

Numerical resolution of McKean-Vlasov FBSDEs using neural networks ^{*}

Maximilien GERMAIN [†], Joseph MIKAEL [‡], Xavier WARIN [§]

February 21, 2022

to appear in Methodology and Computing in Applied Probability

Abstract

We propose several algorithms to solve McKean-Vlasov Forward Backward Stochastic Differential Equations (FBSDEs). Our schemes rely on the approximating power of neural networks to estimate the solution or its gradient through minimization problems. As a consequence, we obtain methods able to tackle both mean-field games and mean-field control problems in moderate dimension. We analyze the numerical behavior of our algorithms on several examples including non linear quadratic models.

Key words: Neural networks, McKean-Vlasov FBSDEs, Deep BSDE, mean-field games, machine learning.

MSC Classification: 65C30, 68T07, 49N80, 35Q89.

1 Introduction

This paper is dedicated to the numerical resolution in moderate dimension of the following McKean-Vlasov Forward Backward Stochastic Differential Equations (MKV FBSDEs)

$$\begin{cases} X_t &= \xi + \int_0^t b(s, X_s, Y_s, Z_s, \mathcal{L}(X_s), \mathcal{L}(Y_s), \mathcal{L}(Z_s)) ds + \int_0^t \sigma(s, X_s, \mathcal{L}(X_s)) dW_s \\ Y_t &= g(X_T, \mathcal{L}(X_T)) + \int_t^T f(s, X_s, Y_s, Z_s, \mathcal{L}(X_s), \mathcal{L}(Y_s), \mathcal{L}(Z_s)) ds - \int_t^T Z_s dW_s \end{cases} \quad (1)$$

with $b : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^k \times \mathbb{R}^{k \times d} \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^k) \times \mathcal{P}_2(\mathbb{R}^{k \times d}) \mapsto \mathbb{R}^d$, $\sigma : \mathbb{R} \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}^d$, $g : \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}^k$, and $f : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^k \times \mathbb{R}^{k \times d} \times \mathcal{P}_2(\mathbb{R}^d) \times \mathcal{P}_2(\mathbb{R}^k) \times \mathcal{P}_2(\mathbb{R}^{k \times d}) \mapsto \mathbb{R}^k$.

W_t is a d -dimensional \mathcal{F}_t -Brownian motion where $(\Omega, \mathcal{A}, \mathcal{F}_t, \mathbb{P})$ is a given filtered probability space and $T > 0$.

ξ is a given random variable in $L^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^d)$ and $\mathcal{P}_2(\mathbb{R}^n)$ stands for the space of square integrable probability measures over \mathbb{R}^n endowed with the 2-Wasserstein distance

$$\mathcal{W}_2(\mu, \nu) = \inf \left\{ \sqrt{\mathbb{E}[(X - X')^2]} \mid X, X' \in L^2(\Omega, \mathcal{F}, \mathbb{P}; \mathbb{R}^n), \mathcal{L}(X) = \mu, \mathcal{L}(X') = \nu \right\}.$$

At last $\mathcal{L}(\cdot)$ is a generic notation for the law of a random variable.

^{*}This work is supported by FiME, Laboratoire de Finance des Marchés de l'Energie.

[†]EDF R&D, Université de Paris, LPSM mgermain@lpsm.fr

[‡]EDF R&D joseph.mikael@edf.fr

[§]EDF R&D & FiME xavier.warin@edf.fr

This kind of equation is linked to non local PDEs known as master equations. We refer to [CD18], chapter 4 and 5 of volume 2, for an introduction on the subject. In [CCD15], it is shown for example that under regularity conditions, when the drift b is independent of time and the law $\mathcal{L}(Z_t)$, the driver f does not depend on time and the law $\mathcal{L}(Z_t)$, σ does not depend on time, then the resolution of equation (1) provides a way to estimate the solution of the equation

$$\begin{aligned} & \partial_t \mathcal{U}(t, x, \mu) + b(x, \mathcal{U}(t, x, \mu), \partial_x \mathcal{U}(t, x, \mu) \sigma(x, \mu), \mu, \eta) \cdot \partial_x \mathcal{U}(t, x, \mu) \\ & + \frac{1}{2} \text{Tr}[\partial_{xx}^2 \mathcal{U}(t, x, \mu) \sigma \sigma^\top(x, \mu)] + f(x, \mathcal{U}(t, x, \mu), \partial_x \mathcal{U}(t, x, \mu) \sigma(x, \mu), \mu, \eta) \\ & + \int_{\mathbb{R}^d} \partial_\mu \mathcal{U}(t, x, \mu)(y) \cdot b(y, \mathcal{U}(t, y, \mu), \partial_x \mathcal{U}(t, x, \mu) \sigma(x, \mu), \mu, \eta) \, d\mu(y) \\ & + \int_{\mathbb{R}^d} \frac{1}{2} \text{Tr}[\partial_x \partial_\mu \mathcal{U}(t, x, \mu)(y) \sigma \sigma^\top(y, \mu)] \, d\mu(y) = 0 \text{ on } [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \quad (2) \end{aligned}$$

with initial condition $\mathcal{U}(0, x, \mu) = g(x, \mu)$ on $\mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d)$, and where η is a notation for the image of the probability measure μ by the mapping $x \in \mathbb{R}^d \rightarrow \mathcal{U}(t, x, \mu) \in \mathbb{R}^k$. Under suitable assumptions, [CCD15] proves that the solution to (1) admits the so-called decoupling field representation $Y_s = U(s, X_s, \mathcal{L}(X_s))$ where $(t, x, \mu) \in [0, T] \times \mathbb{R}^d \times \mathcal{P}_2(\mathbb{R}^d) \mapsto U(T - t, x, \mu)$ is a classical solution to (2). See Theorem 2.9 and equation (2.12) in [CCD15]. Their results are stated in the more general of dynamics depending in the joint law of (X_t, Y_t) but we state them here with the particular case of dependence in the marginal laws $\mathcal{L}(X_t), \mathcal{L}(Y_t)$ to be coherent with (1). This MKV FBSDE representation is used in [CCD19] to build a numerical scheme for the approximation of (2) when the drift b is independent of Z . The equation above is non local due to the integral terms and the term $\partial_\mu \mathcal{U}(t, x, \mu)(y)$ stands for the Wasserstein derivative of \mathcal{U} in the direction of the measure at point (t, x, μ) and evaluated at the continuous coordinate y .

Equations (1) appear as well as probabilistic formulations of mean-field games or mean-field controls characterizing the value function V of the game. Mean-field games are introduced by [LL06a] and [LL06b] to model games with interactions between many similar players. In this theory, each player's dynamics and cost take into account the empirical distribution of all agents. At the limit of an infinite number of players, the search for a Nash equilibrium with close loop controls boils down to a control problem concerning a representative player whose law enters in the cost and dynamics. Two probabilistic approaches based on Forward Backward Stochastic Differential Equations can be used to solve these problems:

- A first approach called the **Pontryagin approach** consists as shown in [CD13] in applying the strong Pontryagin principle to these control problems. Under regularity and convexity conditions, Y_t appears to be a stochastic representation of the gradient of the value function V . In this case the coefficients b, f of the related MKV FBSDE (1) do not depend on $Z_t, \mathcal{L}(Y_t), \mathcal{L}(Z_t)$ and $k = d$.
- Another approach called the **Weak approach** permits to solve the optimization problem by estimating directly Y_t as the value function V of the problem as shown in [CL15]. In this case the coefficients b, f of the related MKV FBSDE (1) do not depend on $Y_t, \mathcal{L}(Y_t), \mathcal{L}(Z_t)$ and $k = 1$.

The numerical resolution of equations (1) is rather difficult since:

- The dynamics are coupled through both the drift and the driver of the BSDE.
- The McKean-Vlasov structure of the problem requires to solve a fixed point in probability spaces.

In the linear-quadratic setting (with quadratic cost to minimize but linear dynamics) the weak approach applied to mean-field problems gives a problem in low dimension but with quadratic coupling in Z_t appearing in the backward dynamic. In contrast, the Pontryagin approach exhibits a problem in potentially high dimension (the Z component is a $d \times d$ matrix in this case) but with a linear coupling in Y_t which is easier to solve numerically.

In the case of mean-field games, only the law of X_t is present in the dynamic of (1). In mean-field games of controls (also called extended mean-field games, see [AK20] and the references therein), individuals interact through their controls instead of their states as in the model of trade crowding in [CL18]. The law of the control thus appears in the dynamic of (1) and may give rise to some FBSDE depending on the law of Z_t in the weak approach or the law of Y_t in the Pontryagin approach.

Existence and uniqueness of a solution to the fully coupled system (1) are studied by [CD18] when the drift does not depend on the law of Z . Their Theorem 4.29 gives existence of a solution under a non-degeneracy condition. However, uniqueness is a priori only expected to hold in small time, as stated in Theorem 4.24 from [CD18]. In the latter we will assume that existence and uniqueness hold for the MKV FBSDE we aim to solve.

In [CCD19] and [Ang+19], tree and grid algorithms are proposed and tested in dimension 1. It is worth mentioning that these techniques suffer from the so-called curse of dimensionality and cannot be applied when the dimension describing a player state is high (typically greater than 3 or 4). This is due to the discretization of the state space.

However, new approaches using machine learning are developed since 2017 to solve non linear parabolic PDEs through a BSDE representation. Two kinds of methods have emerged:

- The first to appear are **global methods** first proposed in [HJE17] to solve semi-linear PDEs. They rely on a single high-dimensional optimization problem whose resolution is difficult. It consists in the training of as many neural networks as time steps by solving in a forward way the backward representation of the PDE solution. The Z_t process is represented by a different neural network Z_i^θ with parameters θ at each date t_i . Instead of solving the BSDE starting from the terminal condition, the method writes it down as a forward equation and an optimization problem aiming to reach the terminal condition $g(X_T)$ by minimizing a mean squared error $\mathbb{E}|Y_T - g(X_T)|^2$. The approach is extended to fully nonlinear equations (nonlinear in the solution, its gradient and hessian) in [BEJ19] and the authors show that the methodology can solve some equations in high dimension. [CWNMW19] showed that it is more effective to use a single network for all dates and besides proposed an original fixed point algorithm to solve semilinear PDEs.
- A second kind of **local methods** first proposed in [HPW20] is based on local optimization problems solved at each time step in a backward way. Contrarily to the global method, the successive optimization problems are here in moderate dimension. Each optimization step at date t_i consists in the training of only two local neural networks Y_i^θ, Z_i^θ with parameters θ . For instance, instead of solving 1 optimization problem with N neural networks in the global method, [HPW20] solves N learning problems with 2 neural networks. Moreover the resolution is simplified by the initialization of the neural networks at time t_i to their previously computed values at time t_{i+1} , which provides a good approximation for the current value. This strategy is inspired by the standard backward resolution of BSDE with conditional expectations from [BT04] and [GLW05]. The methodology is extended to the much more challenging case of fully nonlinear PDEs in [PWG21] by combining it with some ideas proposed in [Bec+19]. Extensive tests performed in [HPW20] show that the local method gives better results than the global

one, such as [PWG21] in the case of fully nonlinear dynamics. Especially, these papers show that local methods can be used with a larger time horizon T than the global method.

Machine learning techniques to solve coupled FBSDEs are investigated by several authors in [HL20] and [Ji+20], and a first method for McKean-Vlasov FBSDEs with delay is studied by [FZ20] for a linear quadratic equation in dimension one. Similar and more general ideas are presented and tested in dimension one in [CL19] alongside convergence results, together with an additional method directly solving mean-field control problems by minimizing the cost without writing down optimality conditions. The resulting algorithms proposed all rely on the global approach first initiated in [HJE17].

Our paper aims to extend these methods and to propose new ones for the resolution of McKean-Vlasov FBSDEs in moderate dimension, and go beyond one dimensional examples for which standard methods are already available (see [ACD10; Lau21]). In fact, one major advantage for the use of neural networks for solving control problems is their ability to efficiently represent high-dimensional functions without using space grids. We also study the influence of the maturity T on the algorithms. We first propose to modify the previously proposed algorithm to stabilize its convergence. Our modification allows us to reduce the variance of the estimators used in the dynamic of X_t and Y_t . Then we propose a second algorithm relaxing the fixed point iteration algorithm by adding a neural network learning the distribution of the solution thanks to a penalization in the loss function. At last we propose a resolution scheme based on some local resolution as in [HPW20].

To simplify the presentation, we consider first order interaction models, that is to say that the dependency of the drift and cost function with respect to the laws $\mathcal{L}(X_t), \mathcal{L}(Y_t), \mathcal{L}(Z_t)$ only concerns expectations in the form

$$u_t = (u_t^X, u_t^Y, u_t^Z) := (\mathbb{E}[\varphi_1(X_t)], \mathbb{E}[\varphi_2(Y_t)], \mathbb{E}[\varphi_3(Z_t)]), \quad (3)$$

for some continuous functions $\varphi_1, \varphi_2, \varphi_3$ with adequate domains and codomains. In this framework we can rewrite by abuse of notation

$$\begin{aligned} b(s, X_s, Y_s, Z_s, \mathcal{L}(X_s), \mathcal{L}(Y_s), \mathcal{L}(Z_s)) &= b(s, X_s, Y_s, Z_s, u_s) \\ \sigma(s, X_s, \mathcal{L}(X_s)) &= \sigma(s, X_s, u_s^X) \\ g(X_T, \mathcal{L}(X_T)) &= g(X_T, u_T^X) \\ f(s, X_s, Y_s, Z_s, \mathcal{L}(X_s), \mathcal{L}(Y_s), \mathcal{L}(Z_s)) &= f(s, X_s, Y_s, Z_s, u_s). \end{aligned}$$

For instance when $\varphi_1, \varphi_2, \varphi_3$ are power functions with positive integers as exponents we recover probability distribution moments. See Remark 1 for more general cases beyond first order interaction.

We provide multidimensional tests to show how these machine learning approaches can overcome the curse of dimensionality on some test cases first coming from a mean-field game of controls: we solve the FBSDE derived from the weak approach and the Pontryagin approach. We also consider an example arising from a non linear quadratic mean-field game. Then we compare all the methods on some general test cases of FBSDE involving linear or quadratic dependence on the processes X_t, Y_t, Z_t and on their distributions.

The structure of the paper is the following: in sections 2 and 3 we describe the proposed schemes, and in section 4 we provide a numerical study of our methods in dimension 10 (except for the one-dimensional Example of Section 4.2). We show that our algorithms can solve non linear-quadratic models with small maturities.

2 Machine learning global solvers

In this section we propose three **global algorithms** based on the approach in [HJE17].

2.1 Algorithm principle

We propose a generalized and refined version of the Algorithm 2 from [CL19]. We recall that a similar technique with additional networks is used in [FZ20] for delayed McKean-Vlasov equations but is tested only on a one dimensional linear quadratic example. Our methods also take advantage of different expectation computation methods, introduced in section 2.2. We present in section 4 several tests in dimension 10 where the laws of X, Y, Z are involved.

We consider the Euler-Maruyama discretized FBSDE system (1) on a regular time grid $t_k = \frac{kT}{N}$ for $k \in \llbracket 0, N \rrbracket$:

$$\begin{cases} X_{t_{i+1}} &= X_{t_i} + b(t_i, X_{t_i}, Y_{t_i}, Z_{t_i}, u_{t_i}) \Delta t + \sigma(t_i, X_{t_i}, u_{t_i}^X) \Delta W_i \\ Y_{t_{i+1}} &= Y_{t_i} - f(t_i, X_{t_i}, Y_{t_i}, Z_{t_i}, u_{t_i}) \Delta t + Z_{t_i} \Delta W_i, \end{cases} \quad (4)$$

with terminal condition $Y_{t_N} = g(X_{t_N}, u_{t_N}^X)$ and initial condition $X_0 = \xi$. We recall that u_t is defined in (3). We note $\Delta t := t_{i+1} - t_i = \frac{T}{N}$ and $(\Delta W_i)_{i=0, \dots, N-1} := (W_{t_{i+1}} - W_{t_i})_{i=0, \dots, N-1}$ the Brownian increments. In the FBSDE theory, one requires the processes $(X_{t_i}, Y_{t_i}, Z_{t_i})$ to be \mathcal{F}_{t_i} -adapted. Therefore the backward part of the system can also be written in the conditional expectation form

$$\begin{cases} Y_{t_i} = \mathbb{E}[Y_{t_{i+1}} + f(t_i, X_{t_i}, Y_{t_i}, Z_{t_i}, u_{t_i}) \Delta t | \mathcal{F}_{t_i}] \\ Z_{t_i} = \mathbb{E}[Y_{t_{i+1}} \frac{\Delta W_i}{\Delta t} | \mathcal{F}_{t_i}], \end{cases} \quad (5)$$

where we see how the process Z is defined. This process, specific to the stochastic case, allows the Y component to be \mathcal{F}_t -adapted, even though we fix its terminal condition. It is a major difference between backward ordinary differential equations and backward stochastic differential equations. We see that the whole system is coupled therefore we need to design a method allowing to solve simultaneously both equations of (4).

We solve the system by the Merged Deep BSDE method introduced in [CWNMW19]. Z_{t_i} is approximated by a single feedforward neural network $\mathcal{Z}_{\theta^z}(t_i, X_{t_i})$ and Y_0 by a neural network $\mathcal{Y}_{\theta^y}(X_0)$ with parameters $\theta = (\theta^y, \theta^z)$. With this point of view, the discretized Brownian motion W_t acts as training data in the language of machine learning, so that we can generate a training set as large as desired. Extensive tests conducted in [CWNMW19] show that the use of a Merged network improves the training in comparison with the Deep BSDE method of [HJE17]. Indeed it lowers the number of parameters to learn hence reduces the complexity of the problem. It empirically improves the accuracy of the method but also makes the training faster. That is why we focus on this architecture. It is also used by [CL19] which considers controls in the form $(\zeta(t_i, X_i))_{i=1, \dots, N}$ for a neural network ζ . The use of recurrent networks such as Long Short Term Memory networks as in [FZ20] is possible but tests achieved in [CWNMW19] seem to show that it does not bring more accuracy on Markovian problems. Other alternatives may include the GroupSort network [ALG19] for a better control of the Lipschitz constant of the approximation, or some special networks preserving some properties of the solution but they have not been tested.

The motivation for such an approximation comes from the notion of decoupling field, also used for numerical purposes in [Ang+19] or [CCD19], which gives the existence of functions u, v (see the paragraph below (2)) such that

$$Y_t = u(t, X_t, \mathcal{L}(X_t)), \quad Z_t = v(t, X_t, \mathcal{L}(X_t)). \quad (6)$$

Numerically, it is enough to consider Y and Z as a function of the couple (t, X_t) . In fact, the law of the solution (and therefore its moments) can be seen as a function of t . That's why we search for a representation

$$Y_t = \tilde{u}(t, X_t), \quad Z_t = \tilde{v}(t, X_t). \quad (7)$$

The forward-backward system is transformed into a forward system and an optimization problem aiming to satisfy the terminal condition of the BSDE through the loss function $\mathbb{E}[(Y_T - g(X_T, u_{t_N}^X))^2]$. To simplify notations, $X_i := X_{t_i}$ and similarly for Y and Z .

In practice the loss function is minimized with the Adam gradient descent method [KB14]. In any case, the goal of our scheme is to learn both the optimal control and the distribution of X_t, Y_t, Z_t . In the following, B is the batch size, that is the number of particles we will sample to approximate the loss function in expectation form, see for instance (10). N is the number of time steps and M is the number of previous batches expectations to keep in memory.

We use for \mathcal{Z} and \mathcal{Y} feedforward neural networks with 3 hidden layers ($d+10$ neurons in each) with hyperbolic tangent function as activation functions and an output layer with identity as activation. It is worth noticing that because the merged neural network takes the couple (t, X) as inputs, we cannot use batch normalization since the distribution of X_i is not stationary over time.

2.2 Estimation of the expectation

A key step for the methods is to estimate the mean-field parameter u . It has a significant effect on the algorithms performances. We note $\theta_m = (\theta_m^y, \theta_m^z)$ the neural network parameters at optimization iteration m and $u_i = (u_i^X, u_i^Y, u_i^Z)$ the estimation of u_{t_i} . In the algorithms described below, the approximated processes are considered as functions of the parameters θ of the neural network. Several methods can be used to approximate the moments of the solution involved in the stochastic McKean-Vlasov dynamics:

- **Direct:** use the empirical mean of the current batch of particles

$$u_i = \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_m)), \sum_{j=1}^B \varphi_3(Z_i^j(\theta_m)) \right), \quad i = 0, \dots, N-1. \quad (8)$$

Alternatively one could use instead the last batch particles to estimate the law

$$u_i = \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_{m-1})), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_{m-1})), \sum_{j=1}^B \varphi_3(Z_i^j(\theta_{m-1})) \right), \quad i = 0, \dots, N-1. \quad (9)$$

The difference lies in the fact that in one case the optimization of the parameters θ_m at iteration m modifies the current estimation of the law whereas using the previously computed parameters θ_{m-1} fixes the law and simplifies the optimization problem. In practice, for the numerical tests of Section 4 we use the formula (8). This approach requires to handle very large batches, typically of the order of $B = 10,000$ sample paths get a reasonable approximation of the laws. This is the approach used by [FZ20] and [CL19].

We solve the following optimization problem

$$\begin{aligned}
& \min_{\theta=(\theta^y, \theta^z)} \frac{1}{B} \sum_{k=1}^B \left| Y_N^k(\theta) - g\left(X_N^k(\theta), \frac{1}{B} \sum_{j=1}^B \varphi_1(X_N^j(\theta))\right) \right|^2 \tag{10} \\
& X_{i+1}^j(\theta) = X_i^j(\theta) + b\left(t_i, X_i^j(\theta), Y_i^j(\theta), \mathcal{Z}_{\theta^z}\left(t_i, X_i^j(\theta)\right), u_i\right) \Delta t + \sigma\left(t_i, X_i^j(\theta), u_i^X\right) \Delta W_i^j \\
& Y_{i+1}^j(\theta) = Y_i^j(\theta) - f\left(t_i, X_i^j(\theta), Y_i^j(\theta), \mathcal{Z}_{\theta^z}\left(t_i, X_i^j(\theta)\right), u_i\right) \Delta t + \mathcal{Z}_{\theta^z}\left(t_i, X_i^j(\theta)\right) \Delta W_i^j \\
& u_i = \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta)), \sum_{j=1}^B \varphi_3\left(\mathcal{Z}_{\theta^z}\left(t_i, X_i^j(\theta)\right)\right) \right) \\
& X_0^j = \xi^j \sim \xi, \quad j = 1, \dots, B \\
& Y_0^j(\theta) = \mathcal{Y}_{\theta^y}(X_0^j), \\
& i = 1, \dots, N-1.
\end{aligned}$$

The **Global Direct solver** leads to algorithm 1.

Algorithm 1 Global Direct solver

- 1: Let $\mathcal{Y}_{\theta^y}(\cdot)$ be a neural network with parameter θ^y , defined on \mathbb{R}^d and valued in \mathbb{R}^k , $\mathcal{Z}_{\theta^z}(\cdot, \cdot)$ be a neural network with parameter θ^z , defined on $\mathbb{R}^+ \times \mathbb{R}^d$ and valued in $\mathbb{R}^{k \times d}$, so that $\theta = (\theta^y, \theta^z)$ is initialized with value $\theta_0 = (\theta_0^y, \theta_0^z)$.
 - 2: **for** m from 0 to K **do** ▷ Stochastic gradient iterations
 - 3: Sample $(\xi^j)_{j=1, \dots, B}$ from B independent copies of the initial condition ξ .
 - 4: Set $\forall j \in \llbracket 1, B \rrbracket$, $X_0^j(\theta_m) = \xi^j \in \mathbb{R}^d$, $Y_0^j(\theta_m) = \mathcal{Y}_{\theta_m^y}(\xi^j) \in \mathbb{R}^k$.
 - 5: **for** i from 0 to $N-1$ **do**
 - 6: $u_i = (u_i^X, u_i^Y, u_i^Z) = \frac{1}{B} \sum_{j=1}^B \left(\varphi_1(X_i^j(\theta_m)), \varphi_2(Y_i^j(\theta_m)), \varphi_3\left(\mathcal{Z}_{\theta_m^z}\left(t_i, X_i^j(\theta_m)\right)\right) \right)$
 - 7: **for** j from 1 to B **do**
 - 8: Sample δ_i^j from a d -dimensional standard Gaussian vector.
 - 9: $X_{i+1}^j(\theta_m) = X_i^j(\theta_m) + b\left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z}\left(t_i, X_i^j(\theta_m)\right), u_i\right) \Delta t + \sqrt{\Delta t} \sigma\left(t_i, X_i^j(\theta_m), u_i^X\right) \delta_i^j$
 - 10: $Y_{i+1}^j(\theta_m) = Y_i^j(\theta_m) - f\left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z}\left(t_i, X_i^j(\theta_m)\right), u_i\right) \Delta t + \sqrt{\Delta t} \mathcal{Z}_{\theta_m^z}\left(t_i, X_i^j(\theta_m)\right) \delta_i^j$
 - 11: **end for**
 - 12: **end for**
 - 13: $\overline{X}_N(\theta_m) = \frac{1}{B} \sum_{j=1}^B \varphi_1(X_N^j(\theta_m))$,
 - 14: $J(\theta_m) = \frac{1}{B} \sum_{j=1}^B \left(Y_N^j(\theta_m) - g\left(X_N^j(\theta_m), \overline{X}_N(\theta_m)\right) \right)^2$
 - 15: Calculate $\nabla J(\theta_m)$ by back-propagation.
 - 16: Update $\theta_{m+1} = \theta_m - \rho_m \nabla J(\theta_m)$.
 - 17: **end for**
-

- **Dynamic:** a method which dynamically updates the estimation on $(M+1)B$ samples. The expectations from the last M batches are kept in memory in an array

$$\begin{aligned}
& (\zeta_{i,r}) \quad i = 0, \dots, N-1, \\
& \quad \quad r = 0, \dots, M-1
\end{aligned}$$

initialized with values $(\mathbb{E}[\varphi_1(\xi)], \varphi_2(0), \varphi_3(0))^{N \times M}$.

At iteration $m - 1$, $\nu_i^{(m-1)}$ is defined as the empirical mean on these previous sample paths. On a new batch, the expectation is computed by averaging the previous estimation $\nu_i^{(m-1)}$ and the current batch empirical mean by the following algorithm used for $i = 0, \dots, N - 1$:

$$\begin{aligned} \nu_i^{(m-1)} &= \frac{1}{M} \sum_{r=0}^{M-1} \zeta_{i,r}, \\ u_i &= \frac{M\nu_i^{(m-1)} + \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_m)), \sum_{j=1}^B \varphi_3(Z_i^j(\theta_m)) \right)}{M+1}, \\ \zeta_{i,m\%M} &= \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_m)), \sum_{j=1}^B \varphi_3(Z_i^j(\theta_m)) \right). \end{aligned} \quad (11)$$

The notation $m\%M$ refers to the remainder of the Euclidian division of m by M . This technique allows to use smaller batches of size 100 or 1000. Thus it is more efficient in terms of convergence speed in comparison with the direct approach. This method can be seen as a dynamic fixed point approach.

The idea behind this update rule comes from online learning in machine learning. $1/M$ can be interpreted as a learning rate quantifying the updating speed. From the current estimation of the particles law, we introduce a small correction related to the new observed samples. Therefore the estimation is much more stable through iterations compared to the instantaneous update of the law used by the Direct method. After M batches, the older samples are forgotten, since they don't represent anymore the current law. Indeed we expect the convergence for a good choice of M . If this parameter is too small the stabilization would be inefficient and on the contrary a too large M would slow down the learning process by introducing a bias in the law. For instance in our numerical experiments of Section 4 we use $M = 100$ for a total of 2000 gradient descent iterations.

Remark 1. *If the law dependence is more general than a first order interaction and is given by a continuous function $F : \mu \in \mathcal{P}_2(\mathbb{R}^d) \mapsto \mathbb{R}^k$ then the Direct method can be straightforwardly applied to the equation by estimating $F(\mathcal{L}(X_t))$ by the so-called empirical projection $F(\frac{1}{B} \sum_{j=1}^B \delta_{X_{t_i}^j})$ for identically distributed particles $(X_{t_i}^j)_{j=1, \dots, B}$ on a time grid t_0, \dots, t_N . Concerning the Dynamic approach, it would require to keep in memory the previously computed particles from the last M batches which is costly.*

Remark 2. *The fixed point approach is known to be convergent theoretically only for small maturities. In practice, the theoretical bound on the maturity found on the simple example given for example in paragraph 3.1 in [Ang+19] is far too pessimistic. We will see that the restriction is not relevant on all our test cases.*

For a given iteration m , given the estimations $(\zeta_{i,r})_{\substack{i=0, \dots, N-1, \\ r=0, \dots, M-1}}$ of u_i on the last M

iterations, we perform one gradient descent step for the following optimization problem

$$\begin{aligned}
& \min_{\theta_m=(\theta_m^y, \theta_m^z)} \frac{1}{B} \sum_{k=1}^B \left| Y_N^k(\theta_m) - g \left(X_N^k(\theta_m), \frac{1}{B} \sum_{j=1}^B \varphi_1(X_N^j(\theta_m)) \right) \right|^2 \\
X_{i+1}^j(\theta_m) &= X_i^j(\theta_m) + b \left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right), \tilde{u}_i \right) \Delta t \\
&\quad + \sigma \left(t_i, X_i^j(\theta_m), \tilde{u}_i^X(\theta_m) \right) \Delta W_i^j \\
Y_{i+1}^j(\theta_m) &= Y_i^j(\theta_m) - f \left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right), \tilde{u}_i \right) \Delta t \\
&\quad + \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right) \Delta W_i^j \\
u_i &= \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_m)), \sum_{j=1}^B \varphi_3 \left(\mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right) \right) \right) \\
\tilde{u}_i &= \frac{\sum_{r=0}^{M-1} \zeta_{i,r} + u_i}{M+1} \\
X_0^j &= \xi^j \sim \xi, \quad j = 1, \dots, B \\
Y_0^j(\theta_m) &= \mathcal{Y}_{\theta_m^y}^j(X_0^j) \\
i &= 1, \dots, N-1.
\end{aligned}$$

Then we update $(\xi_i)_i$ by forgetting the oldest estimation and keeping in memory the new one, $(u_i)_i$ (see (11)). The **Global Dynamic solver** is given more explicitly in algorithm 2.

Algorithm 2 Global Dynamic solver

- 1: Let $\mathcal{Y}_{\theta^y}(\cdot)$ be a neural network with parameter θ^y , defined on \mathbb{R}^d and valued in \mathbb{R}^k , $\mathcal{Z}_{\theta^z}(\cdot, \cdot)$ be a neural network with parameter θ^z , defined on $\mathbb{R}^+ \times \mathbb{R}^d$ and valued in $\mathbb{R}^{k \times d}$, so that $\theta = (\theta^y, \theta^z)$ is initialized with value $\theta_0 = (\theta_0^y, \theta_0^z)$.
 - 2: Set $\forall i \in \llbracket 0, N-1 \rrbracket, \forall r \in \llbracket 0, M-1 \rrbracket, \zeta_{i,r} = (\mathbb{E}[\varphi_1(\xi)], \varphi_2(0), \varphi_3(0))$.
 - 3: **for** m from 0 to K **do**
 - 4: Sample $(\xi^j)_{j=1, \dots, B}$ from B independent copies of the initial condition ξ .
 - 5: Set $\forall j \in \llbracket 1, B \rrbracket, X_0^j(\theta_m) = \xi^j \in \mathbb{R}^d, Y_0^j(\theta_m) = \mathcal{Y}_{\theta_m^y}(\xi^j) \in \mathbb{R}^k$.
 - 6: **for** i from 0 to $N-1$ **do**
 - 7: $u_i = \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_m)), \sum_{j=1}^B \varphi_3 \left(\mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right) \right) \right)$
 - 8: $\tilde{u}_i = (\tilde{u}_i^X, \tilde{u}_i^Y, \tilde{u}_i^Z) = \frac{\sum_{r=0}^{M-1} \zeta_{i,r} + u_i}{M+1}$
 - 9: **for** j from 1 to B **do**
 - 10: Sample δ_i^j from a d -dimensional standard Gaussian vector.
 - 11: $X_{i+1}^j(\theta_m) = X_i^j(\theta_m) + b \left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right), \tilde{u}_i \right) \Delta t + \sqrt{\Delta t} \sigma \left(t_i, X_i^j(\theta_m), \tilde{u}_i^X \right) \delta_i^j$
 - 12: $Y_{i+1}^j(\theta_m) = Y_i^j(\theta_m) - f \left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right), \tilde{u}_i \right) \Delta t + \sqrt{\Delta t} \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right) \delta_i^j$
 - 13: **end for**
 - 14: $\zeta_{i, m \% M} = u_i$
 - 15: **end for**
 - 16: $\overline{X}_N(\theta_m) = \frac{1}{B} \sum_{j=1}^B \varphi_1(X_N^j(\theta_m)),$
 - 17: $J(\theta_m) = \frac{1}{B} \sum_{j=1}^B \left(Y_N^j(\theta_m) - g \left(X_N^j(\theta_m), \overline{X}_N(\theta_m) \right) \right)^2$
 - 18: Calculate $\nabla J(\theta_m)$ by back-propagation.
 - 19: Update $\theta_{m+1} = \theta_m - \rho_m \nabla J(\theta_m)$.
 - 20: **end for**
-

- **Expectation:** estimate u_t by a neural network Ψ_{θ^Ψ} with input t and parameters θ^Ψ .

$$u_i(\theta^\Psi) = \Psi_{\theta^\Psi}(t_i) = (\Psi_{\theta^\Psi}^X(t_i), \Psi_{\theta^\Psi}^Y(t_i), \Psi_{\theta^\Psi}^Z(t_i)), \quad i = 0, \dots, N. \quad (12)$$

A penalization term

$$\mathbb{E} \left[\frac{\lambda}{N} \sum_{i=0}^{N-1} \left\| \Psi_{\theta^\Psi}(t_i) - \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta)), \sum_{j=1}^B \varphi_3(Z_i^j(\theta)) \right) \right\|_2^2 \right],$$

is added to the loss function. We will see that in practice this method is quite involved to use because the performances heavily depend upon the choice of the parameter λ . This approach provides a relaxation of the fixed point method.

We solve the following optimization problem

$$\begin{aligned}
& \min_{\theta=(\theta^y, \theta^z, \theta^\Psi)} \frac{1}{B} \sum_{k=1}^B \left| Y_N^k(\theta) - g \left(X_N^k(\theta), \frac{1}{B} \sum_{j=1}^B \varphi_1(X_N^j(\theta)) \right) \right|^2 \\
& \quad + \frac{\lambda}{N} \sum_{i=0}^{N-1} \left\| u_i(\theta^\Psi) - \Psi_{\theta^\Psi}(t_i) \right\|^2 \\
X_{i+1}^j(\theta) &= X_i^j(\theta) + b \left(t_i, X_i^j(\theta), Y_i^j(\theta), \mathcal{Z}_{\theta^z} \left(t_i, X_i^j(\theta) \right), \Psi_{\theta^\Psi}(t_i) \right) \Delta t \\
& \quad + \sigma \left(t_i, X_i^j(\theta), \Psi_{\theta^\Psi}(t_i)^X \right) \Delta W_i^j \\
Y_{i+1}^j(\theta) &= Y_i^j(\theta) - f \left(t_i, X_i^j(\theta), Y_i^j(\theta), \mathcal{Z}_{\theta^z} \left(t_i, X_i^j(\theta) \right), \Psi_{\theta^\Psi}(t_i) \right) \Delta t \\
& \quad + \mathcal{Z}_{\theta^z} \left(t_i, X_i^j(\theta) \right) \Delta W_i^j \\
u_i &= \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta)), \sum_{j=1}^B \varphi_3 \left(\mathcal{Z}_{\theta^z} \left(t_i, X_i^j(\theta) \right) \right) \right) \\
X_0^j &= \xi^j \sim \xi, \quad j = 1, \dots, B \\
Y_0^j(\theta) &= \mathcal{Y}_{\theta^y}(X_0^j) \\
i &= 1, \dots, N-1.
\end{aligned}$$

The **Global Expectation solver** is described in algorithm 3. The parameter λ is chosen by trial and error.

Algorithm 3 Global Expectation solver

- 1: Let $\mathcal{Y}_{\theta^y}(\cdot)$ be a neural network with parameter θ^y , defined on \mathbb{R}^d and valued in \mathbb{R}^k , $\mathcal{Z}_{\theta^z}(\cdot, \cdot)$ defined on $\mathbb{R}^+ \times \mathbb{R}^d$, $\Psi_{\theta^\Psi}(\cdot) = (\Psi_{\theta^\Psi}^X(\cdot), \Psi_{\theta^\Psi}^Y(\cdot), \Psi_{\theta^\Psi}^Z(\cdot))$ defined on \mathbb{R}^+ be neural networks with parameters θ^z, θ^Ψ , taking values respectively in $\mathbb{R}^{k \times d}$ and $\mathbb{R}^d \times \mathbb{R}^k \times \mathbb{R}^{k \times d}$, so that $\theta = (\theta^y, \theta^z, \theta^\Psi)$ is initialized with value $\theta_0 = (\theta_0^y, \theta_0^z, \theta_0^\Psi)$.
 - 2: **for** m from 0 to K **do**
 - 3: Sample $(\xi^j)_{j=1, \dots, B}$ from B independent copies of the initial condition ξ .
 - 4: Set $\forall j \in \llbracket 1, B \rrbracket$, $X_0^j(\theta_m) = \xi^j \in \mathbb{R}^d, Y_0^j(\theta_m) = \mathcal{Y}_{\theta_m^y}(\xi^j) \in \mathbb{R}^k$.
 - 5: **for** i from 0 to $N - 1$ **do**
 - 6: $u_i = \frac{1}{B} \left(\sum_{j=1}^B \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^B \varphi_2(Y_i^j(\theta_m)), \sum_{j=1}^B \varphi_3 \left(\mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right) \right) \right)$
 - 7: **for** j from 1 to B **do**
 - 8: Sample δ_i^j from a d -dimensional Gaussian vector.
 - 9: $X_{i+1}^j(\theta_m) = X_i^j(\theta_m) + b \left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right), \Psi_{\theta_m^\Psi}(t_i) \right) \Delta t + \sqrt{\Delta t} \sigma \left(t_i, X_i^j(\theta_m), \Psi_{\theta_m^\Psi}^X(t_i) \right) \delta_i^j$
 - 10: $Y_{i+1}^j(\theta_m) = Y_i^j(\theta_m) - f \left(t_i, X_i^j(\theta_m), Y_i^j(\theta_m), \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right), \Psi_{\theta_m^\Psi}(t_i) \right) \Delta t + \sqrt{\Delta t} \mathcal{Z}_{\theta_m^z} \left(t_i, X_i^j(\theta_m) \right) \delta_i^j$
 - 11: **end for**
 - 12: **end for**
 - 13: $\overline{X}_N(\theta_m) = \frac{1}{B} \sum_{j=1}^B \varphi_1(X_N^j(\theta_m))$,
 - 14: $J(\theta_m) = \frac{1}{B} \sum_{j=1}^B \left(Y_N^j(\theta_m) - g \left(X_N^j(\theta_m), \overline{X}_N(\theta_m) \right) \right)^2 + \frac{\lambda}{N} \sum_{i=0}^{N-1} \left(u_i - \Psi_{\theta_m^\Psi}(t_i) \right)^2$
 - 15: Calculate $\nabla J(\theta_m)$ by back-propagation.
 - 16: Update $\theta_{m+1} = \theta_m - \rho_m \nabla J(\theta_m)$.
 - 17: **end for**
-

We will compare the performances of these techniques on several examples in section 4.

3 A local solver

We also propose a local method inspired by the Deep Backward Dynamic Programming introduced by [HPW20] and [PWG21]. It considers local minimization problems between contiguous time steps. In this case there are as many networks as time steps. We replace a global optimization setting by a set of smaller problems.

In this method for $i \in \llbracket 0, N - 1 \rrbracket$, Z_i and Y_i are approximated by a neural network $(\mathcal{Z}_{\theta^z}^i(\cdot), \mathcal{Y}_{\theta^y}^i(\cdot))$ with parameters $\theta = (\theta_0^y, \theta_0^z, \dots, \theta_{N-1}^y, \theta_{N-1}^z)$. At iteration m , with $\theta_m = (\theta_{m,0}^y, \theta_{m,0}^z, \dots, \theta_{m,N-1}^y, \theta_{m,N-1}^z)$, we simulate $X_i(\theta_m)$ with the previously computed parameters θ_m .

$$\begin{aligned} X_{i+1}^j(\theta_m) &= X_i^j(\theta_m) + b \left(t_i, X_i^j(\theta_m), \mathcal{Y}_{\theta_m^y}^i \left(X_i^j(\theta_m) \right), \mathcal{Z}_{\theta_m^z}^i \left(X_i^j(\theta_m) \right), \tilde{u}_i \right) \Delta t \\ &+ \sigma \left(t_i, X_i^j(\theta_m), \tilde{u}_i^X \right) \Delta W_i^j. \end{aligned} \quad (13)$$

This first step allows to find the areas visited by the controlled process. Using R samples, we compute the empirical mean $\overline{m}_i^m = \frac{1}{R} \left(\sum_{j=1}^R X_i^j(\theta_m) \right)$ and variance $V_i^m = \frac{1}{R} \sum_{j=1}^R \left(X_i^j(\theta_m) \right)^2 -$

$\frac{1}{R^2} \left(\sum_{j=1}^R X_i^j(\theta_m) \right)^2$ of $X_i(\theta_m)$. We estimate u_t as in the Dynamic method (11):

$$u_i = \frac{1}{R} \left(\sum_{j=1}^R \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^R \varphi_2 \left(\mathcal{Y}_{\theta_{m,i}}^i \left(t_i, X_i^j(\theta_m) \right) \right), \sum_{j=1}^R \varphi_3 \left(\mathcal{Z}_{\theta_{m,i}}^i \left(t_i, X_i^j(\theta_m) \right) \right) \right)$$

$$\tilde{u}_i = \frac{\sum_{r=0}^{M-1} \zeta_{i,r} + u_i}{M+1}.$$

A priori R can be different from the batch size B . It is the batch size used for the estimation of the law with the previously computed parameters. In the numerical tests of Section 4 we use $B = 100$ or $B = 300$ for the backward optimization and a larger value $R = 50000$ for the Monte-Carlo forward estimation of the law. Then we solve backward problems to find the θ_{m+1} by sampling B independent copies of X_i through a Gaussian distribution $\mathcal{N}(\overline{m}_i^m, V_i^m)$ with frozen parameters θ_m :

- First sample B independent copies X_N^1, \dots, X_N^B of X_N following a Gaussian distribution $\mathcal{N}(\overline{m}_N^m, V_N^m)$. $Y_{\theta_{m+1,N}}^N(X_N^j)$ is set to the terminal condition $g(X_N^j, u_N^X)$.
- For i from $N-1$ to 0 :
Sample B independent copies X_i^1, \dots, X_i^B of X_i following a Gaussian distribution $\mathcal{N}(\overline{m}_i^m, V_i^m)$. Diffuse according to the dynamics (13) of X , starting from X_i^1, \dots, X_i^B , to obtain $X_{i+1}^1, \dots, X_{i+1}^B$.
Solve the local optimization problem

$$\min_{\theta=(\theta^y, \theta^z)} \frac{1}{B} \sum_{j=1}^B \left| Y_{\theta_{m+1,i+1}}^{i+1} \left(X_{i+1}^j \right) - Y_{\theta^y}^i \left(X_i^j \right) + f \left(t_i, X_i^j, \mathcal{Y}_{\theta^y}^i \left(X_i^j \right), \mathcal{Z}_{\theta^z}^i \left(X_i^j \right), \tilde{u}_i \right) \Delta t - \mathcal{Z}_{\theta^z}^i \left(X_i^j \right) \Delta W_i^j \right|^2,$$

starting from the parameter value $\theta_{m+1,i+1}$. We can then update the θ value by denoting as $\theta_{m+1,i}$ the argmin value of the minimization problem.

- Repeat the previous steps for the iteration $m+1$ until reaching K iterations.

In the version of the **Local Dynamic solver** given in algorithm 4, we use the dynamic update of the expectations introduced previously in the dynamic solver of section 2. In this algorithm H stands for the number of gradient steps to perform at each step of the algorithm and R is the number of samples for the laws estimation.

Remark 3. *Because we have to learn the dynamic of the forward process, the use of a backward resolution is not as obvious as in [Hur+21; Bac+21]. We have to alternate between forward dynamic estimations and backward resolutions.*

More precisely, here we solve fully coupled FBSDEs, whereas the works [HPW20; PWG21] consider decoupled FBSDEs, where the forward process X can be simulated independently of Y, Z . Estimating the law of X and sampling from a normal distribution allows us to decouple and solve locally the FBSDEs in areas visited by X and its law. However, because of this freezing of the forward dynamics, another fixed point problem has to be solved. Other approaches for fully coupled FBSDEs like [HL20] and [Ji+20] rely on the global machine learning method initiated by [HJE17].

Notice that we estimate the law dependent part of the dynamics in the first loop but for the resolution loop, we use a normal distribution as training measure. Another natural training measure could be the empirical law of the process. Here, we employ an hybrid method between the two methods proposed in Remark 2.1 from [Bac+21]. We exploit the approximated control to estimate the law

of X and use a Gaussian law with the same two first moments as X as training measure. In the numerical examples of [Bac+21], the authors also use Gaussian distributions as training measures, but with fixed parameters through the iterations. Here we adapt the training measures during the training to have the same expectation and variance as the controlled process thanks to an exploration step.

Algorithm 4 Local Dynamic solver

- 1: Let $(\mathcal{Y}_{\theta^y}^i(\cdot), \mathcal{Z}_{\theta^z}^i(\cdot))$ be some neural networks defined on \mathbb{R}^d with values in $\mathbb{R}^k \times \mathbb{R}^{k \times d}$ for $i = 0, \dots, N-1$ and parameters $\theta = (\theta_0^y, \theta_0^z, \dots, \theta_{N-1}^y, \theta_{N-1}^z)$ initialized with values $\theta_0 = (\theta_{0,0}^y, \theta_{0,0}^z, \dots, \theta_{0,N-1}^y, \theta_{0,N-1}^z)$.
 - 2: Set $\forall i \in \llbracket 0, N \rrbracket, \forall r \in \llbracket 0, M-1 \rrbracket, \zeta_{i,r} = (\mathbb{E}[\xi], 0, 0)$.
 - 3: **for** m from 0 to K **do**
 - 4: Sample δ_i^j from a d -dimensional standard Gaussian vector, $i = 0, \dots, N, j = 1, \dots, R$.
 - 5: Sample $(\xi^j)_{j=1, \dots, B}$ from R independent copies of the initial condition ξ .
 - 6: Set $\forall j \in \llbracket 1, R \rrbracket, X_0^j(\theta_m) = \xi^j \in \mathbb{R}^d$.
 - 7: **for** i from 0 to N **do** ▷ Forward estimation of the laws
 - 8: $l_i^m = \frac{1}{R} \left(\sum_{j=1}^R X_i^j(\theta_m) \right)$
 - 9: $u_i = \frac{1}{R} \left(\sum_{j=1}^R \varphi_1(X_i^j(\theta_m)), \sum_{j=1}^R \varphi_2(\mathcal{Y}_{\theta_{m,i}^y}^i(X_i^j(\theta_m))), \sum_{j=1}^R \varphi_3(\mathcal{Z}_{\theta_{m,i}^z}^i(X_i^j(\theta_m))) \right)$
 - 10: $V_i^m = \frac{1}{R} \sum_{j=1}^R \left(X_i^j(\theta_m) \right)^2 - \frac{1}{R^2} \left(\sum_{j=1}^R X_i^j(\theta_m) \right)^2$
 - 11: $\tilde{u}_i = (\tilde{u}_i^X, \tilde{u}_i^Y, \tilde{u}_i^Z) = \frac{\sum_{r=0}^{M-1} \zeta_{i,r} + u_i}{M+1}$
 - 12: $\zeta_{i,m \% M} = u_i(\theta_m)$
 - 13: **for** j from 1 to R **do**
 - 14: $X_{i+1}^j(\theta_m) = X_i^j(\theta_m) + b \left(t_i, X_i^j(\theta_m), \mathcal{Y}_{\theta_{m,i}^y}^i(X_i^j(\theta_m)), \mathcal{Z}_{\theta_{m,i}^z}^i(X_i^j(\theta_m)), \tilde{u}_i \right) \Delta t + \sqrt{\Delta t} \sigma \left(t_i, X_i^j(\theta_m), \tilde{u}_i^X \right) \delta_i^j$
 - 15: **end for**
 - 16: **end for**
 - 17: **for** i from $N-1$ to 0 **do** ▷ Backward resolution
 - 18: $\hat{\theta}_0 = \theta_{m,i}$
 - 19: **for** h from 0 to $H-1$ **do** ▷ Gradient descent with simulated data for X
 - 20: **for** j from 1 to B **do**
 - 21: Sample Ξ_i^j, Θ_i^j from d -dimensional standard Gaussian vectors.
 - 22: $x_i^j = l_i^m + \sqrt{V_i^m} \Theta_i^j$
 - 23: $x_{i+1}^j = x_i^j + b \left(t_i, x_i^j, \mathcal{Y}_{\hat{\theta}_h^y}^i(x_i^j), \mathcal{Z}_{\hat{\theta}_h^z}^i(x_i^j), \tilde{u}_i \right) \Delta t + \sqrt{\Delta t} \sigma \left(t_i, x_i^j, \tilde{u}_i^X \right) \Xi_i^j$
 - 24: **if** $i = N-1$ **then**
 - 25: $Y_{i+1}^j = g \left(x_N^j, \tilde{u}_N^X \right)$
 - 26: **else**
 - 27: $Y_{i+1}^j = \mathcal{Y}_{\hat{\theta}_{m+1,i+1}^y}^{i+1} \left(x_{i+1}^j \right)$
 - 28: **end if**
 - 29: **end for**
 - 30: $J^i(\hat{\theta}_h) = \frac{1}{B} \sum_{j=1}^B \left(f \left(t_i, x_i^j, \mathcal{Y}_{\hat{\theta}_h^y}^i(x_i^j), \mathcal{Z}_{\hat{\theta}_h^z}^i(x_i^j), \tilde{u}_i \right) \Delta t + Y_{i+1}^j - \mathcal{Y}_{\hat{\theta}_h^y}^i(x_i^j) - \sqrt{\Delta t} \mathcal{Z}_{\hat{\theta}_h^z}^i(x_i^j) \Xi_i^j \right)^2$
 - 31: Calculate $\nabla J^i(\hat{\theta}_h)$ by back-propagation.
 - 32: Update $\hat{\theta}_{h+1} = \hat{\theta}_h - \rho_h \nabla J^i(\hat{\theta}_h)$.
 - 33: **end for**
 - 34: $\theta_{m+1,i} = \hat{\theta}_H$
 - 35: **end for**
 - 36: **end for**
-

4 Numerical results

The algorithms are implemented in Python with the Tensorflow library [Aba+16]. Each numerical experiment is conducted using a node composed of 2 Intel® Xeon® Gold 5122 Processors, 192 Go of RAM, and 2 GPU nVidia® Tesla® V100 16Go. The multi-GPU parallelization on the global solver is conducted using the Horovod library [SDB18]. The methods we test are:

- **Global Direct:** algorithm 1 at page 7. Batch size $B = 10000$.
- **Global Dynamic:** algorithm 2 at page 10. Batch size $B = 200$ and $M = 100$.
- **Global Expectation:** algorithm 3 at page 12. Batch size $B = 2000$.
- **Local Dynamic:** algorithm 4 at page 15. Batch size $B = 300$ (Weak), $B = 100$ (Pontryagin) and $M = 20$, $R = 50000$.

The name depends on the two choices made regarding the algorithm: the use a Global or Local method but also what law estimation technique is applied. If the algorithm is applied to equations coming from the Pontryagin (abbreviated in Pont.) or the Weak approach, it is specified in its name. In the following tables, the incorrect results are highlighted in red. In Section 4.1 the three best results are highlighted in green whereas in Section 4.3 only the best value is.

4.1 Linear price impact model

We use a linear-quadratic mean-field game of controls model studied in [Ang+19] and [CD18] for comparison. This model is useful for numerical tests since the analytic solution is known. The MFG of controls model for the representative player is given by:

$$\begin{aligned} \min_{\alpha \in \mathbb{A}} \quad & \mathbb{E} \left[\int_0^T \left(\frac{c_\alpha}{2} \|\alpha_t\|^2 + \frac{c_X}{2} \|X_t\|^2 - \gamma X_t \cdot u_t \right) dt + \frac{c_g}{2} \|X_T\|^2 \right] \\ \text{subject to} \quad & X_t = x_0 + \int_0^t \alpha_s ds + \sigma W_t \end{aligned} \quad (14)$$

and the fixed point $\mathbb{E}[\alpha_t] = u_t$. In this case, the mean-field interaction is exerted through the law of the control process.

The Pontryagin optimality principle gives the system:

$$\begin{cases} dX_t &= -\frac{1}{c_\alpha} Y_t dt + \sigma dW_t \\ X_0 &= x_0 \\ dY_t &= -(c_X X_t + \frac{\gamma}{c_\alpha} \mathbb{E}[Y_t]) dt + Z_t dW_t \\ Y_T &= c_g X_T. \end{cases} \quad (15)$$

In this case, the output Z of the neural network is a matrix of size $d \times d$ and Y is a vector of size d . The weak representation of the value function gives:

$$\begin{cases} dX_t &= -\frac{1}{c_\alpha} \sigma^{-1} Z_t dt + \sigma dW_t \\ X_0 &= x_0 \\ dY_t &= -\left(\frac{c_X}{2} \|X_t\|^2 + \frac{\gamma}{c_\alpha} X_t \cdot \sigma^{-1} \mathbb{E}[Z_t] + \frac{1}{2c_\alpha} \|\sigma^{-1} Z_t\|^2 \right) dt + Z_t dW_t \\ Y_T &= \frac{c_g}{2} \|X_T\|^2. \end{cases} \quad (16)$$

In this case, the output Z of the neural network is a vector of size d and Y is a scalar. Therefore we may be able work in higher dimensions.

Remark 4. *With LQ models, the dynamics of Y is linear in the Pontryagin approach and quadratic in the Weak approach. Thus the potentially high dimension of one method is counterbalanced by the complex dynamics of the other technique.*

For our numerical experiments we take $c_X = 2, x_0 = 1, \sigma = 0.7, \gamma = 2, c_\alpha = 2/3, c_g = 0.3$. If not stated otherwise, the simulations are conducted with $T = 1, d = 10, \Delta t = 0.01$.

Method \ T	0.25	0.75	1.0	1.5
Reference	0.7709	0.1978	0.0811	0.0125
Glob. Direct Pontryagin	0.763 (1.3e-03)	0.187 (2.5e-03)	0.075 (2.7e-03)	0.012 (5.0e-03)
Glob. Dyn. Pont.	0.762 (2.3e-03)	0.189 (4.0e-03)	0.078 (5.5e-03)	0.013 (6.7e-03)
Glob. Exp. Pont. (0.1)	0.763 (1.6e-03)	0.604 (1.1e-01)	0.729 (1.1e-01)	0.803 (1.5e-01)
Glob. Exp. Pont. (1.)	0.762 (1.4e-03)	0.251 (2.7e-02)	0.467 (7.6e-02)	0.639 (1.1e-01)
Glob. Exp. Pont. (10.)	0.763 (1.5e-03)	0.216 (1.7e-02)	0.275 (3.7e-02)	0.574 (1.7e-01)
Glob. Exp. Pont. (100.)	0.776 (8.4e-03)	0.797 (1.1e-01)	1.042 (1.3e-01)	1.613 (2.6e-01)
Glob. Direct Weak	0.778 (2.0e-03)	0.200 (1.4e-02)	0.092 (2.9e-02)	0.025 (2.0e-02)
Glob. Dyn. Weak	0.775 (4.4e-03)	0.212 (2.0e-02)	0.083 (4.1e-02)	0.016 (5.6e-02)
Glob. Exp. Weak (0.1)	0.877 (1.9e-02)	0.654 (9.9e-02)	0.595 (2.3e-01)	0.28 (6.0e-01)
Glob. Exp. Weak (1.)	0.901 (2.2e-03)	0.664 (9.8e-02)	0.617 (1.0e-01)	0.507 (1.9e-01)
Glob. Exp. Weak (10.)	0.887 (1.1e-02)	0.698 (7.3e-02)	0.6541 (6.3e-02)	0.49 (2.3e-01)
Glob. Exp. Weak (100.)	0.887 (2.0e-03)	0.650 (9.4e-02)	0.602 (9.1e-02)	0.492 (2.4e-01)
Loc. Dyn. Pont	0.767 (3.5e-04)	0.189 (6.3e-04)	0.076 (7.6e-04)	0.011 (7.5e-04)
Loc. Dyn. Weak	0.944 (8.7e-04)	0.740 (2.6e-02)	0.692 (1.6e-02)	0.625 (2.2e-02)

Table 1: Mean of $\mathbb{E}[X_T]$ over the 10 dimensions (and standard deviation) for several maturities T on the price impact model (14). We perform 2000 iterations for global (glob.) methods and 20000 iterations for local (loc.) methods. For the expectation method, the value of the λ penalization parameter is given under parenthesis.

Global Direct Pontryagin	1877 s.
Global Dyn. Pontryagin	1336 s.
Global Exp. Pontryagin	1562 s.
Global Direct Weak	2205s.
Global Dyn. Weak	1605 s.
Global Exp. Weak	1670 s.
Local Dynamic Pontryagin	11627 s.
Local Dynamic Weak	12689 s.

Table 2: Duration times of the methods (2000 iterations for global methods, 20000 iterations for local methods) on the price impact model (14) with $T = 1$. on one run

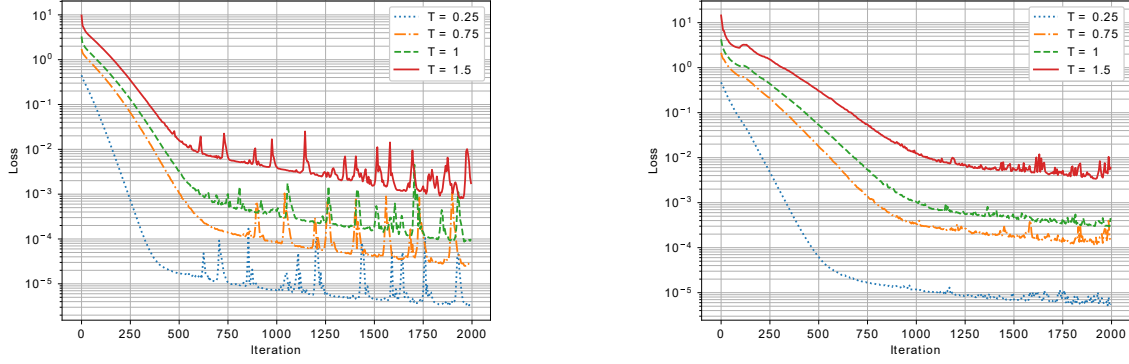


Figure 1: Learning curves for Global Direct (left) and Global Dynamic (right) Pontryagin methods on the price impact model (14). The loss is the L^2 error between Y_T and the terminal condition of the backward equation.

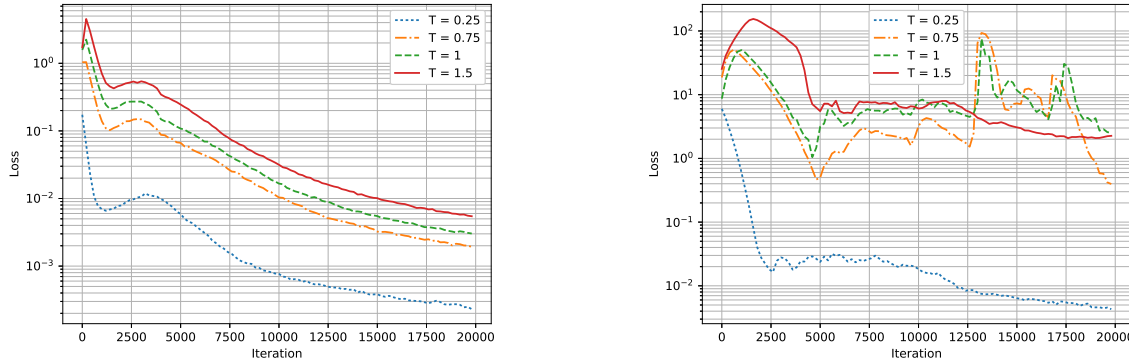


Figure 2: Learning curves for Local Dynamic Pontryagin (left) and Local Dynamic Weak (right) methods on the price impact model (14). The loss is the sum of the local L^2 errors between the neural network Y and the Euler discretization for all time steps.




Figure 3: Learning curves for Global Direct Weak (left) and Global Dynamic Weak (right) methods on the price impact model (14). The loss is the L^2 error between Y_T and the terminal condition of the backward equation.




Figure 4: First coordinate of the optimal control evaluated on a sample path for Global Direct Pontryagin (left) and Global Dynamic (right) Pontryagin methods after 2000 iterations on the price impact model (14) with $T = 1$.




Figure 5: First coordinate of the optimal control evaluated on a sample path for Local Dynamic Pontryagin (left) and Local Dynamic Weak (right) methods after 20000 iterations on the price impact model (14) with $T = 1$.

All methods except Global Expectation Weak and Local Dynamic Weak converge to the exact solution for small maturities. These two solvers do not converge to the right solution for any time horizon.

The choice of the parameter λ influences a lot the output of the Global Expectation Pontryagin scheme, as observed in Table 1. The best results are obtained when λ is of order 10 but we notice that a large range of values seems to work fine for very small time horizon T . However the Global Expectation Weak scheme never works. We will not test the expectation methods on the other test cases since they are less efficient than the other methods.

We see in Figure 2 that the Local Dynamic Pontryagin method needs more iterations for the loss to stabilize than the Global methods. We cannot hope for more iterations to help the convergence in the Global Weak methods since the loss in the learning curves of Figure 3 reaches a plateau. The algorithms solving the system coming from the Pontryagin principle perform better than the others. The dynamic estimation of the expectation allows to gain training speed and to smooth the loss, as seen in Figure 1 and Table 2. As another accuracy test, we can also plot the optimal control for which we have an analytical expression. We see in Figures 4, 5 that Global Pontryagin and Local Dynamic Pontryagin methods perform well but that the Local Dynamic Weak method does not seem to converge, which confirms what is observed in Table 1.

Global Direct	535 s.
Global Dynamic	425 s.
Local Dynamic	10900 s.

Table 3: Duration times of the algorithms for model (17) on one run (2000 iterations for global methods, 20000 iterations for the local method)

4.2 A one-dimensional mixed model

We consider the following one-dimensional example from [CL19; Ang+19]:

$$\begin{cases} dX_t &= -\rho Y_t dt + \sigma dW_t, X_0 = x_0 \\ dY_t &= \arctan(\mathbb{E}[X_t]) dt + Z_t dW_t, Y_T = \arctan(X_T). \end{cases} \quad (17)$$

This model comes from the Pontryagin principle applied to the mean-field game problem

$$\begin{aligned} \min_{\alpha} \mathbb{E} \left[\int_0^T \left(\frac{1}{2\rho} \alpha_s^2 - X_s \arctan(u_s) \right) ds + g(X_T) \right] \\ dX_t = \alpha_t dt + dW_t, X_0 = x_0, \end{aligned}$$

with the fixed point $u_s = \mathbb{E}[X_s]$, and where g is an antiderivative of \arctan . We take the same model parameters as in [CL19] ($T = 1$ and $x_0 = 1$) and obtain in Figure 6 with all our methods the same results as in their Figure 4. For the numerical resolution we choose 100 time steps. Notice that we use 3 hidden layers with 11 neurons in each when [CL19] uses 100 neurons by layer. We see that our smaller number of neurons is enough for this example resolution.

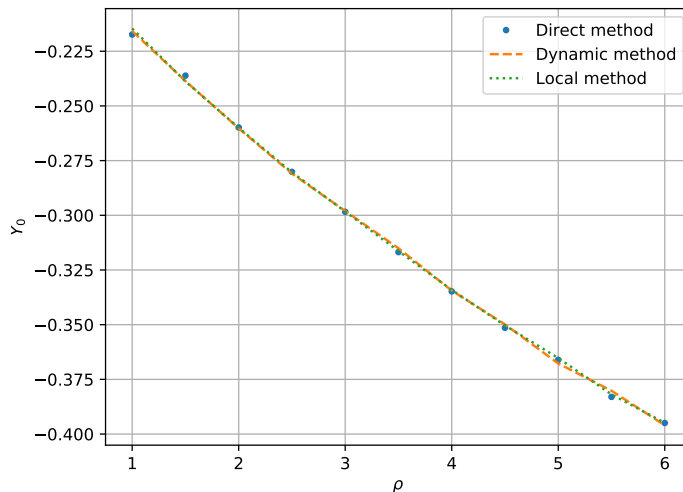


Figure 6: Value of Y_0 as a function of parameter ρ for the model (17)

The duration of the methods are given in Table 3 which illustrates again the speed gain in using the dynamic method.

4.3 Beyond the mean-field games case

In this Section we design non Linear Quadratic models in order to test the limitations of our methods. We construct general MKV FBSDES with explicit solutions following a log-normal distribution. Let X_t^i be defined by

$$dX_t^i = aX_t^i dt + \sigma X_t^i dW_t^i, \quad (18)$$

$$X_0^i = \xi, \quad (19)$$

with $a, \sigma, \xi \in \mathbb{R}$. We obtain explicitly

$$\begin{aligned} X_t^i &= \xi e^{(a - \frac{\sigma^2}{2})t + \sigma W_t^i}, \\ h_t^i &:= \mathbb{E}[X_t^i] = \xi e^{at}, \\ k_t^i &:= \mathbb{E}[(X_t^i)^2] = \xi e^{(2a + \sigma^2)t}. \end{aligned}$$

We choose $u : (t, x) \mapsto e^{\alpha t} \log(\prod_{i=1}^n x^i)$ and the following dynamic for Y_t

$$Y_t = e^{\alpha t} \log\left(\prod_i X_t^i\right) = e^{\alpha t} \sum_i \left[\log(\xi) + \left(a - \frac{\sigma^2}{2}\right)t + \sigma W_t^i \right],$$

such that

$$\begin{aligned} c_t &:= \mathbb{E}[Y_t] = e^{\alpha t} \sum_i \left[\log(\xi) + \left(a - \frac{\sigma^2}{2}\right)t \right], \\ d_t &:= \mathbb{E}[Y_t^2] = e^{2\alpha t} \left(\left[\sum_i \left(\log(\xi) + a - \frac{\sigma^2}{2} \right)t \right]^2 + \sum_i \sigma^2 t \right). \end{aligned}$$

As we have $Y_t = u(t, X_t)$, we obtain $Z_t^i = \sigma X_t^i \partial_{x_i} u(t, X_t)$ following:

$$\begin{aligned} Z_t^i &= \sigma e^{\alpha t} \\ e_t^i &:= \mathbb{E}[Z_t^i] = \sigma e^{\alpha t}, \\ f_t^i &:= \mathbb{E}[(Z_t^i)^2] = \sigma^2 e^{2\alpha t}. \end{aligned}$$

Introducing

$$\begin{aligned} \phi(t, x) &:= \partial_t u + \sum_i a x_i \partial_{x_i} u + \sum_i \frac{(\sigma x_i)^2}{2} \partial_{x_i^2}^2 u \\ &= e^{\alpha t} \left(\alpha \log\left(\prod_i x^i\right) + \sum_i \left(a - \frac{\sigma^2}{2}\right) \right), \end{aligned}$$

$u(t, X_t)$ solves the PDE

$$\partial_t u + \sum_i a x^i \partial_{x_i} u + \sum_i \frac{\sigma^2}{2} \partial_{x_i^2}^2 u - \phi(t, x) = 0.$$

This semilinear PDE is related to the BSDE associated with the driver $f(t, x) = -\phi(t, x)$ for forward dynamics (18).

Using some chosen \mathbb{R}^d valued functions ψ and \mathbb{R}^k valued functions κ , we express all dynamics in a McKean-Vlasov setting:

$$\begin{cases} dX_t^i &= (aX_t^i + \psi(Y_t, Z_t^i, \mathbb{E}[X_t^i], \mathbb{E}[(X_t^i)^2], \mathbb{E}[Y_t], \mathbb{E}[Y_t^2], \mathbb{E}[Z_t^i], \mathbb{E}[(Z_t^i)^2]) \\ &\quad - \psi(e^{\alpha t} \log(\prod_i X_t^i), \sigma e^{\alpha t}, h_t^i, k_t^i, c_t, d_t, e_t^i, f_t^i)) dt + \sigma X_t^i dW_t^i \\ X_0^i &= \xi \\ dY_t &= -f(t, X_t, Y_t, Z_t, \mathbb{E}[X_t], \mathbb{E}[X_t^2], \mathbb{E}[Y_t], \mathbb{E}[Y_t^2], \mathbb{E}[Z_t], \mathbb{E}[Z_t^2]) dt \\ &\quad + Z_t dW_t \\ Y_T &= e^{\alpha T} \log(\prod_i X_T^i) \end{cases} \quad (20)$$

Method \ T	0.25	0.75	1.0	1.5
Reference	1.0253	1.0779	1.1052	1.1618
Global Direct	1.025 (2e-03)	1.076 (3e-03)	1.095 (4e-03)	1.162 (7e-03)
Global Dynamic	1.026 (2e-03)	1.077 (4e-03)	1.105 (3e-03)	1.163 (5e-03)
Local Dynamic	1.025 (2e-04)	1.092 (5e-04)	1.146 (8e-04)	1.28 (1e-03)

Table 4: Mean of $\mathbb{E}[X_T]$ over the 10 dimensions (and standard deviation) for several maturities T (2000 iterations for global methods, 20000 iterations for local method) on the fully coupled linear model (21).

Global Direct	Global Dynamic	Local Dynamic
2081 s.	1308 s.	14811 s.

Table 5: Duration times of the methods (2000 iterations for global methods, 20000 iterations for local method) on the fully coupled linear model (21) for $T = 1$.

with

$$\begin{aligned}
& f(t, X_t, Y_t, Z_t, x_1, x_2, y_1, y_2, z_1, z_2) \\
&= -\phi(t, x) + \kappa(Y_t, Z_t, x_1, x_2, y_1, y_2, z_1, z_2) \\
& - \kappa \left(e^{\alpha t} \log \left(\prod_i X_t^i \right), \sigma e^{\alpha t}, h_t^i, k_t^i, c_t, d_t, e_t^i, f_t^i \right).
\end{aligned}$$

and $f : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^{1 \times d} \times \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}^{1 \times d} \times \mathbb{R}^{1 \times d} \mapsto \mathbb{R}$.

We consider two models of this kind for numerical tests.

4.3.1 A fully coupled linear example

We consider a linear McKean-Vlasov FBSDE in Y_t, Z_t and their law dynamics for X_t and Y_t :

$$\left\{ \begin{array}{l}
dX_t^i = (aX_t^i + b(Y_t + Z_t^i + \mathbb{E}[X_t^i] + \mathbb{E}[Y_t] + \mathbb{E}[Z_t^i]) \\
\quad - b \left(e^{\alpha t} \log \left(\prod_{i=1}^d X_t^i \right) + \sigma e^{\alpha t} + h_t^i + c_t + e_t^i \right) dt + \sigma X_t^i dW_t^i \\
X_0^i = \xi \\
dY_t = \left(\phi(t, X_t) + b(Y_t + \frac{1}{d} \sum_{i=1}^d Z_t^i + \frac{1}{d} \sum_{i=1}^d \mathbb{E}[X_t^i] + \mathbb{E}[Y_t] \right. \\
\quad \left. + \frac{1}{d} \sum_{i=1}^d \mathbb{E}[Z_t^i]) - b \left(e^{\alpha t} \log \left(\prod_{i=1}^d X_t^i \right) + \frac{1}{d} \sum_{i=1}^d \sigma e^{\alpha t} + \frac{1}{d} \sum_{i=1}^d h_t^i \right. \right. \\
\quad \left. \left. + c_t + \frac{1}{d} \sum_{i=1}^d e_t^i \right) \right) dt + Z_t dW_t \\
Y_T = e^{\alpha T} \log \left(\prod_{i=1}^d X_T^i \right).
\end{array} \right. \quad (21)$$




Figure 7: Learning curves for Local Dynamic method on the fully coupled linear model (21). The loss is the sum of the local L^2 errors between the neural network Y and the Euler discretization for all time steps.

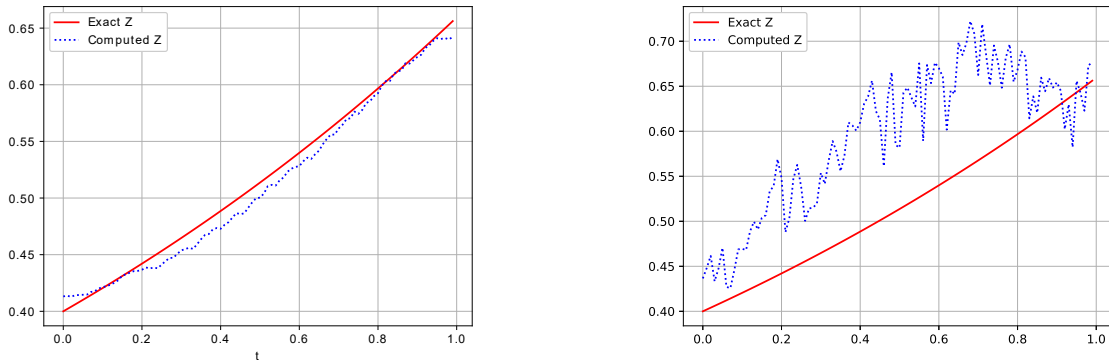


Figure 8: First coordinate of Z_t evaluated on a sample path for Global Dynamic (left) and Local Dynamic (right) methods after 2000 iterations (respectively 20000) on the fully coupled linear model (21).




Figure 9: First coordinate of Y_t evaluated on a sample path for Global Direct (left) and Local Dynamic methods (right) after 2000 iterations (respectively 20000) on the fully coupled linear model (21).

Method \ T	0.25	0.75	1.0	1.5
Reference	1.0253	1.0779	1.1052	1.1618
Global Direct	1.024 (1.8e-03)	1.065 (4.3e-03)	12.776 (3.3e-02)	DV
Global Dynamic	1.025 (2.1e-03)	1.072 (3.1e-03)	0.961 (7.0e-03)	DV
Local Dynamic	1.024 (1.6e-04)	-7.180 (9.0e-04)	0.411 (1.1e-03)	DV

Table 6: Mean of $\mathbb{E}[X_T]$ over the 10 dimensions (and standard deviation) for several maturities T (2000 iterations for global methods, 20000 iterations for Local method) on the fully coupled quadratic model (22).

Global Direct	Global Dynamic	Local Dynamic
2072 s.	1309 s.	14823 s.

Table 7: Duration times of the methods (2000 iterations for global methods, 20000 iterations for local method) on the fully coupled quadratic model (22) for $T = 1$.

We take $a = b = 0.1$, $\alpha = 0.5$, $\sigma = 0.4$, $\xi = 1$. The three algorithms demonstrate good performances on this test case. Both processes Y, Z are well represented by the neural network. However the Local Dynamic method is less precise than the global methods when the maturity grows. We see in Table 4, Figure 8, Figure 9 that the Local Dynamic method is biased when $T = 1$ when the Global methods achieve a great accuracy. It looks like the results of the Local Dynamic method cannot be improved since the loss flattens in Figure 7.

4.3.2 A fully coupled quadratic example

We consider a quadratic McKean-Vlasov FBSDE in Y_t, Z_t and their law dynamics for X_t and Y_t :
We consider a quadratic McKean-Vlasov FBSDE in Y_t, Z_t and their law dynamics for X_t and Y_t :

$$\left\{ \begin{array}{l}
dX_t^i = (aX_t^i + b(Y_t + Z_t^i + \mathbb{E}[X_t^i] + \mathbb{E}[Y_t] + \mathbb{E}[Z_t^i]) + \sigma X_t^i dW_t^i \\
\quad - b \left(e^{\alpha t} \log \left(\prod_{i=1}^d X_t^i \right) + \sigma e^{\alpha t} + h_t^i + c_t + e_t^i \right) \\
\quad + c \left(Y_t^2 + (Z_t^i)^2 + \mathbb{E}[(X_t^i)^2] + \mathbb{E}[Y_t^2] + \mathbb{E}[(Z_t^i)^2] \right) \\
\quad - c \left(e^{2\alpha t} \log \left(\prod_{i=1}^d X_t^i \right)^2 + \sigma^2 e^{2\alpha t} + (h_t^i)^2 + c_t^2 + (e_t^i)^2 \right) \right) dt \\
X_0^i = \xi \\
dY_t = \left(\phi(t, X_t) + b(Y_t + \frac{1}{d} \sum_{i=1}^d Z_t^i + \frac{1}{d} \sum_{i=1}^d \mathbb{E}[X_t^i] + \mathbb{E}[Y_t]) \right. \\
\quad + \frac{b}{d} \sum_{i=1}^d \mathbb{E}[Z_t^i] - b \left(e^{\alpha t} \log \left(\prod_{i=1}^d X_t^i \right) + \frac{1}{d} \sum_{i=1}^d \sigma e^{\alpha t} \right) + c \mathbb{E}[Y_t^2] \\
\quad - b \left(\frac{1}{d} \sum_{i=1}^d h_t^i + c_t + \frac{1}{d} \sum_{i=1}^d e_t^i \right) + c(Y_t^2 + \frac{1}{d} \sum_{i=1}^d (Z_t^i)^2 + \frac{1}{d} \sum_{i=1}^d \mathbb{E}[(X_t^i)^2]) \\
\quad + \frac{c}{d} \sum_{i=1}^d \mathbb{E}[(Z_t^i)^2]) - c \left(e^{2\alpha t} \log \left(\prod_{i=1}^d X_t^i \right)^2 + \frac{1}{d} \sum_{i=1}^d \sigma^2 e^{2\alpha t} \right) \\
\quad \left. - c \left(\frac{1}{d} \sum_{i=1}^d (h_t^i)^2 + c_t^2 + \frac{1}{d} \sum_{i=1}^d (e_t^i)^2 \right) \right) dt + Z_t dW_t \\
Y_T = e^{\alpha T} \log \left(\prod_{i=1}^d X_T^i \right).
\end{array} \right. \quad (22)$$




Figure 10: Learning curves for Global Direct (left) and Global Dynamic (right) methods on the fully coupled quadratic model (22). The loss is the L^2 error between Y_T and the terminal condition of the backward equation.

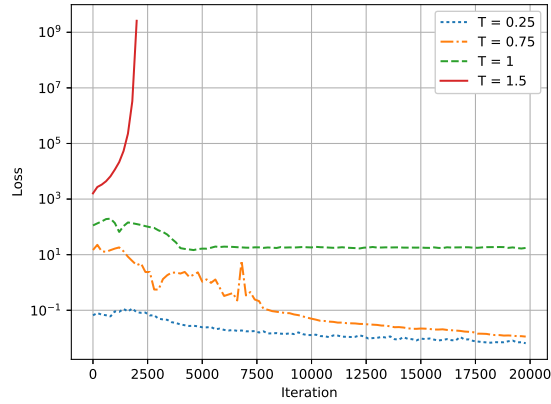


Figure 11: Learning curves for the Local Dynamic methods on the fully coupled quadratic model (22). The loss is the sum of the local L^2 errors between the neural network Y and the Euler discretization for all time steps.

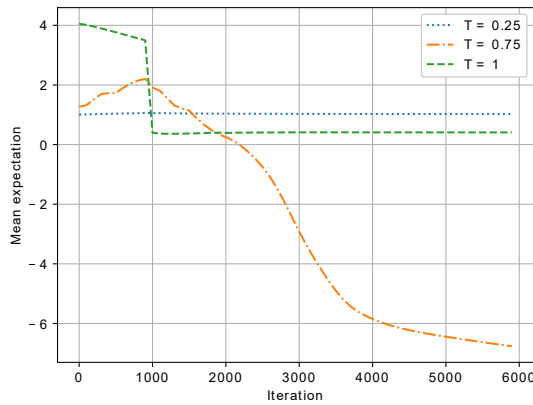


Figure 12: $\mathbb{E}[X_T]$ for Local Dynamic method on the fully coupled quadratic model (22).




Figure 13: First coordinate of Z_t evaluated on a sample path for Global Direct (left) ($T = 0.75$) and Local Dynamic methods (right) ($T = 0.25$) methods after 2000 iterations (respectively 20000) on the fully coupled quadratic model (22).




Figure 14: First coordinate of Y_t evaluated on a sample path for Global Dynamic (left) method ($T = 0.75$) and Local Dynamic method (right) ($T = 0.25$) after 2000 iterations (respectively 20000 for the local method) on the fully coupled quadratic model (22).

We take $a = b = c = 0.1$, $\alpha = 0.5$, $\sigma = 0.4$, $\xi = 1$. We observe in Table 6 convergence of the methods for small maturities and divergence beyond $T = 1$. Note that the dynamic estimation of the expectation prevents the algorithm to explode for $T = 1$, contrarily to the Global Direct method. However, it does not converge to the true solution in this case. Indeed the loss plateaus at the value 2 in Figure 10 (right), so the terminal condition of the BSDE is not properly respected. The Global Dynamic method also produces better result than the Global Direct one (see Figure 14 and Table 6). We notice from Figure 13 and Figure 14 that the estimated Y, Z processes have the good shape but some errors are still present after convergence.

Concerning the Local Dynamic method, we see in Figure 12 that the estimated expectations are stable around zero for a few iterations but then become negative. It may be due to the lack of a contraction for the fixed point problem. The loss explodes for $T = 1.5$, as seen on the learning curve from Figure 11. For $T = 1$, we see that it stays above 10.

5 Conclusion

We have shown that neural network methods can solve some moderate dimensional FBSDE of McKean-Vlasov type. Comparing the different algorithms we find out that:

- The dynamic update of the expectation is efficient in terms of computation speed (about 30% faster than direct method) and seems to smooth the learning curve.
- For the mean-field games of controls example, the Pontryagin approach performs better than the Weak one for large maturities. On the contrary, the Weak approach is the best for small maturities.
- For the fully coupled linear model we observe no convergence problem whereas for the fully coupled quadratic one we can solve only the equation on a small time horizon. However the Local Dynamic method is not very accurate for larger maturities.
- The Local Dynamic method faces more difficulties for quadratic problems than the global methods do. It also requires more iterations, hence more time, to converge.
- The expectation law estimation method does not work well and requires to empirically choose a proper penalization parameter, which is troublesome.
- The methods can be used in dimension 10, thus applied to more realistic problems than usually. For instance, in the price impact model, the number of dimensions corresponds to the number of assets involved in the trading. Thus, developing methods able to deal with problems in high dimensions can help us to handle large portfolios.

We recommend the use of the global method combined with dynamic moment estimation which offers the best accuracy and training speed among all the tested methods. For linear quadratic mean-field games of controls, it appears to be better to use the Weak approach for small maturities and the Pontryagin method for larger time horizons. The use of a Local Dynamic method is possible but requires too many iterations to converge hence it is not competitive in terms of computation time.

References

- [Aba+16] M. Abadi et al. “TensorFlow: A System for Large-scale Machine Learning”. In: *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*. OSDI’16. Savannah, GA, USA: USENIX Association, 2016, pp. 265–283. ISBN: 978-1-931971-33-1. URL: <http://dl.acm.org/citation.cfm?id=3026877.3026899>.
- [ACD10] Y. Achdou and I. Capuzzo-Dolcetta. “Mean Field Games: Numerical Methods”. In: *SIAM Journal on Numerical Analysis* 48 (Jan. 2010). DOI: 10.1137/090758477.
- [AK20] Y. Achdou and Z. Kobeissi. “Mean Field Games of Controls: Finite Difference Approximations”. In: *arXiv:2003.03968* (2020).
- [ALG19] C. Anil, J. Lucas, and R. Grosse. “Sorting Out Lipschitz Function Approximation”. In: *Proceedings of the 36th ICML*. Ed. by K. Chaudhuri and R. Salakhutdinov. Vol. 97. 2019, pp. 291–301.
- [Ang+19] A. Angiuli et al. “Cemracs 2017: numerical probabilistic approach to MFG”. In: *ESAIM: Proceedings and Surveys* 65 (Jan. 2019), pp. 84–113. DOI: 10.1051/proc/201965084.
- [Bac+21] A. Bachouch et al. “Deep neural networks algorithms for stochastic control problems on finite horizon: numerical computations”. In: *Methodol. Comput. Appl. Probab* (2021).

- [Bec+19] C. Beck et al. “Deep splitting method for parabolic PDEs”. In: *arXiv preprint arXiv:1907.03452* (2019).
- [BEJ19] C. Beck, W. E, and A. Jentzen. “Machine Learning Approximation Algorithms for High-Dimensional Fully Nonlinear Partial Differential Equations and Second-order Backward Stochastic Differential Equations”. In: *J. Nonlinear Sci.* 29.4 (Aug. 2019), pp. 1563–1619. ISSN: 1432-1467. DOI: 10.1007/s00332-018-9525-3.
- [BT04] B. Bouchard and N. Touzi. “Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations”. In: *Stochastic Process. Appl.* 111.2 (2004), pp. 175–206.
- [CCD15] J.-F. Chassagneux, D. Crisan, and F. Delarue. “A probabilistic approach to classical solutions of the master equation for large population equilibria”. In: *to appear in Memoirs of the AMS* (2015).
- [CCD19] J.-F. Chassagneux, D. Crisan, and F. Delarue. “Numerical method for FBSDEs of McKean-Vlasov type”. In: *Annals of Applied Probability* 29 (2019). DOI: 10.1214/18-AAP1429.
- [CD13] R. Carmona and F. Delarue. “Probabilistic Analysis of Mean-Field Games”. In: *SIAM Journal on Control and Optimization* 51.4 (2013), pp. 2705–2734. DOI: 10.1137/120883499. URL: <https://doi.org/10.1137/120883499>.
- [CD18] R. Carmona and F. Delarue. *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer, 2018.
- [CL15] R. Carmona and D. Lacker. “A probabilistic weak formulation of mean field games and applications”. In: *Annals of Applied Probability* 25.3 (June 2015), pp. 1189–1231. DOI: 10.1214/14-AAP1020. URL: <https://doi.org/10.1214/14-AAP1020>.
- [CL18] P. Cardaliaguet and C.-A. Lehalle. “Mean field game of controls and an application to trade crowding”. In: *Mathematics and Financial Economics* 12.3 (2018), pp. 335–363. DOI: <https://doi.org/10.1007/s11579-017-0206-z>.
- [CL19] 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”. In: *arXiv preprint arXiv:1908.01613, to appear in The Annals of Applied Probability* (Aug. 2019).
- [CWNMW19] Q. Chan-Wai-Nam, J. Mikael, and X. Warin. “Machine Learning for Semi Linear PDEs”. In: *Journal of Scientific Computing* 79 (Feb. 2019), 1667–1712. DOI: 10.1007/s10915-019-00908-3.
- [FZ20] J.-P. Fouque and Z. Zhang. “Deep Learning Methods for Mean Field Control Problems with Delay”. In: *Frontiers in Applied Mathematics and Statistics* 6 (2020). DOI: <https://doi.org/10.3389/fams.2020.00011>.
- [GLW05] E. Gobet, J-P. Lemor, and X. Warin. “A regression-based Monte Carlo method to solve backward stochastic differential equations”. In: *Ann. Appl. Probab.* 15.3 (2005), pp. 2172–2202.

- [HJE17] J. Han, A. Jentzen, and W. E. “Solving high-dimensional partial differential equations using deep learning”. In: *Proceedings of the National Academy of Sciences* 115 (July 2017). DOI: 10.1073/pnas.1718942115.
- [HL20] J. Han and J. Long. “Convergence of the Deep BSDE Method for Coupled FBSDEs”. In: *Probability, Uncertainty and Quantitative Risk* 5.1 (2020), pp. 1–33. DOI: <https://doi.org/10.1186/s41546-020-00047-w>.
- [HPW20] C. Huré, H. Pham, and X. Warin. “Deep backward schemes for high-dimensional nonlinear PDEs”. In: *Mathematics of Computation* 89.324 (2020), pp. 1547–1579. DOI: <https://doi.org/10.1090/mcom/3514>.
- [Hur+21] C. Huré et al. “Deep neural networks algorithms for stochastic control problems on finite horizon: convergence analysis”. In: *SIAM J. Numer. Anal.* 59.1 (2021), 525–557.
- [Ji+20] S. Ji et al. “Three Algorithms for Solving High-Dimensional Fully Coupled FBSDEs Through Deep Learning”. In: *IEEE Intelligent Systems* 35.3 (2020), pp. 71–84. DOI: 10.1109/MIS.2020.2971597.
- [KB14] D. Kingma and J. Ba. “Adam: A Method for Stochastic Optimization”. In: *International Conference on Learning Representations* (Dec. 2014).
- [Lau21] M. Lauriere. “Numerical Methods for Mean Field Games and Mean Field Type Control”. In: *arXiv:2106.06231* (2021).
- [LL06a] J.-M. Lasry and P.-L. Lions. “Jeux à champ moyen. I – Le cas stationnaire”. In: *Comptes Rendus Mathématique - C R MATH* 343 (Nov. 2006), pp. 619–625. DOI: 10.1016/j.crma.2006.09.019.
- [LL06b] J.-M. Lasry and P.-L. Lions. “Jeux à champ moyen. II – Horizon fini et contrôle optimal”. In: *Comptes Rendus. Mathématique. Académie des Sciences, Paris* 10 (Nov. 2006). DOI: 10.1016/j.crma.2006.09.018.
- [PWG21] H. Pham, X. Warin, and M. Germain. “Neural networks-based backward scheme for fully nonlinear PDEs”. In: *SN Partial Differential Equations and Applications* 2.16 (2021).
- [SDB18] A. Sergeev and M. Del Balso. *Horovod: fast and easy distributed deep learning in TensorFlow*. Feb. 2018.