta^{(T),\Delta} J(\Theta) = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} J^{(T),\Delta}(\Theta_i) U_i \quad \text{and} \quad \hat{V}_\zeta^{(T)} J(\Theta) = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} J^{(T),\Delta}(\Theta_i) V_i.$$

- The output of Algorithm 1 averaged over the randomness of the state trajectories:

$$\begin{aligned} \hat{V}_\theta^{\Delta,N,pop} J(\Theta) &= \mathbb{E}[\tilde{V}_\theta^{\Delta,N,pop} J(\Theta) | \underline{U}, \underline{V}], \\ \hat{V}_\zeta^{\Delta,N,pop} J(\Theta) &= \mathbb{E}[\tilde{V}_\zeta^{\Delta,N,pop} J(\Theta) | \underline{U}, \underline{V}], \end{aligned} \tag{D.3}$$

where $\underline{U} = (U_1, \dots, U_{\tilde{N}})$ and $\underline{V} = (V_1, \dots, V_{\tilde{N}})$.

In the above notations, given a finite horizon $T > 0$ (to be chosen later on) and $\Theta = (\theta, \zeta) \in H$, we introduced the following time-truncated expected cost functional

$$\begin{aligned} J^{(T)}(\Theta) &= \mathbb{E} \left[\int_0^T e^{-\beta t} ((Y_t^\Theta)^\top Q Y_t^\Theta + (Z_t^\Theta)^\top \hat{Q} Z_t^\Theta \right. \\ &\quad \left. + \int (a^\top R a + \lambda \log p^\Theta(Y_t^\Theta, Z_t^\Theta, a)) p^\Theta(Y_t^\Theta, Z_t^\Theta, a) da) dt \right], \end{aligned}$$

and also introduced the corresponding time discretized cost functional, with action execution noise

$$\begin{aligned} J^{(T),\Delta}(\Theta) &= h \mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\mathbf{x}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}])^\top Q (\mathbf{x}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}]) + \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}]^\top \hat{Q} \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}] \right. \right. \\ &\quad \left. \left. + (\alpha_{t_l}^{\Theta,\Delta})^\top R \alpha_{t_l}^{\Theta,\Delta} + \lambda \log p^\Theta(\mathbf{x}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}], \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}], \alpha_{t_l}^{\Theta,\Delta}) \right) \right], \end{aligned} \tag{D.4}$$

where $(\mathbf{x}_{t_l}^{\Theta,\Delta})_{0 \leq l \leq n}$ is the time-discretization scheme of $(X_t^\Theta)_{t \in [0, T]}$, over the same time grid Δ as the interacting agents dynamics (4.4), with dynamics

$$\begin{cases} \mathbf{x}_{t_{l+1}}^{\Theta,\Delta} &= \mathbf{x}_{t_l}^{\Theta,\Delta} + (B \mathbf{x}_{t_l}^{\Theta,\Delta} + \bar{B} \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}] + D \alpha_{t_l}^{\Theta,\Delta}) h \\ &\quad + \gamma \sqrt{h} \mathbf{w}_l + \gamma_0 \sqrt{h} \mathbf{w}_l^0, \\ \mathbf{x}_0^{\Theta,\Delta} &= X_0 \end{cases} \tag{D.5}$$

with $\alpha_{t_l}^{\Theta,\Delta} = \theta(\mathbf{x}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}]) + \zeta \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}] + \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}$, $(\boldsymbol{\xi}_{t_l})_{0 \leq l \leq n-1}$ being *i.i.d.* random variables with law $\mathcal{N}(0, \mathbb{I}_m)$.

Step 2: Denoting by $\tilde{\varepsilon} = \frac{\varepsilon'}{5\sqrt{2}} = \frac{1}{10\sqrt{2}} \sqrt{\frac{\varepsilon}{\kappa}}$, choosing $r < \min(\tilde{r}(b), \frac{1}{h_r(b, 1/\tilde{\varepsilon})}, \frac{1}{h_r(b, 1/\tilde{\varepsilon})})$ and $\tilde{N} > h_{\tilde{N}}(b, \frac{1}{\tilde{\varepsilon}})$, Proposition D.2 guarantees that with probability $1 - 4(\frac{d}{\tilde{\varepsilon}})^{-d}$

$$\|\mathbb{E}_\theta^5\|_{\mathbb{F}} \leq \tilde{\varepsilon} \quad \text{and} \quad \|\mathbb{E}_\zeta^5\|_{\mathbb{F}} \leq \tilde{\varepsilon}.$$

Choosing $T = \check{T}(\varepsilon)$ as in the statement of Theorem 4.1, according to Proposition D.3, it holds

$$\|\mathbf{E}_\theta^4\|_{\mathbb{F}} \leq \tilde{\varepsilon} \quad \text{and} \quad \|\mathbf{E}_\zeta^4\|_{\mathbb{F}} \leq \tilde{\varepsilon}, \quad \mathbb{P} - \text{a.s.}$$

Next choosing $n = \check{n}(\varepsilon)$ as in the statement of Theorem 4.1, it follows from Proposition D.4 that

$$\|\mathbf{E}_\theta^3\|_{\mathbb{F}} \leq \tilde{\varepsilon} \quad \text{and} \quad \|\mathbf{E}_\zeta^3\|_{\mathbb{F}} \leq \tilde{\varepsilon}, \quad \mathbb{P} - \text{a.s.}$$

Choosing $N = \check{N}(\varepsilon)$ large enough (again as in the statement of Theorem 4.1), according to estimates (D.9) in Proposition D.6, it holds

$$\|\mathbf{E}_\theta^2\|_{\mathbb{F}} \leq \tilde{\varepsilon} \quad \text{and} \quad \|\mathbf{E}_\zeta^2\|_{\mathbb{F}} \leq \tilde{\varepsilon}, \quad \mathbb{P} - \text{a.s.}$$

Finally, according to Proposition D.5, up to a modification of \tilde{N} , namely taking \tilde{N} large enough such that

$$\tilde{N} > \max\{h_{\tilde{N}}(b, \frac{1}{\check{r}(\varepsilon)}, \frac{1}{\varepsilon}), h'_{\tilde{N}}(b, \frac{1}{\check{r}(\varepsilon)}, \check{T}(\varepsilon), \check{n}(\varepsilon), \check{N}(\varepsilon), 1/\tilde{\varepsilon})\}$$

one has

$$\|\mathbf{E}_\theta^1\|_{\mathbb{F}} \leq \tilde{\varepsilon} \quad \text{and} \quad \|\mathbf{E}_\zeta^1\|_{\mathbb{F}} \leq \tilde{\varepsilon}$$

with probability $1 - (\frac{d}{\tilde{\varepsilon}})^{-d}$.

Putting the above estimates together yields

$$\|(\tilde{\nabla}_\theta^{\Delta, N, pop} - \nabla_\theta)J(\Theta)\|_{\mathbb{F}} \leq \frac{\varepsilon'}{\sqrt{2}} \quad \text{and} \quad \|(\tilde{\nabla}_\zeta^{\Delta, N, pop} - \nabla_\zeta)J(\Theta)\|_{\mathbb{F}} \leq \frac{\varepsilon'}{\sqrt{2}}$$

so that

$$\|(\tilde{\nabla}^{\Delta, N, pop} - \nabla)J(\Theta)\|_H \leq \varepsilon'$$

with probability at least $1 - 4(\frac{d}{\tilde{\varepsilon}})^{-d}$.

Step 3: As discussed in Subsection D.1, the previous inequality implies that the probability of $\mathbb{M}_1 \cap \mathbb{M}_2$ defined in (D.1) with $\nu_1 = \frac{1}{2}$ and $\nu_2 = \frac{9}{4}$ is not smaller than $1 - 4(\frac{d}{\tilde{\varepsilon}})^{-d}$.

It then follows from Lemma D.1 that for any $\rho \in (0, \frac{4}{9L(b)})$, still with the probability not smaller than $1 - 4(\frac{d}{\tilde{\varepsilon}})^{-d}$, one has

$$\Theta_{[1]} = \Theta_{[0]} - \rho \tilde{\nabla}^{\Delta, N, pop} J(\Theta_{[0]}) \in \mathcal{R}(b)$$

and

$$J(\Theta_{[1]}) - J(\Theta^*) \leq \left(1 - \frac{\rho(4 - 9\rho\check{L}(b))}{8\bar{\kappa}(b)}\right)(J(\Theta_{[0]}) - J(\Theta^*)).$$

By induction, we eventually deduce that

$$J(\Theta_{[k]}) - J(\Theta^*) \leq \left(1 - \frac{\rho(4 - 9\rho\check{L}(b))}{8\bar{\kappa}(b)}\right)^k (J(\Theta_{[0]}) - J(\Theta^*))$$

with probability at least $1 - 4k(\frac{d}{\tilde{\varepsilon}})^{-d}$. The conclusion of Theorem 4.1 easily follows from the previous inequality.

D.3 Analysis of the single step GD algorithm

In this section, we analyze the convergence of a general single-step GD algorithm. This part serves as an important building block in the proof of Theorem 4.1. Our main result is Proposition D.1. We first prove some technical results.

Lemma D.2. *The map $(\theta, \zeta) \mapsto \nabla J(\Theta)$ is $\check{L}(b)$ -Lipschitz continuous on $\mathcal{R}(b)$, that is, for all $\Theta = (\theta, \zeta), \Theta' = (\theta', \zeta') \in \mathcal{R}(b)$,*

$$\|\nabla J(\Theta) - \nabla J(\Theta')\|_H \leq \check{L}(b)\|\Theta - \Theta'\|_H,$$

recalling that $\check{L}(b) = \max\{L(b), \hat{L}(b)\}$, where $L(b)$ and $\hat{L}(b)$ are the Lipschitz constants of $\nabla J_1(\theta)$ and $\nabla J_2(\zeta)$ on $\mathcal{S}(b)$ and $\hat{\mathcal{S}}(b)$ respectively, defined in (3.2) of Proposition 3.3.

Proof. It is clear that $\Theta \in \mathcal{R}(b)$ implies $(\theta, \zeta) \in \mathcal{S}(b) \times \hat{\mathcal{S}}(b)$ so that

$$\|\nabla J_1(\theta) - \nabla J_1(\theta')\|_F \leq L(b)\|\theta - \theta'\|_F \leq \check{L}(b)\|\theta - \theta'\|_F$$

$$\|\nabla J_2(\zeta) - \nabla J_2(\zeta')\|_F \leq \hat{L}(b)\|\zeta - \zeta'\|_F \leq \check{L}(b)\|\zeta - \zeta'\|_F$$

which in turn clearly yields

$$\begin{aligned} \|\nabla J(\Theta) - \nabla J(\Theta')\|_H^2 &= \|\nabla_\theta J(\theta, \zeta) - \nabla_\theta J(\theta', \zeta')\|_F^2 + \|\nabla_\zeta J(\theta, \zeta) - \nabla_\zeta J(\theta', \zeta')\|_F^2 \\ &= \|\nabla J_1(\theta) - \nabla J_1(\theta')\|_F^2 + \|\nabla J_2(\zeta) - \nabla J_2(\zeta')\|_F^2 \\ &\leq \check{L}^2(b)(\|\theta - \theta'\|_F^2 + \|\zeta - \zeta'\|_F^2) = \check{L}^2(b)\|\Theta - \Theta'\|_H^2. \end{aligned}$$

The proof is now complete. \square

For $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$ and $\mathbf{G} = (G_1, G_2) \in H$ satisfying

$$\langle \mathbf{G}, \nabla J(\Theta) \rangle_H \geq \nu_1 \|\nabla J(\Theta)\|_H^2, \quad \|\mathbf{G}\|_H^2 \leq \nu_2 \|\nabla J(\Theta)\|_H^2, \quad (\text{D.6})$$

we set

$$\Theta_\rho := \Theta - \rho \mathbf{G} = (\theta - \rho G_1, \zeta - \rho G_2). \quad (\text{D.7})$$

Lemma D.3. *For any $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$, any $\mathbf{G} \in H$ satisfying (D.6) and any $\rho \in (0, \frac{2\nu_1}{\nu_2 \check{L}(b)})$, Θ_ρ defined by (D.7) satisfies*

$$\Theta_\rho \in \mathcal{R}(b).$$

Proof. We let

$$\rho_{\max} = \sup\{\rho' \geq 0 \mid \Theta_{\rho'} := \Theta - \rho' \mathbf{G} = (\theta - \rho' G_1, \zeta - \rho' G_2) \in \mathcal{R}(b), \forall \rho \in [0, \rho']\}.$$

Firstly, from the first condition of \mathbf{G} , \mathbf{G} is also a descent direction of the function J , we have $\rho_{\max} > 0$. Next, $\mathcal{R}(b)$ is compact thus bounded, we have $\rho_{\max} < +\infty$.

Assume that $\rho_{\max} < \frac{2\nu_1}{\nu_2\check{L}(b)}$. By the continuity of $\rho \mapsto J(\theta - \rho\mathbf{G})$ and the definition of $\mathcal{R}(b)$, we have $J(\Theta - \rho_{\max}\mathbf{G}) = \check{J}(\Theta - \rho_{\max}\mathbf{G}) + v(\lambda) = b + v(\lambda)$. Therefore, for any $\rho \in (0, \rho_{\max}]$, we have $\Theta - \rho\mathbf{G} \in \mathcal{R}(b)$. A second order Taylor's expansion together with the Lipschitz continuity of ∇J and (D.6) gives

$$\begin{aligned} J(\Theta - \rho\mathbf{G}) &\leq J(\Theta) - \rho\langle \nabla J(\Theta), \mathbf{G} \rangle_H + \frac{\rho^2\check{L}(b)}{2}\|\mathbf{G}\|_H^2 \\ &\leq J(\Theta) - \rho\nu_1\|\nabla J(\Theta)\|_H^2 + \frac{\rho^2\nu_2\check{L}(b)}{2}\|\nabla J(\Theta)\|_H^2 \end{aligned}$$

so that

$$J(\Theta - \rho\mathbf{G}) - J(\Theta) \leq \frac{-\rho(2\nu_1 - \rho\nu_2\check{L}(b))}{2}\|\nabla J(\Theta)\|_H^2 < 0$$

which in turn clearly gives

$$J(\Theta - \rho\mathbf{G}) < J(\Theta) = \check{J}(\Theta) + v(\lambda) \leq b + v(\lambda).$$

This last inequality contradicts the fact that $J(\Theta - \rho_{\max}\mathbf{G}) = b + v(\lambda)$. We thus conclude that $\rho_{\max} \geq \frac{2\nu_1}{\nu_2\check{L}(b)}$ and $\Theta_\rho = \Theta - \rho\mathbf{G} \in \mathcal{R}(b)$ for all $\rho \in (0, \frac{2\nu_1}{\nu_2\check{L}(b)})$. \square

Proposition D.1. *For any $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$, any $G \in H$ satisfying (D.6) and any $\rho \in (0, \frac{2\nu_1}{\nu_2\check{L}(b)})$, Θ_ρ defined by (D.7) satisfies*

$$J(\Theta_\rho) - J(\Theta^*) \leq \left(1 - \frac{\rho(2\nu_1 - \rho\check{L}(b)\nu_2)}{2\bar{\kappa}(b)}\right)(J(\Theta) - J(\Theta^*)),$$

where $\bar{\kappa}(b) := \max\left(\kappa_1, \kappa_2, \frac{\nu_1^2}{2\nu_2\check{L}(b)}\right) + \frac{1}{2}$.

Proof. Writing again a second order Taylor expansion and using (D.6), we get

$$\begin{aligned} J(\Theta - \rho\mathbf{G}) &\leq J(\Theta) - \rho\langle \nabla J(\Theta), \mathbf{G} \rangle_H + \frac{\rho^2\check{L}(b)}{2}\|\mathbf{G}\|_H^2 \\ &\leq J(\Theta) - \rho\nu_1\|\nabla J(\Theta)\|_H^2 + \frac{\rho^2\nu_2\check{L}(b)}{2}\|\nabla J(\Theta)\|_H^2 \\ &\leq J(\Theta) + \left(\frac{\rho^2\nu_2\check{L}(b)}{2} - \rho\nu_1\right)\|\nabla J(\Theta)\|_H^2. \end{aligned}$$

Note that the gradient domination inequality (4.3) is satisfied with $\bar{\kappa}(b)$ instead of κ and since

$$\frac{\rho(2\nu_1 - \rho\check{L}(b)\nu_2)}{2\bar{\kappa}(b)} \in (0, 1),$$

we get

$$J(\Theta - \rho\mathbf{G}) - J(\Theta) \leq \left(\frac{\rho^2\nu_2\check{L}(b)}{2} - \rho\nu_1\right)\|\nabla J(\Theta)\|_H^2 \leq \frac{1}{\bar{\kappa}(b)}\left(\frac{\rho^2\nu_2\check{L}(b)}{2} - \rho\nu_1\right)(J(\Theta) - J(\Theta^*))$$

so that

$$J(\Theta - \rho \mathbf{G}) - J(\Theta^*) \leq \left(1 - \frac{\rho(2\nu_1 - \rho\nu_2\check{L}(b))}{2\bar{\kappa}(b)}\right)(J(\Theta) - J(\Theta^*)).$$

□

D.4 Error analysis of J and its gradient

We here study the five terms appearing in the decomposition of the error of $(\tilde{\nabla}_\theta^{\Delta, N, pop} - \nabla_\theta)J$ and $(\tilde{\nabla}_\zeta^{\Delta, N, pop} - \nabla_\zeta)J$ introduced in Step 1 of Section D.2.

- **Bounding $\|\hat{\nabla} - \nabla\|_F$ (Error of perturbation with the exact expected functional cost).**

The proof of the following result is postponed to Section D.4.1.

Proposition D.2. *For all $\varepsilon > 0$, there exist $h_r(b, \frac{1}{\varepsilon})$ with at most polynomial growth in b and $\frac{1}{\varepsilon}$ and $h_{\tilde{N}}(b, \frac{1}{r}, \frac{1}{\varepsilon})$ with at most polynomial growth in b , $\frac{1}{r}$ and $\frac{1}{\varepsilon}$ such that for all $r < \min(\check{r}(b), \frac{1}{h_r(b, \frac{1}{\varepsilon})})$, all $\tilde{N} > h_{\tilde{N}}(b, \frac{1}{r}, \frac{1}{\varepsilon})$ and all $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$,*

$$\|(\hat{\nabla}_\theta - \nabla_\theta)J(\Theta)\|_F \leq \varepsilon \quad \text{and} \quad \|(\hat{\nabla}_\zeta - \nabla_\zeta)J(\Theta)\|_F \leq \varepsilon,$$

with probability at least $1 - (d/\varepsilon)^{-d}$.

- **Bounding $\|\hat{\nabla}^{(T)} - \hat{\nabla}\|_F$ (Horizon truncation error).**

The proof of the following result is postponed to Section D.4.2.

Proposition D.3. \mathbb{P} -a.s., for all $\Theta \in \mathcal{R}(b)$, one has

$$\begin{aligned} \|(\hat{\nabla}_\theta - \hat{\nabla}_\theta^{(T)})J(\Theta)\|_F &\leq \frac{d}{r}c_1(2b)e^{-c_2(2b)T}, \\ \|(\hat{\nabla}_\zeta - \hat{\nabla}_\zeta^{(T)})J(\Theta)\|_F &\leq \frac{d}{r}c_1(2b)e^{-c_2(2b)T}. \end{aligned}$$

- **Bounding $\|\hat{\nabla}^{(T)} - \hat{\nabla}^{(T), \Delta}\|_F$ (Time discretization error).**

The proof of the following result is postponed to Section D.4.3.

Proposition D.4. *There exists a constant $c_3 = c_3(2b) > 0$ (non-decreasing with respect to b) such that \mathbb{P} -a.s., for all $\Theta \in \mathcal{R}(b)$, it holds*

$$\begin{aligned} \|(\hat{\nabla}_\theta^{(T)} - \hat{\nabla}_\theta^{(T), \Delta})J(\Theta)\|_F &\leq \frac{d}{r}c_3(2b)\frac{T}{n}, \\ \|(\hat{\nabla}_\zeta^{(T)} - \hat{\nabla}_\zeta^{(T), \Delta})J(\Theta)\|_F &\leq \frac{d}{r}c_3(2b)\frac{T}{n}. \end{aligned}$$

- **Bounding $\|\tilde{\nabla}^{\Delta,N,pop} - \hat{\nabla}^{\Delta,N,pop}\|_F$ (Statistical error).**

The proof of the following result is postponed to Section D.4.4.

Proposition D.5. *(Statistical error on gradient estimators) For all $\varepsilon > 0$, there exist $h'_r(b, \frac{1}{\varepsilon})$ and $h'_{\tilde{N}}(b, \frac{1}{r}, T, n, N, \frac{1}{\varepsilon})$ with at most polynomial growth in $\frac{1}{r}$, b , $\frac{1}{\varepsilon}$, $\log N$, $\log n$, $\log T$ such that for all $r < \min\{\tilde{r}(b), 1/h'_r(b, \frac{1}{\varepsilon})\}$ and all $\tilde{N} \geq h'_{\tilde{N}}(b, \frac{1}{r}, T, n, N, \frac{1}{\varepsilon})$, with probability at least $1 - (d/\varepsilon)^{-d}$, for all $\Theta \in \mathcal{R}(b)$, it holds*

$$\begin{aligned} \|(\tilde{\nabla}_{\theta}^{\Delta,N,pop} - \hat{\nabla}_{\theta}^{\Delta,N,pop})J(\Theta)\|_F &\leq \varepsilon, \\ \|(\tilde{\nabla}_{\zeta}^{\Delta,N,pop} - \hat{\nabla}_{\zeta}^{\Delta,N,pop})J(\Theta)\|_F &\leq \varepsilon. \end{aligned}$$

- **Bounding $\|\hat{\nabla}^{\Delta,N,pop} - \hat{\nabla}^{(T),\Delta}\|_F$ (Particle discretization error).**

The proof of the following result is postponed to Section D.4.5.

Proposition D.6. *Let $\bar{\mathcal{J}}_{pop}^{\Delta,N}(\Theta) := \mathbb{E}[\mathcal{J}_{pop}^{\Delta,N}(\Theta)]$, $\Theta \in \mathcal{R}(b)$, recalling that $\mathcal{J}_{pop}^{\Delta,N}$ is defined by (4.7). There exists $c_4 = c_4(b) > 0$ such that for all $\Theta \in \mathcal{R}(b)$,*

$$|(\bar{\mathcal{J}}_{pop}^{\Delta,N} - J^{(T),\Delta})(\Theta)| \leq \frac{c_4}{N}. \quad (\text{D.8})$$

Moreover, \mathbb{P} -a.s., for all $\Theta \in \mathcal{R}(b)$, it holds

$$\begin{aligned} \|(\hat{\nabla}_{\theta}^{\Delta,N,pop} - \hat{\nabla}_{\theta}^{(T),\Delta})J(\Theta)\|_F &\leq \frac{d}{r} \frac{c_4(2b)}{N}, \\ \|(\hat{\nabla}_{\zeta}^{\Delta,N,pop} - \hat{\nabla}_{\zeta}^{(T),\Delta})J(\Theta)\|_F &\leq \frac{d}{r} \frac{c_4(2b)}{N}. \end{aligned} \quad (\text{D.9})$$

D.4.1 Proof of Proposition D.2

The proof is reminiscent of Lemma 31 [6] or Lemma 30 [8]. For any $\Theta = (\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$ and $r > 0$, let us introduce the following smooth approximation of J defined by

$$J_r(\theta, \zeta) = \mathbb{E}_{(P_r, \hat{P}_r)}[J(\theta + P_r, \zeta + \hat{P}_r)] = \mathbb{E}_{(P_r, \hat{P}_r)}[J_1(\theta + P_r) + J_2(\zeta + \hat{P}_r)] + v(\lambda) \quad (\text{D.10})$$

where P_r and \hat{P}_r are two independent random variables uniformly distributed in \mathbb{S}_r and $\mathbb{E}_{(P_r, \hat{P}_r)}$ stands for the expectation taken with respect to the random variables P_r, \hat{P}_r .

The following result is directly taken from Lemma 26 in [8]. For any $(\theta, \zeta) \in \mathcal{S} \times \hat{\mathcal{S}}$, it holds

$$\begin{aligned} \nabla_{\theta} J_r(\theta, \zeta) &= \frac{d}{r^2} \mathbb{E}_{P_r}[J_1(\theta + P_r)P_r] = \frac{d}{r^2} \mathbb{E}_{(P_r, \hat{P}_r)}[J(\theta + P_r, \zeta + \hat{P}_r)P_r] \\ \nabla_{\zeta} J_r(\theta, \zeta) &= \frac{d}{r^2} \mathbb{E}_{\hat{P}_r}[J_2(\zeta + \hat{P}_r)\hat{P}_r] = \frac{d}{r^2} \mathbb{E}_{(P_r, \hat{P}_r)}[J(\theta + P_r, \zeta + \hat{P}_r)\hat{P}_r]. \end{aligned} \quad (\text{D.11})$$

We then write

$$\|(\hat{\nabla}_\theta - \nabla_\theta)J(\Theta)\|_F \leq \|\hat{\nabla}_\theta J(\Theta) - \nabla_\theta J_r(\Theta)\|_F + \|\nabla_\theta J_r(\Theta) - \nabla_\theta J(\Theta)\|_F.$$

For the second term $\|\nabla_\theta J_r(\Theta) - \nabla_\theta J(\Theta)\|_F$, using the Lipschitz property of the gradient $\nabla J(\Theta)$ shown in the Section C.3, there exists $\tilde{h}_r(b, \frac{1}{\varepsilon})$ with polynomial growth in b and $\frac{1}{\varepsilon}$ such that for all $r < 1/\tilde{h}_r(b, \frac{1}{\varepsilon})$, for all $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$ and $\Theta' = (\theta', \zeta')$ such that $\|\theta' - \theta\|_F \leq r$ and $\|\zeta' - \zeta\|_F \leq r$, we have $\|\nabla_\theta J(\Theta') - \nabla_\theta J(\Theta)\|_F \leq \frac{\varepsilon}{2}$. Then according to (D.10), noticing that $\nabla_\theta J_r(\Theta) = \mathbb{E}_{(P_r, \hat{P}_r)}[\nabla_\theta J(\theta + P_r, \zeta + \hat{P}_r)]$, we have

$$\|\nabla_\theta J_r(\Theta) - \nabla_\theta J(\Theta)\|_F \leq \mathbb{E}_{(P_r, \hat{P}_r)}[\|\nabla_\theta J(\theta + P_r, \zeta + \hat{P}_r) - \nabla_\theta J(\Theta)\|_F] \leq \frac{\varepsilon}{2}.$$

Then, for the first term $\|\hat{\nabla}_\theta J(\Theta) - \nabla_\theta J_r(\Theta)\|_F$, using (D.11), we notice that $\hat{\nabla}_\theta J(\Theta)$ is exactly the Monte-Carlo approximation for the expectation $\nabla_\theta J_r(\Theta)$. Assuming that $r < \check{r}(b)$, according to Lemma D.1, for all $\Theta \in \mathcal{R}(b)$, $i = 1, \dots, \tilde{N}$, $(\theta + U_i, \zeta + V_i) \in \mathcal{R}(2b)$ and each individual sample has the norm bounded by $\frac{d}{r}(2b + v(\lambda))$ according to (4.1)-(4.2). Thus we can apply vector Bernstein's Inequality to deduce that there exists $h_{sample}(b, \frac{1}{r}, \frac{1}{\varepsilon})$ with polynomial growth in b , $\frac{1}{r}$ and $\frac{1}{\varepsilon}$ such that for all $\tilde{N} \geq h_{sample}(b, \frac{1}{r}, \frac{1}{\varepsilon})$, for all $\Theta \in \mathcal{R}(b)$, we have $\|\hat{\nabla}_\theta J(\Theta) - \nabla_\theta J_r(\Theta)\|_F \leq \frac{\varepsilon}{2}$ with probability at least $1 - (d/\varepsilon)^{-d}$.

The bound for $\|(\hat{\nabla}_\zeta - \nabla_\zeta)J(\Theta)\|_F$ is derived in a similar manner, and thus its proof is omitted.

D.4.2 Proof of Proposition D.3

In order to establish our result on $\|(\hat{\nabla}_\theta - \hat{\nabla}_\theta^{(T)})J(\Theta)\|_F$ and $\|(\hat{\nabla}_\zeta - \hat{\nabla}_\zeta^{(T)})J(\Theta)\|_F$, we first need to derive an error bound on $|(J^{(T)} - J)(\Theta)|$.

Lemma D.4. *There exist two positive constants $c_1 = c_1(b)$ and $c_2 = c_2(b)$ depending only upon $b \in \mathbb{R}_+$, the model parameters and λ such that for all $\Theta \in \mathcal{R}(b)$*

$$|(J^{(T)} - J)(\Theta)| \leq c_1(b)e^{-c_2(b)T}.$$

Besides, $c_1(\cdot)$ and $c_2(\cdot)$ are non-decreasing functions.

Before proving the above lemma, we will need the following technical result which is directly taken from Lemma 12 [16]. Its proof is thus omitted.

Lemma D.5. *Let the matrices $F, X \succ 0$ and $\Omega \succ 0$ satisfy*

$$FX + XF^\top + \Omega = 0$$

then for any $t \geq 0$,

$$\|e^{Ft}\|_2^2 \leq \frac{\|X\|_2}{\sigma_{\min}(X)} e^{-\frac{\sigma_{\min}(\Omega)}{\|X\|_2} t}$$

where $\|\cdot\|_2$ denotes the largest singular value of the matrices.

Proof of Lemma D.4. From the definition of J and $J^{(T)}$, we get

$$\begin{aligned}
& (J - J^{(T)})(\Theta) \\
&= \mathbb{E} \left[\int_T^\infty e^{-\beta t} ((Y_t^\Theta)^\top Q Y_t^\Theta + (Z_t^\Theta)^\top \hat{Q} Z_t^\Theta + \int_{\mathbb{R}^m} (a^\top R a + \lambda \log p^\Theta(Y_t^\Theta, Z_t^\Theta, a)) p^\Theta(Y_t^\Theta, Z_t^\Theta, a) da) dt \right] \\
&= \mathbb{E} \left[\int_T^\infty e^{-\beta t} ((Y_t^\Theta)^\top (Q + \theta^\top R \theta) Y_t^\Theta + (Z_t^\Theta)^\top (\hat{Q} + \zeta^\top R \zeta) Z_t^\Theta) dt \right] + v^{(T)}(\lambda)
\end{aligned}$$

where

$$v^{(T)}(\lambda) = \left(-\frac{\lambda m}{2} \log(\pi \lambda) + \frac{\lambda}{2} \log |\det(R)| \right) \int_T^\infty e^{-\beta t} dt = \frac{e^{-\beta T}}{\beta} \left(-\frac{\lambda m}{2} \log(\pi \lambda) + \frac{\lambda}{2} \log |\det(R)| \right).$$

In particular, note that $v(\lambda)$ defined in (2.7) is in fact $v^{(0)}(\lambda)$ for $T = 0$.

Denoting for simplicity $\mathbf{C}_t = \mathbb{E}[Y_t^\theta (Y_t^\theta)^\top]$ and $\hat{\mathbf{C}}_t = \mathbb{E}[Z_t^\zeta (Z_t^\zeta)^\top]$, one has

$$(J - J^{(T)})(\Theta) = \langle Q + \theta^\top R \theta, \int_T^\infty e^{-\beta t} \mathbf{C}_t dt \rangle + \langle \hat{Q} + \zeta^\top R \zeta, \int_T^\infty e^{-\beta t} \hat{\mathbf{C}}_t dt \rangle + v^{(T)}(\lambda) \quad (\text{D.12})$$

As proved in Lemma B.2 (using Itô's formula),

$$\begin{aligned}
\mathbf{C}_t &= \exp(t(B + D\theta)) \mathbf{C}_0 \exp(t(B + D\theta)^\top) + \int_0^t \exp(-(s-t)(B + D\theta)) \gamma \gamma^\top \exp(-(s-t)(B + D\theta)^\top) ds, \\
\hat{\mathbf{C}}_t &= \exp(t(\hat{B} + D\zeta)) \hat{\mathbf{C}}_0 \exp(t(\hat{B} + D\zeta)^\top) + \int_0^t \exp(-(s-t)(\hat{B} + D\zeta)) \gamma_0 \gamma_0^\top \exp(-(s-t)(\hat{B} + D\zeta)^\top) ds.
\end{aligned}$$

The equations (2.1) satisfied by Σ_θ and $\hat{\Sigma}_\zeta$ together with Lemma D.5 guarantee that

$$\begin{aligned}
\|e^{(B - \frac{\beta}{2} I_d + D\theta)t}\|_2^2 &\leq \frac{\|\Sigma_\theta\|_2}{\sigma_{\min}(\Sigma_\theta)} e^{-\frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2} t}, \\
\|e^{(\hat{B} - \frac{\beta}{2} I_d + D\zeta)t}\|_2^2 &\leq \frac{\|\hat{\Sigma}_\zeta\|_2}{\sigma_{\min}(\hat{\Sigma}_\zeta)} e^{-\frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2} t},
\end{aligned}$$

so that

$$\begin{aligned}
\|e^{(B+D\theta)t}\|_2^2 &\leq \|e^{(B - \frac{\beta}{2} I_d + D\theta)t}\|_2^2 \|e^{\frac{\beta}{2} I_d t}\|_2^2 \leq \frac{\|\Sigma_\theta\|_2}{\sigma_{\min}(\Sigma_\theta)} e^{-\frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2} t} e^{2\|\frac{\beta}{2} I_d\|_2 t} = \frac{\|\Sigma_\theta\|_2}{\sigma_{\min}(\Sigma_\theta)} e^{(\beta - \frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2})t}, \\
\|e^{(\hat{B}+D\zeta)t}\|_2^2 &\leq \|e^{(\hat{B} - \frac{\beta}{2} I_d + D\zeta)t}\|_2^2 \|e^{\frac{\beta}{2} I_d t}\|_2^2 \leq \frac{\|\hat{\Sigma}_\zeta\|_2}{\sigma_{\min}(\hat{\Sigma}_\zeta)} e^{-\frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2} t} e^{2\|\frac{\beta}{2} I_d\|_2 t} = \frac{\|\hat{\Sigma}_\zeta\|_2}{\sigma_{\min}(\hat{\Sigma}_\zeta)} e^{(\beta - \frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2})t}.
\end{aligned}$$

Hence,

$$\begin{aligned}
\|\mathbf{C}_t\|_F &\leq \|\mathbf{C}_0\|_F \|e^{(B+D\theta)t}\|_2^2 + \|\gamma\|^2 \int_0^t \|e^{(B+D\theta)v}\|_2^2 dv \\
&\leq \frac{\|\Sigma_\theta\|_2}{\sigma_{\min}(\Sigma_\theta)} \left(\|\mathbf{C}_0\|_F + \frac{\|\gamma\|_F^2}{\beta - \frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2}} e^{(\beta - \frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2})t} - \frac{\|\gamma\|_F^2}{\beta - \frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2}} \right)
\end{aligned}$$

and similarly

$$\|\hat{\mathbf{C}}_t\|_F \leq \frac{\|\hat{\Sigma}_\zeta\|_2}{\sigma_{\min}(\hat{\Sigma}_\zeta)} \left((\|\hat{\mathbf{C}}_0\|_F + \frac{\|\gamma_0\|_F^2}{\beta - \frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2}}) e^{(\beta - \frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2})t} - \frac{\|\gamma_0\|_F^2}{\beta - \frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2}} \right).$$

Since $\Theta = (\theta, \zeta) \in \mathcal{R}(b) \subset \mathcal{S}(b) \times \hat{\mathcal{S}}(b)$, recalling that $\mathcal{S}(b)$ as well as $\hat{\mathcal{S}}(b)$ are compact sets, using the continuity to $\theta \mapsto \|\Sigma_\theta\|_2$ and $\zeta \mapsto \|\hat{\Sigma}_\zeta\|_2$, there exist $\tilde{c}_1(b), \tilde{c}_2(b) > 0$ such that for all $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$

$$\tilde{c}_2(b) \leq \|\Sigma_\theta\|_2 \leq \tilde{c}_1(b), \quad \tilde{c}_2(b) \leq \|\hat{\Sigma}_\zeta\|_2 \leq \tilde{c}_1(b)$$

and using the fact that $\Sigma_\theta \succeq M, \hat{\Sigma}_\theta \succeq \hat{M}$, we get

$$\begin{aligned} e^{-\beta t} \|\mathbf{C}_t\|_F &\leq \frac{\|\Sigma_\theta\|_2}{\sigma_{\min}(\Sigma_\theta)} \left((\|\mathbf{C}_0\|_F + \frac{\|\gamma\|_F^2}{\beta - \frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2}}) e^{-\frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2} t} - \frac{\|\gamma\|_F^2 e^{-\beta t}}{\beta - \frac{\sigma_{\min}(M)}{\|\Sigma_\theta\|_2}} \right) \\ &\leq \frac{\tilde{c}_1(b)}{\sigma_{\min}(M)} \left((\|\mathbf{C}_0\|_F + \frac{\|\gamma\|_F^2}{\beta - \frac{\sigma_{\min}(M)}{\tilde{c}_2(b)}}) e^{-\frac{\sigma_{\min}(M)}{\tilde{c}_2(b)} t} - \frac{\|\gamma\|_F^2 e^{-\beta t}}{\beta - \frac{\sigma_{\min}(M)}{\tilde{c}_1(b)}} \right), \end{aligned}$$

and

$$\begin{aligned} e^{-\beta t} \|\hat{\mathbf{C}}_t\|_F &\leq \frac{\|\hat{\Sigma}_\zeta\|_2}{\sigma_{\min}(\hat{\Sigma}_\zeta)} \left((\|\hat{\mathbf{C}}_0\|_F + \frac{\|\gamma_0\|_F^2}{\beta - \frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2}}) e^{-\frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2} t} - \frac{\|\gamma_0\|_F^2 e^{-\beta t}}{\beta - \frac{\sigma_{\min}(\hat{M})}{\|\hat{\Sigma}_\zeta\|_2}} \right) \\ &\leq \frac{\tilde{c}_1(b)}{\sigma_{\min}(\hat{M})} \left((\|\hat{\mathbf{C}}_0\|_F + \frac{\|\gamma_0\|_F^2}{\beta - \frac{\sigma_{\min}(\hat{M})}{\tilde{c}_2(b)}}) e^{-\frac{\sigma_{\min}(\hat{M})}{\tilde{c}_1(b)} t} - \frac{\|\gamma_0\|_F^2 e^{-\beta t}}{\beta - \frac{\sigma_{\min}(\hat{M})}{\tilde{c}_1(b)}} \right). \end{aligned}$$

By using again the fact that $\mathcal{R}(b)$ is compact, there exists a constant $\tilde{c}_3(b)$ such that for all $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$:

$$\max\{\|Q + \theta^\top R\theta\|_F, \|\hat{Q} + \zeta^\top R\zeta\|_F\} \leq \tilde{c}_3(b).$$

Combing back to (D.12) and plugging the above estimates, we obtain for all $\Theta \in \mathcal{R}(b)$

$$\begin{aligned} |(J^{(T)} - J)(\Theta)| &\leq \|Q + \theta^\top R\theta\|_F \int_T^\infty e^{-\beta t} \|\mathbf{C}_t\|_F dt + \|\hat{Q} + \zeta^\top R\zeta\|_F \int_T^\infty e^{-\beta t} \|\hat{\mathbf{C}}_t\|_F dt + |v^{(T)}(\lambda)| \\ &\leq \tilde{c}_3(b) C_1(b) \int_T^\infty e^{-C_2(b)t} dt + \tilde{c}_3(b) C_3(b) \int_T^\infty e^{-C_4(b)t} dt + O_\lambda(e^{-\beta T}) \\ &\leq c_1(b) e^{-c_2(b)T}, \end{aligned}$$

for some $c_1(b) > 0$ and $c_2(b) > 0$ depending only upon b and λ .

Finally, the monotonicity of c_1 and c_2 with respect to b is a consequence of the fact that if $b_1 \leq b_2$ then $\mathcal{R}(b_1) \subset \mathcal{R}(b_2)$. \square

Proof of Proposition D.3. Noticing that \mathbb{P} -a.s. $(\theta + U_i, \zeta + V_i) \in \mathcal{R}(2b)$, from Lemma D.4,

$$|(J^{(T)} - J)(\Theta_i)| \leq c_1(2b)e^{-c_2(2b)T}, \quad \mathbb{P}\text{-a.s.}$$

and, recalling that $\|U_i\|_F = \|V_i\|_F = r$, \mathbb{P} -a.s. it holds

$$\begin{aligned} \|(\hat{\nabla}_\theta - \hat{\nabla}_\theta^{(T)})J(\Theta_i)\|_F &\leq \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} |(J^{(T)} - J)(\Theta_i)| \|U_i\|_F \leq \frac{d}{r} c_1(2b)e^{-c_2(2b)T}, \\ \|(\hat{\nabla}_\zeta - \hat{\nabla}_\zeta^{(T)})J(\Theta_i)\|_F &\leq \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} |(J^{(T)} - J)(\Theta_i)| \|V_i\|_F \leq \frac{d}{r} c_1(2b)e^{-c_2(2b)T}. \end{aligned}$$

□

D.4.3 Proof of Proposition D.4

We start with the following technical result related to the weak discretization error on the cost value function.

Lemma D.6. *There exists a constant $c_3 = c_3(b) > 0$ (non-decreasing with respect to b) such that for all $\Theta \in \mathcal{R}(b)$, it holds*

$$|(J^{(T),\Delta} - J^{(T)})(\Theta)| \leq c_3(b)h = c_3(b)\frac{T}{n}.$$

Proof. Step 1: We introduce the two processes $(\mathbf{y}_{t_l}^{\Theta,\Delta})_{0 \leq l \leq n}$ and $(\mathbf{z}_{t_l}^{\Theta,\Delta})_{0 \leq l \leq n}$ defined by

$$\mathbf{y}_{t_l}^{\Theta,\Delta} = \mathbf{x}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}] \text{ and } \mathbf{z}_{t_l}^{\Theta,\Delta} = \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}],$$

with dynamics

$$\begin{aligned} \mathbf{y}_{t_{l+1}}^{\Theta,\Delta} &= \mathbf{y}_{t_l}^{\Theta,\Delta} + ((B + D\theta)\mathbf{y}_{t_l}^{\Theta,\Delta} + D\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l})h + \sqrt{h}\gamma\mathbf{w}_l, \quad \mathbf{y}_0^{\Theta,\Delta} = X_0 - \mathbb{E}_0[X_0], \\ \mathbf{z}_{t_{l+1}}^{\Theta,\Delta} &= \mathbf{z}_{t_l}^{\Theta,\Delta} + (\hat{B} + D\zeta)\mathbf{z}_{t_l}^{\Theta,\Delta}h + \sqrt{h}\gamma_0\mathbf{w}_l^0, \quad \mathbf{z}_0^{\Theta,\Delta} = \mathbb{E}_0[X_0]. \end{aligned} \quad (\text{D.13})$$

Recalling (D.4) together with the dynamics defined in (D.5), one has

$$J^{(T),\Delta}(\Theta) = h\mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\mathbf{y}_{t_l}^{\Theta,\Delta})^\top (Q + \theta^\top R\theta)\mathbf{y}_{t_l}^{\Theta,\Delta} + (\mathbf{z}_{t_l}^{\Theta,\Delta})^\top (\hat{Q} + \zeta^\top R\zeta)\mathbf{z}_{t_l}^{\Theta,\Delta} + \beta v(\lambda) \right) \right].$$

We also define another auxiliary functional cost:

$$\begin{aligned} \tilde{J}^{(T),\Delta}(\Theta) &= h\mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}])^\top Q (\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}]) + \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}]^\top \hat{Q} \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}] \right. \right. \\ &\quad \left. \left. + \int (a^\top R a + \lambda \log p^\Theta(\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}], \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}], a)) p^\Theta(\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}], \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}], a) da \right) \right], \end{aligned} \quad (\text{D.14})$$

where $(\tilde{\mathbf{x}}_t^{\Theta,\Delta})_{0 \leq l \leq n}$ is the time-discretization scheme of $(X_t^\Theta)_{t \in [0, T]}$, over the same time grid Δ as the interacting agents dynamics (4.4), with dynamics

$$\begin{cases} \tilde{\mathbf{x}}_{t_{l+1}}^{\Theta,\Delta} &= \tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} + (B\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} + \bar{B}\mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}] + D \int a\pi^\Theta(\mathrm{d}\mathbf{a}|\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} - \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}], \mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}]))h \\ &\quad + \sqrt{h}\gamma\mathbf{w}_l + \sqrt{h}\gamma_0\mathbf{w}_l^0; \\ &= \tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} + ((B + D\theta)\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta} + (\bar{B} - D\theta + D\zeta)\mathbb{E}_0[\tilde{\mathbf{x}}_{t_l}^{\Theta,\Delta}])h + \sqrt{h}\gamma\mathbf{w}_l + \sqrt{h}\gamma_0\mathbf{w}_l^0 \\ \tilde{\mathbf{x}}_0^{\Theta,\Delta} &= X_0 \end{cases}$$

with $\Theta = (\theta, \zeta) \in \mathcal{R}(b)$.

Let us also introduce the two processes $(\tilde{\mathbf{y}}_t^{\Theta,\Delta})_{l=0,\dots,n}$ and $(\tilde{\mathbf{z}}_t^{\Theta,\Delta})_{l=0,\dots,n}$ defined by

$$\tilde{\mathbf{y}}_t^{\Theta,\Delta} = \tilde{\mathbf{x}}_t^{\Theta,\Delta} - \mathbb{E}_0[\tilde{\mathbf{x}}_t^{\Theta,\Delta}], \quad \tilde{\mathbf{z}}_t^{\Theta,\Delta} = \mathbb{E}_0[\tilde{\mathbf{x}}_t^{\Theta,\Delta}], \quad (\text{D.15})$$

with dynamics

$$\begin{aligned} \tilde{\mathbf{y}}_{t_{l+1}}^{\Theta,\Delta} &= \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta} + (B + D\theta)\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta}h + \sqrt{h}\gamma\mathbf{w}_l, & \tilde{\mathbf{y}}_0^{\Theta,\Delta} &= X_0 - \mathbb{E}_0[X_0], \\ \tilde{\mathbf{z}}_{t_{l+1}}^{\Theta,\Delta} &= \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta} + (\hat{B} + D\zeta)\tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta}h + \sqrt{h}\gamma_0\mathbf{w}_l^0, & \tilde{\mathbf{z}}_0^{\Theta,\Delta} &= \mathbb{E}_0[X_0]. \end{aligned}$$

Note that $\tilde{J}^{(T),\Delta}$ in (D.14), can be written using the dynamics of $(\tilde{\mathbf{y}}_t^{\Theta,\Delta})_{l=0,\dots,n}$ and $(\tilde{\mathbf{z}}_t^{\Theta,\Delta})_{l=0,\dots,n}$. Namely, one has

$$\begin{aligned} \tilde{J}^{(T),\Delta}(\Theta) &= h\mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta})^\top Q \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta} + (\tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta})^\top \hat{Q} \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta} \right. \right. \\ &\quad \left. \left. + \int_{\mathcal{A}} (a^\top R a + \lambda \log p^\Theta(\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta}, \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta}, z)) p^\Theta(\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta}, \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta}, z) \mathrm{d}z \right) \right] \\ &= h\mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta})^\top (Q + \theta^\top R \theta) \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta} + (\tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta})^\top (\hat{Q} + \zeta^\top R \zeta) \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta} \right. \right. \\ &\quad \left. \left. + \left(-\frac{\lambda m}{2} \log(\pi \lambda) + \frac{\lambda}{2} \log |\det(R)| \right) \right) \right] \\ &= h\mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta})^\top (Q + \theta^\top R \theta) \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta} + (\mathbf{z}_{t_l}^{\Theta,\Delta})^\top (\hat{Q} + \zeta^\top R \zeta) \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta} \right) \right] \\ &\quad + \beta v(\lambda) h \sum_{l=0}^{n-1} e^{-\beta t_l}. \end{aligned}$$

Step 2: We first prove an upper-bound on $|(\tilde{J}^{(T),\Delta} - J^{(T),\Delta})(\Theta)|$. From (D.13)-(D.15), we get that $\tilde{\mathbf{z}}_t^{\Theta,\Delta} = \mathbf{z}_t^{\Theta,\Delta}$, for all $l = 0, \dots, n$ and

$$(\mathbf{y}_{t_{l+1}}^{\Theta,\Delta} - \tilde{\mathbf{y}}_{t_{l+1}}^{\Theta,\Delta}) = (\mathbf{y}_{t_l}^{\Theta,\Delta} - \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta}) + ((B + D\theta)(\mathbf{y}_{t_l}^{\Theta,\Delta} - \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta}) + D\sqrt{\frac{\lambda}{2}}R^{-1}\xi_{t_l})h, \quad \mathbf{y}_0^{\Theta,\Delta} - \tilde{\mathbf{y}}_0^{\Theta,\Delta} = 0,$$

so that

$$\mathbf{y}_{t_l}^{\Theta, \Delta} - \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta} = \sum_{l'=0}^{l-1} (I_d + h(B + D\theta))^{l-1-l'} hD \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_{l'}}.$$

Therefore,

$$\begin{aligned} & \mathbb{E}[(\mathbf{y}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \mathbf{y}_{t_l}^{\Theta, \Delta} - (\tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta}] \\ &= 2\mathbb{E}[(\mathbf{y}_{t_l}^{\Theta, \Delta} - \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta}] + \mathbb{E}[(\mathbf{y}_{t_l}^{\Theta, \Delta} - \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) (\mathbf{y}_{t_l}^{\Theta, \Delta} - \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})] \\ &= 2\mathbb{E}\left[\left(\sum_{l'=0}^{l-1} (I_d + h(B + D\theta))^{l-1-l'} hD \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_{l'}}\right)^\top (Q + \theta^\top R\theta) \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta}\right] \\ &+ \mathbb{E}\left[\left(\sum_{l'=0}^{l-1} (I_d + h(B + D\theta))^{l-1-l'} hD \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_{l'}}\right)^\top (Q + \theta^\top R\theta) \left(\sum_{l'=0}^{l-1} (I_d + h(B + D\theta))^{l-1-l'} hD \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_{l'}}\right)\right] \\ &=: \mathbf{A}_1 + \mathbf{A}_2. \end{aligned}$$

Since $(\boldsymbol{\xi}_{t_l})_{0 \leq l \leq n}$ is independent of $(\tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})_{0 \leq l \leq n}$ which has zero-mean, we have $\mathbf{A}_1 = 0$. As for \mathbf{A}_2 , also using the independence of $(\boldsymbol{\xi}_{t_l})_{l=0, \dots, n}$, we have for any $S \in \mathbb{S}_{\geq 0}^m$, $\mathbb{E}[\boldsymbol{\xi}_{t_i}^\top S \boldsymbol{\xi}_{t_j}] = \delta_{i,j} \text{tr}(S)$, $i, j = 0, \dots, n$, $\delta_{i,j} = 1_{i=j}$, so that

$$\begin{aligned} \mathbf{A}_2 &= \sum_{i=0}^{l-1} \sum_{j=0}^{l-1} \mathbb{E}\left[\left((I_d + h(B + D\theta))^{l-1-i} hD \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_i}\right)^\top \right. \\ &\quad \left. (Q + \theta^\top R\theta) \left((I_d + h(B + D\theta))^{l-1-j} hD \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_j}\right)\right] \\ &= h^2 \sum_{l'=0}^{l-1} \text{tr}\left(\left((I_d + h(B + D\theta))^{l-1-l'} D \sqrt{\frac{\lambda}{2} R^{-1}}\right)^\top (Q + \theta^\top R\theta) \left((I_d + h(B + D\theta))^{l-1-l'} D \sqrt{\frac{\lambda}{2} R^{-1}}\right)\right). \end{aligned}$$

Consequently, there exists a constant $C_1 > 0$ depending only upon T, Q, D, λ, R and Θ (with at most of polynomial growth in $\|\Theta\|$) such that $|\mathbf{A}_2| \leq \frac{C_1}{n}$ for all $l = 0, \dots, n$. Hence, up to a modification of C_1 , for all $\Theta \in \mathcal{R}(b)$, it holds

$$\begin{aligned} |(\tilde{J}^{(T), \Delta} - J^{(T), \Delta})(\Theta)| &= \left| h \mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} \left((\mathbf{y}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \mathbf{y}_{t_l}^{\Theta, \Delta} - (\tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta} \right) \right] \right| \\ &= \left| h \sum_{l=0}^{n-1} e^{-\beta t_l} \mathbb{E} \left[(\mathbf{y}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \mathbf{y}_{t_l}^{\Theta, \Delta} - (\tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta})^\top (Q + \theta^\top R\theta) \tilde{\mathbf{y}}_{t_l}^{\Theta, \Delta} \right] \right| \\ &\leq h \sum_{l=0}^{n-1} \frac{C_1(\theta)}{n} \\ &\leq \frac{C_1(\theta)}{n}. \end{aligned} \tag{D.16}$$

Step 3: We now establish an upper-bound for $|(J^{(T)} - \tilde{J}^{(T),\Delta})(\Theta)|$. Let us recall that

$$J^{(T)}(\Theta) = \mathbb{E} \left[\int_0^T e^{-\beta t} ((Y_t^\Theta)^\top (Q + \theta^\top R \theta) Y_t^\Theta + (Z_t^\Theta)^\top (\hat{Q} + \zeta^\top R \zeta) Z_t^\Theta) dt \right] + \beta v(\lambda) \int_0^T e^{-\beta t} dt.$$

Note carefully that $(\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta})_{0 \leq l \leq n}$ (resp. $(\tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta})_{0 \leq l \leq n}$) is the time discretized version of the process $(Y_t^\Theta)_{t \in [0, T]}$ (resp. $(Z_t^\Theta)_{t \in [0, T]}$) with linear drift and constant diffusion coefficients. Hence, it follows from standard results on the weak approximation error (see e.g. [17], [18]) that

$$\begin{aligned} & \left| \mathbb{E} \left[\int_0^T e^{-\beta t} ((Y_t^\Theta)^\top (Q + \theta^\top R \theta) Y_t^\Theta + (Z_t^\Theta)^\top (\hat{Q} + \zeta^\top R \zeta) Z_t^\Theta) dt \right] \right. \\ & \left. - \mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} h ((\tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta})^\top (Q + \theta^\top R \theta) \tilde{\mathbf{y}}_{t_l}^{\Theta,\Delta} + (\tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta})^\top (\hat{Q} + \zeta^\top R \zeta) \tilde{\mathbf{z}}_{t_l}^{\Theta,\Delta}) \right] \right| \leq c_3(\Theta) h, \end{aligned}$$

for some $\Theta \mapsto c_3(\Theta)$ with at most polynomial growth. Moreover, standard results on Riemann integrals give $|h \sum_{l=0}^{n-1} e^{-\beta t_l} - \int_0^T e^{-\beta t} dt| \leq Ch$.

Combining the two previous bounds with (D.16), we conclude that there exists $c_3 = c_3(\Theta) > 0$ (with at most polynomial growth in Θ) such that

$$|(J^{(T),\Delta} - J^{(T)})(\Theta)| \leq |(J^{(T),\Delta} - \tilde{J}^{(T),\Delta})(\Theta)| + |(\tilde{J}^{(T),\Delta} - J^{(T)})(\Theta)| \leq c_3 h.$$

Since $\mathcal{R}(b)$ is compact, c_3 may be considered to depend on Θ only through b . \square

Proof of Proposition D.4 Recalling that \mathbb{P} -a.s. $(\theta + U_i, \zeta + V_i) \in \mathcal{R}(2b)$, it follows from Lemma D.6 and the fact that $\|U_i\|_F = \|V_i\|_F = r$ that \mathbb{P} -a.s.

$$\begin{aligned} \|\hat{\nabla}_\theta^{(T)} - \hat{\nabla}_\theta^{(T),\Delta} J(\Theta)\|_F &\leq \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} |(J^{(T)} - J^{(T),\Delta})(\Theta_i)| \|U_i\|_F \leq \frac{d}{r} c_3(2b) \frac{T}{n}, \\ \|\hat{\nabla}_\zeta^{(T)} - \hat{\nabla}_\zeta^{(T),\Delta} J(\Theta)\|_F &\leq \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} |(J^{(T)} - J^{(T),\Delta})(\Theta_i)| \|V_i\|_F \leq \frac{d}{r} c_3(2b) \frac{T}{n}. \end{aligned}$$

\square

D.4.4 Proof of Proposition D.5

Taking the empirical average over the particles in both sides of (4.6), we deduce that the dynamics of the process $(\hat{\mu}_{t_l}^{\Theta,\Delta,N} := \frac{1}{N} \sum_{j=1}^N X_{t_l}^{\Theta,\Delta,(j)})_{0 \leq l \leq N}$ is given by

$$\begin{aligned} \hat{\mu}_{t_{l+1}}^{\Theta,\Delta,N} &= \hat{\mu}_{t_l}^{\Theta,\Delta,N} + ((\hat{B} + D\zeta) \hat{\mu}_{t_l}^{\Theta,\Delta,N} + D \sqrt{\frac{\lambda}{2}} R^{-1} \bar{\boldsymbol{\xi}}_{t_l}^N) h + \sqrt{h} \gamma \bar{\mathbf{w}}_l^N + \sqrt{h} \gamma_0 \mathbf{w}_l^0, \\ \hat{\mu}_0^{\Theta,\Delta,N} &= \hat{\mu}_0^N := \frac{1}{N} \sum_{j=1}^N X_0^{(j)}, \end{aligned} \tag{D.17}$$

where $\bar{\boldsymbol{\xi}}_t^N = \frac{1}{N} \sum_{j=1}^N \boldsymbol{\xi}_t^{(j)}$ and $\bar{\mathbf{w}}_t^N = \frac{1}{N} \sum_{j=1}^N \mathbf{w}_t^{(j)}$. Then, let us introduce the auxiliary process $(Y_t^{\Theta, \Delta, (j)})$ defined by

$$Y_t^{\Theta, \Delta, (j)} = X_t^{\Theta, \Delta, (j)} - \hat{\mu}_t^{\Theta, \Delta, N}.$$

From (4.6) and (D.17), we get

$$\begin{aligned} Y_{t+1}^{\Theta, \Delta, (j)} &= Y_t^{\Theta, \Delta, (j)} + ((B + D\theta)Y_t^{\Theta, \Delta, (j)} + D\sqrt{\frac{\lambda}{2}}R^{-1}(\boldsymbol{\xi}_t^{(j)} - \bar{\boldsymbol{\xi}}_t^N))h + \sqrt{h}\gamma(\mathbf{w}_t^{(j)} - \bar{\mathbf{w}}_t^N), \\ Y_0^{\Theta, \Delta, (j)} &= Y_0^{(j)} := X_0^{(j)} - \hat{\mu}_0^N. \end{aligned} \tag{D.18}$$

Note that $(Y_t^{\Theta, \Delta, (j)})_{1 \leq j \leq N}$ (resp. $\hat{\mu}_t^{\Theta, \Delta, N}$) depends on Θ only through θ (resp. ζ). In order to simplify the notation, from now on, we will write $(Y_t^{\theta, \Delta, (j)})_{1 \leq j \leq N}$ and $\hat{\mu}_t^{\zeta, \Delta, N}$.

The average cost (4.7) can be decomposed as follows

$$\begin{aligned} \mathcal{J}_{pop}^{\Delta, N}(\Theta) &:= \frac{h}{N} \sum_{j=1}^N \left(\sum_{l=0}^{n-1} e^{-\beta t_l} ((X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top (Q + \theta^\top R\theta)(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}) \right. \\ &\quad + (\hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top (\hat{Q} + \zeta^\top R\zeta)\hat{\mu}_{t_l}^{\Theta, \Delta, N} + 2(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top \theta^\top R\zeta\hat{\mu}_{t_l}^{\Theta, \Delta, N} \\ &\quad + 2(\theta(X_{t_l}^{\Theta, \Delta, (j)} - \hat{\mu}_{t_l}^{\Theta, \Delta, N}) + \zeta\hat{\mu}_{t_l}^{\Theta, \Delta, N})^\top R\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)} + (\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)})^\top R\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)} \\ &\quad \left. - (\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)})^\top R\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)} - \frac{\lambda}{2} \log\left(\frac{(\pi\lambda)^m}{\det(R)}\right) \right) \\ &= \frac{h}{N} \sum_{j=1}^N \left(\sum_{l=0}^{n-1} e^{-\beta t_l} ((Y_{t_l}^{\theta, \Delta, (j)})^\top (Q + \theta^\top R\theta)Y_{t_l}^{\theta, \Delta, (j)} + (\hat{\mu}_{t_l}^{\zeta, \Delta, N})^\top (\hat{Q} + \zeta^\top R\zeta)\hat{\mu}_{t_l}^{\zeta, \Delta, N} \right. \\ &\quad \left. + 2(\theta Y_{t_l}^{\theta, \Delta, (j)} + \zeta\hat{\mu}_{t_l}^{\zeta, \Delta, N})^\top R\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)} + \beta v(\lambda) \right) \\ &= \mathcal{J}_1(\theta) + \mathcal{J}_2(\zeta) + \beta v(\lambda)h \sum_{l=0}^{n-1} e^{-\beta t_l}, \end{aligned} \tag{D.19}$$

where

$$\begin{aligned} \mathcal{J}_1(\theta) &:= \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} e^{-\beta t_l} ((Y_{t_l}^{\theta, \Delta, (j)})^\top (Q + \theta^\top R\theta)Y_{t_l}^{\theta, \Delta, (j)} + 2(Y_{t_l}^{\theta, \Delta, (j)})^\top \theta^\top R\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)}), \\ \mathcal{J}_2(\zeta) &:= h \sum_{l=0}^{n-1} e^{-\beta t_l} ((\hat{\mu}_{t_l}^{\zeta, \Delta, N})^\top (\hat{Q} + \zeta^\top R\zeta)\hat{\mu}_{t_l}^{\zeta, \Delta, N} + 2(\hat{\mu}_{t_l}^{\zeta, \Delta, N})^\top \zeta^\top R\sqrt{\frac{\lambda}{2}}R^{-1}\bar{\boldsymbol{\xi}}_t^N). \end{aligned} \tag{D.20}$$

Step 1. Recall that, for $\Theta = (\theta, \zeta)$, Algorithm 1 gives

$$\begin{aligned}\tilde{\nabla}_{\theta}^{\Delta, N, pop} J(\Theta) &= \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathcal{J}_{pop}^{\Delta, N, i} U_i = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathcal{J}_{pop}^{\Delta, N, i} U_i, \\ \tilde{\nabla}_{\zeta}^{\Delta, N, pop} J(\Theta) &= \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathcal{J}_{pop}^{\Delta, N, i} V_i = \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \mathcal{J}_{pop}^{\Delta, N, i} V_i,\end{aligned}$$

and observe that $\mathcal{J}_{pop}^{\Delta, N, i}$ can be decomposed as follows

$$\mathcal{J}_{pop}^{\Delta, N, i} = \mathcal{J}_{pop,1}^{\Delta, N, i} + \mathcal{J}_{pop,2}^{\Delta, N, i} + \beta v(\lambda) h \sum_{l=1}^n e^{-\beta t_l}$$

with

$$\begin{aligned}\mathcal{J}_{pop,1}^{\Delta, N, i} &:= \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^n e^{-\beta t_l} \left((Y_{t_l}^{\theta_i, \Delta, (j)})^\top (Q + \theta_i^\top R \theta_i) Y_{t_l}^{\theta_i, \Delta, (j)} + 2(Y_{t_l}^{\theta_i, \Delta, (j)})^\top \theta_i^\top R \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}^{i, (j)} \right), \\ \mathcal{J}_{pop,2}^{\Delta, N, i} &:= h \sum_{l=0}^{n-1} e^{-\beta t_l} \left((\hat{\mu}_{t_l}^{\zeta_i, \Delta, N})^\top (\hat{Q} + \zeta_i^\top R \zeta_i) \hat{\mu}_{t_l}^{\zeta_i, \Delta, N} + 2(\hat{\mu}_{t_l}^{\zeta_i, \Delta, N})^\top \zeta_i^\top R \sqrt{\frac{\lambda}{2} R^{-1}} \bar{\boldsymbol{\xi}}_{t_l}^{i, N} \right),\end{aligned}\tag{D.21}$$

where $(\theta_i, \zeta_i) = (\theta + U_i, \zeta + V_i)$ and the two processes $(Y_{t_l}^{\theta_i, \Delta, (j)})_{0 \leq l \leq n}$, $(\hat{\mu}_{t_l}^{\zeta_i, \Delta, N})_{0 \leq l \leq n}$ satisfy

$$\begin{aligned}Y_{t_{l+1}}^{\theta_i, \Delta, (j)} &= Y_{t_l}^{\theta_i, \Delta, (j)} + ((B + D\theta_i) Y_{t_l}^{\theta_i, \Delta, (j)} + D \sqrt{\frac{\lambda}{2} R^{-1}} (\boldsymbol{\xi}_{t_l}^{i, (j)} - \bar{\boldsymbol{\xi}}_{t_l}^{i, N})) h + \sqrt{h} \gamma (\mathbf{w}_l^{i, (j)} - \bar{\mathbf{w}}_l^{i, N}), \\ Y_0^{\theta_i, \Delta, (j)} &= Y_0^{i, (j)} := X_0^{i, (j)} - \hat{\mu}_0^{i, N}, \\ \hat{\mu}_{t_{l+1}}^{\zeta_i, \Delta, N} &= \hat{\mu}_{t_l}^{\zeta_i, \Delta, N} + ((\hat{B} + D\zeta_i) \hat{\mu}_{t_l}^{\zeta_i, \Delta, N} + D \sqrt{\frac{\lambda}{2} R^{-1}} \bar{\boldsymbol{\xi}}_{t_l}^{i, N}) h + \sqrt{h} \gamma \bar{\mathbf{w}}_l^{i, N} + \sqrt{h} \gamma_0 \mathbf{w}_l^{i, 0}, \\ \hat{\mu}_0^{\zeta_i, \Delta, N} &= \hat{\mu}_0^{i, N} := \frac{1}{N} \sum_{j=1}^N X_0^{i, (j)}.\end{aligned}\tag{D.22}$$

Here, $\left\{ (\boldsymbol{\xi}^{i, (j)}, \mathbf{w}^{i, (j)}, X_0^{i, (j)})_{1 \leq j \leq N}, \mathbf{w}^{i, 0}; i = 1, \dots, \tilde{N} \right\}$ are *i.i.d.* copies of $(\boldsymbol{\xi}^{(j)}, \mathbf{w}^{(j)}, X_0^{(j)})_{1 \leq j \leq N}$, \mathbf{w}^0 and $\bar{\boldsymbol{\xi}}^{i, N} = \frac{1}{N} \sum_{j=1}^N \boldsymbol{\xi}^j$, $\bar{\mathbf{w}}^{i, N} = \frac{1}{N} \sum_{j=1}^N \mathbf{w}^j$.

Step 2: We adapt the arguments of Lemma 35 and Lemma 44 in [6]. We introduce the sub-exponential norm $\|\cdot\|_{\psi_1}$ and sub-Gaussian norm $\|\cdot\|_{\psi_2}$ of the random vector X :

$$\|X\|_{\psi_1} = \inf \left\{ t > 0, \mathbb{E} \left[\exp \left(\frac{|X|}{t} \right) \right] \leq 2 \right\}; \quad \|X\|_{\psi_2} = \inf \left\{ t > 0, \mathbb{E} \left[\exp \left(\frac{|X|^2}{t^2} \right) \right] \leq 2 \right\}.$$

Recalling (D.3), one has

$$\begin{aligned}
& \|(\tilde{\nabla}_\theta^{\Delta,N,pop} - \hat{\nabla}_\theta^{\Delta,N,pop})J(\Theta)\|_F \\
&= \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop}^{\Delta,N,i}|U_i, V_i])U_i \right\|_F \\
&\leq \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta,N,i}|U_i])U_i \right\|_F + \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,2}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,2}^{\Delta,N,i}|U_i])U_i \right\|_F,
\end{aligned}$$

and

$$\begin{aligned}
& \|(\tilde{\nabla}_\zeta^{\Delta,N,pop} - \hat{\nabla}_\zeta^{\Delta,N,pop})J(\Theta)\|_F \\
&= \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop}^{\Delta,N,i}|U_i, V_i])V_i \right\|_F \\
&\leq \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta,N,i}|V_i])V_i \right\|_F + \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,2}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,2}^{\Delta,N,i}|V_i])V_i \right\|_F.
\end{aligned}$$

We now provide some upper-estimate on the probabilities $\mathbb{P}(\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,\ell}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,\ell}^{\Delta,N,i}|U_i])U_i\|_F \geq \varepsilon)$ and $\mathbb{P}(\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,\ell}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,\ell}^{\Delta,N,i}|V_i])V_i\|_F \geq \varepsilon)$ for $\ell = 1, 2$.

We only deal with $\mathbb{P}(\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta,N,i}|U_i])U_i\|_F \geq \varepsilon)$ inasmuch the proof for the other terms are similar.

For $i = 1, \dots, \tilde{N}$, $j = 1, \dots, N$ and $l = 0, \dots, n-1$, we let

$$\begin{aligned}
\eta_l^{i,j} &= (Y_{t_l}^{\theta_i, \Delta, (j)})^\top (Q + \theta_i^\top R \theta_i) Y_{t_l}^{\theta_i, \Delta, (j)} + 2(Y_{t_l}^{\theta_i, \Delta, (j)})^\top \theta_i^\top R \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}^{i, (j)} \\
&\quad - \mathbb{E}[(Y_{t_l}^{\theta_i, \Delta, (j)})^\top (Q + \theta_i^\top R \theta_i) Y_{t_l}^{\theta_i, \Delta, (j)} + 2(Y_{t_l}^{\theta_i, \Delta, (j)})^\top \theta_i^\top R \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}^{i, (j)} | U_i] \\
&= (Y_{t_l}^{\theta_i, \Delta, (j)})^\top (Q + \theta_i^\top R \theta_i) Y_{t_l}^{\theta_i, \Delta, (j)} + 2(Y_{t_l}^{\theta_i, \Delta, (j)})^\top \theta_i^\top R \sqrt{\frac{\lambda}{2} R^{-1}} \boldsymbol{\xi}_{t_l}^{i, (j)} \\
&\quad - \mathbb{E}[(Y_{t_l}^{\theta_i, \Delta, (j)})^\top (Q + \theta_i^\top R \theta_i) Y_{t_l}^{\theta_i, \Delta, (j)} | U_i].
\end{aligned}$$

where the second equality come from the fact that, conditionally on U_i , the two random variables $Y_{t_l}^{\theta_i, \Delta, (j)}$ and $\boldsymbol{\xi}_{t_l}^{i, (j)}$ are independent and that $\mathbb{E}[Y_{t_l}^{\theta_i, \Delta, (j)} | U_i] = \mathbb{E}[Y_{t_l}^{\theta_i, \Delta, (j)}]_{|\theta=\theta_i} = 0$. From (D.21) and the definition of $\eta_l^{i,j}$, we get

$$\mathcal{J}_{pop,1}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta,N,i}|U_i] = \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} e^{-\beta t_l} \eta_l^{i,j}$$

so that

$$\begin{aligned}
& \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta,N,i} | U_i]) U_i \right\|_F \\
&= \left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \left(\frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} e^{-\beta t_l} \eta_l^{i,j} \right) U_i \right\|_F \\
&= \frac{hd}{r^2} \left\| \left(\frac{1}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} \frac{e^{-\beta t_l}}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \eta_l^{i,j} \right) U_i \right\|_F \\
&\leq \frac{Td}{r^2} \sup_{1 \leq j \leq N, 0 \leq l \leq n-1} \left\| \frac{e^{-\beta t_l}}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \eta_l^{i,j} U_i \right\|_F.
\end{aligned}$$

Hence, noting that the random variables $(\eta^{i,j})_{1 \leq j \leq N}$ have the same law and letting $\delta = \frac{r^2 \varepsilon}{2Td}$, the previous estimate yields

$$\begin{aligned}
& \mathbb{P} \left(\left\| \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta,N,i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta,N,i} | U_i]) U_i \right\|_F \geq \frac{\varepsilon}{2} \right) \\
&\leq \mathbb{P} \left(\frac{Td}{r^2} \sup_{j=1, \dots, N} \sup_{l=0, \dots, n-1} \left\| \frac{e^{-\beta t_l}}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \eta_l^{i,j} U_i \right\|_F \geq \frac{\varepsilon}{2} \right) \\
&= \mathbb{P} \left(\sup_{j=1, \dots, N} \sup_{l=0, \dots, n-1} \left\| \frac{e^{-\beta t_l}}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \eta_l^{i,j} U_i \right\|_F \geq \delta \right) \tag{D.23} \\
&\leq \sum_{j=1}^N \sum_{l=0}^{n-1} \mathbb{P} \left(\left\| \frac{e^{-\beta t_l}}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \eta_l^{i,j} U_i \right\|_F \geq \delta \right) \\
&= N \sum_{l=0}^{n-1} \mathbb{P} \left(\left\| \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (e^{-\beta t_l} \eta_l^{i,1}) U_i \right\|_F \geq \delta \right).
\end{aligned}$$

Now, it easily follows from standard computations based on the dynamics (D.22) of $Y^{\theta_i, \Delta, (1)}$ and since the θ_i are bounded that the random variables $Y_{t_l}^{\theta_i, \Delta, (1)}$ have finite sub-Gaussian norm $\|Y_{t_l}^{\theta_i, \Delta, (1)}\|_{\Psi_2}$ and since $\|\xi_{t_l}^{i,(j)}\|_{\Psi_2} < \infty$ we deduce that $\|\eta_l^{i,1}\|_{\Psi_1} < \infty$ uniformly in (l, i) . By Lemma 37 in [6], we get

$$\mathbb{P} \left(\left\| \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \eta_l^{i,1} U_i \right\|_F \geq \delta \right) \leq 2d \exp \left(- \frac{\tilde{N} \delta^2}{2(\chi_l^2 + \chi_l \delta)} \right), \tag{D.24}$$

where $\chi_l := C_1 r \sup_{\theta' \in \mathbb{B}(\theta, r)} \|\eta_l^{\theta', 1}\|_{\psi_1}$ and C_1 is a universal constant¹. Hence, combining (D.23) and (D.24), we obtain

$$\mathbb{P}\left(\left\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta, N, i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta, N, i} | U_i]) U_i\right\|_F \geq \frac{\varepsilon}{2}\right) \leq 2dNn \exp\left(-\frac{\tilde{N}\delta^2}{2(\tilde{\chi}^2 + \tilde{\chi}\delta)}\right) \quad (\text{D.25})$$

with $\tilde{\chi} = \sup_{0 \leq l \leq n-1} (e^{-\beta t_l} \chi_l)$.

Then, from Lemma D.7 (see below), taking $r \leq \min\{\tilde{r}(b), \inf_{\Theta \in \mathcal{R}(b)} \frac{1}{h_1(\Theta)}\}$ (recalling that $\mathcal{R}(b) \subseteq \mathcal{S}(b) \times \hat{\mathcal{S}}(b)$ is compact and that according to Proposition C.1, $\|\theta\|_F \leq \mathbf{Bd}_\theta(b)$) and letting

$$\hat{\chi}(b) = C\tilde{r}(b)(1 + \mathbf{Bd}_\theta(b)^2 + \tilde{r}(b)^2),$$

one has $\tilde{\chi} \leq \hat{\chi}(b)$ uniformly for all $\Theta \in \mathcal{R}(b)$.

The previous bound on $\tilde{\chi}$ together with (D.25) implies that if

$$\tilde{N} \geq h_{\tilde{N}}^{(1,1)}(r, b) := 2 \frac{(\hat{\chi}(b)^2 + \hat{\chi}(b)\delta)}{\delta^2} \left((d+1) \log\left(\frac{d}{\varepsilon}\right) + \log N + \log n + \log(4\varepsilon) \right),$$

then, for all $\Theta \in \mathcal{R}(b)$, one has

$$\mathbb{P}\left(\left\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta, N, i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta, N, i} | U_i]) U_i\right\|_F \geq \frac{\varepsilon}{2}\right) \leq \frac{1}{2} \left(\frac{d}{\varepsilon}\right)^{-d}.$$

Similarly, there exists $h_{\tilde{N}}^{(1,2)}(r, b)$ such that for any $\tilde{N} \geq h_{\tilde{N}}^{(1,2)}(r, b)$ and any $\Theta \in \mathcal{R}(b)$,

$$\mathbb{P}\left(\left\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,2}^{\Delta, N, i} - \mathbb{E}[\mathcal{J}_{pop,2}^{\Delta, N, i} | U_i]) U_i\right\|_F \geq \frac{\varepsilon}{2}\right) \leq \frac{1}{2} \left(\frac{d}{\varepsilon}\right)^{-d}.$$

Hence, choosing $\tilde{N} \geq \max\{h_{\tilde{N}}^{(1,1)}(r, b), h_{\tilde{N}}^{(1,2)}(r, b)\}$, we get

$$\begin{aligned} \mathbb{P}\left(\left\|\tilde{\nabla}_\theta^{\Delta, N, pop} J(\Theta) - \hat{\nabla}_\theta^{\Delta, N, pop} J(\Theta)\right\|_F \geq \varepsilon\right) &\leq \mathbb{P}\left(\left\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,1}^{\Delta, N, i} - \mathbb{E}[\mathcal{J}_{pop,1}^{\Delta, N, i} | U_i]) U_i\right\|_F \geq \frac{\varepsilon}{2}\right) \\ &\quad + \mathbb{P}\left(\left\|\frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} (\mathcal{J}_{pop,2}^{\Delta, N, i} - \mathbb{E}[\mathcal{J}_{pop,2}^{\Delta, N, i} | U_i]) U_i\right\|_F \geq \frac{\varepsilon}{2}\right) \\ &\leq \left(\frac{d}{\varepsilon}\right)^{-d}. \end{aligned}$$

The proof of the upper-bound on $\mathbb{P}(\|\tilde{\nabla}_\zeta^{\Delta, N, pop} J(\Theta) - \hat{\nabla}_\zeta^{\Delta, N, pop} J(\Theta)\|_F \geq \varepsilon)$ being similar is omitted. The proof of Proposition D.5 is now complete. \square

¹Referring to [6], one would use the standard operator norm $\|\cdot\|$ in (D.24). However, due to the equivalence of the matrix norms, we can replace $\|\cdot\|$ by the Frobenius norm $\|\cdot\|_F$, up to a modification of the constant $C_1 > 0$ which may depend on the dimensions m and d .

Lemma D.7. *There exists $h_1(\|\theta\|)$ with at most of polynomial growth in $\|\theta\|$ and a constant $C < \infty$ such that for all $r \leq \frac{1}{h_1(\Theta)}$, it holds*

$$\check{\chi} \leq Cr(1 + \|\theta\|_F^2 + r^2).$$

Proof. Let

$$\Psi_l^{\theta'} = (Y_{t_l}^{\theta', \Delta, (1)})^\top (Q + (\theta')^\top R \theta') Y_{t_l}^{\theta', \Delta, (1)} + 2(Y_{t_l}^{\theta', \Delta, (1)})^\top (\theta')^\top R \sqrt{\frac{\lambda}{2} R^{-1} \xi_{t_l}^{(1)}}, \quad \theta' \in \mathbb{B}(\theta, r),$$

recalling that $Y^{\theta', \Delta, (1)}$ is given by (D.18) with $\theta = \theta'$.

Since

$$\|\Psi_l^{\theta'} - \mathbb{E}[\Psi_l^{\theta'}]\|_{\psi_1} \leq 2\|\Psi_l^{\theta'}\|_{\psi_1},$$

and $\eta_l^{\theta', 1} = \Psi_l^{\theta'} - \mathbb{E}[\Psi_l^{\theta'}]$, we get

$$\check{\chi} \leq 2C_1 r \sup_{0 \leq l \leq n-1, \theta' \in \mathbb{B}(\theta, r)} (e^{-\beta t_l} \|\Psi_l^{\theta'}\|_{\psi_1}).$$

The triangle inequality for the sub-exponential norm $\|\cdot\|_{\Psi_1}$ yields

$$\|\Psi_l^{\theta'}\|_{\psi_1} \leq \|(Y_{t_l}^{\theta', \Delta, (1)})^\top (Q + (\theta')^\top R \theta') Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_1} + 2\|(Y_{t_l}^{\theta', \Delta, (1)})^\top (\theta')^\top R \sqrt{\frac{\lambda}{2} R^{-1} \xi_{t_l}^{(1)}}\|_{\psi_1}. \quad (\text{D.26})$$

From Proposition 2.4 [24],

$$\begin{aligned} \|(Y_{t_l}^{\theta', \Delta, (1)})^\top (Q + (\theta')^\top R \theta') Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_1} &\leq \text{Tr}(Q + (\theta')^\top R \theta') \|Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_2}^2 \\ &\leq Cd(\|Q\|_F + \|R\|_F(\|\theta\|_F + r)^2) \|Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_2}^2, \end{aligned}$$

and, from the inequality $\|XY\|_{\Psi_1} \leq \|X\|_{\Psi_2} \|Y\|_{\Psi_2}$ stemming from standard Young's inequality,

$$\begin{aligned} \|(Y_{t_l}^{\theta', \Delta, (1)})^\top (\theta')^\top R \sqrt{\frac{\lambda}{2} R^{-1} \xi_{t_l}^{(1)}}\|_{\psi_1} &= \|(\theta' Y_{t_l}^{\theta', \Delta, (1)}) \cdot (R \sqrt{\frac{\lambda}{2} R^{-1} \xi_{t_l}^{(1)}})\|_{\psi_1} \\ &\leq \|\theta' Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_2} \|R \sqrt{\frac{\lambda}{2} R^{-1} \xi_{t_l}^{(1)}}\|_{\psi_2} \\ &\leq C \|\theta'\|_F \|Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_2} \|R \sqrt{\frac{\lambda}{2} R^{-1} \xi_{t_l}^{(1)}}\|_{\psi_2} \\ &\leq C(\|\theta\|_F + r) \|Y_{t_l}^{\theta', \Delta, (1)}\|_{\psi_2} \end{aligned}$$

for some constant $C = C(\lambda, \|R\|_F)$.

Similarly to Lemma 36 [6], we can establish an universal upper-bound for the sub-Gaussian norm of $Y_{t_l}^{\theta', \Delta, (j)}$ for all $\theta' \in \mathbb{B}(\theta, r)$, where $r \leq 1/h_1(\|\theta\|)$, $h_1(\|\theta\|)$ being at most of polynomial growth in $\|\theta\|$. Namely, it holds

$$\sup_{0 \leq l \leq n-1} e^{-\beta t_l} \|Y_{t_l}^{\theta', \Delta, (j)}\|_{\psi_2} \leq C \sup\{\|Y_0^{(j)}\|_{\psi_2}, \sup_{0 \leq l \leq n-1} \|\mathbf{w}_l^{(j)} - \bar{\mathbf{w}}_l^N\|_{\psi_2}, \sup_{0 \leq l \leq n-1} \|\xi_l^{(j)} - \bar{\xi}_l^N\|_{\psi_2}\}.$$

Plugging the three previous inequalities into (D.26), we get

$$\|\Psi_l^{\theta'}\|_{\psi_1} \leq C(1 + \|\theta\|_F^2 + r^2)$$

which eventually yields

$$\tilde{\chi} \leq 2C_1 r \sup_{l=1,\dots,n} \sup_{\theta' \in \mathbb{B}(\theta,r)} (e^{-\beta t_l} \|\Psi_l^{\theta'}\|_{\psi_1}) \leq Cr(1 + \|\theta\|_F^2 + r^2),$$

for all $r \leq 1/h_1(\|\theta\|)$. □

D.4.5 Proof of Proposition D.6.

Recalling that $\tilde{\mathcal{J}}_{pop}^{\Delta,N}(\Theta) = \mathbb{E}[\mathcal{J}_{pop}^{\Delta,N}(\Theta)]$ and using the fact that $(U_i, V_i)_{1 \leq i \leq N}$ is independent of $((X_0^{(j)}, (\xi_{t_l}^{(j)})_{0 \leq l \leq n-1}, W^{(j)})_{j=1,\dots,N}, W^0)$, we get $\mathbb{E}[\mathcal{J}_{pop}^{\Delta,N,i} | U_i, V_i] = \tilde{\mathcal{J}}_{pop}^{\Delta,N}(\Theta_i)$ and

$$\begin{aligned} \hat{\nabla}_{\theta}^{\Delta,N,pop} J(\Theta) &= \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \tilde{\mathcal{J}}_{pop}^{\Delta,N}(\theta + U_i, \zeta + V_i) U_i, \\ \hat{\nabla}_{\zeta}^{\Delta,N,pop} J(\Theta) &= \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} \tilde{\mathcal{J}}_{pop}^{\Delta,N}(\theta + U_i, \zeta + V_i) V_i. \end{aligned}$$

Step 1: We introduce the process $(\mathbf{x}^{\Theta,\Delta,(j)})_{j=1,\dots,N}$ with dynamics

$$\begin{aligned} \mathbf{x}_{t_{l+1}}^{\Theta,\Delta,(j)} &= \mathbf{x}_{t_l}^{\Theta,\Delta,(j)} + (B\mathbf{x}_{t_l}^{\Theta,\Delta,(j)} + \bar{B}\mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta,(j)}] + D\boldsymbol{\alpha}_{t_l}^{\Theta,\Delta,(j)})h + \sqrt{h}\gamma\mathbf{w}_l^{(j)} + \sqrt{h}\gamma_0\mathbf{w}_l^0, \\ \mathbf{x}_0^{\Theta,\Delta,(j)} &= X_0^{(j)}, \end{aligned}$$

with

$$\boldsymbol{\alpha}_{t_l}^{\Theta,\Delta,(j)} = \theta(\mathbf{x}_{t_l}^{\Theta,\Delta,(j)} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta,(j)}]) + \zeta\mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta,(j)}] + \sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)},$$

and where the random variables $(X_0^{(j)})_{1 \leq j \leq N}$, $(\boldsymbol{\xi}^{(j)} := (\xi_{t_l}^{(j)})_{0 \leq l \leq n-1})_{1 \leq j \leq N}$, $(\mathbf{w}^{(j)})_{1 \leq j \leq N}$ and \mathbf{w}^0 are the same as those employed in (4.5)-(4.6).

We then define the two processes $\mathbf{y}^{\theta,\Delta,(j)}$ and $\mathbf{z}^{\zeta,\Delta}$ by

$$\mathbf{y}_{t_l}^{\theta,\Delta,(j)} = \mathbf{x}_{t_l}^{\Theta,\Delta,(j)} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta,(j)}], \quad \mathbf{z}_{t_l}^{\zeta,\Delta} = \mathbb{E}_0[\mathbf{x}_{t_l}^{\Theta,\Delta}],$$

with $\mathbf{y}_0^{\theta,\Delta,(j)} = X_0^{(j)} - \mathbb{E}[X_0]$ and $\mathbf{z}_0^{\zeta,\Delta} = \mathbb{E}[X_0]$. Their dynamics are given by

$$\begin{aligned} \mathbf{y}_{t_{l+1}}^{\theta,\Delta,(j)} &= \mathbf{y}_{t_l}^{\theta,\Delta,(j)} + ((B + D\theta)\mathbf{y}_{t_l}^{\theta,\Delta,(j)} + D\sqrt{\frac{\lambda}{2}}R^{-1}\boldsymbol{\xi}_{t_l}^{(j)})h + \gamma\sqrt{h}\mathbf{w}_l^{(j)}, \\ \mathbf{z}_{t_{l+1}}^{\zeta,\Delta} &= \mathbf{z}_{t_l}^{\zeta,\Delta} + (\hat{B} + D\zeta)\mathbf{z}_{t_l}^{\zeta,\Delta}h + \gamma_0\sqrt{h}\mathbf{w}_l^0. \end{aligned} \tag{D.27}$$

We again notice that $\mathbf{y}^{\theta,\Delta,(j)}$ (resp. $\mathbf{z}^{\theta,\Delta}$) depends only on θ (resp. ζ). Hence, in order to simplify the notation, from now on, we will write $\mathbf{y}^{\theta,\Delta,(j)}$ and $\mathbf{z}^{\zeta,\Delta}$. Note that $(\mathbf{y}_{t_l}^{\theta,\Delta,(j)})_{j=1,\dots,N}$ are *i.i.d.* copies of $\mathbf{y}^{\theta,\Delta}$ and $\mathbf{z}^{\zeta,\Delta} = \mathbf{z}^{\theta,\Delta}$ where $(\mathbf{y}^{\theta,\Delta}, \mathbf{z}^{\theta,\Delta})$ are defined by (D.13). Hence,

$$\begin{aligned} J^{(T),\Delta,N}(\Theta) &:= \frac{h}{N} \sum_{j=1}^N \mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} ((\mathbf{x}_{t_l}^{\theta,\Delta,(j)} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\theta,\Delta}])^\top Q (\mathbf{x}_{t_l}^{\theta,\Delta,(j)} - \mathbb{E}_0[\mathbf{x}_{t_l}^{\theta,\Delta}]) \right. \\ &\quad \left. + \mathbb{E}_0[\mathbf{x}_{t_l}^{\theta,\Delta,(j)}]^\top \hat{Q} \mathbb{E}_0[\mathbf{x}_{t_l}^{\theta,\Delta}] + (\boldsymbol{\alpha}_{t_l}^{\theta,\Delta,(j)})^\top R \boldsymbol{\alpha}_{t_l}^{\theta,\Delta,(j)} + \lambda \log \pi^\Theta(\boldsymbol{\alpha}_{t_l}^{\theta,\Delta,(j)} | \mathbf{x}_{t_l}^{\theta,\Delta,(j)}, \mathbb{E}_0[\mathbf{x}_{t_l}^{\theta,\Delta}]) \right] \\ &= \frac{h}{N} \sum_{j=1}^N \mathbb{E} \left[\sum_{l=0}^{n-1} e^{-\beta t_l} ((\mathbf{y}_{t_l}^{\theta,\Delta,(j)})^\top (Q + \theta^\top R \theta) \mathbf{y}_{t_l}^{\theta,\Delta,(j)} + (\mathbf{z}_{t_l}^{\zeta,\Delta})^\top (\hat{Q} + \zeta^\top R \zeta) \mathbf{z}_{t_l}^{\zeta,\Delta} + \beta v(\lambda)) \right] \\ &= J^{(T),\Delta}(\Theta), \end{aligned}$$

recalling that $J^{(T),\Delta}(\Theta)$ is defined by (D.4). We will also use the decomposition $J^{(T),\Delta,N} = J_1^{(T),\Delta} + J_2^{(T),\Delta} + \beta v(\lambda) h \sum_{l=1}^n e^{-\beta t_l}$ with

$$\begin{aligned} J_1^{(T),\Delta}(\theta) &:= \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} e^{-\beta t_l} \mathbb{E} \left[(\mathbf{y}_{t_l}^{\theta,\Delta,(j)})^\top (Q + \theta^\top R \theta) \mathbf{y}_{t_l}^{\theta,\Delta,(j)} \right], \\ J_2^{(T),\Delta}(\zeta) &:= h \sum_{l=0}^{n-1} e^{-\beta t_l} \mathbb{E} \left[(\mathbf{z}_{t_l}^{\zeta,\Delta})^\top (\hat{Q} + \zeta^\top R \zeta) \mathbf{z}_{t_l}^{\zeta,\Delta} \right]. \end{aligned}$$

Then, according to (D.19) and (D.20),

$$\bar{J}_{pop}^{\Delta,N}(\Theta) = \mathbb{E}[\mathcal{J}_{pop}^{\Delta,N}(\Theta)] = \bar{J}_1(\theta) + \bar{J}_2(\zeta) + \beta v(\lambda) h \sum_{l=1}^n e^{-\beta t_l}$$

where

$$\begin{aligned} \bar{J}_1(\theta) &= \mathbb{E}[\mathcal{J}_1(\theta)] = \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} e^{-\beta t_l} \mathbb{E}[(Y_{t_l}^{\theta,\Delta,(j)})^\top (Q + \theta^\top R \theta) Y_{t_l}^{\theta,\Delta,(j)}] \\ \bar{J}_2(\zeta) &= \mathbb{E}[\mathcal{J}_2(\zeta)] = h \sum_{l=0}^{n-1} e^{-\beta t_l} \mathbb{E}[(\hat{\mu}_{t_l}^{\zeta,\Delta,N})^\top (\hat{Q} + \zeta^\top R \zeta) \hat{\mu}_{t_l}^{\zeta,\Delta,N}]. \end{aligned}$$

Hence, in order to establish an upper-bound for $(J^{(T),\Delta,N} - \bar{J}_{pop}^{\Delta,N})(\Theta) = (J_1^{(T),\Delta} - \bar{J}_1)(\theta) + (J_2^{(T),\Delta} - \bar{J}_2)(\zeta)$, we need to quantify the $L^p(\mathbb{P})$ error for the difference $\mathbf{y}_{t_l}^{\theta,\Delta,(j)} - Y_{t_l}^{\theta,\Delta,(j)}$ and $\mathbf{z}_{t_l}^{\zeta,\Delta} - \hat{\mu}_{t_l}^{\zeta,\Delta,N}$, $j = 1, \dots, N$ and $l = 0, \dots, n-1$.

Step 2: Let

$$\begin{aligned} \mathcal{Q}_{t_l}^{\theta,\Delta,(j)} &= \mathbf{y}_{t_l}^{\theta,\Delta,(j)} - Y_{t_l}^{\theta,\Delta,(j)}, \quad j = 1, \dots, N, \\ \hat{\mathcal{Q}}_{t_l}^{\zeta,\Delta,N} &= \mathbf{z}_{t_l}^{\zeta,\Delta} - \hat{\mu}_{t_l}^{\zeta,\Delta,N}. \end{aligned}$$

From (D.17)-(D.18) and (D.27), we easily get

$$\begin{aligned}\mathcal{Q}_{t_{l+1}}^{\theta,\Delta,(j)} &= \mathcal{Q}_{t_l}^{\theta,\Delta,(j)} + h((B + D\theta)\mathcal{Q}_{t_l}^{\theta,\Delta,(j)} + D\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} + \sqrt{h}\gamma\bar{\mathbf{w}}_l^N), \\ \hat{\mathcal{Q}}_{t_{l+1}}^{\zeta,\Delta,N} &= \hat{\mathcal{Q}}_{t_l}^{\zeta,\Delta,N} + h((\hat{B} + D\zeta)\hat{\mathcal{Q}}_{t_l}^{\zeta,\Delta,N} - D\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} - \sqrt{h}\gamma\bar{\mathbf{w}}_l^N),\end{aligned}\tag{D.28}$$

with $\mathcal{Q}_0^{\theta,\Delta,(j)} = \mathcal{Q}_0 := \hat{\mu}_0^N - \mathbb{E}[X_0]$, $\hat{\mathcal{Q}}_0^{\zeta,\Delta} = \hat{\mathcal{Q}}_0 := \mathbb{E}[X_0] - \hat{\mu}_0^N$. Note carefully that for any $l = 0, \dots, n$, the random variables $(\mathcal{Q}_{t_l}^{\theta,\Delta,(j)})_{1 \leq j \leq N}$ are equal. We thus omit the superscript j and simply write $\mathcal{Q}_{t_l}^{\theta,\Delta}$ in what follows.

From (D.28), we get

$$\begin{aligned}|\mathcal{Q}_{t_{l+1}}^{\theta,\Delta,(j)}|^2 &= |\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}|^2 + 2\langle \mathcal{Q}_{t_l}^{\theta,\Delta,(j)}, h(B + D\theta)\mathcal{Q}_{t_l}^{\theta,\Delta,(j)} + hD\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} + \sqrt{h}\gamma\bar{\mathbf{w}}_l^N \rangle \\ &\quad + \left| h(B + D\theta)\mathcal{Q}_{t_l}^{\theta,\Delta,(j)} + hD\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} + \sqrt{h}\gamma\bar{\mathbf{w}}_l^N \right|^2 \\ &\leq (1 + hC_1(\theta))|\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}|^2 + 2\langle \mathcal{Q}_{t_l}^{\theta,\Delta,(j)}, hD\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} + \sqrt{h}\gamma\bar{\mathbf{w}}_l^N \rangle \\ &\quad + 2\left| hD\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} + \sqrt{h}\gamma\bar{\mathbf{w}}_l^N \right|^2,\end{aligned}$$

for some $\theta \mapsto C_1(\theta)$ with at most of quadratic growth in $\|\theta\|$. Taking expectation in both sides of the previous inequality and recalling that $\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}$ is independent of $\bar{\xi}^N$ and $\bar{\mathbf{w}}^N$, we get

$$\begin{aligned}\mathbb{E}[|\mathcal{Q}_{t_{l+1}}^{\theta,\Delta,(j)}|^2] &\leq (1 + hC_1(\theta))\mathbb{E}[|\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}|^2] + 2\mathbb{E}\left| hD\sqrt{\frac{\lambda}{2}R^{-1}\bar{\xi}_{t_l}^N} + \sqrt{h}\gamma\bar{\mathbf{w}}_l^N \right|^2 \\ &\leq (1 + hC_1(\theta))\mathbb{E}[|\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}|^2] + C_2\frac{h}{N},\end{aligned}$$

for some constant $C_2 < \infty$ and where, for the last inequality, we used the fact that $\mathbb{E}|\bar{\xi}_{t_l}^N|^2 = m/N$ and $\mathbb{E}|\bar{\mathbf{w}}_l^N|^2 = d/N$, $l = 0, \dots, n-1$. From the discrete Grönwall lemma, up to a modification of C_2 , we get

$$\mathbb{E}[|\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}|^2] \leq (1 + hC_1(\theta))^l \left(\mathbb{E}[|\mathcal{Q}_0^{\theta,\Delta,(j)}|^2] + \frac{C_2}{C_1(\theta)N} \right).$$

Now, since $\mathbb{E}[|\mathcal{Q}_0^{\theta,\Delta,(j)}|^2] = \text{Var}(X_0)/N$, using the standard inequality $1 + x \leq \exp(x)$, $x \in \mathbb{R}$, the previous inequality eventually implies

$$\sup_{0 \leq l \leq n} \mathbb{E}[|\mathcal{Q}_{t_l}^{\theta,\Delta,(j)}|^2] \leq \frac{\exp(C_1(\theta))}{N},$$

up to a modification of $C_1(\theta)$ (with at most of quadratic growth). Similar arguments yield

$$\sup_{0 \leq l \leq n} \mathbb{E}[|\hat{\mathcal{Q}}_{t_l}^{\zeta,\Delta,N}|^2] \leq \frac{\exp(C_2(\zeta))}{N}.$$

The two previous upper-bounds in turn imply

$$\begin{aligned}
|(J_1^{(T),\Delta} - \bar{J}_1)(\theta)| &\leq C(\|Q\|_F + \|R\|_F\|\theta\|_F^2) \frac{h}{N} \sum_{j=1}^N \sum_{l=0}^{n-1} e^{-\beta t_l} \mathbb{E}[|Y_{t_l}^{\theta,\Delta,(j)} - \mathbf{y}_{t_l}^{\theta,\Delta,(j)}|^2] \\
&\leq CT(\|Q\|_F + \|R\|_F\|\theta\|_F^2) \sup_{l=0,\dots,n-1} \mathbb{E}[|Q_{t_l}^{\theta,\Delta}|^2] \\
&\leq \frac{\exp(C_1(\theta))}{N} (\|Q\|_F + \|R\|_F\|\theta\|_F^2),
\end{aligned}$$

and similarly

$$|(J_2^{(T),\Delta} - \bar{J}_2)(\zeta)| \leq \frac{\exp(C_2(\zeta))}{N} (\|Q\|_F + \|R\|_F\|\zeta\|_F^2),$$

up to a modification of $C_1(\theta)$ and $C_2(\zeta)$. Recalling that $\mathcal{R}(b)$ is compact, the right-hand side of the two previous inequalities can be uniformly bounded by a constant depending only upon b . The proof of (D.8) is now complete.

Step 3: After recalling that $\Theta_i \in \mathcal{R}(2b)$, it directly follows from (D.8) that

$$|(\bar{J}_{pop}^{\Delta,N} - J^{(T),\Delta})(\Theta_i)| \leq \frac{c_4(2b)}{N}$$

and, since $\|U_i\|_F = \|V_i\| = r$, \mathbb{P} - a.s

$$\begin{aligned}
\|(\hat{V}_\theta^{\Delta,N,pop} - \hat{V}_\theta^{(T),\Delta})J(\Theta)\|_F &\leq \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} |(\bar{J}_{pop}^{\Delta,N} - J^{(T),\Delta})(\Theta_i)| \|U_i\|_F \leq \frac{d}{r} \frac{c_4(2b)}{N}, \\
\|(\hat{V}_\zeta^{\Delta,N,pop} - \hat{V}_\zeta^{(T),\Delta})J(\Theta)\|_F &\leq \frac{d}{r^2} \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} |(\bar{J}_{pop}^{\Delta,N} - J^{(T),\Delta})(\Theta_i)| \|V_i\|_F \leq \frac{d}{r} \frac{c_4(2b)}{N},
\end{aligned}$$

which ends the proof. \square

References

- [1] A. Angiuli, J.-P. Fouque, and M. Laurière. Unified reinforcement Q-learning for mean field game and control problems. *Mathematics of Control, Signals and Systems*, 34:217–271, 2022.
- [2] M. Basei and H. Pham. A weak martingale approach to linear-quadratic mckean-vlasov stochastic control problems. *Journal of Optimization Theory and Applications*, 181(2):347–382, 2019.

- [3] J. Bu, A. Mesbahi, and M. Mesbahi. Policy gradient-based algorithms for continuous-time linear quadratic control. *arXiv:2006.09178*, 2020.
- [4] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games: vol. I, Mean Field FBSDEs, Control, and Games*. Springer, 2018.
- [5] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games: vol. II, Mean Field game with common noise and Master equations*. Springer, 2018.
- [6] R. Carmona, M. Laurière, and Z. Tan. Linear-quadratic mean-field reinforcement learning: convergence of policy gradient methods. *arXiv:1910.04295*, 2019.
- [7] R. Carmona, M. Laurière, and Z. Tan. Model-free mean-field reinforcement learning: Mean-field MDP and mean-field Q-learning. *The Annals of Applied Probability*, 33(6B):5334 – 5381, 2023.
- [8] M. Fazel, R. Ge, S.M. Kakade, and M. Mesbahi. Global convergence of policy gradient methods for the linear quadratic regulator. *Proceedings of the 35th International Conference on Machine Learning*, pages 1467–1476, 2018.
- [9] N. Frikha, M. Germain, M. Laurière, H. Pham, and X. Song. Actor-critic learning for mean-field control in continuous time. *arXiv: 2303.06993*, 2023.
- [10] M. Giegrich, C. Reisinger, and Y. Zhang. Convergence of policy gradient methods for finite-horizon exploratory linear-quadratic control problems. *SIAM Journal on Control and Optimization*, 62(2):1060–1092, 2024.
- [11] H. Gu, X. Guo, X. Wei, and R. Xu. Mean field controls with Q-learning for cooperative MARL: convergence and complexity analysis. *SIAM Journal on Mathematics of Data Science*, 3(4), 2021.
- [12] X. Guo, X. Li, and R. Xu. Fast policy learning for linear-quadratic control with entropy regularization. *arXiv:2311.14168v3*, 2023.
- [13] B. Hambly, R. Xu, and H. Yang. Policy gradient methods for the noisy linear quadratic regulator over a finite horizon. *SIAM Journal on Control and Optimization*, 59(5):3359–3391, 2021.
- [14] Y. Jia and X.Y. Zhou. Policy gradient and actor–critic learning in continuous time and space: Theory and algorithms. *Journal of Machine Learning Research*, 2021.
- [15] Y. Jia and X.Y. Zhou. q learning in continuous time. *Journal of Machine Learning Research*, 2023.

- [16] H. Mohammadi, A. Zare, M. Soltanolkotabi, and M. R. Jovanović. Convergence and sample complexity of gradient methods for the model-free linear–quadratic regulator problem. *IEEE Transactions on Automatic Control*, 67(5):2435–2450, 2022.
- [17] Rémi Munos. Policy gradient in continuous time. *Journal of Machine Learning Research*, 7(27):771–791, 2006.
- [18] Eckhard Platen Peter E. Kloeden. *Numerical Solution of Stochastic Differential Equations*. Springer Berlin, Heidelberg, 1992.
- [19] H. Pham and X. Warin. Actor-critic learning algorithms for mean-field control with moment neural networks. arXiv: 2309.04317, 2023.
- [20] R. Sutton and A. Barto. *Reinforcement Learning: An Introduction*. Cambridge, MA:MIT, 2018.
- [21] L. Szpruch, T. Treetanthiploet, and Y. Zhang. Optimal scheduling of entropy regularizer for continuous-time linear-quadratic reinforcement learning. *SIAM Journal on Control and Optimization*, 62(1):135–166, 2024.
- [22] H. Wang, T. Zariphopoulou, and X.Y. Zhou. Reinforcement learning in continuous time and space: A stochastic control approach. *Journal of Machine Learning Research*, 21(198):1–34,, 2020.
- [23] W. Wang, J. Han, Z. Yang, and Z. Wang. Global convergence of policy gradient for linear-quadratic mean-field control/game in continuous time. *Proceedings of Machine Learning Research*, 2021.
- [24] K. Zajkowski. Bounds on tail probabilities for quadratic forms in dependent sub-gaussian random variables. *Statistics & Probability Letters*, 167:108898, 2020.