

Extended mean field control: a global numerical solution via finite-dimensional approximation

Athena Picarelli*

Marco Scaratti *

Jonathan Tam †

March 18, 2026

Abstract

We investigate the global numerical approximation of a class of extended mean field control problems (MFC), where the dynamics and costs depend on the joint distribution of the state and the control. We propose a framework to approximate the value function globally over the Wasserstein space, moving beyond the restriction of fixed initial conditions. Our approach exploits the propagation of chaos by approximating the infinite-dimensional MFC problem by an N -player cooperative game, together with the usage of finite-dimensional solvers. This method avoids the need to parametrise functions on an infinite-dimensional space, offering a balance between probabilistic rigor and computational efficiency.

1 Introduction

In this article, we consider a finite-dimensional approximation for a class of extended mean field control problems. Our method provides a global approximation of the value function on the infinite-dimensional Wasserstein space, as opposed to a fixed initial condition. We appeal to the propagation of chaos [15, 25] and consider an approximating finite-player cooperative game, which allows us to fit a finite-dimensional function by proxy.

The study of mean field control (MFC) problems, arising from the control of SDEs of McKean-Vlasov type (MKV SDE), is closely related to the theory of mean field games (MFG), both of which act as an asymptotic approximation of stochastic games between a large number of players. In particular, MFC problems correspond to symmetric cooperative games, where the role of the limiting SDE is that of a central planner seeking a social optimum. MFC and MFG problems appear in many economics and financial applications, including planning in energy markets [2], as well as optimal liquidation and price impact problems with aggregate flow [3, 16, 22]. An extensive treatment of MFC and MFG problems can be found in the two-volume monograph [5]. From a control perspective, MFC problems can be viewed in the lens of a single-agent optimisation problem in which the dynamics and cost depend on the law of the controlled state process. Such formulations can arise in contexts such as optimal transport and target constraint problems, where related Schrödinger bridge problems have recently attracted attention in generative modelling in finance [14, 32, 33].

Due to the nature of the MKV SDE, the state variable in the value function for an MFC problem is a probability measure in the Wasserstein space. Whilst the value function can be considered as the unique viscosity solution to a class of Master Bellman PDEs under certain conditions [9, 10], it is typical in practice to numerically approximate the value function by particle methods. The finite-dimensional reduction of problems on infinite-dimensional spaces of probability measures have been considered in theoretical works [17, 31], with particle approximation methods being considered in [7, 13, 18, 30]. It is typical to parametrise the control with a neural network, and minimising a loss function based on either a dynamic programming or BSDE characterisation of the control problem. This method has demonstrated empirical effectiveness at overcoming the curse of dimensionality associated with more traditional PDE-based methods [21].

*Department of Economics, University of Verona,
{athena.picarelli, marco.scaratti}@univr.it.

†Mathematical Institute, University of Oxford, tam@maths.ox.ac.uk.

The main contribution of this paper is to provide a numerical approximation of the MFC value function on the Wasserstein space globally. Given a value function v for the MFC problem, the propagation of chaos yields an approximation of the form

$$v(t, \mu) \approx \int_{(\mathbb{R}^d)^N} \bar{v}_N(t, x_1, \dots, x_N) d\mu^{\otimes N}, \quad (1.1)$$

where \bar{v}_N is the value function of an associated N -player cooperative game. We therefore approximate \bar{v}_N globally, which is a finite-dimensional function. Afterwards, $v(t, \mu)$ can be computed by sampling \bar{v}_N according to the measure μ . This contrasts with the aforementioned works, where the MFC dynamics are simulated directly via particle methods, and solutions are only given for fixed initial conditions. To the best of our knowledge, [28, 29] provide a global approximation of the value function by another neural network parametrisation, followed by a fitting with regression. However, learning a function on the Wasserstein space comes at a large computational cost, and the estimation of the joint law $\mathcal{L}(X_t, \alpha_t)$, where X and α denote the state and control processes respectively, becomes an extra source of approximation error. By first passing to the finite-dimensional setting, and learning instead the associated (finite-dimensional) value function, we benefit from the following:

1. *More efficient training of the neural network.* The law of the state process $\mathcal{L}(X_t)$ appear in MFC problems in either the coefficients of the state dynamics, the cost functions, or as an input in feedback controls of the form $a(t, X_t, \mathcal{L}(X_t))$. A large number of trajectories (usually in the order of 10^6) would be required for an accurate Monte Carlo estimate of the term $\mathcal{L}(X_t)$. In the finite-dimensional problem, the mean field terms are replaced with the empirical distributions $\frac{1}{N} \sum_{n=1}^N \delta_{X_t^n}$. In this case, N Brownian trajectories are sufficient for each estimate of the value function. This gives a higher flexibility on the batch size for training: a smaller batch size is useful for quicker gradient descent iterates, whilst a larger batch size can reduce the variance in the gradient updates for a more stable training process.
2. *Avoiding an a priori discretisation of the Wasserstein space.* Training the finite-dimensional value function \bar{v}_N does not require a discretisation of the underlying state space. Moreover, the evaluation of the integral in (1.1) only requires evaluating sampled inputs through the trained neural network, so the Monte Carlo error can be reduced to the desired accuracy with modest memory capacity constraints. Our method avoids accumulating errors associated with the quantisation or truncation of moments over multiple time steps that would occur when simulating the mean-field system.
3. *Flexible choice of finite-dimensional solvers.* Our method is independent of the choice of the finite-dimensional solver. Depending on the problem assumptions, dynamic programming or BSDE characterisations of the value function can be used as a basis for the algorithm [18, 21, 26]. In contrast, BSDE characterisations are much harder to obtain in the mean field regime when the coefficients depend on the law of the controls $\mathcal{L}(\alpha_t)$, as in [1], which restricts the possible choice of algorithms.
4. *Approximating non-feedback controls.* In the spirit of the approach in [9], we show that under mild assumptions, the MFC problem can always be approximated with a regularised problem through the addition of independent Brownian noise. The regularised problem then admits an optimal feedback control. We show in a numerical example in Section 5 that our method is able to approximate a non-feedback control that depends only on its random initial condition. This justifies a neural network parametrisation of the control, which are implicitly feedback in nature.

The rest of the paper is organised as follows. In Section 2 we outline the probabilistic setting for the extended MFC problem and the properties associated with the value function. In Section 3 we give the finite-dimensional approximation for the value function of the mean field problem, which forms the basis for our numerical scheme. In Section 4 we present the numerical scheme and the computational steps in detail. Finally we conclude in Section 5 with three numerical experiments.

2 Extended mean field control problem

2.1 Probabilistic framework

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space supporting an m -dimensional Brownian motion $W = (W_t)_{t \geq 0}$. Let $\mathbb{F}^W = (\mathcal{F}_t^W)_{t \geq 0}$ denote the \mathbb{P} -completion of the natural filtration generated by W , which satisfies the usual conditions. We assume there exists a sub- σ -algebra of $\mathcal{G} \subseteq \mathcal{F}$ such that \mathcal{G} and \mathcal{F}_∞^W are independent, and that there exists an \mathcal{G} -measurable random variable $U_G : \Omega \rightarrow \mathbb{R}$ with uniform distribution on $[0, 1]$. The assumptions on \mathcal{G} imply that for every $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, there exists a random variable $\xi \in L^2(\Omega, \mathcal{G}, \mathbb{P}; \mathbb{R}^d)$ such that $\mathcal{L}(\xi) = \mu$ (cf. [10, Lemma 2.1]). Finally, denote by $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ the filtration with $\mathcal{F}_t = \mathcal{G} \vee \mathcal{F}_t^W$, which satisfies the usual conditions.

2.2 The control problem

We consider a finite horizon optimal control problem. Given a non-empty Polish space (A, d_A) , let \mathcal{A} denote the set of \mathbb{F} -progressively measurable processes $\alpha : [0, T] \times \Omega \rightarrow A$. We write $\mathcal{P}_2^{X, A} := \mathcal{P}_2(\mathbb{R}^d \times A)$. Let $[b, \sigma] : [0, T] \times \mathbb{R}^d \times A \times \mathcal{P}_2^{X, A} \rightarrow [\mathbb{R}^d, \mathbb{R}^{d \times m}]$. For an initial time $t \in [0, T]$, we consider the state equation governed by a stochastic differential equation of McKean-Vlasov type, with the additional dependence on the *joint law of the state and the control*, given by:

$$dX_s = b(s, X_s, \alpha_s, \mathcal{L}(X_s, \alpha_s)) ds + \sigma(s, X_s, \alpha_s, \mathcal{L}(X_s, \alpha_s)) dW_s, \quad X_t = \xi, \quad (2.1)$$

for $s \in (t, T]$, where $\xi \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$ and $\alpha \in \mathcal{A}$. To emphasise the dependence of X on (t, ξ, α) , we will write $X^{t, \xi, \alpha}$ when needed.

Remark 2.1. *The space $\mathcal{P}_2(\mathbb{R}^d)$, as well as the counterpart $L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$ for the initial condition, can be generalized to $\mathcal{P}_q(\mathbb{R}^d)$, resp. $L^q(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$, for any $q \geq 1$ and the results of the current section still hold true.*

We impose the following assumptions on the coefficients to guarantee the existence and uniqueness of solutions to Equation (2.1). For simplicity, we use $|\cdot|$ to denote both the Euclidean norm on \mathbb{R}^d as well as the Frobenius norm of an $\mathbb{R}^{d \times m}$ matrix.

Assumption 2.2.

(i) b, σ are bounded and continuous.

(ii) For all $(t, (x, x'), (\nu, \nu'), a) \in [0, T] \times (\mathbb{R}^d)^2 \times (\mathcal{P}_2^{X, A})^2 \times A$, there exists a constant $C_{\text{Lip}} \geq 0$ such that

$$|[b, \sigma](t, x, a, \nu) - [b, \sigma](t, x', a, \nu')| \leq C_{\text{Lip}}(|x - x'| + \mathcal{W}_2(\nu, \nu')).$$

Proposition 2.3. *Given Assumption 2.2, for any $t \in [0, T]$, $\alpha \in \mathcal{A}$, and $\xi, \xi' \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$, there exists a unique, indistinguishable solution $(X_r^{t, \xi, \alpha})_{r \in [t, T]}$ to (2.1), satisfying*

$$\mathbb{E} \left[\sup_{r \in [t, T]} |X_r^{t, \xi, \alpha}|^2 \right] \leq C_1 (1 + \mathbb{E}[|\xi|^2]),$$

for a positive constant C_1 independent of t, ξ, α . Also, for all $s \in [t, T]$,

$$\mathbb{E} \left[\sup_{r \in [s, T]} |X_r^{t, \xi, \alpha} - X_r^{s, \xi', \alpha}|^2 \right] \leq C_2 \left(\mathbb{E}[|\xi - \xi'|^2] + (1 + \mathbb{E}[|\xi|^2] + \mathbb{E}[|\xi'|^2])|s - t| \right), \quad (2.2)$$

for a positive constant C_2 independent of t, s, ξ, ξ', α . Moreover, the flow property holds, that is, for every $t \leq r \leq s \leq T$,

$$X_r^{t, \xi, \alpha} = X_r^{s, X_s^{t, \xi, \alpha}}, \quad \mathbb{P}\text{-a.s.}$$

Proof. The proof relies on standard arguments; see e.g. [11, Lemma 2.1, Remark 2.3, Proposition 3.3] and [10, Proposition 2.10]. \square

Let $f : [0, T] \times \mathbb{R}^d \times A \times \mathcal{P}_2^{X, \mathcal{A}} \rightarrow \mathbb{R}$ and $g : \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$. To define the mean field control problem, for a given admissible control $\alpha \in \mathcal{A}$ we consider the functional

$$J(t, \xi, \alpha) = \mathbb{E} \left[\int_t^T f(s, X_s, \alpha_s, \mathcal{L}(X_s, \alpha_s)) ds + g(X_T, \mathcal{L}(X_T)) \right].$$

We assume the following conditions apply to the running and terminal payoff functionals.

Assumption 2.4.

(i) f, g are bounded and continuous.

(ii) For all $(t, (x, x'), (\mu, \mu'), (\nu, \nu'), a) \in [0, T] \times (\mathbb{R}^d)^2 \times (\mathcal{P}_2(\mathbb{R}^d))^2 \times (\mathcal{P}_2^{X, \mathcal{A}})^2 \times A$, there exists a constant $C'_{\text{Lip}} \geq 0$ such that

$$\begin{aligned} |f(t, x, a, \nu) - f(t, x', a, \nu')| + |g(x, \mu) - g(x', \mu')| \\ \leq C'_{\text{Lip}} (|x - x'| + \mathcal{W}_2(\nu, \nu') + \mathcal{W}_2(\mu, \mu')). \end{aligned}$$

Remark 2.5. The boundedness conditions on b, σ, f, g required in Assumptions 2.2 and 2.4 can be relaxed with some standard growing conditions; see for example [10, Assumptions $(\mathbf{A}_{A, b, \sigma}), (\mathbf{A}_{f, g}), (\mathbf{A}_{f, g})_{\text{cont}}$] or [11, Assumption $(\mathbf{A1})$]. In addition, the results of the current section hold true by considering a weaker Lipschitz assumption for the measure argument, that is, for all $(t, (x, x'), (\nu, \nu'), a) \in [0, T] \times (\mathbb{R}^d)^2 \times (\mathcal{P}_2^{X, \mathcal{A}})^2 \times A$, there exists a constant $\bar{C}_{\text{Lip}} \geq 0$ such that

$$|[b, \sigma, f](t, x, a, \nu) - [b, \sigma, f](t, x', a, \nu')| \leq \bar{C}_{\text{Lip}} (|x - x'| + \mathcal{W}_2(\nu_x, \nu'_x)),$$

where $\nu_x \in \mathcal{P}_2(\mathbb{R}^d)$ denotes the marginal of ν on the state, i.e., $\nu_x(\cdot) = \nu(\cdot \times A)$. The intention is to retain these items to facilitate the exposition as a whole.

Thanks to the assumptions above, the so-called *lifted value function*, given by

$$V(t, \xi) := \inf_{\alpha \in \mathcal{A}} J(t, \xi, \alpha)$$

is well-defined. We now recall some important properties of the lifted value function, which will be employed in Section 3.

Proposition 2.6. Under Assumptions 2.2 and 2.4 the lifted value function V satisfies the following.

1. V is bounded.
2. V is jointly continuous, that is $V(t_n, \xi_n) \rightarrow V(t, \xi)$ as $n \rightarrow \infty$ for every couple $(t, \xi) \in [0, T] \times L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$ and sequence $\{(t_n, \xi_n)\}_{n \in \mathbb{N}_+}$ where $t_n \in [0, T]$ and $\xi_n \in L^2(\Omega, \mathcal{F}_{t_n}, \mathbb{P}; \mathbb{R}^d)$ for all $n \in \mathbb{N}_+$, such that $|t_n - t| + \mathbb{E}[|\xi_n - \xi|^2] \rightarrow 0$ as $n \rightarrow \infty$.
3. V is Lipschitz in the second argument, i.e., there exists $L \geq 0$, independent of t, ξ, ξ' , such that

$$|V(t, \xi) - V(t, \xi')| \leq L \mathbb{E}[|\xi - \xi'|^2]^{1/2}, \quad \forall t \in [0, T], \xi, \xi' \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d).$$

Proof. The boundedness of V follows from the boundedness assumption of the coefficients f and g . To prove the second claim, we first note that the Lipschitz continuity of f, g implies the following joint continuity property: for any sequence $\{(x_m, \zeta_m)\}_{m \in \mathbb{N}} \subset \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$ converging to $(x, \zeta) \in \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, as $m \rightarrow \infty$, we have

$$\sup_{\substack{t \in [0, T], a \in A \\ \eta \in \mathcal{P}_2(A)}} |f(t, x_m, a, \nu_m) - f(t, x, a, \nu)| + |g(x_m, \zeta_m) - g(x, \zeta)| \xrightarrow{m \rightarrow \infty} 0, \quad (2.3)$$

where $\nu_m, \nu \in \mathcal{P}_2^{X, \mathcal{A}}$ have marginals (ζ_m, η) and (ζ, η) , respectively. Let $\alpha \in \mathcal{A}$, and let $K \geq 0$ such that $\|f\|_\infty \leq K$. Assume without loss of generality $t_n \leq t$ (an analogous argument holds for $t_n > t$). Then, we

have

$$\begin{aligned}
& |J(t_n, \xi_n, \alpha) - J(t, \xi, \alpha)| \\
&= \left| \mathbb{E} \left[\int_{t_n}^T f(r, X_r^{t_n, \xi_n, \alpha}, \alpha_r, \mathcal{L}(X_r^{t_n, \xi_n, \alpha}, \alpha_r)) dr - \int_t^T f(r, X_r^{t, \xi, \alpha}, \alpha_r, \mathcal{L}(X_r^{t, \xi, \alpha}, \alpha_r)) dr \right. \right. \\
&\quad \left. \left. + g(X_T^{t_n, \xi_n, \alpha}, \mathcal{L}(X_T^{t_n, \xi_n, \alpha})) - g(X_T^{t, \xi, \alpha}, \mathcal{L}(X_T^{t, \xi, \alpha})) \right] \right| \\
&\leq \mathbb{E} \left[\int_{t_n}^t |f(r, X_r^{t_n, \xi_n, \alpha}, \alpha_r, \mathcal{L}(X_r^{t_n, \xi_n, \alpha}, \alpha_r))| dr \right. \\
&\quad \left. + \int_t^T |f(r, X_r^{t_n, \xi_n, \alpha}, \alpha_r, \mathcal{L}(X_r^{t_n, \xi_n, \alpha}, \alpha_r)) - f(r, X_r^{t, \xi, \alpha}, \alpha_r, \mathcal{L}(X_r^{t, \xi, \alpha}, \alpha_r))| dr \right. \\
&\quad \left. + |g(X_T^{t_n, \xi_n, \alpha}, \mathcal{L}(X_T^{t_n, \xi_n, \alpha})) - g(X_T^{t, \xi, \alpha}, \mathcal{L}(X_T^{t, \xi, \alpha}))| \right] \\
&\leq K|t - t_n| + \mathbb{E} \left[|g(X_T^{t_n, \xi_n, \alpha}, \mathcal{L}(X_T^{t_n, \xi_n, \alpha})) - g(X_T^{t, \xi, \alpha}, \mathcal{L}(X_T^{t, \xi, \alpha}))| \right] \\
&\quad + \mathbb{E} \left[\int_t^T \sup_{(r, a, \eta) \in [t, T] \times A \times \mathcal{P}_2(A)} |f_r^{t_n, \xi_n}(a, \eta) - f_r^{t, \xi}(a, \eta)| dr \right],
\end{aligned}$$

where $f_r^{t, \xi}(a, \eta) := f(r, X_r^{t, \xi, \alpha}, a, \mathcal{L}(X_r^{t, \xi, \alpha}, \eta))$. Observe that by taking the product measure \otimes as a particular coupling, one has

$$\sup_{r \in [t, T]} \mathcal{W}_2^2 \left(\mathcal{L}(X_r^{t_n, \xi_n, \alpha}), \mathcal{L}(X_r^{t, \xi, \alpha}) \right) \leq \mathbb{E} \left[\sup_{r \in [t, T]} |X_r^{t_n, \xi_n, \alpha} - X_r^{t, \xi, \alpha}|^2 \right],$$

and by applying the estimate (2.2) from Proposition 2.3, it follows by the joint continuity property of f and g that

$$|J(t_n, \xi_n, \alpha) - J(t, \xi, \alpha)| \xrightarrow{n \rightarrow \infty} 0, \quad \forall \alpha \in \mathcal{A}.$$

Finally, taking the supremum over the control gives

$$|V(t_n, \xi_n) - V(t, \xi)| \leq \sup_{\alpha \in \mathcal{A}} |J(t_n, \xi_n, \alpha) - J(t, \xi, \alpha)| \xrightarrow{n \rightarrow \infty} 0.$$

For the third claim, the proof relies on similar arguments using the estimate (2.2) from Proposition 2.3; see e.g. [9, proof of Proposition 2.3]. \square

Remark 2.7. *The joint continuity of f, g stated in (2.3) is the minimal assumption to get the conclusions in Proposition 2.6; note that the stronger Lipschitz continuity property we required in Assumption 2.4 will be needed for the finite-dimensional approximation in Section 3.*

Looking at the lifted value function, it turns out that it only depends on the law of ξ ; such a property is referred to as the *law invariance property*.

Theorem 2.8. *Given Assumptions 2.2 and 2.4, for any $t \in [0, T]$ and $\xi, \xi' \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$ such that $\mathcal{L}(\xi) = \mathcal{L}(\xi')$, it holds*

$$V(t, \xi) = V(t, \xi').$$

Therefore, one can define the intrinsic value function $v : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$ by

$$v(t, \mu) := V(t, \xi), \quad \forall (t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d), \tag{2.4}$$

where $\xi \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^d)$, such that $\mathcal{L}(\xi) = \mu$.

Proof. The proof of the result adapts the arguments presented in [10, Theorem 3.6] to our setting. We sketch here the main steps and differences, referring to [10, Section 3.3 and Appendix B] for technical details.

Step 1. Canonical representation. Suppose that there exist two \mathcal{F}_t -measurable random variables $U_\xi, U_{\xi'}$ with uniform distribution on $[0, 1]$, being independent on ξ, ξ' respectively. From [10, Lemma B.2], for any $\alpha \in \mathcal{A}$ there exists a $Prog(\mathbb{F}^{W,t}) \otimes \mathcal{B}(\mathbb{R}^d) \otimes \mathcal{B}([0, 1])$ -measurable function, denoted by \mathbf{a} , such that

$$\begin{aligned} \mathbf{a} : [0, T] \times \Omega \times \mathbb{R}^d \times [0, 1] &\rightarrow A, \quad (s, \omega, x, y) \mapsto \mathbf{a}_s(x, y)(\omega), \\ (\xi, (\alpha_s)_{s \in [t, T]}, (W_s - W_t)_{s \in [t, T]}) &\stackrel{\mathcal{L}}{=} (\xi, \mathbf{a}_s(\xi, U_\xi)_{s \in [t, T]}, (W_s - W_t)_{s \in [t, T]}), \end{aligned}$$

where the notation $\mathbf{a}_s(\xi, U_\xi)$ refers to the function $(s, \omega) \mapsto \mathbf{a}_s(\xi(\omega), U_\xi(\omega))(\omega)$, $\stackrel{\mathcal{L}}{=}$ stands for equality in law on the space $(\Omega, \mathcal{F}, \mathbb{P})$, and $\mathbb{F}^{W,t} = (\mathcal{F}_s^{W,t})_{s \geq t}$ is the \mathbb{P} -completion of the filtration generated by Brownian increments $(W_s - W_t)_{s \geq t}$ from time t on, so that $Prog(\mathbb{F}^{W,t})$ denotes the progressive σ -field on $\Omega \times [0, T]$, relative to the filtration $\mathbb{F}^{W,t}$. Similarly, we have

$$(\xi, (\alpha_s)_{s \in [t, T]}, (W_s - W_t)_{s \in [t, T]}) \stackrel{\mathcal{L}}{=} (\xi', (\mathbf{a}_s(\xi', U_{\xi'}))_{s \in [t, T]}, (W_s - W_t)_{s \in [t, T]}),$$

and also

$$((X_s^{t, \xi, \alpha})_{s \in [t, T]}, (\alpha_s)_{s \in [t, T]}) \stackrel{\mathcal{L}}{=} ((X_s^{t, \xi', \mathbf{a}(\xi', U_{\xi'})})_{s \in [t, T]}, (\mathbf{a}_s(\xi', U_{\xi'}))_{s \in [t, T]}).$$

Step 2. Conclusions. As a consequence of Step 1, we get $J(t, \xi, \alpha) = J(t, \xi', \mathbf{a}(\xi', U_{\xi'})) \geq V(t, \xi')$ for every admissible control α . By inverting the roles of ξ and ξ' , we get the converse inequality, from which $V(t, \xi) = V(t, \xi')$ and the intrinsic value function v is well-defined. It remains to show that $U_\xi, U_{\xi'}$ exist; the discrete case where $\mathcal{L}(\xi) = \sum_{i=1}^m p_i \delta_{x_i}$, for $\{x_1, \dots, x_m\} \subset \mathbb{R}^d$, $x_i \neq x_j$ for $i \neq j$ and $p_i > 0$ with $\sum_{i=1}^m p_i = 1$ follows from [10, Lemma B.3]. Given the continuity of $\xi \mapsto V(t, \xi)$ from Proposition 2.6, the case for arbitrary random variables ξ and ξ' follows from a monotone convergence argument. \square

The results from Proposition 2.6 and the law invariance property Theorem 2.8, ensure that the intrinsic value function v holds the same properties for V :

Corollary 2.9. *Under Assumptions 2.2 and 2.4 the value function v satisfies the following.*

1. v is bounded.
2. v is jointly continuous (in the sense stated in Proposition 2.6).
3. v is Lipschitz in the second argument, i.e. there exists $L > 0$ such that

$$|v(t, \mu) - v(t, \mu')| \leq L \mathcal{W}_2(\mu, \mu'), \quad \forall t \in [0, T], \mu, \mu' \in \mathcal{P}_2(\mathbb{R}^d).$$

3 Approximation of the value function

As the law μ of the initial state belongs to $\mathcal{P}_2(\mathbb{R}^d)$, the value function v is defined on an infinite-dimensional measure space. The main idea of the following section is to consider the extended mean field control problem as the limit, for $N \rightarrow \infty$, of a cooperative N -player stochastic differential game; then, in Section 4, the resulting approximation will be employed to construct a numerical scheme. In particular, the approximation we propose exploits the link between the intrinsic value function v and the value function of a finite-dimensional cooperative differential game, denoted in the following by \bar{v}_N or $\bar{v}_{\varepsilon, N}$. This link has been theoretically established in [9, Appendix A], in a more restrictive framework. If compared with [9], we allow ourselves to work under weaker assumptions on the coefficients, as we are not concerned with the formulation of the Master Bellman equation and the uniqueness of its viscosity solutions. Moreover, we extend the result in [9] by considering the additional dependence on the joint law of state and control in both drift and diffusion.

We present the approximation result considering separately the case of non-degenerate and degenerate diffusion σ . By *non-degenerate* we mean that there exists a constant $\theta > 0$ such that for all $(t, x, a, \nu) \in [0, T] \times \mathbb{R}^d \times A \times \mathcal{P}_2^{X, A}$ one has

$$\sigma \sigma^\top(t, x, a, \nu) \geq \theta \mathbf{I}_d,$$

where \mathbf{I}_d represents the d -dimensional identity matrix. In the case of a non-degenerate diffusion, we directly approximate by means of a cooperative game between a finite but large number of players. This relies on a

propagation of chaos argument which requires non-degeneracy in the diffusion term. When the diffusion is possibly degenerate, a regularisation argument can be employed to approximate the original problem with a non-degenerate one. Therefore, in the general case, our approximation can be summarised into two main steps:

- I. (if needed) Approximation by a non-degenerate control problem.
- II. Approximation by means of a cooperative N -player stochastic differential game.

Regarding the first step, the approximation of the value function still remains defined on an infinite-dimensional space, but allows us in the second step to apply a propagation of chaos result; this motivates the general Assumption 2.2 to work with possibly degenerate diffusion σ . For non-degenerate diffusions, we only need the second approximation step.

3.1 Case 1: σ non-degenerate

Given the mean field control problem, we introduce the corresponding finite-player game. Let $N \in \mathbb{N}_+$ and let $(\Omega, \mathcal{F}, \mathbb{P})$ be an (auxiliary) complete probability space, supporting independent m -dimensional Brownian motions $\bar{W}^1, \dots, \bar{W}^N$. Consider the corresponding \mathbb{P} -completion of the natural filtration generated by $\bar{W}^1, \dots, \bar{W}^N$, denoted by $\bar{\mathbb{F}}^W = (\bar{\mathcal{F}}_t^W)_{t \geq 0}$. In analogy with the mean field control problem, we assume the existence of a sub- σ -algebra $\bar{\mathcal{G}} \subset \bar{\mathcal{F}}$ such that $\bar{\mathcal{G}}$ and $\bar{\mathcal{F}}_\infty^W$ are independent and $\mathcal{P}_2(\mathbb{R}^{Nd}) = \{\mathcal{L}(\bar{\xi}) \mid \bar{\xi} \in L^2(\Omega, \bar{\mathcal{G}}, \mathbb{P}; \mathbb{R}^{Nd})\}$; finally, let $\bar{\mathbb{F}} = (\bar{\mathcal{F}}_t)_{t \geq 0}$ where $\bar{\mathcal{F}}_t = \bar{\mathcal{F}}_t^W \vee \bar{\mathcal{G}}$.

To consider the control problem for each player, we introduce $\bar{\mathcal{A}}^N$, the family of $\bar{\mathbb{F}}$ -progressively measurable processes $\bar{\alpha} : [0, T] \times \Omega \rightarrow A^N$, where $\bar{\alpha} = (\bar{\alpha}^1, \dots, \bar{\alpha}^N)$. For every initial time $t \in [0, T]$, control $\bar{\alpha} \in \bar{\mathcal{A}}^N$ and initial conditions $\bar{\xi} = (\bar{\xi}^1, \dots, \bar{\xi}^N) \in L^2(\Omega, \bar{\mathcal{F}}_t, \mathbb{P}; \mathbb{R}^{Nd})$, we consider the system of fully coupled controlled stochastic differential equations

$$d\bar{X}_s^n = b_N^n(s, \bar{X}_s, \bar{\alpha}_s) dt + \sigma_N^n(s, \bar{X}_s, \bar{\alpha}_s) d\bar{W}_s^n, \quad \bar{X}_t^n = \bar{\xi}^n, \quad (3.1)$$

for $s \in [t, T]$ and $n = 1, \dots, N$, where $\bar{X} = \bar{X}^{t, \bar{\xi}, \bar{\alpha}} = (\bar{X}^{1, t, \bar{\xi}, \bar{\alpha}}, \dots, \bar{X}^{N, t, \bar{\xi}, \bar{\alpha}})$, and b_N^n, σ_N^n, f_N^n and g_N^n are defined, for $\mathbf{x} = (x_1, \dots, x_N)$ and $\mathbf{a} = (a_1, \dots, a_N)$:

$$\begin{aligned} [b_N^n, \sigma_N^n, f_N^n] : [0, T] \times \mathbb{R}^{Nd} \times A^N &\rightarrow [\mathbb{R}^d, \mathbb{R}^{d \times m}, \mathbb{R}], & g_N^n : \mathbb{R}^{Nd} &\rightarrow \mathbb{R} \\ (t, \mathbf{x}, \mathbf{a}) &\mapsto [b, \sigma, f] \left(t, x_n, a_n, \frac{1}{N} \sum_{j=1}^N \delta_{(x_j, a_j)} \right), & \mathbf{x} &\mapsto g \left(x_n, \frac{1}{N} \sum_{j=1}^N \delta_{x_j} \right) \end{aligned}$$

for every $t \in [0, T]$ and $\bar{\mu} \in \mathcal{P}_2(\mathbb{R}^{Nd})$, with $\bar{\xi} = (\bar{\xi}^1, \dots, \bar{\xi}^N) \sim \bar{\mu} = \mu^1 \otimes \dots \otimes \mu^N$. Then, the map

$$v_N(t, \bar{\mu}) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}^N} \frac{1}{N} \sum_{n=1}^N \bar{\mathbb{E}} \left[\int_t^T f_N^n(s, \bar{X}_s, \bar{\alpha}_s) ds + g_N^n(\bar{X}_T) \right] \quad (3.2)$$

is well-defined (cf. [15]).

Theorem 3.1. *Under Assumptions 2.2 and 2.4, and assuming A is compact, let $t \in [0, T]$, $q > 2$ and $\{\mu^n\}_{n \in \mathbb{N}_+} \subset \mathcal{P}_q(\mathbb{R}^d)$ such that $\sup_{n \geq 1} \frac{1}{N} \sum_{n=1}^N \int_{\mathbb{R}^d} |x|^q \mu^n(dx) < \infty$. Then*

$$\lim_{N \rightarrow \infty} \left| v_N(t, \bar{\mu}) - v \left(t, \frac{1}{N} \sum_{n=1}^N \mu^n \right) \right| = 0.$$

Proof. See [15, Theorem 3.3]. □

Corollary 3.2. *Under Assumptions 2.2 and 2.4 and A compact, let $q > 2$, $(t, \mu) \in [0, T] \times \mathcal{P}_q(\mathbb{R}^d)$, and let $\mu^{\otimes N} := \mu \otimes \dots \otimes \mu$. It holds that*

$$\lim_{N \rightarrow \infty} v_N(t, \mu^{\otimes N}) = v(t, \mu).$$

Note that the value function v_N is still defined on an infinite-dimensional space but comes from a cooperative N -player game, where the coefficients of the problem take values from finite-dimensional spaces. In the following result, we characterise the value function v_N by averaging the players' (deterministic) positions, thereby obtaining the desired finite-dimensional approximation.

Theorem 3.3. *Under Assumptions 2.2 and 2.4, let $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$ and*

$$v_N : [0, T] \times \mathbb{R}^{Nd} \rightarrow \mathbb{R}, \quad \bar{v}_N(t, x_1, \dots, x_N) := v_N(t, \delta_{x_1} \otimes \dots \otimes \delta_{x_N}). \quad (3.3)$$

It holds that

$$v_N(t, \mu^{\otimes N}) = \int_{\mathbb{R}^{Nd}} \bar{v}_N(t, x_1, \dots, x_N) \mu(dx_1) \dots \mu(dx_N). \quad (3.4)$$

Remark 3.4. *Equation (3.4) simply stands for*

$$v_N(t, \mu^{\otimes N}) = \int_{\mathbb{R}^{Nd}} v_N(t, \delta_{x_1} \otimes \dots \otimes \delta_{x_N}) \mu(dx_1) \dots \mu(dx_N),$$

and the choice to write $v_N(t, \delta_{x_1} \otimes \dots \otimes \delta_{x_N})$ as $\bar{v}_N(t, x_1, \dots, x_N)$ is made to emphasise the fact that \bar{v}_N lives on a finite-dimensional space.

Proof of Theorem 3.3. We report here the main steps of the proof, referring to [9, equality (A.10)] for additional technical details. We first observe that \bar{v}_N corresponds to the value function of the cooperative N -player game presented in the second approximation step, with *deterministic* initial state $\bar{x} = (x_1, \dots, x_N)$ instead of the random vector $\bar{\xi} = (\xi^1, \dots, \xi^N)$. In addition, the optimal control problem involves the coefficients $b_N^n, \sigma_N^n, f_N^n, g_N^n$, which have been defined directly from b, σ, f, g , hence they satisfy Assumptions 2.2 and 2.4. By employing similar computations as in Proposition 2.6 and observing that

$$\mathcal{W}_2^2 \left(\frac{1}{N} \sum_{n=1}^N \delta_{X^n}, \frac{1}{N} \sum_{n=1}^N \delta_{Y^n} \right) \leq \bar{\mathbb{E}} \left[\frac{1}{N} \sum_{n=1}^N |X^n - Y^n|^2 \right],$$

for $\{X^n\}_{n=1}^N, \{Y^n\}_{n=1}^N \subset L^2(\bar{\Omega}, \bar{\mathcal{F}}, \bar{\mathbb{P}}; \mathbb{R}^d)$, it follows that \bar{v}_N , and clearly v_N , satisfy the properties of Corollary 2.9, suitably adapted to the context. Identity (3.4) follows as a special case of the more general equality

$$v_N(t, \bar{\mu}) = \int_{\mathbb{R}^{Nd}} \bar{v}_N(t, x^1, \dots, x^N) \bar{\mu}(dx^1, \dots, dx^N), \quad \forall (t, \bar{\mu}) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^{Nd}),$$

from which (3.4) follows upon choosing $\bar{\mu} = \mu^{\otimes N}$; the latter is equivalent to proving

$$v_N(t, \bar{\mu}) = \bar{\mathbb{E}}[\bar{v}_N(t, \bar{\xi})], \quad \forall (t, \bar{\xi}) \in [0, T] \times L^2(\bar{\Omega}, \bar{\mathcal{F}}_t, \bar{\mathbb{P}}; \mathbb{R}^{Nd}) \text{ s.t. } \mathcal{L}(\bar{\xi}) = \bar{\mu}. \quad (3.5)$$

The proof of (3.5) consists of showing the double inequality $v_N(t, \bar{\mu}) \leq \bar{\mathbb{E}}[\bar{v}_N(t, \bar{\xi})]$ and $v_N(t, \bar{\mu}) \geq \bar{\mathbb{E}}[\bar{v}_N(t, \bar{\xi})]$ (cf. [9, Theorem A.7]). This is possible first considering $\bar{\xi} \in L^2(\bar{\Omega}, \bar{\mathcal{F}}_t, \bar{\mathbb{P}}; \mathbb{R}^{Nd})$ taking only a finite number of values, that is $\bar{\xi} = \sum_{k=1}^K \bar{x}_k 1_{\bar{E}_k}$, for some $K \in \mathbb{N}_+, \bar{x}_k \in \mathbb{R}^{Nd}$ and $\bar{E}_k \subset \sigma(\bar{\xi})$ with $\{\bar{E}_k\}_{k=1}^K$ being a partition of $\bar{\Omega}$. The general case $\bar{\xi} \in L^2(\bar{\Omega}, \bar{\mathcal{F}}_t, \bar{\mathbb{P}}; \mathbb{R}^{Nd})$ can be deduced from approximation arguments, recalling that \bar{v}_N satisfies the properties of Corollary 2.9, i.e. is bounded, jointly continuous and Lipschitz in \bar{x} . \square

Finally, we state the main result.

Theorem 3.5. *Under Assumptions 2.2 and 2.4, and assuming A is compact, it holds that, for every $(t, \mu) \in [0, T] \times \mathcal{P}_q(\mathbb{R}^d)$, $q > 2$,*

$$v(t, \mu) = \lim_{N \rightarrow \infty} \int_{\mathbb{R}^{Nd}} \bar{v}_N(t, x_1, \dots, x_N) \mu(dx_1) \dots \mu(dx_N). \quad (3.6)$$

Proof. The proof follows directly from Corollary 3.2 and Theorem 3.3. \square

Remark 3.6. *We recall an additional important approximation result. We presented an approximation of the value function, through its finite-dimensional counterpart, which corresponds to the value function \bar{v}_N of a cooperative N -player stochastic differential game. In Section 4, we approximate \bar{v}_N by learning δ_N -optimal controls of the large population cooperative control problem, for $\delta_N \searrow 0$, which can be shown to converge towards an optimal control of the original mean field control problem. Note that the convergence of the control is provided in the "weak sense", given by measure-valued rules. We refer to the most recent result in literature [15, Proposition 3.4].*

3.2 Case 2: σ degenerate

As remarked before, if σ is degenerate, we first need an approximation with a non-degenerate problem in order to apply Theorem 3.1. We present in detail the first approximation step and briefly sketch the second one, which is a straightforward adaptation of Section 3.1.

3.2.1 Approximation by a non-degenerate control problem

We consider an auxiliary complete probability space $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{P}})$ supporting a m -dimensional Brownian motion $\hat{W} = (\hat{W}_t)_{t \geq 0}$ and a new additional d -dimensional Brownian motion $\hat{B} = (\hat{B}_t)_{t \geq 0}$, independent of \hat{W} . Denote by $\hat{\mathcal{A}}$ the control process space and by $\hat{\mathbb{F}} = (\hat{\mathcal{F}}_t)_{t \geq 0}$ the standard filtration, both defined in analogy to Section 2. Now, consider $\varepsilon > 0$, $\hat{\xi} \in L^2(\hat{\Omega}, \hat{\mathcal{F}}_t, \hat{\mathbb{P}}; \mathbb{R}^d)$, $\hat{\alpha} \in \hat{\mathcal{A}}$ and denote by $\hat{X}^{\varepsilon, t, \hat{\xi}, \hat{\alpha}}$ the unique solution of the non-degenerate McKean-Vlasov stochastic differential equation

$$d\hat{X}_s = b(s, \hat{X}_s, \hat{\alpha}_s, \mathcal{L}(\hat{X}_s, \hat{\alpha}_s)) ds + \sigma(s, \hat{X}_s, \hat{\alpha}_s, \mathcal{L}(\hat{X}_s, \hat{\alpha}_s)) d\hat{W}_s + \varepsilon d\hat{B}_s, \quad \hat{X}_t = \hat{\xi}.$$

We then define the lifted valued function

$$V_\varepsilon(t, \hat{\xi}) := \inf_{\hat{\alpha} \in \hat{\mathcal{A}}} J(t, \hat{\xi}, \hat{\alpha}) = \inf_{\hat{\alpha} \in \hat{\mathcal{A}}} \hat{\mathbb{E}} \left[\int_t^T f(s, \hat{X}_s, \hat{\alpha}_s, \mathcal{L}(\hat{X}_s, \hat{\alpha}_s)) ds + g(\hat{X}_T, \mathcal{L}(\hat{X}_T)) \right].$$

The law invariance property, as stated in Theorem 2.8 holds, and we can define the intrinsic value function $v_\varepsilon : [0, T] \times \mathcal{P}_2(\mathbb{R}^d) \rightarrow \mathbb{R}$. Clearly, the properties of Corollary 2.9 hold the same for v_ε . In addition, since \hat{B} is independent of \hat{W} , we have that $v_0(t, \mu) = v(t, \mu)$, for all $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$. We then formulate the first limit result.

Lemma 3.7. *Under Assumptions 2.2 and 2.4, there exists a non-negative constant \hat{C} , depending only on $C_{\text{Lip}}, C'_{\text{Lip}}$ and T , such that, for every $\varepsilon \geq 0$ we have*

$$|v_\varepsilon(t, \mu) - v_0(t, \mu)| \leq \hat{C}\varepsilon,$$

for all $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$.

Proof. By employing similar arguments as in [24, Theorem 2.5.9] and [9, Lemma A.5], one gets

$$\hat{\mathbb{E}} \left[\sup_{s \in [t, T]} |X_s^{\varepsilon, t, \hat{\xi}, \hat{\alpha}} - X_s^{0, t, \hat{\xi}, \hat{\alpha}}|^2 \right] \leq C\varepsilon^2,$$

where $C \equiv C(C_{\text{Lip}}, T) \geq 0$, for every $(\hat{\xi}, \hat{\alpha}) \in L^2(\hat{\Omega}, \hat{\mathcal{F}}_t, \hat{\mathbb{P}}; \mathbb{R}^d) \times \hat{\mathcal{A}}$. Finally, with standard techniques as the ones shown in Proposition 2.6, we have

$$|v_\varepsilon(t, \mu) - v_0(t, \mu)| \leq 2C'_{\text{Lip}}(T+1)\sqrt{C\varepsilon^2} =: \hat{C}\varepsilon. \quad \square$$

3.2.2 Approximation by a cooperative stochastic differential game of finitely many players

Let $N \in \mathbb{N}_+$ and let $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ an auxiliary complete probability space, supporting $\tilde{W}^1, \dots, \tilde{W}^N$ independent m -dimensional Brownian motions, $\tilde{B}^1, \dots, \tilde{B}^N$ independent d -dimensional Brownian motions. Finally define $\tilde{\mathbb{F}} = (\tilde{\mathcal{F}}_t)_{t \geq 0}$ according to Section 3.1.

We introduce $\tilde{\mathcal{A}}^N$ the family of $\tilde{\mathbb{F}}$ -progressively measurable processes $\tilde{\alpha} : [0, T] \times \tilde{\Omega} \rightarrow A^N$, where $\tilde{\alpha} = (\tilde{\alpha}^1, \dots, \tilde{\alpha}^N)$. For every initial time $t \in [0, T]$, control $\tilde{\alpha} \in \tilde{\mathcal{A}}^N$ and initial conditions $\tilde{\xi} = (\tilde{\xi}^1, \dots, \tilde{\xi}^N) \in L^2(\tilde{\Omega}, \tilde{\mathcal{F}}_t, \tilde{\mathbb{P}}; \mathbb{R}^{Nd})$, we consider the system of fully coupled controlled stochastic differential equations

$$d\tilde{X}_s^i = b_N^n(s, \tilde{X}_s, \tilde{\alpha}_s) dt + \sigma_N^n(s, \tilde{X}_s, \tilde{\alpha}_s) d\tilde{W}_s^n + \varepsilon d\tilde{B}_s^n, \quad \tilde{X}_t^i = \tilde{\xi}^i,$$

for $s \in [t, T]$ and $n = 1, \dots, N$, where $\tilde{X} = \tilde{X}^{t, \tilde{\xi}, \tilde{\alpha}} = (\tilde{X}^{1, t, \tilde{\xi}, \tilde{\alpha}}, \dots, \tilde{X}^{N, t, \tilde{\xi}, \tilde{\alpha}})$, and the coefficients $b_N^n, \sigma_N^n, f_N^n, g_N^n$ are defined as in Section 3.1. Again, for every $t \in [0, T]$ and $\tilde{\mu} \in \mathcal{P}_2(\mathbb{R}^{Nd})$, with $\tilde{\xi} = (\tilde{\xi}^1, \dots, \tilde{\xi}^N) \sim \tilde{\mu} = \tilde{\mu}^1 \otimes \dots \otimes \tilde{\mu}^N$, it is well defined

$$v_{\varepsilon, N}(t, \tilde{\mu}) = \inf_{\tilde{\alpha} \in \tilde{\mathcal{A}}^N} \frac{1}{N} \sum_{n=1}^N \tilde{\mathbb{E}} \left[\int_t^T f_N^n(s, \tilde{X}_s, \tilde{\alpha}_s) ds + g_N^n(\tilde{X}_T) \right]$$

According to Theorem 3.1 the following lemma holds.

Lemma 3.8. *Under Assumptions 2.2 and 2.4 and A compact, let $(t, \mu) \in [0, T] \times \mathcal{P}_q(\mathbb{R}^d)$, $q > 2$, then it holds that*

$$\lim_{n \rightarrow \infty} v_{\varepsilon, N}(t, \mu^{\otimes N}) = v_{\varepsilon}(t, \mu).$$

Moreover, analogously to Theorem 3.3, one has the following representation for $v_{\varepsilon, N}$.

Theorem 3.9. *Under Assumptions 2.2 and 2.4, let $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$ and consider the map*

$$\bar{v}_{\varepsilon, N} : [0, T] \times \mathbb{R}^{Nd}, \quad \bar{v}_{\varepsilon, N}(t, x_1, \dots, x_N) := v_{\varepsilon, N}(t, \delta_{x_1} \otimes \dots \otimes \delta_{x_N}).$$

Then it holds that

$$v_{\varepsilon, N}(t, \mu^{\otimes N}) = \int_{\mathbb{R}^{Nd}} \bar{v}_{\varepsilon, N}(t, x_1, \dots, x_N) \mu(dx_1) \dots \mu(dx_N).$$

3.2.3 Approximation result

By Lemma 3.8, 3.7, and Theorem 3.9, we can finally state the main approximation result for diffusions that are possibly degenerate.

Theorem 3.10. *Under Assumptions 2.2 and 2.4, and assuming A is compact, it holds that, for every $(t, \mu) \in [0, T] \times \mathcal{P}_q(\mathbb{R}^d)$, $q > 2$,*

$$v(t, \mu) = \lim_{\varepsilon \rightarrow 0} \lim_{N \rightarrow \infty} \int_{\mathbb{R}^{Nd}} \bar{v}_{\varepsilon, N}(t, x_1, \dots, x_N) \mu(dx_1) \dots \mu(dx_N).$$

Proof. The proof follows from Lemma 3.7, Lemma 3.8 and Theorem 3.9. □

Remark 3.11. *Similarly to Remark 3.6, we stress here that $\bar{v}_{\varepsilon, N}$ is approximated by learning δ_n^ε -optimal controls of the large population cooperative control problem, for $\delta_n^\varepsilon \searrow 0$, converging in the sense of measure-valued rules towards an optimal control of the starting mean field control problem. This convergence is only provided for every fixed $\varepsilon > 0$, and is not uniform in ε .*

The result of Theorem 3.3 (resp. Theorem 3.10) implies that for a sufficiently large N (and sufficiently small ε), a simple integration of the finite dimensional value function \bar{v}_N (resp. $\bar{v}_{\varepsilon, N}$) can be used as an approximation of the value function of the extended mean-field optimal control problem. We provide an implementable algorithm to obtain this approximation in the next section.

Remark 3.12. *Under stronger assumptions, the value function $v(t, \mu)$ is a viscosity solution of the Master Bellman equation:*

$$\begin{cases} \partial_t v(t, \mu) = F(t, \mu, v(t, \mu), \partial_\mu v(t, \mu)(\cdot), \partial_x \partial_\mu v(t, \mu)(\cdot)), & (t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d), \\ v(T, \mu) = \int_{\mathbb{R}^d} g(x, \mu) \mu(dx), & \mu \in \mathcal{P}_2(\mathbb{R}^d), \end{cases}$$

where F is the associated maximised Hamiltonian for the functions f, b, σ , see e.g. [9, Equation (3.3)] or [10, Equation (5.1)]. In this case, the numerical method of Section 4 can be employed to approximate the possibly unique solution of the Master Bellman equation.

4 A computational scheme with neural networks

We detail here the precise steps for the numerical computation of the mean field value function (2.4). As outlined in the introduction, we avoid learning directly the value function in the Wasserstein space, but instead we first learn an approximating finite-dimensional value function \bar{v}_N (or $\bar{v}_{\varepsilon, N}$) before evaluating a corresponding integral for the input measure. We point out that \bar{v}_N (resp. $\bar{v}_{\varepsilon, N}$) is the value function associated to a classical optimal control problem in dimension Nd and is the unique viscosity solution to a Hamilton-Jacobi-Bellman (HJB) equation under our assumptions.

For simplicity, we fix a time horizon of $[0, T]$ and large $N \in \mathbb{N}_+$ (the choice of N depends on the specific problem, see Section 5). Due to the potentially high dimensionality of the problem, classical numerical

methods for HJB equations, such as finite differences, finite elements and semi-Lagrangian schemes, suffer from the well known ‘‘curse of dimensionality’’. For this reason, the algorithm we propose to numerically compute \bar{v}_N (resp. $\bar{v}_{\varepsilon,N}$) can make use of neural networks so that it is implementable in high dimension. As we already mentioned in the introduction, the choice of the numerical method used to approximate the finite dimensional value function is not restrictive and the one proposed below (Sections 4.1 and 4.2) is just a possible one having the advantage of working in the most general setting.

For the purposes of numerical computation, we shall assume that the control space A is a subset of some Euclidean space \mathbb{R}^{d_A} . For our experiments, we will parametrise both the control and value functions with neural networks. We adopt a similar network structure to that used in [26]. Firstly, the output of two sub-networks, taking the time t and the particles (x_1, \dots, x_N) as inputs respectively, are concatenated. Then, this concatenated output is passed as an input to one final sub-network. Each sub-network is of feedforward form, comprising of one hidden layer. By a feedforward network we mean a map $\mathcal{NN} : \mathbb{R}^{w_0} \rightarrow \mathbb{R}^{w_l}$, specified by the number of layers l , the widths $\mathbf{w} = (w_0, \dots, w_l)$ of each layer and an activation function $\sigma_{\text{act}} : \mathbb{R} \rightarrow \mathbb{R}$. For an input $x \in \mathbb{R}^{w_0}$ we have

$$\mathcal{NN}(x) = \sigma_{\text{out}}(W_l(\sigma_{\text{act}}(W_{l-1}(\dots\sigma_{\text{act}}(W_1x + b_1)) + b_{l-1})) + b_l)$$

where $W_j \in \mathbb{R}^{w_j \times w_{j-1}}$, $b_j \in \mathbb{R}^{w_j}$ for $1 \leq j \leq l$ are known as the weights and biases of the network, and the activation function σ_{act} is applied component-wise to vectors. The function $\sigma_{\text{out}} : \mathbb{R} \rightarrow \mathbb{R}$, applied component-wise on the last layer, serves to modify the range of the neural network. In contrast to [26], the width of the hidden layer to grows linearly with respect to the dimension of the players; in our experiments this is taken to be $10 + Nd$. This is done to allow for a sufficient number of parameters in the network to avoid underfitting.

4.1 Obtaining an approximately optimal policy

The first step is to obtain an approximately optimal policy for the finite N -player game through a policy gradient iteration. To do so, we first approximate the system of N -players $(\bar{X}^1, \dots, \bar{X}^N)$ in (3.1) with the Euler-Maruyama scheme as follows. We discretise the time horizon $[0, T]$ into equal subintervals of length $\Delta t > 0$, and let $\Delta W_s^n := W_{s+\Delta t}^n - W_s^n$ denote the Brownian increments. The approximating time-discretised system is denoted by $(\hat{X}^1, \dots, \hat{X}^N)$, which for each $n \in \{1, \dots, N\}$ and $s \in \{0, \Delta t, 2\Delta t, \dots, T - \Delta t\}$, satisfies the dynamics

$$\hat{X}_{s+\Delta t}^n = \hat{X}_s^n + b_N^n(s, \hat{X}_s, \alpha_s)\Delta t + \sigma_N^n(s, \hat{X}_s, \alpha_s)\Delta W_s^n, \quad (4.1)$$

where the functions b_N^n , σ_N^n , f_N^n and g_N^n are as given in Section 3. If σ is possibly degenerate, we simulate additional Brownian increments $\Delta B_t^n = B_{t+\Delta t}^n - B_t^n$, independent of ΔW_t , and for some small $\varepsilon > 0$, consider instead

$$\hat{X}_{s+\Delta t}^n = \hat{X}_s^n + b_N^n(s, \hat{X}_s, \alpha_s)\Delta t + \sigma_N^n(s, \hat{X}_s, \alpha_s)\Delta W_s^n + \varepsilon \Delta B_s^n.$$

We now parameterise the control by a neural network with parameters θ , and denote this control by α_θ . Choose a measure $\mu_0 \in \mathcal{P}_2(\mathbb{R}^d)$ to act as a sampling measure for \hat{X}_0 (see Remark 4.1 on the choice of such a sampling measure), and consider the (discretised) cost functional

$$J_N(\theta) = \frac{1}{N} \sum_{n=1}^N \mathbb{E}_{\hat{X}_0 \sim \mu_0} \left[\sum_{s=0}^{T-\Delta t} f_N^n(s, \hat{X}_s, \alpha_s)\Delta t + g_N^n(\hat{X}_T) \right].$$

The gradient of the cost functional with respect to the parameters θ is given by

$$\nabla_\theta J_N(\theta) = \frac{1}{N} \sum_{n=1}^N \mathbb{E}_{\hat{X}_0 \sim \mu_0} \left[\sum_{s=0}^{T-\Delta t} \nabla_\theta f_N^n(s, \hat{X}_s, \alpha_\theta(s, \hat{X}_s))\Delta t + \nabla_\theta g_N^n(\hat{X}_T) \right].$$

We then estimate $\nabla_\theta J_N(\theta)$ with Monte Carlo simulation. Let $M > 0$ denote the total number of sample trajectories of the system \hat{X} , given by (4.1). For each $1 \leq m \leq M$ and $1 \leq n \leq N$, let $\hat{X}_s^{n,m}$ denote the m^{th} sample of the n^{th} player at time s . The Monte Carlo estimator of $\nabla_\theta J_N(\theta)$ is then given by

$$G_\theta := \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \sum_{s=0}^{T-\Delta t} \nabla_\theta f_N^n(s, \hat{X}_s^{n,m}, \alpha_\theta(s, \hat{X}_s^{n,m}))\Delta t + \nabla_\theta g_N^n(\hat{X}_T^{n,m}).$$

To apply the policy iteration, first initialise the parameters of the neural network by some θ_0 , as well as a learning rate $\gamma > 0$. Then the successive iterations for the parameters are given by the update rule

$$\theta_{j+1} := \theta_j - \gamma G_{\theta_j}, \quad j \geq 0.$$

At each iteration j , the cost functional $J_N(\theta_j)$ is evaluated. The iteration procedure terminates when no further minimisation in J_N occurs after a prespecified number of iterations. We take the neural network after termination as an approximately optimal control, and denote this by α^* . We note that although each $\hat{X}_t^{n,m}$ depend implicitly on the parameters θ , the terms $\nabla_{\theta} g_N^n(\hat{X}_T^{:,m})$ and $\nabla_{\theta} f_N^n(s, \hat{X}_s^{:,m}, \alpha_{\theta}(s, \hat{X}_s^{:,m}))$ can be computed efficiently in practice with automatic differentiation tools.

Remark 4.1. *The sampling measure μ_0 should have support on the domain of X_0 and should be chosen such that there is a sufficient sampling of the domain. Alternatively, one can also truncate the domain and restrict the initial conditions within some compact set $\mathcal{K} \subset \mathbb{R}^d$, sampling \mathcal{K} according to the uniform distribution.*

4.2 Computing the finite-dimensional value function

For the rest of this section, we will assume that the diffusion coefficient σ is non-degenerate and apply the results from Section 3.1. For the case that σ is possibly degenerate, it is sufficient to refer to Section 3.2, and substitute the value function \bar{v}_N with $\bar{v}_{\varepsilon,N}$ (for ε small).

Recall from (3.2) and (3.3) that \bar{v}_N is the value function for the N -player game, and that for any $t \in [0, T]$, $(x_1, \dots, x_N) \in (\mathbb{R}^d)^N$, we have

$$\bar{v}_N(t, x_1, \dots, x_N) = \inf_{\bar{\alpha} \in \bar{\mathcal{A}}^N} \frac{1}{N} \sum_{n=1}^N \mathbb{E} \left[\int_t^T f_N^n(s, \bar{X}_s, \bar{\alpha}_s) ds + g_N^n(\bar{X}_T) \right].$$

By using the discretised dynamics \hat{X} and approximately optimal control α^* as obtained in the previous section, given an initial condition $\hat{X}_t^{:,m} = (x_1, \dots, x_N)$, an estimate of $\bar{v}_N(t, x_1, \dots, x_N)$ is given by

$$\bar{v}_N^{MC}(t, x_1, \dots, x_N) := \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \sum_{s=t}^{T-\Delta t} f_N^n(s, \hat{X}_s^{:,m}, \alpha^*(s, \hat{X}_s^{:,m})) \Delta t + g_N^n(\hat{X}_T^{:,m}). \quad (4.2)$$

Then, in virtue of Theorem 3.3, given $\mu \in \mathcal{P}_2(\mathbb{R}^d)$, an approximation of the (mean-field) value function $v(t, \mu)$ can be obtained by evaluating the integral

$$\int_{\mathbb{R}^{Nd}} \bar{v}_N^{MC}(t, x_1, \dots, x_N) \mu(dx_1) \dots \mu(dx_N). \quad (4.3)$$

This provides an approximation for $v(t, \mu)$ for each fixed $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^d)$. To obtain an approximation of v across the whole domain $[0, T] \times \mathcal{P}_2(\mathbb{R}^d)$, we employ an additional regression step to approximate \bar{v}_N , using samples of (4.2) over the finite-dimensional domain $[0, T] \times \mathbb{R}^{Nd}$. Since we do not need to approximate the distribution of \hat{X} in (4.2), the number of Monte Carlo samples M can be chosen flexibly to accommodate for individual computational resources. This differs from the approach of [28, 29], where the regression is performed directly for the mean-field value function v on the infinite-dimensional domain $[0, T] \times \mathcal{P}_2(\mathbb{R}^d)$, and a high number of Monte Carlo samples per iteration is required to approximate the distribution of \hat{X} in (2.1).

In the case of regression, consider a parametrised function $\bar{v}_{N,\eta} : [0, T] \times \mathbb{R}^{Nd} \rightarrow \mathbb{R}$ with parameters η . Let $K > 0$ denote the number of samples for regression, and let $(t^k, \mathbf{x}^k) \in [0, T] \times \mathbb{R}^{Nd}$, $1 \leq k \leq K$. The samples are denoted by $y_1, \dots, y_K \in \mathbb{R}$, which correspond to the values obtained from the expression (4.2) with the initial conditions $(t^1, \mathbf{x}^1), \dots, (t^K, \mathbf{x}^K)$ respectively. Consider a mean-squared error (MSE) loss function

$$L(\eta) := \frac{1}{K} \sum_{k=1}^K |\bar{v}_{N,\eta}(t^k, \mathbf{x}^k) - y_k|^2.$$

Let $\gamma' > 0$ be the regression learning rate, and let η_j be the function parameters at the j^{th} iteration, which are updated according to $\eta_{j+1} := \eta_j - \gamma' \nabla_{\eta} L(\eta)$, for $j \geq 0$. As in the case of the policy iteration step, the iteration terminates when $L(\eta)$ does not decrease after a prespecified number of iterations. The parameters

after termination are taking to be approximately optimal and are denoted by η^* . For simplicity, we denote the resulting function \bar{v}_{N,η^*} by \bar{v}_N^* .

In practice, we observe a slow convergence rate during the training of $\bar{v}_{N,\eta}$ when minimising for $L(\eta)$. To accelerate the training process, we employ the differential regression method as outlined in [26], which learns in addition the gradient of the value function during training. In particular, each y_k can be considered as a function of the initial condition (t^k, \mathbf{x}^k) , so that the derivatives $\partial_t y_k$ and $\nabla_{\mathbf{x}} y_k$ are unbiased estimators of $\partial_t \bar{v}_N^{MC}(t, x_1, \dots, x_N)$ and $\nabla_{\mathbf{x}} \bar{v}_N^{MC}(t, x_1, \dots, x_N)$ respectively. The gradient $\nabla_{\mathbf{x}} y_k$ can in turn be evaluated by computing the partial derivatives of y_k with respect to each simulated trajectory for each player across each $s \in \{0, \Delta t, \dots, T - \Delta t\}$ via automatic differentiation. Then, $\bar{v}_{N,\eta}$ is trained on the weighted loss function

$$L_w(\eta) := \frac{1}{K} \sum_{k=1}^K \left(w_{\text{val}} |\bar{v}_{N,\eta}(t^k, \mathbf{x}^k) - y_k|^2 + w_{\text{der}} \left[\|\nabla_{\mathbf{x}} \bar{v}_{N,\eta}(t^k, \mathbf{x}^k) - \nabla_{\mathbf{x}} y_k\|_2^2 + |\partial_t \bar{v}_{N,\eta}(t^k, \mathbf{x}^k) - \partial_t y_k|^2 \right] \right),$$

for some constants $w_{\text{val}}, w_{\text{der}} > 0$. We find that for our experiments, optimising over L_w was more effective with shorter experiment times, compared to the approach in [26] of alternating between two separate loss functions of the value function and its derivatives.

Remark 4.2. *When the diffusion coefficient σ_N^n is independent of the control α , it is also possible to learn an optimal control via a backwards stochastic differential equation (BSDE) characterisation of the value function. In this case, the learning procedure coincides with the one of a BSDE solvers such as in [21, 23].*

Remark 4.3. *The Lipschitz assumptions Assumptions 2.2 and 2.4 on the coefficients of b, σ, f, g are sufficient conditions for the obtaining an unbiased estimator of $\nabla_{\mathbf{x}} \bar{v}_{N,\eta}(t^k, \mathbf{x}^k)$. We refer to the textbook [20, Ch. 7] and the references within for a more detailed treatment of pathwise derivative estimators.*

4.3 Recovering the mean field value function

Given, for $N \in \mathbb{N}_+$, the approximation \bar{v}_N^* of the finite-dimensional value function \bar{v}_N and recall that the mean field value function is given by the limit in (3.6), we obtain its approximation by evaluating, for N big enough, the integral in (3.6). In other words we have, for N sufficiently large,

$$v(t, \mu) \approx I_N := \int_{\mathbb{R}^{Nd}} \bar{v}_N^*(t, x_1, \dots, x_N) \mu(dx_1) \dots \mu(dx_N).$$

In the case where the process X is real-valued, the integral I_N can be evaluated by using the inverse transform sampling. Given $\mu \in \mathcal{P}(\mathbb{R})$, let $F_\mu : \mathbb{R} \rightarrow [0, 1]$ be the cumulative distribution function associated to μ . Then one obtains

$$I_N = \mathbb{E} [\bar{v}_N^*(t, F_\mu^{-1}(U_1), \dots, F_\mu^{-1}(U_N))],$$

where U_1, \dots, U_N are independent uniformly distributed random variables over $[0, 1]$, so that evaluating I_N is reduced to sampling over the unit hypercube. This allows for the required number of samples to be reduced by the use of quasi-Monte Carlo methods, for example employing the use of Sobol sequences. We summarise the overall algorithm in this section in pseudo-code in Line 1 for the MSE loss function L . Observe that the first “while” loop in Line 1 learns the optimal control, while the second loop learns the finite-dimensional value function.

5 Numerical experiments

In this section we consider three experiments for our proposed numerical algorithm. The rate of convergence of the finite dimensional value functions towards the mean field value function is only proven for a few specific cases [4, 6, 8, 19], and to our best knowledge, the general case remains an open problem. Our numerical tests suggests that an upper bound of $N = 500$ players can approximate the mean field value function up to a 0.1% degree of accuracy. In the case of a weaker mean field coupling, it maybe be possible to consider a lower dimension for a similar level of accuracy, such as the case in Example 3, where we perform the regression

Algorithm 1: Finite-dimensional approximation for the MFC problem

```
1: Input: Dynamics  $b_N^n, \sigma_N^n$ ; costs  $f_N^n, g_N^n$ ; stepsize  $\Delta t$ ; max patience  $P_{\max}$ 
2: Input: Control  $\alpha_\theta$ ; learning rate  $\gamma$ 
3:  $v_{\text{best}} \leftarrow +\infty, \text{patience} \leftarrow 0$ 
4: while  $\text{patience} < P_{\max}$  do
5:   Initialise  $v = 0$ , sample  $x = (x_1, \dots, x_N)$  uniformly on the domain
6:   for  $t = 0, \Delta t, \dots, T - \Delta t$  do
7:      $a \leftarrow \alpha_\theta(t, x), v \leftarrow v + \sum_{n=1}^N (f_N^n(t, x, a)/N)\Delta t$ 
8:     Generate Brownian increments  $\Delta W_t^n, n \in \{1, \dots, N\}$ 
9:      $x_n \leftarrow x_n + b_N^n(t, x, a)\Delta t + \sigma_N^n(t, x, a)\Delta W_t^n, n \in \{1, \dots, N\}$ 
10:     $v \leftarrow v + \sum_{n=1}^N g_N^n(x)/N, \theta \leftarrow \theta - \gamma \nabla_\theta v$ 
11:    if  $v < v_{\text{best}}$  then
12:       $v_{\text{best}} \leftarrow v, \text{patience} \leftarrow 0$ 
13:    else
14:       $\text{patience} \leftarrow \text{patience} + 1$ 
15:  return  $\alpha_{\theta^*}$ 
16: Input: Learning rate  $\gamma'$ ; value function  $\bar{v}_{N,\eta}$ ; measure  $\mu$ 
17:  $L_{\text{best}} \leftarrow +\infty, \text{patience} \leftarrow 0$ 
18: while  $\text{patience} < P_{\max}$  do
19:   Initialise loss  $L = 0$ , sample  $x = (x_1, \dots, x_N)$  uniformly on the domain
20:   Repeat Lines 7 to 10 to compute  $v$ , but replace  $\alpha_\theta$  with  $\alpha_{\theta^*}$ .
21:    $L \leftarrow L + |\bar{v}_{N,\eta}(0, x_1, \dots, x_N) - v|^2, \eta \leftarrow \eta - \gamma' \nabla_\eta L$ 
22:   if  $L < L_{\text{best}}$  then
23:      $L_{\text{best}} \leftarrow L, \text{patience} \leftarrow 0$ 
24:   else
25:      $\text{patience} \leftarrow \text{patience} + 1$ 
26: return  $\int_{x_1, \dots, x_N} \bar{v}_N^*(0, x_1, \dots, x_N) \mu(x_1) \dots \mu(x_N)$ 
```

step with 100 players. This reduces the complexity of the problem compared to the sampling and regression of the value function in the Wasserstein space.

In all experiments, the training is performed on the Google Colab platform, using the NVIDIA T4 GPU. We choose our examples such that explicit solutions are available for benchmarking purposes. In these cases, we evaluate the performance of the method by computing a residual loss for the mean field control problem. Let $P > 0$ be the number of samples in a test set of quantized measures, at which that value functions will be evaluated, and let $L > 0$ be the number of sample points on the underlying domain \mathbb{R}^{Nd} . The quantized measures can be generated as follows. Sample a number of points $x_1, \dots, x_L \in \mathbb{R}^{Nd}$, and sample the exponential random variables $E_1, \dots, E_L \sim \text{Exp}(1)$, set $\pi_l = E_l / (\sum_j E_j)$. For $1 \leq p \leq P$, define the quantized measures $\mu^p = \sum_{l=1}^L \pi_l \delta_{x_l}$. Then, we compute the residual loss function

$$\mathcal{L}_{res} := \frac{1}{P} \sum_{p=1}^P \left| v(t, \mu^p) - \int_{\mathbb{R}^{Nd}} \bar{v}_N^*(t, x_1, \dots, x_N) \mu^p(dx_1) \dots \mu^p(dx_N) \right|^2.$$

When an analytical expression for $v(t, \mu^p)$ is not available, we can still evaluate the performance of the method by considering the accuracy of the approximation \bar{v}_N^* for the finite-dimensional value function. In this case, we compute the loss associated to the finite-dimensional Hamilton-Jacobi-Bellman (HJB) equation as follows. Let $\Delta t, \Delta x \in \mathbb{R}_+$ be given step sizes in the time and spatial directions respectively. Let $\mathcal{T} = \{0, \Delta t, \dots, T - \Delta t, T\}$, and let \mathcal{X} be a discretised grid with step Δx on the domain \mathbb{R}^d , on which the computation is performed. For a sufficiently smooth function $\phi : [0, T] \times \mathbb{R}^{Nd} \rightarrow \mathbb{R}$, we denote by $\mathcal{D}_t(\phi)$ the approximation of the partial derivative $\partial_t \phi$ obtained by automatic differentiation. Similarly, we denote by $\mathcal{D}_x(\phi)$ and $\mathcal{D}_x^2(\phi)$ the approximation of the gradient vector $D_x(\phi)$ and Hessian matrix $D_x^2(\phi)$ respectively.

Then, consider the HJB loss

$$\begin{aligned} \mathcal{L}_{\text{HJB}} := & \frac{1}{|\mathcal{T}||\mathcal{X}|} \sum_{\substack{t \in \mathcal{T} \\ \mathbf{x} \in \mathcal{X}}} \left| \mathcal{D}_t(\bar{v}_N^*)(t, \mathbf{x}) + H(t, \mathbf{x}, \mathcal{D}_{\mathbf{x}}(\bar{v}_N^*)(t, \mathbf{x}), \mathcal{D}_{\mathbf{x}}^2(\bar{v}_N^*)(t, \mathbf{x})) \right|^2 \\ & + \frac{1}{|\mathcal{X}|} \sum_{\mathbf{x} \in \mathcal{X}} \left| \bar{v}_N^*(T, \mathbf{x}) - \frac{g_N^n(\mathbf{x})}{N} \right|^2, \end{aligned}$$

where $H : [0, T] \times \mathbb{R}^{Nd} \times \mathbb{R}^{Nd} \times \mathbb{R}^{(Nd \times Nd)} \rightarrow \mathbb{R}$ is the Hamiltonian for the associated finite-dimensional PDE, i.e.,

$$\begin{aligned} H(t, \mathbf{x}, p, M) := & \frac{1}{N} \sum_{n=1}^N f_N^n(t, \mathbf{x}, \alpha_\theta^*(t, \mathbf{x})) + \sum_{n=1}^N \langle b_N^n(t, \mathbf{x}, \alpha_\theta^*(t, \mathbf{x})), p^n \rangle \\ & + \frac{1}{2} \sum_{n=1}^N \sum_{n'=1}^N \text{tr} \left(\sigma_N^n(t, \mathbf{x}, \alpha_\theta^*(t, \mathbf{x})) \sigma_N^{n'}(t, \mathbf{x}, \alpha_\theta^*(t, \mathbf{x}))^\top M_{n', n} \right) \end{aligned}$$

where p^n contains the entries (p_n, \dots, p_{n+d-1}) of the vector p , and $M_{n', n}$ is the submatrix of M containing the rows from n' to $n'+d-1$ and the columns from n to $n+d-1$ inclusive. Finally, for illustration purposes, we will also plot the values of \bar{v}_N^* against Monte Carlo estimates, generated by the learned control α^* , along cross-sections or curves in the computed domain.

5.1 Example 1

Suppose $W = (W_t)_{t \in [0, T]}$ is a real-valued Brownian motion on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $X^{0, \xi, \alpha} = (X_t^{0, \xi, \alpha})_{t \in [0, T]}$ be a real-valued controlled stochastic process, solution to the SDE

$$dX_t = \alpha_t dt + dW_t, \quad X_0 = \xi \sim \mu,$$

where α is a real-valued progressively measurable process valued in $A \subseteq \mathbb{R}$. Let $m_0 \in \mathbb{R}$, $\sigma_0 \in \mathbb{R}_+$ and consider the following value function

$$v(0, \mu) := \inf_{\alpha \in \mathcal{A}} \left\{ \left(\mathbb{E}[X_T^{0, \xi, \alpha}] - m_0 \right)^2 + \left(\text{Var}(X_T^{0, \xi, \alpha}) - \sigma_0^2 \right)^2 \right\}. \quad (5.1)$$

Remark 5.1. Notice that $v(0, \mu) \geq 0$ and a sufficient condition for the existence of an optimal control is that there exists a control $\alpha \in \mathcal{A}$ such that $\mathbb{E}[X_T^{0, \xi, \alpha}] = m_0$ and $\text{Var}(X_T^{0, \xi, \alpha}) = \sigma_0^2$. Consequently, when looking at solving the minimisation problem (5.1), m_0 and σ_0 can be interpreted, respectively, as a target mean and standard deviation for the terminal state. In this context, the cost function plays a role as a penalty function on the distribution of the state at terminal time. See also, for example, the soft-constrained Schrödinger bridge problem in [27].

Let $m_0 = 0$ and $\sigma_0 = 1$. Having $\mathbb{E}[X_T^{0, \xi, \alpha}] = \mathbb{E}[\xi] + \mathbb{E}[\int_0^T \alpha_t dt]$, one easily gets

$$v(0, \mu) \geq \inf_{\alpha \in \mathcal{A}} \left\{ \mathbb{E} \left[\xi + \int_0^T \alpha_t dt \right]^2 \right\} \geq \mathbb{E}[\xi]^2 + 2 \inf_{\alpha \in \mathcal{A}} \left\{ \mathbb{E}[\xi] \mathbb{E} \left[\int_0^T \alpha_t dt \right] \right\}. \quad (5.2)$$

Assuming, for instance, $\mathbb{E}[\xi] \geq 0$ and $A \subseteq [0, +\infty)$ we deduce

$$v(0, \mu) \geq \mathbb{E}[\xi]^2. \quad (5.3)$$

It can be easily checked that, if $T = 1$ and $\xi \sim \delta_{x_0}$, with $x_0 \geq 0$, taking $\hat{\alpha} \equiv 0$, the lower bound $\mathbb{E}[\xi]^2$ is achieved. Indeed, one has $\mathbb{E}[X_T^{0, \xi, \hat{\alpha}}] = x_0$ and $\text{Var}(X_T^{0, \xi, \hat{\alpha}}) = 1$, ending with $J(0, \xi, \hat{\alpha}) = x_0^2 = \mathbb{E}[\xi]^2$ and then $v(0, \delta_{x_0}) = x_0^2$.

We consider hereafter the control constrained into $A = [0, 1]$. From the numerical point of view, the constraint has been introduced with the `sigmoid` as activation function $\sigma_{\text{out}} : \mathbb{R} \rightarrow (0, 1)$ for the last layer of the neural network learning the control α^* , i.e. $\sigma_{\text{out}}(x) = 1/(1 + e^{-x})$. It easy to check that minimizing over




Figure 5.1: (Example 1) Top: empirical convergence of the finite-dimensional value function as the number of particles $N \rightarrow \infty$. Bottom: exact mean field solution $v(0, \delta_{x_0})$ versus generated Monte Carlo estimates, for $x_0 \in \{0.000, 0.125, 0.250, 0.375, 0.500\}$, number of players N increasing.

controls valued on $[0, 1]$ in (5.2) gives the same value as taking $(0, 1)$ -valued controls. It is indeed sufficient to consider the sequence of controls $\alpha_j \equiv \frac{1}{j}$, for $j \in \mathbb{N}_+$, to get

$$\mathbb{E}[X_T^{0, \xi, \alpha_j}] = \mathbb{E}[\xi] + \frac{T}{j} = x_0 + \frac{1}{j} \quad \text{and} \quad \text{Var}(X_T^{0, \xi, \alpha}) = \text{Var}(\xi) + T = 1,$$

then

$$J\left(0, \delta_{x_0}, \alpha_j\right) = \left(x_0 + \frac{1}{j}\right)^2 + (1 - 1)^2 = x_0^2 + \frac{1}{j^2} + 2x_0 \frac{1}{j} \xrightarrow{j \rightarrow +\infty} x_0^2 = \mathbb{E}[\xi]^2,$$

which combined with Equation (5.3), implies $v(0, \delta_{x_0}) = \mathbb{E}[\xi]^2$.

For our experiments, we trained the neural network for $\bar{v}_N(\cdot, 0)$ for points uniformly distributed in $[-1, 1]$. The above benchmark of the method is provided to be $v(0, \delta_{x_0}) = x_0^2$, and we plot the absolute relative error for different values of N within 10 and 10^3 and using 10,000 trajectories for the simulation. The top graph of Figure 5.1 demonstrates empirical convergence of our method through the exact mean field value function, for two different initial points. Moreover, as previously shown in Equation (5.3), $\mathbb{E}[\xi]^2$ represents a lower bound for $J(0, \xi, \alpha)$, for every admissible constrained control α in $[0, 1]$. A graphical representation for $\xi \sim \delta_{x_0}$, i.e. $J(0, \delta_{x_0}, \alpha) \geq x_0^2$, is provided in the bottom graph of Figure 5.1 (unless the Monte Carlo error is involved), along with approximation improvement with increasing N . Finally, in Figure 5.2 we plot the quantiles of the distribution of the trajectory X_t^n for a single player. The plot resembles a Gaussian distribution, for each time step, and this has been verified by testing the assumption with the Pearson normality test (`scipy.stats.normaltest`, returning an average p-value = 0.5753 over time, with a minimum value above




Figure 5.2: (Example 1) Plot of the distribution quantiles over time, for one simulated player of the process X_t^n , starting at $\xi \sim \delta_{0.5}$. The histogram of X_T^n at the terminal time $T = 1$ compared with theoretical density $\mathcal{N}(0.5, 1)$ is plotted simultaneously on the right vertical axis.

0.05, sampling 100,000 trajectories). This is what we expect, taking the optimal control $\hat{\alpha} \equiv 0$, as for all $s \in [0, 1]$, $X_s^{0, x_0, \hat{\alpha}} = x_0 + W_s \sim \mathcal{N}(x_0, s)$.

5.2 Example 2

We consider an optimal liquidation problem with permanent price impact as described in [1]. A group of investors aim to liquidate their position by controlling their trading speed over time. This generates a permanent market impact as the price is assumed to depend on the average trading speed of the investors. By considering a continuum of investors, the inventory process is given by

$$dQ_t = \alpha_t dt, \quad Q_0 = \xi \sim \mu_0 \in \mathcal{P}_2(\mathbb{R}).$$

Here α_t generally takes on negative values as investors are liquidating their portfolio. The wealth process is given by

$$dX_t = -\alpha_t(S_t + k\alpha_t) dt, \quad X_0 = 0,$$

where the parameter k represents the temporary market impact which only affects the individual investor. The price process S follows the dynamics

$$dS_t = \lambda \mathbb{E}[\alpha_t] dt + \sigma dW_t, \quad S_0 = s_0,$$

where the parameter $\lambda \geq 0$ represents the permanent price impact. To ease the notation we will write $(S, Q, X) = (S^{0, s_0, \alpha}, Q^{0, \xi, \alpha}, X^{0, 0, \alpha})$. The cost functional is

$$\mathbb{E} \left[-X_T - Q_T(S_T - \psi Q_T) + \phi \int_0^T Q_t^2 dt \right],$$

where ψ is a penalty parameter for excess inventory at terminal time, and ϕ is an ‘urgency’ parameter for the running cost. By rewriting the cost functional as

$$\mathbb{E} \left[\int_0^T (\alpha_t S_t + k\alpha_t^2 + \phi Q_t^2) dt - Q_T(S_T - \psi Q_T) \right],$$

the problem can then be considered as an extended MFC problem, with a two-dimensional state process (S, Q) and value function

$$v(0, \mu) = \inf_{\alpha \in \mathcal{A}} \mathbb{E} \left[\int_0^T (\alpha_t S_t + k\alpha_t^2 + \phi Q_t^2) dt - Q_T(S_T - \psi Q_T) \right], \quad \mu = \delta_{s_0} \otimes \mu_0.$$




Figure 5.3: (Example 2) Left: The absolute relative error of the expected stock price, comparing the explicit solution with the finite-dimensional approximation. Right: empirical convergence of the finite-dimensional value function as the number of particles $N \rightarrow \infty$.

Note that the diffusion coefficient of (S, Q) in this problem is degenerate. It is shown in [1] that the explicit solution of the optimal control is

$$\hat{\alpha}_t = Q_0 r \frac{d_1 e^{-r(T-t)} - d_2 e^{r(T-t)}}{d_1 e^{-rT} + d_2 e^{rT}} + \mathbb{E}[Q_0] \frac{2\lambda\phi(e^{-rt} + e^{rt})}{(d_1 e^{-rT} + d_2 e^{rT})(c_1 e^{-rT} + c_2 e^{rT})}, \quad (5.4)$$

with constants $d_1 = \sqrt{\phi k} - \psi$, $d_2 = \sqrt{\phi k} + \psi$, $c_1 = 2d_1 + \lambda$, $c_2 = 2d_2 - \lambda$ and $r = \sqrt{\phi/k}$. In particular, when the initial condition Q_0 is non-deterministic, the optimal control $\hat{\alpha}_t$ is not in a closed loop form. As outlined in Section 3.2, we shall regularise the problem by considering the system

$$dS_t = \lambda \mathbb{E}[\alpha_t] dt + \sigma dW_t, \quad dQ_t^\varepsilon = \alpha_t dt + \varepsilon dB_t,$$

where $\varepsilon > 0$ is small and B is independent of W , so that we work towards an approximation of $\bar{v}_{\varepsilon, N}$.

For our experiments, we consider an initial condition of $S_0 = 5$ and $Q_0 \sim \mathcal{N}(10, 1)$. We also take the following parameters: $T = 1$, $k = 0.2$, $\lambda = 0.4$, $\phi = 0.1$, $\psi = 2.5$, $\sigma = 1$, and $\varepsilon = 0.01$ for the regularisation parameter. In order to benchmark our method, we utilise the optimal control $\hat{\alpha}_t$, explicitly given in (5.4), to generate Monte Carlo estimates for the corresponding state process (\hat{S}, \hat{Q}) and the value function v . We first consider the trajectory of the expected stock price $t \mapsto \mathbb{E}[\hat{S}_t]$, and compare it with the trajectory of the empirical mean of the finite dimensional approximation $t \mapsto \frac{1}{N} \sum_{n=1}^N \mathbb{E}[S_t^n]$. We plot the absolute relative error of the two trajectories on the left graph of Figure 5.3 for different values of n , using 500,000 trajectories for the simulation to minimise the Monte Carlo error. On the right graph, we illustrate an instance of the decay of the relative error of the approximated value function against the mean field value function as the value of n increases, for a fixed value of $\mu = \mathcal{N}(10, 1)$. By considering the regularised problem, we obtain a feedback control, parametrised in the form of a neural network, that approximates an optimal control of general open loop form.

5.3 Example 3

Suppose $W = (W_s)_{s \in [0, T]}$ is a real-valued Brownian motion on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $X = (X_s)_{s \in [t, T]}$ satisfy the SDE

$$dX_s = \alpha_s ds + \sqrt{2} dW_s, \quad \mathcal{L}(X_t) = \mu,$$

where α is a real-valued progressively measurable process. Given a terminal cost function $g : \mathbb{R} \rightarrow \mathbb{R}$, consider the value function

$$v(t, \mu) = \inf_{\alpha \in \mathcal{A}} \left(\int_t^T \mathbb{E}^\mu \left[\frac{|\alpha_s|^2}{2} \right] ds + \mathbb{E}^\mu [g(X_T)] \right),$$

where \mathbb{E}^μ denotes the expectation under which $\mathcal{L}(X_t) = \mu$. It can be shown that v satisfies the following master Bellman PDE (see [12, Section 2.4]), on $[0, T] \times \mathcal{P}_2(\mathbb{R})$,

$$-\partial_t v(t, \mu) - \int_{\mathbb{R}^d} \left(\operatorname{div}_x \partial_\mu v(t, \mu)(x) - \frac{1}{2} |\partial_\mu v(t, \mu)(x)|^2 \right) \mu(dx) = 0, \quad (5.5)$$

with terminal condition $v(T, \mu) = \int_{\mathbb{R}} g(x) \mu(dx)$. Such a PDE has a unique, classical solution given by

$$v(t, \mu) = \int_{\mathbb{R}} w(t, x) \mu(dx),$$

where $w : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ solves the following one-dimensional HJB equation:

$$\begin{cases} -\partial_t w(t, x) - \Delta w(t, x) + \frac{1}{2} |\nabla w(t, x)|^2 = 0, & (t, x) \in [0, T] \times \mathbb{R}, \\ w(T, x) = g(x), & x \in \mathbb{R}, \end{cases}$$

which has explicit solution

$$w(t, x) = -2 \log \left(\frac{1}{\sqrt{4\pi(T-t)}} \int_{\mathbb{R}} \exp \left(-\frac{(x-y)^2}{4(T-t)} - \frac{g(y)}{2} \right) dy \right).$$

In this context, the argument of the limit in Theorem 3.5 is constant with respect to the number of players N , and one has the equality for any $N \in \mathbb{N}_+$:

$$v(t, \mu) = \int_{\mathbb{R}^N} \frac{1}{N} \sum_{n=1}^N w(t, x_n) \mu(dx_1) \otimes \dots \otimes \mu(dx_N).$$

In this specific case, due to the integral form of the exact solution of the problem, we can test the validity of our numerical approximation by simply examining the accuracy of the finite-dimensional solver. In fact, it is straightforward to verify that

$$\frac{1}{N} \sum_{n=1}^N w(t, x_n) = v_N(t, \mathbf{x}), \quad \mathbf{x} = (x_1, \dots, x_N) \in \mathbb{R}^N,$$

where $v_N : [0, T] \times \mathbb{R}^N \rightarrow \mathbb{R}$ is the finite-dimensional approximation of (5.5), introduced in Section 3, which is a solution the following HJB equation on $[0, T] \times \mathbb{R}^N$,

$$-\partial_t u(t, \mathbf{x}) - \sum_{n=1}^N \Delta_{x_n} u(t, \mathbf{x}) + \inf_{a \in \mathbb{R}^N} \left\{ \sum_{n=1}^N \left(a_n \cdot \nabla_{x_n} u(t, \mathbf{x}) + \frac{|a_n|^2}{2N} \right) \right\} = 0. \quad (5.6)$$

with terminal condition $u(T, \mathbf{x}) = \bar{g}_N(\mathbf{x}) := \frac{1}{N} \sum_{n=1}^N g(x_n)$. Eq. (5.6) simplifies to

$$-\partial_t u - \Delta u + \frac{N}{2} |\nabla u|^2 = 0, \quad \text{on } [0, T] \times \mathbb{R}^N.$$

In our experiments, we take the terminal condition as the function $g(x) = |x|^2/2$. In order to satisfy the boundedness assumption of g in Assumption 2.4, we truncate g as necessary outside some fixed domain. This aligns with the practical experiments, since the computational step is always confined within a bounded grid.

Figure 5.4 shows our experimental results when applied to a finite-dimensional approximation of 100 players. We compare the trained neural network $\bar{v}_{N,\eta}^*$ with the explicit one-dimensional solution, as well as Monte Carlo estimates along two parameterised curves in the domain. We report a residual loss of $\mathcal{L}_{res} = 2.45e - 3$, averaged over 1000 random samples of quantized measures, as well as an HJB loss of $\mathcal{L}_{HJB} = 3.82e - 4$. The HJB loss demonstrates a good fit of the neural network towards the finite-dimensional value function, and the residual loss indicates a good approximation of the mean field value function v on the domain $\mathcal{P}_2(\mathbb{R})$. In terms of computational performance, the training time of the neural network $\bar{v}_{N,\eta}^*$ was approximately 7 hours. After training, an average of 0.66 seconds is required for the evaluation of $v(t, \mu)$ for a given measure μ . In comparison, generating an Monte Carlo estimate of the value function ‘online’ directly with the control α_θ^* requires approximately 2.67 seconds per evaluation. Therefore, for applications requiring knowledge of the value function across the whole domain $\mathcal{P}_2(\mathbb{R}^d)$, for example the computation of level set $\{\mu \in \mathcal{P}_2(\mathbb{R}) : v(t, \mu) = 0\}$, significant savings in computational time can be achieved.

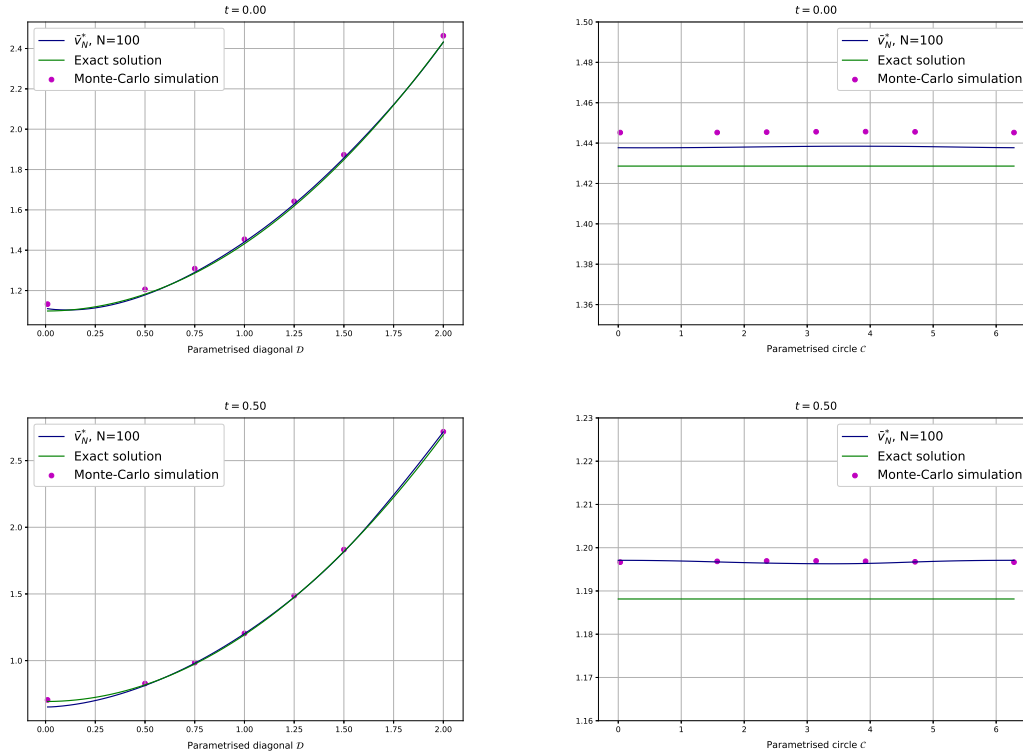


Figure 5.4: (Example 3) Left column: Tracing $\bar{v}_{N,\eta}^*$ along the diagonal of the hypercube $[0, 2]^{100}$. Right column: Tracing $\bar{v}_{N,\eta}^*$ along a circle on the first two dimensions.

Acknowledgments

The authors acknowledge financial support under the National Recovery and Resilience Plan (NRRP): Probabilistic Methods for Energy Transition – CUP G53D23006840001 - Grant Assignment Decree No. 1379 adopted on 01/09/2023 by MUR. Jonathan Tam is also supported by the G-Research research grant.

References

- [1] B. Acciaio, J. Backhoff-Veraguas, and R. Carmona. Extended mean field control problems: stochastic maximum principle and transport perspective. *SIAM J. Control. Optim.*, 57(6):3666–3693, 2019.
- [2] R. Aïd, R. Dumitrescu, and P. Tankov. The entry and exit game in the electricity markets: A mean-field game approach. *J. Dyn. Games*, 8(4):331, 2021.
- [3] P. Cardaliaguet and C.-A. Lehalle. Mean field game of controls and an application to trade crowding. *Math. Financ. Econ.*, 12(3):335–363, 2018.
- [4] P. Cardaliaguet and P. E. Souganidis. Regularity of the value function and quantitative propagation of chaos for mean field control problems. *NoDEA Nonlinear Differential Equations Appl.*, 30(2), 2023.
- [5] R. Carmona, F. Delarue, et al. *Probabilistic theory of mean field games with applications I-II*. Springer, 2018.
- [6] R. Carmona and M. Laurière. Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: II—the finite horizon case. *Ann. Appl. Probab.*, 32(6):4065–4105, 2022.
- [7] R. Carmona and M. Laurière. Deep learning for mean field games and mean field control with applications to finance. *Machine Learning and Data Sciences for Financial Markets: A Guide to Contemporary Practices*, page 369, 2023.

- [8] A. Cecchin, S. Daudin, J. Jackson, and M. Martini. Quantitative convergence for mean field control with common noise and degenerate idiosyncratic noise. *Electron. J. Probab.*, 30:1–62, 2025.
- [9] A. Cosso, F. Gozzi, I. Kharroubi, H. Pham, and M. Rosestolato. Master Bellman equation in the Wasserstein space: Uniqueness of viscosity solutions. *Trans. Amer. Math. Soc.*, 2023.
- [10] A. Cosso, F. Gozzi, I. Kharroubi, H. Pham, and M. Rosestolato. Optimal control of path-dependent McKean-Vlasov SDEs in infinite dimension. *Ann. Appl. Probab.*, 33(4):2863–2918, 2023.
- [11] A. Cosso and H. Pham. Zero-sum stochastic differential games of generalized McKean–Vlasov type. *J. Math. Pures Appl.*, 129:180–212, 2019.
- [12] S. Daudin and B. Seeger. A comparison principle for semilinear Hamilton–Jacobi–Bellman equations in the Wasserstein space. *Calc. Var. Partial Differential Equations*, 63(106), 2024.
- [13] G. Dayanikli, M. Laurière, and J. Zhang. Deep learning for population-dependent controls in mean field control problems with common noise. In *Proc. 23rd Int. Conf. AAMAS*, pages 2231–2233, 2024.
- [14] V. De Bortoli, J. Thornton, J. Heng, and A. Doucet. Diffusion Schrödinger bridge with applications to score-based generative modeling. In *NeurIPS*, volume 34, pages 17695–17709, 2021.
- [15] M. F. Djete. Extended mean field control problem: a propagation of chaos result. *Electron. J. Probab.*, 27:1–53, 2022.
- [16] G. Fu, U. Horst, and X. Xia. A mean-field control problem of optimal portfolio liquidation with semimartingale strategies. *Math. Oper. Res.*, 49(4):2356–2384, 2024.
- [17] W. Gangbo, S. Mayorga, and A. Świąch. Finite dimensional approximations of Hamilton–Jacobi–Bellman equations in spaces of probability measures. *SIAM J. Math. Anal.*, 53(2):1320–1356, Jan. 2021.
- [18] M. Germain, M. Laurière, H. Pham, and X. Warin. DeepSets and their derivative networks for solving symmetric PDEs. *J. Sci. Comput.*, 91(2):63, 2022.
- [19] M. Germain, H. Pham, and X. Warin. Rate of convergence for particle approximation of PDEs in Wasserstein space. *J. Appl. Probab.*, 59(4):992–1008, 2022.
- [20] P. Glasserman. *Monte Carlo methods in financial engineering*, volume 53. Springer, 2004.
- [21] J. Han, A. Jentzen, et al. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Commun. Math. Stat.*, 5(4):349–380, 2017.
- [22] X. Huang, S. Jaimungal, and M. Nourian. Mean-field game strategies for optimal execution. *Appl. Math. Finance*, 26(2):153–185, 2019.
- [23] C. Huré, H. Pham, and X. Warin. Deep backward schemes for high-dimensional nonlinear PDEs. *Math. Comp.*, 89(324):1547–1579, 2020.
- [24] N. V. Krylov. *Controlled Diffusion Processes*. Springer, 1980.
- [25] D. Lacker. Limit theory for controlled McKean–Vlasov dynamics. *SIAM J. Control Optim.*, 55(3):1641–1672, 2017.
- [26] W. Lefebvre, G. Loeper, and H. Pham. Differential learning methods for solving fully nonlinear pdes. *Digital Finance*, 5(1):183–229, 2023.
- [27] J. Ma, Y. Tan, and R. Xu. Schrödinger bridge for generative AI: Soft-constrained formulation and convergence analysis. *arXiv:2510.11829*, 2025.
- [28] H. Pham and X. Warin. Mean-field neural networks: learning mappings on Wasserstein space. *Neural Networks*, 168:380–393, 2023.
- [29] H. Pham and X. Warin. Mean-field neural networks-based algorithms for McKean-Vlasov control problems. *J. Mach. Learn. Res.*, 3(2):176–214, 2024.

- [30] H. M. Soner, J. Teichmann, and Q. Yan. Learning algorithms for mean field optimal control. *arXiv:2503.17869*, 2025.
- [31] M. Talbi. A finite-dimensional approximation for partial differential equations on Wasserstein space. *Stochastic Process. Appl.*, 177:104445, Nov. 2024.
- [32] F. Vargas, P. Thodoroff, A. Lamacraft, and N. Lawrence. Solving Schrödinger bridges via maximum likelihood. *Entropy*, 23(9):1134, 2021.
- [33] G. Wang, Y. Jiao, Q. Xu, Y. Wang, and C. Yang. Deep generative learning via Schrödinger bridge. In *ICML*, pages 10794–10804. PMLR, 2021.