

CONVERGENCE OF A ROBUST DEEP FBSDE METHOD FOR STOCHASTIC CONTROL*

KRISTOFFER ANDERSSON[†], ADAM ANDERSSON[‡], AND C. W. OOSTERLEE[§]

Abstract. In this paper, we propose a deep learning based numerical scheme for strongly coupled forward backward stochastic differential equations (FBSDEs), stemming from stochastic control. It is a modification of the deep BSDE method in which the initial value to the backward equation is not a free parameter, and with a new loss function being the weighted sum of the cost of the control problem, and a variance term which coincides with the mean squared error in the terminal condition. We show by a numerical example that a direct extension of the classical deep BSDE method to FBSDEs fails for a simple linear-quadratic control problem, and we motivate why the new method works. Under regularity and boundedness assumptions on the exact controls of time continuous and time discrete control problems, we provide an error analysis for our method. We show empirically that the method converges for three different problems, one being the one that failed for a direct extension of the deep BSDE method.

Key words. FBSDE, deep learning, stochastic control

MSC codes. 49M25, 60H30, 60H35, 65C30

DOI. 10.1137/22M1478057

1. Introduction. Forward backward stochastic differential equations (FBSDEs) constitute an important family of models with many applications in a wide variety of fields such as finance, physics, chemistry, and engineering. As the name suggests, an FBSDE consists of two stochastic differential equations (SDEs), one forward SDE and one backward SDE, commonly referred to as the forward equation and the backward equation, respectively. The forward equation is a classical SDE with a given initial value, while the backward equation has a given stochastic terminal value and the initial value is part of the solution. In this paper, we are concerned with stochastic control problems, which typically lead to coupled FBSDEs, meaning that the primary stochastic variables in the backward SDE impact the forward SDE and vice versa. Since closed form solutions to FBSDEs are rare, one often has to rely on numerical approximations.

In this paper, we propose a method which falls into the rapidly growing category of neural network based approximation schemes for FBSDEs and PDEs. Even though there are a number of works in this direction, the methods stem from the pioneering work [31] and are similar in spirit. In the present paper, nonconvergence of the method from [31], originally proposed for noncoupled FBSDEs, is identified when applied to

* Submitted to the journal's Methods and Algorithms for Scientific Computing section February 14, 2022; accepted for publication (in revised form) October 6, 2022; published electronically February 27, 2023.

<https://doi.org/10.1137/22M1478057>

Funding: The work of the first and third authors was supported by the European Union under the H2020-EU.1.3.1. MSCA-ITN-2018 scheme, grant 813261.

[†] Research Group of Scientific Computing, Centrum Wiskunde & Informatica, Amsterdam, Netherlands (Kristoffer.Andersson@CWI.nl).

[‡] Research Group of Computational Mathematics, Chalmers University of Technology and the University of Gothenburg, and Saab AB Radar Solutions, Gothenburg, Sweden (adam.andersson@chalmers.se).

[§] Mathematical Institute, Utrecht University, Utrecht, Netherlands (c.w.oosterlee@uu.nl).

strongly coupled FBSDEs. We present a new family of methods and prove analytically and numerically that it overcomes the convergence problem.

Prior to the recent surge in machine learning based algorithms, the task of approximating FBSDEs was an active field of research for several decades. From a general perspective, approximation schemes can be categorized into *backward* and *forward* numerical methods, referring to the order in time in which the methods operate.

A backward numerical method usually relies on an initialization of the backward equation at the known terminal value (or an approximation of the terminal value). The solution is then approximated recursively, backward in time, by approximating conditional expectations. There are several different methods to approximate these conditional expectations such as tree based methods (see [57, 48]), Fourier based methods (see, e.g., [55, 34, 56]), and least-squares Monte Carlo (LSMC) methods (see, e.g., [13, 16, 22, 9, 18, 28]). A general property of backward methods is that the terminal condition of the backward equation depends on the realization of the forward equation. This is not a problem for *decoupled* FBSDEs, but for coupled FBSDEs the method becomes implicit and iterative schemes which may not always converge need to be employed. A second class of algorithms, which is more suitable for coupled FBSDEs, is the class of forward methods such as PDE based methods (see, e.g., [49, 26]) and Picard linearization schemes (see, e.g., [27, 37, 8]). For a summary on forward and backward numerical methods for FBSDEs, we refer to [17]. A typical drawback for the methods mentioned above is that they suffer from the curse of dimensionality (some methods such as LSMC and Picard schemes may overcome this problem to some extent), meaning that the time complexity and the memory requirements increase exponentially with the dimensions of the problem.

In addition to the classical methods described above, a new branch of algorithms based on neural networks has appeared in recent years. In [31], the authors present a method called the *deep BSDE method*, which relies on a neural network parametrization of the control process and the initial condition of the backward equation. Both the forward and the backward equations are then treated as forward equations and are approximated with the Euler–Maruyama scheme. To achieve an accurate approximation, the parameters of the neural networks are optimized such that the terminal condition on the backward equation is (approximately) satisfied in the mean squared sense. The method has proven to be able to approximate a wide class of equations in very high dimensions (at least 100). Since the original deep BSDE method publication, several papers with adjustments of the algorithm as in, e.g., [5, 7, 4, 24, 53, 40, 25, 41, 58, 33], and others with convergence analysis in, e.g., [32, 38, 10, 21, 29, 39, 42], have been published. In addition to being forward methods, these algorithms are *global* in their approximation, meaning that the optimization of all involved neural networks is carried out simultaneously. This implies that they are optimized subject to one single objective function, also called *loss function*, as it is usually referred to in the machine learning literature. There also exists a branch of neural network based algorithms relying on *local* optimization techniques. Typically, these methods are of backward type and of a similar nature as the LSMC algorithms, but instead of polynomials as their basis functions, neural networks are used. As for the LSMC method, these kinds of algorithms are not easily applied to coupled FBSDEs. Algorithms of this type can be found in, e.g., [36, 35, 15] and with error analysis [23, 51]. For an overview of machine learning algorithms for approximation of PDEs, we refer to [6].

As mentioned, we are interested in coupled FBSDEs and the focus is therefore on global algorithms operating forward in time, with a structure similar to the deep BSDE method. We demonstrate that the approach taken in, e.g., [1, 52, 47, 40, 20], where the deep BSDE method is applied to the FBSDEs associated with the

stochastic control problem, is problematic. As we show in the present paper, even though an accurate approximation of the control problem can be achieved, this does not imply an approximation of the FBSDE in general. Moreover, it is observed that the deep BSDE method does not converge for certain problems and the convergence problem is prominent for the strongly coupled FBSDEs. Our proposed method overcomes this problem by employing the equivalence between the stochastic control problem and the FBSDE. To be more precise, we use the fact that the initial value of the BSDE coincides with the value function of the control problem and hence can be expressed as a minimization problem. Moreover, we use the adaptivity property of the BSDE to conclude that a stochastic version of the value function is \mathcal{F}_0 -measurable and therefore has zero variance. These two properties are then combined in the loss function to achieve a robust approximation scheme for coupled FBSDEs. The effectiveness of our algorithm is demonstrated empirically on a collection of problems with different characteristics. In addition, a theoretical error analysis is carried out, in which we provide convergence rates for the initial and terminal conditions of the FBSDE under a mild assumption and strong convergence of the FBSDE under stronger assumptions. Our main result is similar to the a posteriori error bound for the deep BSDE method which was established for weakly coupled FBSDEs in [32] and later extended to non-Lipschitz coefficients (but for less general diffusion coefficients) in [42]. However, these results are unlikely to be valid for strongly coupled FBSDEs and we find several examples in which the discrete terminal condition converges while the FBSDE approximation does not.

This paper is structured as follows. In section 2, we present the stochastic control problem and explain the reformulations to a PDE as well as an FBSDE. Moreover, we introduce reformulations of the FBSDE to different variational problems which are used for the algorithms in later sections. The section concludes by introducing the time discrete counterparts of the reformulations as well as a discussion on when and why the deep BSDE method fails to converge. In section 3, the fully implementable algorithms are presented together with details on the neural networks used. Section 4 is devoted to error analysis of the proposed algorithm. Classical Euler–Maruyama type discretization errors and errors stemming from differences between time discrete and time continuous stochastic control are discussed. Finally, in section 5 numerical approximations are compared with their analytic counterparts when we have such available.

2. The deep FBSDE method and an improved family of methods. This section contains a formal introduction to our proposed method with motivation from stochastic control and FBSDE theory. Section 2.1 introduces the stochastic control problem, the related Hamilton–Jacobi–Bellman (HJB) equation, and the FBSDE. Alternative formulations of the FBSDE are presented in section 2.2. In section 2.3, we motivate by numerical examples why a direct extension of the deep BSDE method to FBSDEs, as in [1, 52, 47], fails for many problems. Finally, in section 2.4, the proposed robust deep FBSDE method is described. In this section, we present formally the method for the sake of clarity, while in section 4, a more rigorous approach is taken.

2.1. Stochastic control and FBSDEs. Our starting point is a controlled SDE and its associated cost functional

$$(2.1) \quad X_t^u = x_0 + \int_0^t \bar{b}(s, X_s^u, u_s) ds + \int_0^t \sigma(s, X_s^u) dW_s,$$

$$(2.2) \quad J^u(t, x) = \mathbb{E} \left[\int_t^T \bar{f}(s, X_s^u, u_s) ds + g(X_T^u) \mid X_t^u = x \right], \quad t \in [0, T].$$

Here $T \in (0, \infty)$, $d, k, \ell \in \mathbb{N}$, $(W_t)_{t \in [0, T]}$ is a k -dimensional standard Brownian motion, the coefficients $\bar{b}: [0, T] \times \mathbb{R}^d \times \mathbb{R}^\ell \rightarrow \mathbb{R}^d$, $\sigma: [0, T] \times \mathbb{R}^d \times \mathbb{R}^\ell \rightarrow \mathbb{R}^{d \times k}$, $\bar{f}: [0, T] \times \mathbb{R}^d \times \mathbb{R}^\ell \rightarrow \mathbb{R}$, and $g: \mathbb{R}^d \rightarrow \mathbb{R}$ satisfy some extra regularity conditions, and the control process $u = (u_t)_{t \in [0, T]}$ belongs to a set \mathcal{U} of admissible controls, taking values in a set $U \subset \mathbb{R}^\ell$. The aim is to find a control process $u^* \in \mathcal{U}$ that minimizes $J^u(0, x_0)$. Assuming the cost to be bounded from below, the *value function* of the control problem is given by $V(t, x) = \inf_{u \in \mathcal{U}} J^u(t, x)$. For the presentation, we assume uniqueness of the infimum. Under appropriate conditions, the value function satisfies an HJB equation, which is a nonlinear parabolic PDE given by

$$(2.3) \quad \begin{cases} \frac{\partial V}{\partial t}(t, x) + \frac{1}{2} \text{Tr}(\sigma \sigma^\top D_x^2 V)(t, x) + \mathcal{H}(t, x, D_x V(t, x)) = 0, & (t, x) \in [0, T) \times \mathbb{R}^d, \\ V(t, x) = g(x), & (t, x) \in \{T\} \times \mathbb{R}^d. \end{cases}$$

Here Tr denotes the trace of a matrix and the *Hamiltonian*, \mathcal{H} , is given by

$$(2.4) \quad \mathcal{H}(t, x, p) = \inf_{v \in U} [\bar{b}(t, x, v)^\top p + \bar{f}(t, x, v)], \quad (t, x) \in [0, T] \times \mathbb{R}^d, \quad p \in \mathbb{R}^d.$$

Under conditions that guarantee a sufficiently regular solution to (2.3), and the infimum in the Hamiltonian to be attained at $v^* = v^*(t, x, p)$, the optimal control is of the feedback form, $u^*(t, X_t) = v^*(t, X_t, D_x V(t, X_t))$, where we have written $X := X^{u^*}$ for the optimally controlled X^u . The feedback map v^* is for many interesting problems easy to derive. Again, under sufficient regularity, Itô's formula applied to $V(t, X_t)$ yields

$$(2.5) \quad \begin{cases} X_t = x_0 + \int_0^t \bar{b}(s, X_s, Z_s) ds + \int_0^t \sigma(s, X_s) dW_s, \\ Y_t = g(X_T) + \int_t^T \bar{f}(s, X_s, Z_s) ds - \int_t^T \langle Z_s, dW_s \rangle, \quad t \in [0, T], \end{cases}$$

where $Y_t = V(t, X_t)$, $Z_t = \sigma^\top(t, X_t) D_x V(t, X_t)$ and, for $\theta \in \{b, f\}$, we have $\theta(t, X_t, Z_t) := \bar{\theta}(t, X_t, v^*(t, X_t, (\sigma(t, X_t) \sigma^\top(t, X_t))^{-1} \sigma(t, X_t) Z_t))$.

In the rest of this section, we assume the existence of a unique solution (X, Y, Z) of (2.5) in appropriate spaces. Given Z or equivalently $D_x V$, we thus have an optimal control $u_t^* = v^*(t, X_t, (\sigma(t, X_t) \sigma^\top(t, X_t))^{-1} \sigma(t, X_t) Z_t)$. This would make efficient numerical FBSDE schemes very useful for solving the control problem. In the other direction, if we have an optimal control u^* , then in general this does not give us Z , unless $p \mapsto v^*(t, x, p)$ is invertible, and only in this case the control problem naturally suggests numerical schemes for FBSDEs. Below, we introduce a family of numerical schemes for FBSDEs that works regardless of invertibility of the feedback map, but reduces to the control problem in the case of invertibility.

2.2. Alternative formulations of FBSDEs. The deep BSDE method proposed in [31] relies on a reformulation of the FBSDE (2.5) into two forward SDEs, one with a priori unknown initial value. It relies moreover on the Markov property of the FBSDE, which guarantees that $Z_t = \zeta^*(t, X_t)$, for some function $\zeta^*: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$, that we refer to as a Markov map, and optimization is done with respect to such functions and initial values y_0 . More precisely, the FBSDE (2.5) is reformulated into the following variational problem:

$$(2.6) \quad \begin{cases} \text{minimize}_{y_0, \zeta} \mathbb{E}|Y_T^{y_0, \zeta} - g(X_T^{y_0, \zeta})|^2, & \text{where} \\ X_t^{y_0, \zeta} = x_0 + \int_0^t b(s, X_s^{y_0, \zeta}, Z_s^{y_0, \zeta})ds + \int_0^t \sigma(s, X_s^{y_0, \zeta})dW_s, \\ Y_t^{y_0, \zeta} = y_0 - \int_0^t f(s, X_s^{y_0, \zeta}, Z_s^{y_0, \zeta})ds + \int_0^t \langle Z_s^{y_0, \zeta}, dW_s \rangle, \\ Z_t^{y_0, \zeta} = \zeta(t, X_t^{y_0, \zeta}), \quad t \in [0, T], \end{cases}$$

where y_0 and ζ are sought in appropriate spaces. From the theory outlined in section 2.1, under well-posedness and sufficient regularity of (2.5), it holds that $Y_0 = V(0, x_0)$ and $\zeta^* = \sigma^\top D_x V$, and thus we have well-posedness of (2.6). While it seems natural to propose a numerical algorithm based on a discrete version of (2.6), we demonstrate below that such an optimization problem, even for many simple problems, does not converge.

In order to introduce numerical schemes that do not suffer under the above problem, we use the following two properties of the initial value Y_0 of (2.5):

- (i) Y_0 coincides with the control problem's value function (property from the control problem);
- (ii) Y_0 is \mathcal{F}_0 -measurable and therefore has zero variance (property from the FB-SDE).

The two properties are both captured in the following variational problem:

$$(2.7) \quad \begin{cases} \text{minimize}_{\zeta} \Phi_\lambda(\zeta) = \mathbb{E}[\mathcal{Y}_0^\zeta] + \lambda \text{Var}[\mathcal{Y}_0^\zeta], & \text{where} \\ \mathcal{Y}_0^\zeta = g(X_T^\zeta) + \int_0^T f(t, X_t^\zeta, Z_t^\zeta)dt - \int_0^T \langle Z_t^\zeta, dW_t \rangle, \\ X_t^\zeta = x_0 + \int_0^t b(s, X_s^\zeta, Z_s^\zeta)ds + \int_0^t \sigma(s, X_s^\zeta)dW_s, \\ Y_t^\zeta = \mathbb{E}[\mathcal{Y}_0^\zeta] - \int_0^t f(s, X_s^\zeta, Z_s^\zeta)ds + \int_0^t \langle Z_s^\zeta, dW_s \rangle, \quad Z_t^\zeta = \zeta(t, X_t^\zeta), \quad t \in [0, T]. \end{cases}$$

We refer to \mathcal{Y}_0^ζ as the *stochastic cost* and notice that $\mathbb{E}[\mathcal{Y}_0^\zeta] = J^{u(\zeta)}(0, x_0)$, where $u(\zeta) \in \mathcal{U}$ is the control generated by ζ . Thus, the first term of the objective function Φ_λ is the cost function of the control problem. In the case of $p \mapsto v^*(t, x, p)$ being invertible, this term alone, i.e., for $\lambda = 0$, offers an equivalent formulation to (2.6). In other cases, uniqueness of minimizers of $\zeta \mapsto \mathbb{E}[\mathcal{Y}_0^\zeta]$ cannot be guaranteed. But among the minimizers, there is only one ζ^* with the property that the variance of the stochastic cost \mathcal{Y}_0^ζ equals zero. The second term of Φ_λ is introduced to penalize nonzero variance and the minimizer for $\lambda > 0$ is unique. Another important feature of the formulation (2.7) is that Y_0^ζ is determined by ζ alone and (2.7) has thus one degree of freedom less than (2.6). A final observation is that

$$(2.8) \quad \begin{aligned} \text{Var}[\mathcal{Y}_0^\zeta] &= \mathbb{E}[|\mathbb{E}[\mathcal{Y}_0^\zeta] - \mathcal{Y}_0^\zeta|^2] \\ &= \mathbb{E}\left[\left|Y_0^\zeta - \int_0^T f(s, X_s^\zeta, Z_s^\zeta)ds + \int_0^T Z_s^\zeta dW_s - g(X_T^\zeta)\right|^2\right] = \mathbb{E}[|Y_T^\zeta - g(X_T^\zeta)|^2]. \end{aligned}$$

This implies that the second term of Φ is, up to the factor λ , the same as that of (2.6), but with Y_0^ζ not being a variable to optimize. Thus there are strong similarities between (2.6) and (2.7), but in the time discrete setting the latter formulation is shown to be advantageous in the sections below.

2.3. A direct extension of the deep BSDE method and why it fails.

In this section, we present the time discrete counterpart of (2.6). We assume an equidistant time grid, $0 = t_0 < t_1 < \dots < t_N = T$, with $h = t_{n+1} - t_n$ and denote the

Brownian increment, $\Delta W_n = W_{n+1} - W_n$. Throughout the paper, we parameterize discretizations by $h \in (0, 1)$ and by this we mean all $h \in (0, 1) \cap \{T/N : N \geq 1\}$.

The time discrete version of (2.6) is given by

$$(2.9) \quad \begin{cases} \text{minimize } \mathbb{E}_{y_0, \zeta} \left[|Y_N^{h, y_0, \zeta} - g(X_N^{h, y_0, \zeta})|^2 \right], & \text{where} \\ X_n^{h, y_0, \zeta} = x_0 + \sum_{k=0}^{n-1} b(t_k, X_k^{h, y_0, \zeta}, Z_k^{h, y_0, \zeta})h + \sum_{k=0}^{n-1} \sigma(t_k, X_k^{h, y_0, \zeta})\Delta W_k, \\ Y_n^{h, y_0, \zeta} = y_0 - \sum_{k=0}^{n-1} f(t_k, X_k^{h, y_0, \zeta}, Z_k^{h, y_0, \zeta})h + \sum_{k=0}^{n-1} \langle Z_k^{h, y_0, \zeta}, \Delta W_k \rangle, \\ Z_k^{h, y_0, \zeta} = \zeta_k(X_k^{h, y_0, \zeta}). \end{cases}$$

It is a direct extension of the deep BSDE method from [31]. In the literature, it was first applied experimentally to FBSDEs in the masters thesis [1] and thereafter in [52], both for inverted pendulums, and in [47] for an application to attitude control of unmanned aerial vehicles. More examples of implementations of the deep FBSDE method are found in [40, 20].

In [32], the authors consider FBSDEs with coefficients b, σ , and f that may take the Y -component, but not the Z -component, as arguments. Under the relatively strict assumption of weak coupling (also called the monotonicity condition; see, e.g., [2]), it is shown that for h small enough there is a constant C , independent of h , such that $\sup_{t \in [0, T]} (\mathbb{E}[|X_t - \hat{X}_t^{h, y_0, \zeta}|^2] + \mathbb{E}[|Y_t - \hat{Y}_t^{h, y_0, \zeta}|^2]) + \int_0^T \mathbb{E}[|Z_t - \hat{Z}_t^{h, y_0, \zeta}|^2] dt \leq C(h + \mathbb{E}[|Y_N^{h, y_0, \zeta} - g(X_N^{h, y_0, \zeta})|^2])$, where for $t \in [t_k, t_{k+1})$, $\hat{X}_t^{h, y_0, \zeta} := X_k^{h, y_0, \zeta}$, $\hat{Y}_t^{h, y_0, \zeta} := Y_k^{h, y_0, \zeta}$, and $\hat{Z}_t^{h, y_0, \zeta} = \zeta_k(X_k^{h, y_0, \zeta})$. Under some additional assumptions on the coefficients b, f, σ , and g (additional smoothness and boundedness of the coefficients to guarantee a bounded and smooth solution of the associated HJB equation), the results in [32] can be extended to the framework of interest in this paper, i.e., coefficients taking the Z -component as an argument. On the other hand, the weak coupling condition is rarely satisfied for FBSDEs stemming from stochastic control problems, and to the best of our knowledge, there is no known way to relax this condition.

To investigate convergence of (2.9) empirically, we first note that if we would know Y_0 a priori, then the variational problem (2.6) would be reduced to finding ζ . We also know that $D_x V^\top \sigma$ is the minimizer, which would make the objective function identical to zero. In the discrete counterpart, we would expect that if (2.9) converges to (2.6), then for sufficiently small h , the objective function would be close to zero if optimizing only ζ and setting $y_0 = Y_0$. Moreover, for a robust algorithm to emerge from (2.9), it is important that $y_0 \neq Y_0$ results in a larger value of the optimal objective function, at least when y_0 and the true initial value, Y_0 , are “far away” from each other. To formalize this, we introduce the mean squared error

$$(2.10) \quad \text{MSE}(y_0) := \inf_{\zeta} \mathbb{E} \left[|Y_N^{h, y_0, \zeta} - g(X_N^{h, y_0, \zeta})|^2 \right].$$

The aim is to investigate whether or not MSE is minimized at, or close to, the true initial condition Y_0 . Moreover, for each y_0 , we want to investigate the Markov map ζ^{y_0} that minimizes $\zeta \mapsto \mathbb{E}[|Y_N^{h, y_0, \zeta} - g(X_N^{h, y_0, \zeta})|^2]$. The discrete costs, associated with (y_0, ζ^{y_0}) and (y_0, ζ) , are given by $J_0^{h, y_0} = J_0^{h, y_0, \zeta^{y_0}}$ and $J_0^{h, y_0, \zeta} = \mathbb{E}[\mathcal{Y}_0^{h, y_0, \zeta}]$. Here, the *discrete stochastic cost* is given by

$$(2.11) \quad \mathcal{Y}_0^{h, y_0, \zeta} = g(X_N^{h, y_0, \zeta}) + \sum_{k=0}^{N-1} f(t_k, X_k^{h, y_0, \zeta}, Z_k^{h, y_0, \zeta})h - \sum_{k=0}^{N-1} \langle Z_k^{h, y_0, \zeta}, \Delta W_k \rangle.$$

Using the stochastic cost, we have by a substitution that

$$(2.12) \quad \text{MSE}(y_0) = \inf_{\zeta} \mathbb{E} [|\mathcal{Y}_0^{h, y_0, \zeta} - y_0|^2].$$




FIG. 1. Demonstration of the performance of the direct extension of the deep BSDE method to FBSDEs corresponding to two LQ control problems. Left: A one-dimensional problem with $N = 10$ time steps. Right: A two-dimensional problem with $N = 100$ time steps.

Figure 1 shows $y_0 \mapsto \text{MSE}(y_0)$ and $y_0 \mapsto J_0^{h,y_0}$ for two different linear-quadratic (LQ) control problems, respectively, a one-dimensional and a two-dimensional problem. The left figure shows that there is a minimum of MSE at the correct Y_0 and for this problem the method converges. For the right figure, it is clear that there is no minimum of MSE around Y_0 , or anywhere in the range. When y_0 and ζ are jointly optimized, the method has no chance to converge for this problem.

Both problems considered in this section are of the form (5.1) with parameters as in section 5.1.1 for the two-dimensional problem and $A = B = C = R_x = R_u = G = 1$ and $\sigma = 0.5$ for the one-dimensional problem.

We identify three distinct cases from Figure 1:

- $y_0 \approx Y_0$. In this case, MSE is the mean squared error of the discretized FBSDE and it is most reasonable that an approximation of (X, Y, Z) is obtained by optimizing over ζ .
- $y_0 > Y_0$. Under this assumption, any ζ attaining the infimum in (2.10) satisfies $J_0^{h,y_0} = y_0$. Thus minimizing (2.10) is the same as finding the ζ that generates discrete cost y_0 and that at the same time minimizes the mean squared error in the terminal condition. A ζ with cost $y_0 > Y_0$ has the possibility to generate a lower $\text{MSE}(y_0) < \text{MSE}(Y_0)$, depending on y_0 and the problem at hand. This is possible only if the coupling of Z in b is strong enough, so that $g(X_T)$ can be efficiently controlled by Z . In the case with no coupling, i.e., for a BSDE, $y_0 > Y_0$ leads to an MSE of the magnitude $y_0 - Y_0$, and thus MSE is increasing in this regime. This is the reason why noncoupled BSDEs, as in [31], or weakly coupled FBSDEs, as in [32], can be approximated with the deep BSDE method. The left plot of Figure 1 shows this favorable behavior, while the right plot has a decreasing MSE and has no chance to converge. We also see that the cost increases linearly for $y_0 > Y_0$ according to $J_0^{h,y_0} = y_0$.
- $y_0 < Y_0$. Since Y_0 is (approximately) a lower bound of $(y_0, \zeta) \mapsto J_0^{h,y_0,\zeta}$, it holds that any ζ attaining the infimum in $\zeta \mapsto \mathbb{E}[|Y_N^{h,y_0,\zeta} - g(X_N^{h,y_0,\zeta})|^2]$ also minimizes the cost functional $\zeta \mapsto J_0^{h,y_0,\zeta}$. But y_0 does not enter $\mathcal{Y}_0^{h,y_0,\zeta}$ explicitly. Therefore, the minimizer of $\zeta \mapsto \mathbb{E}[\mathcal{Y}_0^{h,y_0,\zeta}]$ does not depend on y_0 . Thus, fixing $y_0 < Y_0$ and optimizing ζ approximates a solution to the control problem but not to the FBSDE. This can clearly be seen in Figure 1 from the cost being constant for $y_0 < Y_0$ in both plots. It is also clear that the MSE increases for decreasing $y_0 < Y_0$.

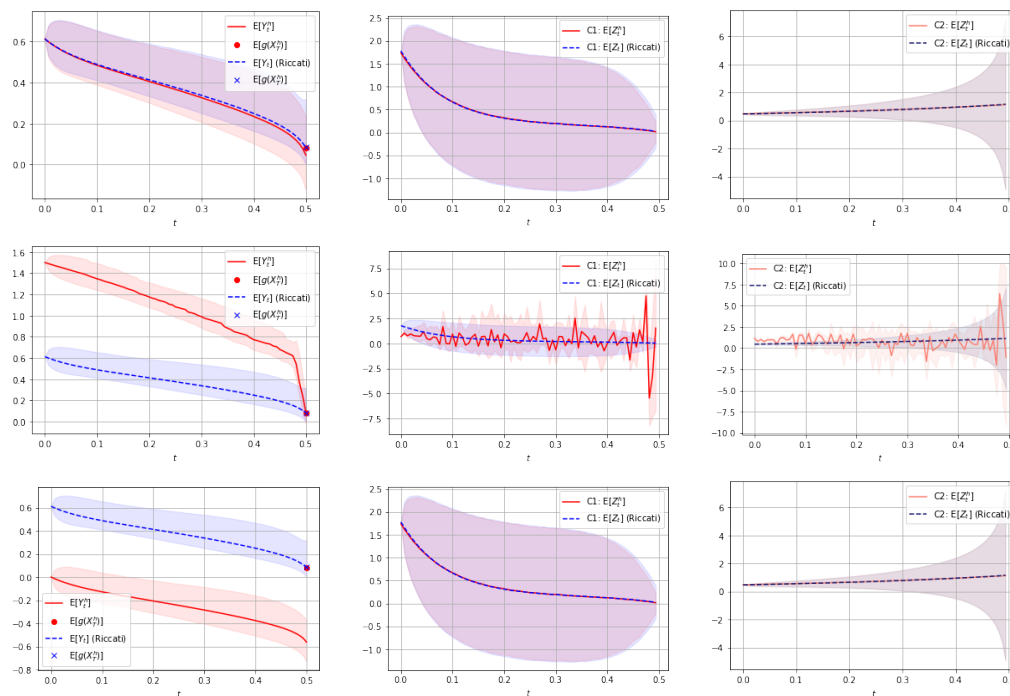


FIG. 2. Demonstration of the performance of the deep FBSDE method with fixed initial condition for $N = 100$ time steps. The shaded areas represent the domain in which 90% of all trajectories lie (the area is bounded by the 5th and the 95th empirical percentile). Left to right: Sample means of the approximate and the semianalytic (Riccati solutions) Y -process and the first and second components of the Z -process (C1 and C2). Top to bottom: Initial conditions $\hat{y}_0 = 0.0 < Y_0$, $\hat{y}_0 = 0.612 = Y_0$, and $\hat{y}_0 = 1.5 > Y_0$.

To further visualize the three cases, Figure 2 shows the empirical means and 90% credible regions (defined as the area between the 5th and the 95th empirical percentiles at each time point) for the true and approximated Y - and Z -processes of the two-dimensional LQ control problem discussed above. In the top row, we see that, in the case $y_0 \approx Y_0$, the two components of the Z -process are very well approximated, but the time discretization error of Y is visible. In the middle row, for $y_0 > Y_0$, we see what is expected based on the discussion above. The Y -process satisfies the terminal value but is otherwise fundamentally distinct from the true Y , and Z is different and oscillating (it is specified to have cost y_0). In the bottom row, the case $y_0 < Y_0$ is shown. Just as explained above, the control problem is solved and therefore the Z -process is well approximated. It should be noted, however, that this is only true since the map $p \mapsto v^*(t, x, p)$, for this specific problem, is invertible. Otherwise, one optimal ζ , in a set of many optimal Markov maps, is approximated. Thus, the control problem is approximately solved; however, the Z -component of the FBSDE is unlikely to be accurate. The Y -process is shifted by $y_0 - Y_0$ and the terminal value is naturally not satisfied.

2.4. A robust deep FBSDE method. Having observed the problems with the direct extension of the deep BSDE method to FBSDEs, we here discretize the alternative formulation (2.7) of the FBSDE to obtain an alternative family of deep FBSDE methods. It reads

(2.13)

$$\begin{cases} \text{minimize } \Phi_{\lambda,h}(\zeta) = \mathbb{E}[\mathcal{Y}_0^{h,\zeta}] + \lambda \text{Var}[\mathcal{Y}_0^{h,\zeta}], & \text{where} \\ \mathcal{Y}_0^{h,\zeta} := g(X_N^{h,\zeta}) + \sum_{k=0}^{N-1} f(t_k, X_k^{h,\zeta}, Z_k^{h,\zeta})h - \sum_{k=0}^{N-1} \langle Z_k^{h,\zeta}, \Delta W_k \rangle, \\ X_n^{h,\zeta} = x_0 + \sum_{k=0}^{n-1} b(t_k, X_k^{h,\zeta}, Z_k^{h,\zeta})h + \sum_{k=0}^{n-1} \sigma(t_k, X_k^{h,\zeta})\Delta W_k, \\ Y_n^{h,\zeta} = \mathbb{E}[\mathcal{Y}_0^{h,\zeta}] - \sum_{k=0}^{n-1} f(t_k, X_k^{h,\zeta}, Z_k^{h,\zeta})h + \sum_{k=0}^{n-1} \langle Z_k^{h,\zeta}, \Delta W_k \rangle, \quad Z_k^{h,\zeta} = \zeta_k(X_k^{h,\zeta}). \end{cases}$$

The main purpose of the current paper is the theoretical and numerical error analysis of (2.13).

2.5. Related methods and comparison. In a comparison with the current literature, we focus on methods for solving stochastic control problems, or the associated FBSDE, by deep learning in a global way, in the sense that only one global optimization problem is solved.

In the early paper [30], time discrete stochastic optimal control problems were solved with deep learning. No explicit connections to FBSDEs, or even to stochastic control in continuous time, were made. The feedback maps for the controls at all time steps were optimized over a family of neural networks, to minimize the discrete cost functional. This methodology is similar to our method when $\lambda = 0$. It only differs in its approximation of the feedback map for u instead of the Markov map for Z . The connection between our proposed method and the deep BSDE method, proposed in [31], is the second term in the loss function of (2.13), corresponding to $\lambda \rightarrow \infty$. From (2.8), it becomes clear that if the driver does not take Y as an input, then this term coincides with the loss function used in the deep BSDE method.

The recent paper [58] is, to our knowledge, the first to introduce the variance penalty term in (2.13) for an FBSDE obtained from the dynamic programming principle. The authors of [59] also use the variance penalty term, but for general decoupled FBSDEs, i.e., not in the context of stochastic control. In [58], the problem is approached differently in that they have one network for the Markov map for Z and one for the feedback map for the control. In [58], the variance term in the loss function is presented as a measurability loss. Their motivation is that if the BSDE is solved, then the initial state of the Y -process is \mathcal{F}_0 -measurable and hence the variance is zero. However, one should bear in mind that this is true only for the continuous BSDE. In the discretized version, the Y -process is not measurable and no arguments for convergence of the discretization are presented in [58]. However, their numerical results are convincing and an error analysis similar to the one presented in section 4 in the present paper could possibly be carried out also in their setting.

Another method which also makes use of the connection to stochastic control (of Hamiltonian systems) was proposed in [41]. In that paper, the stochastic maximum principle approach to stochastic control was used, which results in a different type of FBSDEs, compared to the one obtained with the dynamic programming principle that we consider in this paper. More precisely, it is the Y -process instead of the Z -process which is connected to the control of the SDE. A similarity is that in both papers the algorithms use a method similar to that in [30] to include the cost in the objective function.

Summarizing, only the method in [58] is fully comparable to ours, as it is designed to solve the same problem. It has the variance term but not the mean in its loss

function. By introducing a loss function that includes both we are able to prove convergence and obtain a robust method.

2.6. Decoupled FBSDE and why coupled FBSDEs are important. It is sometimes claimed in the literature that since coupled FBSDEs can be transformed into decoupled BSDE, it is sufficient to have schemes for the latter; see, e.g., [32]. Here, we explain this claim and why we, from a practical and application viewpoint, do not agree. If $\psi: [0, T] \times \mathbb{R}^d \times \mathbb{R}^\ell \rightarrow \mathbb{R}^d$ is sufficiently regular, then it holds by the Itô formula for $Y_t^\psi = V(t, X_t^\psi)$ and $Z_t^\psi = \sigma^\top(t, X_t^\psi) D_x V(t, X_t^\psi)$ that

$$\begin{cases} X_t^\psi = x_0 + \int_0^t (b(s, X_s^\psi, Z_s^\psi) - \psi(s, X_s^\psi, Z_s^\psi)) ds + \int_0^t \sigma(s, X_s^\psi) dW_s, \\ Y_t^\psi = g(X_T^\psi) + \int_t^T (f(s, X_s^\psi, Z_s^\psi) \\ \quad - \langle (\sigma(s, X_s^\psi) \sigma^\top(s, X_s^\psi))^{-1} \sigma(s, X_s^\psi) Z_s, \psi(s, X_s^\psi, Z_s^\psi) \rangle) ds \\ \quad - \int_t^T \langle Z_s^\psi, dW_s \rangle, \quad t \in [0, T]. \end{cases}$$

Thus taking ψ in such a way that $b - \psi$ does not depend on Z decouples the FBSDE, resulting in a BSDE, where the forward equation has no coupling with the backward equation. From this observation, it might be tempting to approximate the optimal Markov map $\sigma^\top(t, x) D_x V(t, x)$ with the deep BSDE method. The problem with this approach is that the deep BSDE method will learn $\sigma^\top(t, x) D_x V(t, x)$ well only around typical trajectories of X^ψ , but not around those of X . While there is empirical evidence that the deep BSDE method overcomes the curse of dimensionality, it does not at all approximate the solution everywhere, but only around the typical solution trajectories of the forward equation. Since X is controlled, it has a different dynamics than X^ψ and this may ruin the applicability of the deep BSDE method for control problems, if the feedback map is desired. If only an approximation of the solution to the HJB equation (the value) is sought, as in [31], then decoupling is feasible.

3. Fully implementable scheme and neural network regression. In this section, we describe how the discrete variational problem (2.13) is approximated with the help of neural network regression. In principle, other function approximators could be used, but neural network regression is arguably one of the most suitable choices due to the ability to approximate complicated high-dimensional functions.

3.1. Fully implementable algorithms. Without further specifications, (2.13) assumes exact optimization over an unspecified set of functions ζ and the exact computation of expectations. To define a fully implementable scheme, the Markov maps $\zeta_0, \dots, \zeta_{N-1}$ are approximated with neural networks $\phi_0^{\theta_0}, \dots, \phi_{N-1}^{\theta_{N-1}}$ with parameters $\theta = (\theta_0, \dots, \theta_{N-1})$ in some parameter space Θ . We specify them in further detail below. Moreover, expectations are approximated with batch Monte Carlo simulation. Let $K_{\text{epochs}} \geq 1, K_{\text{batch}} \geq 1$ be the number of epochs and the number of batches per epoch, respectively. Let further $M_{\text{train}}, M_{\text{batch}} \geq 1$ be the size of the training data set and batch, respectively. We assume that $M_{\text{train}}/(2M_{\text{batch}}) = K_{\text{batch}} \in \mathbb{N}$. Training data are M_{train} independent realizations of the Wiener increments $\Delta W_0, \dots, \Delta W_{N-1} \sim \mathcal{N}(0, h)$ and the training data are reused in K_{epoch} epochs. The training is initialized by random sampling of $\theta^0 \in \Theta$. For each update step in an epoch of the training algorithm, we take $2M_{\text{batch}}$ Wiener increments $\Delta W_0(m), \dots, \Delta W_{N-1}(m)$, $m = 1, 2, \dots, 2M_{\text{batch}}$ from the training data set that were not previously used during the epoch and update θ by approximate optimization of the following problem:

(3.1)

$$\left\{ \begin{array}{l} \text{minimize}_{\theta \in \Theta} \mathcal{L}_{\lambda, h}(\theta) \\ = \frac{1}{M_{\text{batch}}} \left(\sum_{m=1}^{M_{\text{batch}}} \mathcal{Y}_0^{h, \theta}(m) + \lambda \cdot \sum_{m=M_{\text{batch}}+1}^{2M_{\text{batch}}} |g(X_N^{h, \theta}(m)) - Y_N^{h, \theta}(m)|^2 \right), \text{ where} \\ \mathcal{Y}_0^{h, \theta}(m) := g(X_N^{h, \theta}(m)) + \sum_{k=0}^{N-1} f(t_k, X_k^{h, \theta}(m), Z_k^{h, \theta}(m))h \\ \quad - \sum_{k=0}^{N-1} \langle Z_k^{h, \theta}(m), \Delta W_k(m) \rangle, \\ X_n^{h, \theta}(m) = x_0 + \sum_{k=0}^{n-1} b(t_k, X_k^{h, \theta}(m), Z_k^{h, \theta}(m))h + \sum_{k=0}^{n-1} \sigma(t_k, X_k^{h, \theta}(m)) \Delta W_k(m), \\ Y_n^{h, \theta}(m) = \frac{1}{M_{\text{batch}}} \sum_{r=1}^{M_{\text{batch}}} \mathcal{Y}_0^{h, \theta}(r) - \sum_{k=0}^{n-1} f(t_k, X_k^{h, \theta}(m), Z_k^{h, \theta}(m))h \\ \quad + \sum_{k=0}^{n-1} \langle Z_k^{h, \theta}(m), \Delta W_k(m) \rangle, \\ Z_k^{h, \theta}(m) = \phi_k^{\theta_k}(X_k^{h, \theta}(m)). \end{array} \right.$$

When all training data has been used, a new epoch starts. When the K_{epoch} th epoch is finished, the algorithm terminates. The neural network parameters at termination are θ^* . It is an approximation of the parameters θ^{**} that optimize (3.1) in the limit $M_{\text{batch}} \rightarrow \infty$. It should be noted that the expected value of the stochastic sum in $\mathcal{Y}_0^{h, \theta}$ equals zero. Therefore, the algorithm would also work without it, but a practical reason to keep it is that it decreases the variance of $\mathcal{Y}_0^{h, \theta}$ significantly, and this requires fewer Monte Carlo samples to achieve the same accuracy.

3.2. Specification of the neural networks. Here, we introduce the neural networks that we use in our implementations in section 5. The generality is kept to a minimum and more general architectures are of course possible. For each $\phi_k^{\theta_k} : \mathbb{R}^d \rightarrow \mathbb{R}^\ell$, a fully connected neural network with two hidden layers with 20 nodes in each layer and a ReLU activation function $\mathfrak{R}(x) = \max(0, x)$ acting elementwise. More precisely, the $\phi_k^{\theta_k}$ is of the form

$$\phi_k^{\theta_k}(x) = W_k^3 \mathfrak{R}(W_k^2 \mathfrak{R}(W_k^1 x + b_k^1) + b_k^2) + b_k^3,$$

with weight matrices $W_k^1 \in \mathbb{R}^{20 \times d}$, $W_k^2 \in \mathbb{R}^{20 \times 20}$, $W_k^3 \in \mathbb{R}^{\ell \times 20}$, bias vectors $b_k^1, b_k^2 \in \mathbb{R}^{20}$, $b_k^3 \in \mathbb{R}^\ell$, and $\theta_k = (W_k^1, W_k^2, W_k^3, b_k^1, b_k^2, b_k^3)$, where the matrices are assumed vectorized before concatenation.

4. Convergence analysis. In this section, we primarily provide an error analysis for (2.13), i.e., for the semidiscretization in time. In section 4.1, we introduce notation and spaces, and in section 4.2, we present the setting and some further notation. Two technical results on strong and weak convergence are stated and proved in section 4.3. These two results are used in section 4.4 to prove the error in the objective function, in the initial and terminal value of Y , and for the variance of the stochastic cost. The results hold under an assumption on the regularity of the exact continuous and discrete Markov maps. Convergence of the latter to the former is not assumed. Section 4.5 contains strong convergence of the FBSDE under either the stronger assumption of small time T or convergence of the discrete Markov maps. A discussion about a full error analysis for the fully implementable scheme (3.1) is presented in section 4.6.

4.1. Notation and spaces. For Euclidean spaces \mathbb{R}^k , $k \geq 1$, we denote by $\|\cdot\|$ the 2-norm without specifying the dimension. Let $\mathcal{S}^2(\mathbb{R}^k)$ and $\mathcal{H}^2(\mathbb{R}^k)$ be the

spaces of all progressively measurable stochastic processes $y, z: [0, T] \times \Omega \rightarrow \mathbb{R}^k$, for which

$$\|y\|_{\mathcal{S}^2(\mathbb{R}^k)} = \sup_{t \in [0, T]} \left(\mathbb{E}[\|y_t\|^2] \right)^{\frac{1}{2}} < \infty \quad \text{and} \quad \|z\|_{\mathcal{H}^2(\mathbb{R}^k)} = \left(\mathbb{E} \left[\int_0^T \|z_t\|^2 dt \right] \right)^{\frac{1}{2}} < \infty,$$

respectively. For a discretization, $0 = t_0 < t_1 < \dots < t_N = T$ with $t_{n+1} - t_n = h$, for all n , the space $\mathcal{S}_h(\mathbb{R}^k)$ is the space of all \mathcal{F}^h -adapted, square integrable and discrete stochastic processes $y: \{0, 1, \dots, N\} \times \Omega \rightarrow \mathbb{R}^d$, where $\mathcal{F}_n^h = \sigma(\Delta W_m, m = 0, \dots, n-1)$. For $k_1, k_2, k_3, k_4 \geq 1$, $\ell_1, \ell_2, \ell_3 \geq 0$, and regular $S_i \subseteq \mathbb{R}^{k_i}$, $i = 1, 2, 3, 4$, by $C_b^{\ell_1, \ell_2, \ell_3}(S_1 \times S_2 \times S_3; S_4)$, we denote the space of all functions $\phi: S_1 \times S_2 \times S_3 \rightarrow S_4$, whose derivatives up to orders ℓ_1, ℓ_2, ℓ_3 exist and are continuous and bounded. We equip it with the seminorms

$$|\phi|_\gamma = \sup_{x \in S_1 \times S_2 \times S_3} \|\partial^\gamma \phi(x)\|, \quad i \in \{1, \dots, \ell\},$$

where $\gamma = \{(i_1, i_2, i_3) : i_j \in \{0, \dots, \ell_j\}, j = 1, 2, 3\}$ are multi-indices and $\partial^\gamma = \partial_1^{i_1} \partial_2^{i_2} \partial_3^{i_3}$ with $\partial_j^{i_j}$ denoting i_j th partial derivative in variable j . The set $B(\ell_1, \ell_2, \ell_3)$ denotes all multi-indices of length 3 that have exactly one nonzero index. We only use multi-indices in $B(\ell_1, \ell_2, \ell_3)$ and therefore do not impose restrictions on the cross derivatives, as is otherwise common. For functions with fewer than three variables, we reduce the number of indices accordingly, and in the seminorms we write $|\phi|_{i_1, i_2, i_3} = |\phi|_{(i_1, i_2, i_3)}$. For $\alpha \in (0, 1]$, $k_1, k_2, k_3 \geq 1$, $\ell_1, \ell_2 \geq 0$, and regular $S_i \subseteq \mathbb{R}^{k_i}$, $i = 1, 2, 3$, by $C_{H,b}^{\alpha, \ell_1, \ell_2}([0, T] \times S_1 \times S_2; S_3)$, we denote the space of all functions $\phi: [0, T] \times S_1 \times S_2 \rightarrow S_3$ that are α -Hölder continuous in time and whose derivatives of order ℓ_1, ℓ_2 exist, are continuous and bounded, and satisfy the property

$$\begin{aligned} \|\phi\|_{\alpha, \ell_1, \ell_2} &= \sup_{t \in [0, T]} \|\phi(t, 0, 0)\| + \sum_{\gamma \in B} \sup_{t \in [0, T]} |\phi(t, \cdot, \cdot)|_\gamma \\ &\quad + \sup_{(x_1, x_2) \in S_1 \times S_2} \sup_{t_1, t_2 \in [0, T], t_1 \neq t_2} \frac{\mathbf{1}_{(0,1]}(\alpha) \|\phi(t_1, x_1, x_2) - \phi(t_2, x_1, x_2)\|}{(1 + \|x_1\| + \|x_2\|)|t_2 - t_1|^\alpha} < \infty. \end{aligned}$$

Again, when there is only one space variable, we reduce the number of indices. When there is no risk of confusion, we write $|\cdot|_{\alpha, 0, 0}$ to denote the second term defining the $\|\cdot\|_{\alpha, \ell_1, \ell_2}$ -norm and let $|\cdot|_{0, i, 0}$ and $|\cdot|_{0, 0, i}$ coincide with the seminorms on $C_b^{\alpha, \ell_1, \ell_2}([0, T] \times S_1 \times S_2; S_3)$ with the same notation. For $\alpha = 0$, we let $C_{H,b}^{0, \ell_1, \ell_2}([0, T] \times S_1 \times S_2; S_3) = C_b^{0, \ell_1, \ell_2}([0, T] \times S_1 \times S_2; S_3)$.

For any function or process R defined on $[0, T]$, we denote by \check{R} the discrete time function or process defined by $\check{R}_n = R_{t_n}$, $n = 0, \dots, N$. For any discrete function or process R_h defined on $0, \dots, N$, we write \hat{R}_h for the continuous time function or process defined by the piecewise constant interpolation $\hat{R}_{h,t} = R_{h,n}$ for $t \in [t_n, t_{n+1})$ and $\hat{R}_{h,T} = R_{h,N}$.

4.2. Setting and spaces of Markov maps. Let $(W_t)_{t \in [0, T]}$ be a k -dimensional Brownian motion on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$, and $(\alpha, \beta) \in [0, \frac{1}{2}] \times \{1\}$ or $(\alpha, \beta) = (1, 2)$. The coefficients $b: [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \rightarrow \mathbb{R}^d$, $\sigma: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times k}$, $f: [0, T] \times \mathbb{R}^d \times \mathbb{R}^k \rightarrow \mathbb{R}$, and $g: \mathbb{R}^d \rightarrow \mathbb{R}$ satisfy $b \in C_{H,b}^{\alpha, \beta, \beta}([0, T] \times \mathbb{R}^d \times \mathbb{R}^k; \mathbb{R}^d)$, $\sigma \in C_{H,b}^{\alpha, \beta}([0, T] \times \mathbb{R}^d; \mathbb{R}^{d \times k})$, $f \in C_{H,b}^{\alpha, \beta, \beta}([0, T] \times \mathbb{R}^d \times \mathbb{R}^k; \mathbb{R})$, $g \in C_b^1(\mathbb{R}^d; \mathbb{R})$. In the case $\alpha = 1$, we assume that $D_x \sigma = 0$.

We next introduce families of Markov maps. Let $\mathcal{Z} = \mathcal{Z}(b, \sigma, f, g)$ be the collection of all measurable functions $\zeta: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ with the property that the stochastic

processes $(X^\zeta, Y^\zeta, Z^\zeta)_{\zeta \in \mathcal{Z}} \subset \mathcal{S}^2(\mathbb{R}^d) \times \mathcal{S}^2(\mathbb{R}) \times \mathcal{H}^2(\mathbb{R}^d)$ satisfying for all $t \in [0, T]$, \mathbb{P} -almost surely (2.7), are well defined. We write $\mathcal{Z}^{\alpha, \beta} = \mathcal{Z} \cap C_{\text{H}, \text{b}}^{\alpha, \beta}([0, T] \times \mathbb{R}^d; \mathbb{R}^k)$. For the discrete equations, we introduce for every $h \in (0, 1)$ analogously $\mathcal{Z}_h = \mathcal{Z}_h(b, \sigma, f, g)$ to be the collection of all measurable functions $\zeta: \{0, \dots, N_h - 1\} \times \mathbb{R}^d \rightarrow \mathbb{R}^k$ with the property that $(X^{h, \zeta}, Y^{h, \zeta}, Z^{h, \zeta})_{\zeta \in \mathcal{Z}_h} \subset \mathcal{S}_h^2(\mathbb{R}^d) \times \mathcal{S}_h^2(\mathbb{R}) \times \mathcal{S}_h^2(\mathbb{R})$ satisfying for all $n \in \{0, \dots, N_h\}$, \mathbb{P} -almost surely (2.13) are well defined. We write $\mathcal{Z}_h^\beta = \mathcal{Z}_h \cap (C_b^\beta(\mathbb{R}^d; \mathbb{R}^k))^{N+1}$.

By introducing assumptions on the Markov maps, we eliminate the need for assuming smoothness, Lipschitz, polynomial growth, monotonicity, coercivity, or other conditions on the coefficients, for the existence and uniqueness of solutions for (2.7) and (2.13). For unfortunate choices of b, σ, f, g , the spaces \mathcal{Z} and \mathcal{Z}_h might be empty and results hold by default, but given regular b, σ, f, g , it is not hard, using available solution theory, to find $\zeta \in \mathcal{Z}$ and $\zeta_h \in \mathcal{Z}_h$. Still, the entire spaces might be hard to represent but this is not of central importance. In Assumption 4.3 below, the regularities of the optimal ζ^* and ζ_h^* solving (2.7) and (2.13) are of importance, though. Thus classical solution theory for SDE and FBSDE and regularity theory for optimal Markov maps in discrete and continuous time are required for verifying our assumptions for concrete examples. The latter is not well developed; see the discussion prior to Assumption 4.3 below. Thus part of our assumptions require further theoretical development to be verified, but our results show what is required.

4.3. Auxiliary lemmata on strong and weak convergence for SDEs. In the proof of our convergence results in subsection 4.4, we rely on the strong convergence result, stated next. It contains both classical strong convergence of the Euler–Maruyama scheme for Hölder continuous coefficients, including strong order 1 for additive noise, as well as a nonstandard type of strong convergence result for processes that have drift coefficients that for each step size agree between the grid points, but whose coefficients do not necessarily converge as the step size tends to zero.

LEMMA 4.1. *Suppose the setting of subsection 4.2 holds. Let $a \in C_{\text{H}, \text{b}}^{\alpha, \beta}([0, T] \times \mathbb{R}^d; \mathbb{R}^d)$ and $a_h: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, $h \in (0, 1)$, be a family of functions that are constant on each interval $[t_n, t_{n+1})$ and that satisfy $a_h([0, T], \cdot) \subset C^\beta(\mathbb{R}^d; \mathbb{R}^d)$ and $\sup_{h \in (0, 1)} \|a_h\|_{0, \beta} < \infty$, let $\mathcal{X}, \mathcal{X}^{1, h} \in \mathcal{S}^2(\mathbb{R}^d)$, $h \in (0, 1)$ be the unique solutions to*

$$\begin{aligned} d\mathcal{X}_t &= a(t, \mathcal{X}_t)dt + \sigma(t, \mathcal{X}_t)dW_t, \quad t \in (0, T]; \quad \mathcal{X}_0 = x_0, \\ d\mathcal{X}_t^{1, h} &= a_h(t, \mathcal{X}_t^{1, h})dt + \sigma(t, \mathcal{X}_t^{1, h})dW_t, \quad t \in (0, T]; \quad \mathcal{X}_0^{1, h} = x_0, \quad h \in (0, 1), \end{aligned}$$

and let $\mathcal{X}^{2, h}, \mathcal{X}^{3, h} \in \mathcal{S}_h^2(\mathbb{R}^d)$, $h \in (0, 1)$, be the unique solutions to

$$\begin{aligned} \mathcal{X}_{n+1}^{2, h} &= \mathcal{X}_n^{2, h} + a(t_n, \mathcal{X}_n^{2, h})h + \sigma(t_n, \mathcal{X}_n^{2, h})\Delta W_n, \quad n \in \{0, \dots, N-1\}; \\ \mathcal{X}_0^{2, h} &= x_0, \quad h \in (0, 1), \\ \mathcal{X}_{n+1}^{3, h} &= \mathcal{X}_n^{3, h} + a_h(t_n, \mathcal{X}_n^{3, h})h + \sigma(t_n, \mathcal{X}_n^{3, h})\Delta W_n, \quad n \in \{0, \dots, N-1\}; \\ \mathcal{X}_0^{3, h} &= x_0, \quad h \in (0, 1). \end{aligned}$$

It holds that

$$(4.1) \quad \sup_{h \in (0, 1)} \|\mathcal{X}^{1, h}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \sup_{h \in (0, 1)} \|\hat{\mathcal{X}}^{2, h}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \sup_{h \in (0, 1)} \|\hat{\mathcal{X}}^{3, h}\|_{\mathcal{S}^2(\mathbb{R}^d)} < \infty,$$

and there exists a constant C such that

$$\lim_{h \rightarrow 0} h^{-\alpha} \left(\|\mathcal{X} - \hat{\mathcal{X}}^{2,h}\|_{S^2(\mathbb{R}^d)} + \|\mathcal{X}^{1,h} - \hat{\mathcal{X}}^{3,h}\|_{S^2(\mathbb{R}^d)} \right) = \begin{cases} 0 & \text{if } \alpha = 0, \\ C & \text{otherwise.} \end{cases}$$

Proof. Due to the assumption on a_h , there exists a finite constant C such that

$$\|a_h(t, x)\| \leq \sup_{h \in (0,1)} \sup_{t \in [0,T]} (\|a_h(t, 0)\| + |a_h(t, \cdot)|_1) \|x\| \leq C \|x\|.$$

This is enough to show the first assertion, the stability estimate, using a Gronwall argument. For the convergence, we rely on [46], which covers both the cases $\alpha \in [0, \frac{1}{2}]$ and $\alpha = 1$. It is based on the fact that bistability and consistency of order α , in the sense of [46], imply convergence of order α . We start with $\mathcal{X}^{1,h} \rightarrow \mathcal{X}^{3,h}$ and the convergence $\mathcal{X}^{2,h} \rightarrow \mathcal{X}$ is immediate afterward. Since the drift coefficients of $\mathcal{X}^{1,h}$ and $\mathcal{X}^{3,h}$ are both h -dependent and are not assumed to converge, what we want to prove, i.e., $\lim_{h \rightarrow 0} \|\mathcal{X}_t^{1,h} - \hat{\mathcal{X}}_t^{3,h}\|_{S^2(\mathbb{R}^d)} = 0$, does not say anything about convergence of $\mathcal{X}^{3,h}$ and $\mathcal{X}^{1,h}$. In particular, $\mathcal{X}_t^{1,h_1} - \hat{\mathcal{X}}_t^{3,h_2}$ is in general not small for $h_1 \neq h_2$. Due to this special setting, we have to verify that the proofs of [46] are still valid.

We start with proving bistability. First, in [46, Lemma 4.1], it is proven that for small enough h the numerical scheme is bijective. For our proof, it is important that the upper bound for this property does not depend on h . From the proof, it is clear that this bound depends on the reciprocal of $\sup_{h \in (0,1)} \sup_{t \in [0,T]} |a_h(t, \cdot)|_1$ and is therefore bounded away from zero by assumption. By following the proof of [46, Lemma 4.2] line by line, the bistability constants are bounded by $\sup_{h \in (0,1)} \sup_{t \in [0,T]} |a_h(t, \cdot)|_1$ and its reciprocal when Assumptions (S1) and (S2) in [46] hold. It remains to prove (S1) and (S2) with uniform constants. First, (S1) is trivially satisfied with $L = 0$ for any explicit scheme. By inspection of the proof of [46, Theorem 3.3], we see that, for $\alpha \in [0, \frac{1}{2}]$, corresponding to the stochastic θ -method with $\theta = 0$, (S2) is satisfied for all h with $L = 3T \sup_{h \in (0,1)} \sup_{t \in [0,T]} |a_h(t, \cdot)|_1^2 + 12d|\sigma|_1^2$. For $\alpha = 1$, (S2) holds for the same L as the schemes are the same. This concludes the proof of bistability.

Consistency for $\alpha \in [0, \frac{1}{2}]$ is obtained by observing that the proof in [46], referring for details to [11], only relies on (4.1) and the $\frac{1}{2}$ -Hölder continuity of a_h between the grid points, and this is clearly satisfied since a_h is constant on $[t_n, t_{n+1})$. For consistency in case $\alpha = 1$, we rely on the proof of consistency for the Itô–Taylor scheme of order one, i.e., the Milstein scheme, coinciding with our scheme under the assumption of additive noise. It suffices to note that the analysis is conducted per interval, and to check that f_α is globally Lipschitz continuous for $\alpha \in \mathcal{B}(\mathcal{A}_1)$, in the notation of [46]. For this to hold, it suffices that $\sup_{h \in (0,1)} \sup_{t \in [0,T]} |a_h(t, \cdot)|_2 < \infty$. By assumption, this completes the proof for $\mathcal{X}^{1,h}$ and $\mathcal{X}^{3,h}$. For $\mathcal{X}^{2,h}$ and \mathcal{X} , the convergence is included in [46], up to the order of consistency, which in our setting is lower when $\alpha \in [0, \frac{1}{2}]$. It is straightforward to use α -Hölder continuity for $\alpha \in (0, \frac{1}{2})$, exactly as in the case $\alpha = \frac{1}{2}$, and in this way get consistency of order α . For $\alpha = 0$, one uses dominated convergence instead to obtain the convergence without order. This is valid under our assumptions. \square

Our next lemma is a weak convergence type result. With our assumption of low regularity on the FBSDE coefficients, the weak rate does not exceed the strong rate and therefore the proof relies on the strong convergence results of Lemma 4.1. If we were to obtain the weak convergence order $\alpha \in (\frac{1}{2}, 1)$ for multiplicative noise, then this would require a Hölder condition on the third derivative of b, σ, f, g and the case $\alpha = 1$ requires four derivatives [45, 50]. We avoid this generality.

LEMMA 4.2. *Suppose the setting of subsection 4.2 holds. For all functions $\zeta \in \mathcal{Z}^{\alpha,\beta}$, collection of functions $\zeta_h \in \mathcal{Z}_h^\beta$, $h \in (0,1)$, satisfying $\sup_{h \in (0,1)} \|\hat{\zeta}_h\|_{0,\beta} < \infty$, and $\lambda \geq 0$, there exists a constant C such that*

$$\lim_{h \rightarrow 0} h^{-\alpha} \left(|\Phi_\lambda(\zeta) - \Phi_{\lambda,h}(\zeta)| + |\Phi_\lambda(\hat{\zeta}_h) - \Phi_{\lambda,h}(\zeta_h)| \right) = \begin{cases} 0 & \text{if } \alpha = 0, \\ C & \text{otherwise.} \end{cases}$$

Proof. We start with the first term and observe that

$$(4.2) \quad |\Phi_\lambda(\zeta) - \Phi_{\lambda,h}(\zeta)| \leq |\mathbb{E}[\mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta}]| + \lambda |\text{Var}(\mathcal{Y}_0^\zeta) - \text{Var}(\mathcal{Y}_0^{h,\zeta})|.$$

For the proof of both terms of (4.2), we rely on convergence of $(\mathbb{E}[(\mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta})^2])^{\frac{1}{2}}$, which we next prove, beginning with $\alpha > 0$. For $t \in [t_n, t_{n+1})$, denote $\bar{t} = t_n$. By definition,

$$\begin{aligned} \mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta} &= g(X_T^\zeta) - g(X_T^{h,\zeta}) + \int_0^T (f(t, X_t^\zeta, \zeta(t, X_t^\zeta)) - f(\bar{t}, \hat{X}_t^{h,\zeta}, \zeta(\bar{t}, \hat{X}_t^{h,\zeta}))) dt \\ &\quad + \int_0^T \langle \zeta(t, X_t^\zeta) - \zeta(\bar{t}, \hat{X}_t^{h,\zeta}), dW_t \rangle \\ &= g(X_T^\zeta) - g(X_T^{h,\zeta}) + \int_0^T (f(t, X_t^\zeta, \zeta(t, X_t^\zeta)) - f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta}))) dt \\ &\quad + \int_0^T (f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta})) - f(\bar{t}, \hat{X}_t^{h,\zeta}, \zeta(\bar{t}, \hat{X}_t^{h,\zeta}))) dt + \int_0^T \langle \zeta(t, X_t^\zeta) \\ &\quad - \zeta(t, \hat{X}_t^{h,\zeta}), dW_t \rangle + \int_0^T \langle \zeta(t, X_t^\zeta) - \zeta(t, \hat{X}_t^{h,\zeta}), dW_t \rangle \\ &\quad + \int_0^T \langle \zeta(t, \hat{X}_t^{h,\zeta}) - \zeta(\bar{t}, \hat{X}_t^{h,\zeta}), dW_t \rangle. \end{aligned}$$

From the Itô isometry and the triangle inequality, we have

$$\begin{aligned} (\mathbb{E}[(\mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta})^2])^{\frac{1}{2}} &\leq (\mathbb{E}[(g(X_T^\zeta) - g(X_T^{h,\zeta}))^2])^{\frac{1}{2}} \\ &\quad + \int_0^T (\mathbb{E}[(f(t, X_t^\zeta, \zeta(t, X_t^\zeta)) - f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta})))^2])^{\frac{1}{2}} dt \\ &\quad + \int_0^T \mathbb{E}[(f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta})) - f(\bar{t}, \hat{X}_t^{h,\zeta}, \zeta(\bar{t}, \hat{X}_t^{h,\zeta})))^2])^{\frac{1}{2}} dt \\ &\quad + \left(\int_0^T \mathbb{E}[\|\zeta(t, X_t^\zeta) - \zeta(t, \hat{X}_t^{h,\zeta})\|^2] dt \right)^{\frac{1}{2}} \\ &\quad + \left(\int_0^T \mathbb{E}[\|\zeta(t, \hat{X}_t^{h,\zeta}) - \zeta(\bar{t}, \hat{X}_t^{h,\zeta})\|^2] dt \right)^{\frac{1}{2}} = I_1 + I_2 + I_3 + I_4 + I_5. \end{aligned}$$

For I_1 , we use Lipschitz continuity of g and the Cauchy–Schwarz inequality to get

$$(4.3) \quad I_1 \leq |g|_1 \mathbb{E}[\|X_T^\zeta - X_T^{h,\zeta}\|] \leq |g|_1 (\mathbb{E}[\|X_T^\zeta - X_T^{h,\zeta}\|^2])^{\frac{1}{2}} \leq |g|_1 \|X^\zeta - \hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)}.$$

The function $\phi(t, x) = f(t, x, \zeta(t, x))$ is uniformly Lipschitz continuous in x with Lipschitz constant $C_\phi = |f|_{0,1,0} + |f|_{0,0,1}|\zeta|_{0,1}$. This implies, together with the Cauchy–Schwarz inequality,

$$(4.4) \quad I_2 \leq C_\phi \int_0^T (\mathbb{E}[\|X_t^\zeta - \hat{X}_t^{h,\zeta}\|^2])^{\frac{1}{2}} dt \leq C_\phi T \|X^\zeta - \hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)}.$$

Using the assumptions of b and ζ , it is easily verified that ϕ belongs to $C_{H,b}^{\alpha,\beta}([0, T] \times \mathbb{R}^d; \mathbb{R}^d)$ and, therefore, by (4.3), (4.4), and Lemma 4.1, we have $I_1 + I_2 \leq Ch^\alpha$. For the term I_3 , we see that

$$\begin{aligned} I_3 &\leq \left(|f|_{\alpha,0,0} + |f|_{0,0,1} |\zeta|_{\alpha,0} \right) \left(1 + \sup_{h \in (0,1)} \|\hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)} \right) \sum_{n=0}^{N-1} \int_{t_n}^{t_{n+1}} (t - t_n)^\alpha dt \\ &= \frac{|f|_{\alpha,0,0} + |f|_{0,0,1} |\zeta|_{\alpha,0}}{1 + \alpha} \left(1 + \sup_{h \in (0,1)} \|\hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)} \right) Nh^{1+\alpha} \leq Ch^\alpha. \end{aligned}$$

Similarly, for I_4 and I_5 , we have

$$I_4 \leq |\zeta|_{0,1} \left(\int_0^T \mathbb{E} [\|X_t^\zeta - \hat{X}_t^{h,\zeta}\|^2] dt \right)^{\frac{1}{2}} \leq |\zeta|_{0,1} T^{\frac{1}{2}} \|X^\zeta - \hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)} \leq Ch^\alpha$$

and

$$\begin{aligned} I_5 &\leq |\zeta|_{\alpha,0} \left(1 + \sup_{h \in (0,1)} \|\hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)} \right) \left(\sum_{n=0}^{N-1} \int_{t_n}^{t_{n+1}} |t - t_n|^{2\alpha} dt \right)^{\frac{1}{2}} \\ &= \frac{|\zeta|_{\alpha,0}}{1 + 2\alpha} \left(1 + \sup_{h \in (0,1)} \|\hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)} \right) (Nh^{1+2\alpha})^{\frac{1}{2}} \leq Ch^\alpha. \end{aligned}$$

For $\alpha = 0$, the terms I_1, I_2, I_4 need no special treatment. For I_3 , we notice that the function ϕ is only continuous in t and convergence without rate holds by the dominated convergence theorem. This is verified by noting that

$$\begin{aligned} &\sup_{t \in [0, T]} \left(\mathbb{E} \left[\left(f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta})) - f(\bar{t}, \hat{X}_t^{h,\zeta}, \zeta(\bar{t}, \hat{X}_t^{h,\zeta})) \right)^2 \right] \right)^{\frac{1}{2}} \\ &\leq 2 \sup_{t \in [0, T]} \left(\mathbb{E} \left[\left(f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta})) \right)^2 \right] \right)^{\frac{1}{2}} \\ &\leq 2 \sup_{t \in [0, T]} \left(\mathbb{E} \left[\left(f(t, \hat{X}_t^{h,\zeta}, \zeta(t, \hat{X}_t^{h,\zeta})) - f(t, 0, 0) \right)^2 \right] \right)^{\frac{1}{2}} + 2 \sup_{t \in [0, T]} |f(t, 0, 0)| \\ &\leq 2 \sup_{t \in [0, T]} \left(|f(t, \cdot, \cdot)|_{1,0} + |f(t, \cdot, \cdot)|_{0,1} (|\zeta(t, \cdot)|_1 + \|\zeta(t, 0)\|) \right) \sup_{h \in (0,1)} \|\hat{X}^{h,\zeta}\|_{\mathcal{S}^2(\mathbb{R}^d)} \\ &\quad + 2 \sup_{t \in [0, T]} |\mathbb{E}[f(t, 0, 0)]|. \end{aligned}$$

Due to the assumptions, the right-hand side is finite. The term I_5 admits a similar treatment and we refrain from giving details. This proves that

$$(4.5) \quad \lim_{h \rightarrow 0} (\mathbb{E}[(\mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta})^2])^{\frac{1}{2}} = \begin{cases} 0 & \text{if } \alpha = 0, \\ C & \text{otherwise.} \end{cases}$$

For the first term in (4.2), this fact and $|\mathbb{E}[\mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta}]| \leq (\mathbb{E}[(\mathcal{Y}_0^\zeta - \mathcal{Y}_0^{h,\zeta})^2])^{\frac{1}{2}}$ prove the convergence.

For the second term in (4.2), we use the conjugate rule, the Cauchy–Schwarz inequality, and the triangle inequality to get

$$\begin{aligned}
& |\text{Var}(\mathcal{Y}_0^\gamma) - \text{Var}(\mathcal{Y}_0^{h,\tilde{\gamma}})| \\
&= \left| \mathbb{E} \left[|\mathbb{E}[\mathcal{Y}_0^\gamma] - \mathcal{Y}_0^\gamma|^2 - |\mathbb{E}[\mathcal{Y}_0^{h,\tilde{\gamma}}] - \mathcal{Y}_0^{h,\tilde{\gamma}}|^2 \right] \right| \\
&= \left| \mathbb{E} \left[(\mathbb{E}[\mathcal{Y}_0^\gamma - \mathcal{Y}_0^{h,\tilde{\gamma}}] + \mathcal{Y}_0^{h,\tilde{\gamma}} - \mathcal{Y}_0^\gamma) (\mathbb{E}[\mathcal{Y}_0^\gamma + \mathcal{Y}_0^{h,\tilde{\gamma}}] - \mathcal{Y}_0^{h,\tilde{\gamma}} - \mathcal{Y}_0^\gamma) \right] \right| \\
&\leq \left(\mathbb{E} \left[(\mathbb{E}[\mathcal{Y}_0^\gamma - \mathcal{Y}_0^{h,\tilde{\gamma}}] + \mathcal{Y}_0^{h,\tilde{\gamma}} - \mathcal{Y}_0^\gamma)^2 \right] \right)^{\frac{1}{2}} \left(\mathbb{E} \left[(\mathbb{E}[\mathcal{Y}_0^\gamma + \mathcal{Y}_0^{h,\tilde{\gamma}}] - \mathcal{Y}_0^{h,\tilde{\gamma}} - \mathcal{Y}_0^\gamma)^2 \right] \right)^{\frac{1}{2}} \\
&\leq \left(|\mathbb{E}[\mathcal{Y}_0^\gamma - \mathcal{Y}_0^{h,\tilde{\gamma}}]| + (\mathbb{E}[(\mathcal{Y}_0^{h,\tilde{\gamma}} - \mathcal{Y}_0^\gamma)^2])^{\frac{1}{2}} \right) \left(|\mathbb{E}[\mathcal{Y}_0^\gamma + \mathcal{Y}_0^{h,\tilde{\gamma}}]| \right. \\
&\quad \left. + (\mathbb{E}[(\mathcal{Y}_0^{h,\tilde{\gamma}} + \mathcal{Y}_0^\gamma)^2])^{\frac{1}{2}} \right) \\
&\leq 4 \left((\mathbb{E}[(\mathcal{Y}_0^\gamma)^2])^{\frac{1}{2}} + \sup_{h \in (0,1)} (\mathbb{E}[(\mathcal{Y}_0^{h,\tilde{\gamma}})^2])^{\frac{1}{2}} \right) (\mathbb{E}[(\mathcal{Y}_0^{h,\tilde{\gamma}} - \mathcal{Y}_0^\gamma)^2])^{\frac{1}{2}}.
\end{aligned}$$

Together with (4.5), this completes the proof for the term $|\Phi_\lambda(\zeta) - \Phi_{\lambda,h}(\tilde{\zeta})|$.

For the second term $|\Phi_\lambda(\hat{\zeta}_h) - \Phi_{\lambda,h}(\zeta_h)|$, we define the functions $a_h: [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$, $h \in (0, 1)$, by $a_h = \hat{c}_h$, where $c_h: \{0, \dots, N\} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ are given by $c_{h,n}(x) = b(t_n, x, \zeta_{h,n}(x))$. By the assumptions on b and ζ_h , we have that a_h satisfy the assumption of Lemma 4.1. This implies that $\lim_{h \rightarrow 0} \|X^{\zeta_h} - \hat{X}^{h,\zeta_h}\|_{\mathcal{S}^2(\mathbb{R}^d)} = 0$ and $\|X^{\zeta_h} - \hat{X}^{h,\zeta_h}\|_{\mathcal{S}^2(\mathbb{R}^d)} \leq Ch^\alpha$ for $\alpha > 0$. This, with the same calculations as for I_1, I_2, I_4, I_5 , and noting that the analogous for I_3 is zero, completes the proof. \square

4.4. Time discretization error of the initial and terminal values. This and the next section contain our main results. The results are based on an assumption of regularity of the Markov maps $\zeta^* \in \mathcal{Z}$ and $\zeta_{\lambda,h}^* \in \mathcal{Z}_h$, $h \in (0, 1)$, $\lambda \geq 0$, of the FBSDE and its discretizations, respectively. The regularity has to be verified for specific problems and this, in particular for the discrete problem, is a nontrivial and not well-studied problem; see [14, 54, 12]. For the continuous problem, it breaks down into existence and uniqueness of a solution of the FBSDE and regularity of the solution to the HJB equation. In this paper, we provide only the abstract error analysis. Our assumption is stated next.

Assumption 4.3. There exists a unique minimizer $\zeta^* \in \mathcal{Z}$ of $\zeta \mapsto \Phi_1(\zeta)$; it belongs to $\mathcal{Z}^{\alpha,\beta}$ and satisfies $\text{Var}(\mathcal{Y}_0^{\zeta^*}) = 0$. Moreover, for all $\lambda \geq 0$, there exists a collection of minimizers $\zeta_{\lambda,h}^* \in \mathcal{Z}_h$, $h \in (0, 1)$, of $\zeta \mapsto \Phi_{\lambda,h}(\zeta)$ that for each h belongs to \mathcal{Z}_h^β and satisfies $\sup_{h \in (0,1)} \|\hat{\zeta}_{\lambda,h}^*\|_{0,\beta} < \infty$.

Our first result concerns the convergence of the objective function. The proof relies on a split of the error into two error terms, one containing only ζ^* and one containing only $\zeta_{\lambda,h}^*$. This way, we avoid getting an error bound containing the error between the continuous and discrete Markov maps. The result is used throughout our proofs, except in Proposition 4.9, where we have the error of the Markov map in the bound.

THEOREM 4.4. *Suppose the setting of subsection 4.2, let Assumption 4.3 hold and $\lambda \geq 0$. Then, $|\Phi^* - \Phi_{\lambda,h}^*| \rightarrow 0$ as $h \rightarrow 0$, and if $\alpha > 0$, then there exists a constant C , independent of h , such that*

$$|\Phi^* - \Phi_{\lambda,h}^*| \leq Ch^\alpha.$$

Proof. We start by noting that, since $\text{Var}(\mathcal{Y}_0^{\zeta^*}) = 0$, it holds that ζ^* is optimal for all Φ_λ , $\lambda \geq 0$, and $\Phi^* = \Phi_\lambda^* := \Phi_\lambda(\zeta^*)$. Because of optimality of $\zeta^* \in \mathcal{Z}$ and since $\hat{\zeta}_h^* \in \mathcal{Z}$, it holds that $\Phi^* = \Phi_\lambda^* \leq \Phi_\lambda(\hat{\zeta}_h^*)$. This implies

$$\Phi^* \leq \Phi_\lambda(\hat{\zeta}_{\lambda,h}^*) = \Phi_{\lambda,h}^* + \Phi_\lambda(\hat{\zeta}_{\lambda,h}^*) - \Phi_{\lambda,h}^* \leq \Phi_{\lambda,h}^* + |\Phi_\lambda(\hat{\zeta}_{\lambda,h}^*) - \Phi_{\lambda,h}^*|.$$

Similarly, from the optimality of $\zeta_{\lambda,h}^* \in \mathcal{Z}_h$ and since $\check{\zeta}^* \in \mathcal{Z}_h$, it holds that $\Phi_{\lambda,h}^* \leq \Phi_{\lambda,h}(\check{\zeta}^*)$. We have

$$\Phi_{\lambda,h}^* \leq \Phi_{\lambda,h}(\check{\zeta}^*) = \Phi^* + \Phi_{\lambda,h}(\check{\zeta}^*) - \Phi^* \leq \Phi^* + |\Phi_{\lambda,h}(\check{\zeta}^*) - \Phi^*|.$$

The two inequalities equivalently read $\Phi^* - \Phi_{\lambda,h}^* \leq |\Phi(\hat{\zeta}_{\lambda,h}^*) - \Phi_{\lambda,h}^*|$ and $\Phi_{\lambda,h}^* - \Phi^* \leq |\Phi_{\lambda,h}(\check{\zeta}^*) - \Phi^*|$, and thus

$$|\Phi^* - \Phi_{\lambda,h}^*| \leq \max(|\Phi(\hat{\zeta}_{\lambda,h}^*) - \Phi_{\lambda,h}^*|, |\Phi_{\lambda,h}(\check{\zeta}^*) - \Phi^*|) \leq |\Phi(\hat{\zeta}_{\lambda,h}^*) - \Phi_{\lambda,h}^*| + |\Phi_{\lambda,h}(\check{\zeta}^*) - \Phi^*|.$$

The convergence is given by our assumption and Lemma 4.2. This completes the proof. \square

A consequence of the convergence of the objective function is the convergence of the two components of the objective, given that $\lambda > 0$. This is stated in our next theorem.

THEOREM 4.5. *Suppose the setting of subsection 4.2, let Assumption 4.3 hold, $\lambda > 0$, $\mathcal{Y}_0^{h,\lambda} := \mathcal{Y}_0^{h,\zeta_{\lambda,h}^*}$, and $Y_0^{h,\lambda} := \mathbb{E}[\mathcal{Y}_0^{h,\lambda}]$. Then, $|Y_0 - Y_0^{h,\lambda}| + \text{Var}(\mathcal{Y}_0^{h,\lambda}) \rightarrow 0$ as $h \rightarrow 0$, and if $\alpha > 0$, there exists a constant C , independent of h , such that*

$$|Y_0 - Y_0^{h,\lambda}| + \text{Var}(\mathcal{Y}_0^{h,\lambda}) \leq Ch^\alpha.$$

Proof. We start with the proof for the variance and use both the sequences $(\zeta_{0,h}^*)_{h \in (0,1)}$ and $(\zeta_{\lambda,h})_{h \in (0,1)}$. The proof relies on a squeezing argument. Adding and subtracting terms and using the triangle inequality yields

$$(4.6) \quad \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda}) \leq |Y_0 - Y_0^{h,\lambda} - \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda})| + |Y_0 - Y_0^{h,0}| + |Y_0^{h,0} - Y_0^{h,\lambda}|.$$

From the assumption $\text{Var}(\mathcal{Y}_0) = 0$, it holds that $\Phi^* = Y_0$ and Theorem 4.4 gives

$$(4.7) \quad |Y_0 - Y_0^{h,\lambda} - \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda})| + |Y_0 - Y_0^{h,0}| = |\Phi^* - \Phi_{\lambda,h}^*| + |\Phi^* - \Phi_{0,h}^*| \leq Ch^\alpha.$$

For the third term on the right-hand side of (4.6), we first notice that, since $\zeta_{0,h}$ is a minimizer of $\zeta \mapsto Y_0^{h,0}$, it holds that $Y_0^{h,\lambda} - Y_0^{h,0} \geq 0$. This fact, adding $\lambda \text{Var}(\mathcal{Y}_0^{h,\lambda})$, adding and subtracting Y_0 , using the triangle inequality, and using (4.7), gives us

$$(4.8) \quad \begin{aligned} |Y_0^{h,0} - Y_0^{h,\lambda}| &= Y_0^{h,\lambda} - Y_0^{h,0} \leq Y_0^{h,\lambda} + \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda}) - Y_0 - Y_0^{h,0} + Y_0 \\ &\leq |Y_0 - Y_0^{h,\lambda} - \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda})| + |Y_0 - Y_0^{h,\lambda}| \leq Ch^\alpha. \end{aligned}$$

Now, (4.6)–(4.8) complete the proof for the variance. For the convergence of $Y_0^{h,\lambda}$, we conclude

$$|Y_0 - Y_0^{h,\lambda}| \leq |Y_0 - Y_0^{h,\lambda} - \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda})| + \lambda \text{Var}(\mathcal{Y}_0^{h,\lambda}) \leq Ch^\alpha. \quad \square$$

From (2.8) and Theorem 4.5, we directly get convergence in the terminal value and we can also conclude strong convergence of $y_0^{h,\lambda}$. This is stated in the following two corollaries.

COROLLARY 4.6. Suppose the setting of subsection 4.2, let Assumption 4.3 hold, and $\lambda > 0$. Then, $\mathbb{E}[(g(X_N^{h,\lambda}) - Y_N^{h,\lambda})^2] \rightarrow 0$ as $h \rightarrow 0$, and if $\alpha > 0$, then there exists a constant C , independent of h , such that

$$\left(\mathbb{E}\left[(g(X_N^{h,\lambda}) - Y_N^{h,\lambda})^2\right]\right)^{\frac{1}{2}} \leq Ch^{\frac{\alpha}{2}}.$$

COROLLARY 4.7. Suppose the setting of subsection 4.2, let Assumption 4.3 hold, and $\lambda > 0$. Then, $(\mathbb{E}[(\mathcal{Y}_0 - \mathcal{Y}_0^{h,\lambda})^2])^{\frac{1}{2}} \rightarrow 0$ as $h \rightarrow 0$, and if $\alpha > 0$, then there exists a constant C , independent of h , such that

$$\left(\mathbb{E}\left[(\mathcal{Y}_0 - \mathcal{Y}_0^{h,\lambda})^2\right]\right)^{\frac{1}{2}} \leq Ch^{\frac{\alpha}{2}}.$$

Proof. From Theorem 4.5 and the triangle inequality, it holds that

$$\begin{aligned} \left(\mathbb{E}\left[(\mathcal{Y}_0 - \mathcal{Y}_0^{h,\lambda})^2\right]\right)^{\frac{1}{2}} &\leq \left(\mathbb{E}\left[(Y_0 - Y_0^{h,\lambda} + \mathcal{Y}_0 - \mathcal{Y}_0^{h,\lambda})^2\right]\right)^{\frac{1}{2}} + |Y_0 - Y_0^{h,\lambda}| \\ &= \left(\text{Var}(\mathcal{Y}_0^{h,\lambda})\right)^{\frac{1}{2}} + |Y_0 - Y_0^{h,\lambda}| \leq Ch^{\frac{\alpha}{2}}. \end{aligned}$$

This proves the corollary. \square

4.5. Time discretization error of the FBSDE. While the error analysis for the initial and terminal values in the previous subsection required only Assumption 4.3, the strong error analysis of (X, Y, Z) requires more. In Theorem 4.8, we prove strong convergence for small horizon T . This should be compared with the convergence result in [32], which also has a very restrictive assumption on T . Compared to [32], whose bound contains the error in the terminal value, we have no such term. In Proposition 4.9, we prove the strong convergence without a restriction on T with the cost of having a bound containing the error between the Markov maps ζ^* and $\zeta_{\lambda,h}^*$.

THEOREM 4.8. Suppose the setting of subsection 4.2, let Assumption 4.3 hold, $\alpha, \lambda > 0$, and

$$\max\left(T^{\frac{1}{2}}|f|_{0,0,1}, \frac{5T(T|f|_{0,1,0} + |g|_1)|b|_{0,0,1}^2 \exp(5T(|b|_{0,1,0}T + |\sigma|_{0,1}))}{1 - T^{\frac{1}{2}}|f|_{0,0,1}}\right) < 1.$$

Then, there exists a constant C , independent of h , such that

$$\|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \|Y - \hat{Y}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)} \leq Ch^{\frac{\alpha}{2}}.$$

Proof. We start by noting that for $t \in [0, T]$, it holds that

$$(4.9) \quad Y_t - Y_t^{h,\lambda} = Y_0 - Y_0^{h,\lambda} - \int_0^t (f(s, X_s, Z_s) - f(\bar{s}, \hat{X}_s^{h,\lambda}, \hat{Z}_s^{h,\lambda}))ds + \int_0^t (Z_s - \hat{Z}_s^{h,\lambda})dW_s.$$

By the Itô isometry, substitution of (4.9) with $t = T$, and the triangle inequalities, we have

$$\begin{aligned} \|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)} &= \left(\mathbb{E}\left[\int_0^T \|Z_t - \hat{Z}_t^{h,\lambda}\|^2 dt\right]\right)^{\frac{1}{2}} = \left(\mathbb{E}\left[\left(\int_0^T (Z_t - \hat{Z}_t^{h,\lambda})dW_t\right)^2\right]\right)^{\frac{1}{2}} \\ &= \left(\mathbb{E}\left[\left(Y_0 - Y_0^{h,\lambda} + g(X_T) - Y_N^{h,\lambda} - \int_0^T (f(t, X_t, Z_t) - f(\bar{t}, \hat{X}_t^{h,\lambda}, \hat{Z}_t^{h,\lambda}))dt\right)^2\right]\right)^{\frac{1}{2}} \end{aligned}$$

$$\leq |Y_0 - Y_0^{h,\lambda}| + (\mathbb{E}[(g(X_T) - g(X_N^{h,\lambda}))^2])^{\frac{1}{2}} + (\mathbb{E}[(g(X_N^{h,\lambda}) - Y_N^{h,\lambda})^2])^{\frac{1}{2}} \\ + \left(\mathbb{E} \left[\left(\int_0^T (f(t, X_t, Z_t) - f(\bar{t}, \hat{X}_t^{h,\lambda}, \hat{Z}_t^{h,\lambda})) dt \right)^2 \right] \right)^{\frac{1}{2}}.$$

The first three terms are by Theorem 4.5 and Corollary 4.6, bounded from below by $Ch^{\frac{\alpha}{2}} + |g|_1 \|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)}$. By similar arguments as those for I_2, I_3 in the proof of Lemma 4.2, it holds that

$$(4.10) \quad \left(\mathbb{E} \left[\left(\int_0^T (f(t, X_t, Z_t) - f(\bar{t}, \hat{X}_t^{h,\lambda}, \hat{Z}_t^{h,\lambda})) dt \right)^2 \right] \right)^{\frac{1}{2}} \\ \leq T|f|_{\alpha,0,0}(1+\alpha)^{-1}h^\alpha + T|f|_{0,1,0}\|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + T^{\frac{1}{2}}|f|_{0,0,1}\|\hat{Z}^{h,\lambda} - Z\|_{\mathcal{H}^2(\mathbb{R}^k)}.$$

By a kickback argument and by assumption, it holds that

$$(4.11) \quad \|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)} \leq Ch^\alpha + \frac{T|f|_{0,1,0} + |g|_1}{1 - T^{\frac{1}{2}}|f|_{0,0,1}} \|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)}.$$

We next approach the error $\|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)}$. Let $\Xi^{h,\lambda} \in \mathcal{S}^2(\mathbb{R}^d)$, $h \in (0, 1)$, be the family of stochastic processes that for all $t \in [0, T]$, \mathbb{P} -a.s., satisfy

$$\Xi_t^{h,\lambda} = x_0 + \int_0^t b(\bar{s}, \Xi_s^{h,\lambda}, Z_s^{h,\lambda}) ds + \int_0^t \sigma(\bar{s}, \Xi_s^{h,\lambda}) dW_s.$$

Using the triangle inequality, we get

$$\|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} \leq \|X - \Xi^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \|\Xi^{h,\lambda} - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)}.$$

By the arguments used repeatedly in the proof of Lemma 4.1 and by assumptions, it holds that

$$\mathbb{E}[\|X_t - \Xi_t^{h,\lambda}\|^2] \\ \leq Ch^\alpha + 5|b|_{0,0,1}^2 T \|Z - Z^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)}^2 + 5(|b|_{0,1,0} T + |\sigma|_{0,1}) \int_0^t \mathbb{E}[\|X_s - \Xi_s^{h,\lambda}\|^2] ds.$$

From this, it follows by the Gronwall lemma that

$$\mathbb{E}[\|X_t - \Xi_t^{h,\lambda}\|^2] \leq \exp\left(5T(|b|_{0,1,0} T + |\sigma|_{0,1})\right) \left(Ch^\alpha + 5|b|_{0,0,1}^2 T \|Z - Z^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)}^2\right).$$

A use of Lemma 4.1 yields $\|\Xi^{h,\lambda} - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} \leq Ch^\alpha$. We conclude that

$$(4.12) \quad \|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} \leq Ch^\alpha + 5T|b|_{0,0,1}^2 \exp(5T(|b|_{0,1,0} T + |\sigma|_{0,1})) \|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)}.$$

Using (4.12) in (4.11) gives, after a kickback argument, the desired bound $\|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)} \leq Ch^\alpha$.

For $\|Y - \hat{Y}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)}$, we use (4.9) and the standard arguments to get

$$\left(\mathbb{E}[\|Y_t - Y_t^{h,\lambda}\|^2] \right)^{\frac{1}{2}} \leq |Y_0 - Y_0^{h,\lambda}| \\ + \left(\int_0^T \mathbb{E}[\|f(s, X_s, Z_s) - f(\bar{s}, \hat{X}_s^{h,\lambda}, \hat{Z}_s^{h,\lambda})\|^2] ds \right)^{\frac{1}{2}} + \|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)}.$$

Applying Theorem 4.4, (4.10) and the obtained results for $\|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)}$ and $\|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)}$ complete the proof. \square

PROPOSITION 4.9. *Suppose the setting of subsection 4.2, let Assumption 4.3 hold, and $\alpha, \lambda > 0$. Then, there exists a constant C , independent of h , such that*

$$\begin{aligned} & \|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \|Y - \hat{Y}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \|Z - \hat{Z}^{h,\lambda}\|_{\mathcal{H}^2(\mathbb{R}^k)} \\ & \leq C \left(h^{\frac{\alpha}{2}} + \max_{0, \dots, N_h} \left(\mathbb{E} \left[\|\zeta^*(t_n, X_n^{\lambda,h}) - \hat{\zeta}_{h,\lambda}^*(t_n, X_n^{\lambda,h})\|^2 \right] \right)^{\frac{1}{2}} \right). \end{aligned}$$

Proof. This is proved similarly to Theorem 4.8 without the kickback argument and instead of (4.10), using

$$\begin{aligned} & \left(\mathbb{E} \left[\left(\int_0^T (f(t, X_t, Z_t) - f(\bar{t}, \hat{X}_t^h, \hat{Z}_t^h)) dt \right)^2 \right] \right)^{\frac{1}{2}} \\ & \leq C \left(h^{\frac{\alpha}{2}} + \|X - \hat{X}^{h,\lambda}\|_{\mathcal{S}^2(\mathbb{R}^d)} + \max_{0, \dots, N_h} \left(\mathbb{E} \left[\|\zeta^*(t_n, X_n^{\lambda,h}) - \hat{\zeta}_{h,\lambda}^*(t_n, X_n^{\lambda,h})\|^2 \right] \right)^{\frac{1}{2}} \right). \end{aligned}$$

\square

4.6. A discussion on the full error analysis of the robust deep FBSDE method. In subsections 4.4 and 4.5, only the time discretization error is considered, i.e., the error between (2.7) and (2.13). For a full error analysis, the error between the fully implementable scheme (3.1) and (2.7) must be considered. Besides the time discretization error, there are three other sources of error. The first is the error induced by optimizing over the parameters of a neural network, instead of over the vast set \mathcal{Z}_h . By the universal approximation theorem [19], this error can be made arbitrarily small, but this theorem gives no help with suggesting the network architecture that can guarantee a maximal error of desired size. The second error is the Monte Carlo error induced from approximating the expectation in Y_0 by a sample mean. This error allows for a simple error analysis and the Monte Carlo error is of the order $\mathcal{O}(M_{\text{batch}}^{-1/2})$. The final error is the error induced from the inexact optimization procedure of (3.1).

5. Numerical experiments. In this section, we evaluate our algorithm on three different problems. The first two are of LQ type, for which we have access to a semianalytic solution for comparison. The third example uses nonlinear terms in both the drift and diffusion coefficients in the forward equation, and we no longer have access to a reference solution. In the first example, there is a one-to-one map between the feedback control and the Z -process, and we can set $\lambda = 0$ in the loss function. In the second and third examples, this is not the case, and $\lambda > 0$ is necessary for uniqueness of the minimizer to our discrete problem, and convergence to the continuous FBSDE.

In the experimental convergence studies, we approximate $\|\cdot\|_{\mathcal{S}(\mathbb{R}^q)}$ and $\|\cdot\|_{\mathcal{H}(\mathbb{R}^q)}$ with

$$\|A\|_{\mathcal{S}_h^2} = \max_{n \in \{0, 1, \dots, N\}} \left(\frac{1}{M} \sum_{m=1}^M \|A_n(m)\|^2 \right)^{\frac{1}{2}}, \quad \|A\|_{\mathcal{H}_h^2} = \frac{1}{N} \sum_{N=0}^{N-1} \left(\frac{1}{M} \sum_{m=1}^M \|A_n(m)\|^2 \right)^{\frac{1}{2}}.$$

Here, $A(m) = \{A_1(m), A_2(m), \cdot, A_N(m)\}$, $m = 1, 2, \dots, M = 2^{22}$, are independent and identically distributed realizations of some adapted stochastic processes A on the grid.

The norm $\|\cdot\|_{L^2(\Omega;\mathbb{R}^q)}$ is approximated with a sample mean denoted $\|\cdot\|_{L_{h_i}^2}$. For the convergence study, the experimental order of convergence (EOC) is used. It is defined as $\text{EOC}(h_i) = \frac{\log(\text{error}(h_{i+1})) - \log(\text{error}(h_i))}{\log(h_{i+1}) - \log(h_i)}$. In all examples, we use the neural network architecture in section 3.2. We use $M_{\text{train}} = 2^{22}$ training data points and batch size $M_{\text{batch}} = 2^9$ with $K_{\text{epoch}} = 15$ epochs. This gives $K_{\text{batch}} = 2^{12} = 4096$ updates per epoch. For the optimization, the Adam optimizer [44] is used with learning rate 0.1 for the first three epochs, which, after that, is multiplied by a factor of $e^{-0.5}$ for each new epoch. For our use, it was important to choose M_{train} large since in our empirical convergence results we want to isolate the time discretization error. In practice, the method generates acceptable solutions with significantly smaller M_{train} .

5.1. Linear-quadratic control problems. Among all stochastic control problems, the LQ control problem is the most studied and with the most structure; see, e.g., [3]. For our purposes, it has a closed-form analytic solution with which we can compare our numerical approximations.

Let $k = d$, $x_0 \in \mathbb{R}^d$, $A, \sigma \in \mathbb{R}^{d \times d}$, $R_x, G \in \mathbb{S}_+^d$, $R_u \in \mathbb{S}_+^\ell$, and $B \in \mathbb{R}^{d \times \ell}$ be of full rank and $C \in \mathbb{R}^d$. The state equation and cost functional of an LQ-Gaussian control problem are given by

$$\begin{cases} X_t = x_0 + \int_0^t (A(C - X_s) + Bu_s) ds + \int_0^t \sigma dW_s, \\ J^u(t, x) = \mathbb{E}^{t,x} \left[\int_t^T (\langle R_x X_s, X_s \rangle + \langle R_u u_s, u_s \rangle) ds + \langle GX_T, X_T \rangle \right], \quad t \in [0, T]. \end{cases}$$

With the minimizer v^* of the corresponding Hamiltonian, $\inf_{u \in U} \{ \langle D_x V, Bu \rangle + \langle R_u u, u \rangle \}$, we have the optimal feedback control $u_t^* = -\frac{1}{2} R_u^{-1} B^\top D_x V(t, X_t)$. Here, we recall that V is the solution to the associated HJB equation. Its solution is given by $V(t, x) = x^\top P(t)x + x^\top Q(t) + R(t)$, where (P, Q, R) solves the system of ordinary differential equations,

$$\begin{cases} \dot{P}(t) - A^\top P(t) - P(t)A - P(t)BR_u^{-1}B^\top P(t) + R_x = \mathbf{0}_{d \times d}, \\ \dot{Q}(t) + 2P(t)AC - A^\top Q(t) - P(t)BR_u^{-1}B^\top Q(t) = \mathbf{0}_d, \\ \dot{R}(t) + \text{Tr}\{\sigma \sigma^\top P\} + Q(t)^\top AC - \frac{1}{4}Q(t)^\top BR_u^{-1}B^\top Q(t) = 0, \quad t \in [0, T], \\ P(T) = G; \quad Q(T) = \mathbf{0}_d; \quad R(T) = 0. \end{cases}$$

The first equation is a matrix Riccati equation, and we refer to the whole system, slightly inaccurately, as the Riccati equation. The gradient of V satisfies $D_x V(t, x) = 2P(t)x + Q(t)$. The related FBSDE reads

$$(5.1) \quad \begin{cases} X_t = x_0 + \int_0^t [A(C - X_s)BZ_s(t)] ds + \int_0^t \sigma dW_s, \\ Y_t = \langle GX_T, X_T \rangle - \int_t^T (\langle R_x X_s, X_s \rangle - \frac{1}{4} \langle R_u^{-1} B^\top Z_s, B^\top Z_s \rangle) ds \\ \quad + \int_t^T \langle Z_s, \sigma dW_s \rangle, \quad t \in [0, T]. \end{cases}$$

The solution to (5.1) is then given by $Y_t = X_t^\top P(t)X_t + X_t^\top Q(t) + R(t)$ and $Z_t = 2P(t)X_t + Q(t)$. To obtain a benchmark solution for Y and Z in our experiments, the Riccati equation is approximated with the explicit Euler method with 160×2^7 time steps, and X with 160 time steps.

5.1.1. Example with state and control of the same dimension. Our first example concerns a two-dimensional LQ control problem with two-dimensional control. The matrices for the forward equation are given by

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & 0.5 \\ -0.5 & 1 \end{pmatrix}, \quad C = \begin{pmatrix} 0.1 \\ 0.2 \end{pmatrix},$$

$$\sigma = \begin{pmatrix} 0.05 & 0.25 \\ 0.05 & 0.25 \end{pmatrix}, \quad x_0 = \begin{pmatrix} 0.1 \\ 0.1 \end{pmatrix}, \quad T = 0.5,$$

and the penalty matrices for the control problem by

$$R_x = \begin{pmatrix} 100 & 0 \\ 0 & 1 \end{pmatrix}, \quad R_u = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad G = \begin{pmatrix} 1 & 0 \\ 0 & 100 \end{pmatrix}.$$

In the top third of Table 1, we see the convergence rates from the experiment. The regime of the LQ control problem, with, e.g., quadratic dependence in f , does not satisfy the assumptions made in section 4 and a direct comparison cannot be made. Still, we see, for instance, that Y_0^h converges empirically with order 1 while the error in the terminal condition reaches 0.69 and is likely to continue to decrease. In Theorem 4.5 and Corollary 4.6, there is a difference of a factor two between these two errors, which roughly seems to be in line with the rates obtained.

5.1.2. Examples with control in lower dimensions than the state. Our second example concerns a six-dimensional problem with a two-dimensional control. The matrices used for the state equation are

$$A = \text{diag}([1, 2, 3, 1, 2, 3]), \quad B = \begin{pmatrix} 1 & 1 & 0.5 & 1 & 0 & 0 \\ -1 & 1 & 1 & -1 & -1 & 1 \end{pmatrix}^\top,$$

$$C = \text{diag}([-0.2, -0.1, 0, 0, 0.1, 0.2]), \quad \sigma = \text{diag}([0.05, 0.25, 0.05, 0.25, 0.05, 0.25]),$$

$$x_0 = (0.1, 0.1, 0.1, 0.1, 0.1, 0.1)^\top, \quad T = 0.5.$$

The penalty matrices of the control problem are given by

$$R_x = \text{diag}([25, 1, 25, 1, 25, 1]), \quad R_u = \text{diag}([1, 1]), \quad G = \text{diag}([1, 25, 1, 25, 1, 25]).$$

Before we discuss our results, recall that the optimal feedback control at time t is given by $u_t^* = -\frac{1}{2}R_u^{-1}B^\top Z_t$. Since u_t^* takes on values in \mathbb{R}^ℓ and $R_u^{-1}B^\top$ is of rank $\ell < d$ at most, we can conclude that there exist infinitely many processes ζ_t such that $u_t^* = -\frac{1}{2}R_u^{-1}B^\top \zeta_t$. To obtain uniqueness of the control component, we set $\lambda = 1 > 0$.

In Figure 3, the approximations are compared with semianalytic solutions in empirical mean and credible interval and for a representative path realization of X, Y , and Z . Visually, the approximations capture (X, Y, Z) well. The convergence is shown in the middle third of Table 1. The experimental orders decrease below the orders of the previous example. To investigate whether this is the true convergence order, or if other errors are dominating for small time steps, we have done some manual hyperparameter optimization with different training data and batch sizes, learning rates, and neural network architectures, without being able to improve the rates.

The third example aims to demonstrate our methods' ability to deal with high-dimensional problems. Most high-dimensional PDE and BSDE problems in the literature are symmetric (solutions are permutation invariant), and in some cases the solutions can be represented by a one-dimensional BSDE [43, Example 1]. From the parameters below, it becomes evident that the 25-dimensional problem that we choose is highly nonsymmetric and therefore very challenging (arguably more challenging

TABLE 1
Errors and experimental order of convergence for LQ control problems described in sections 5.1.1 and 5.1.2.

| N | $\ X - X^h\ _{S_h^2}$ | | $\ Y - Y^h\ _{S_h^2}$ | | $\ Z - Z^h\ _{\mathcal{H}_h^2}$ | | $\ Y_N^h - g(X_N^h)\ _{L_h^2}$ | | $ Y_0 - Y_0^h $ | | Y_0^h |
|---|-----------------------|------|-----------------------|-------|---------------------------------|------|--------------------------------|-------|-----------------|------|---------|
| | Error | EOC | Error | EOC | Error | EOC | Error | EOC | Error | EOC | Value |
| Problem 1 with $d = 2$ and $\ell = 2$ with analytic initial value $Y_0 = 0.6122$. | | | | | | | | | | | |
| 5 | 6.90e-2 | 1.28 | 1.22 | 1.03 | 5.66e-1 | 1.34 | 9.91e-1 | 0.98 | 7.10e-1 | 1.13 | 1.32 |
| 10 | 2.85e-2 | 1.11 | 6.02e-1 | 0.98 | 2.22e-1 | 1.16 | 5.04e-1 | 0.92 | 3.26e-1 | 1.16 | 0.937 |
| 20 | 1.32e-2 | 1.09 | 3.05e-1 | 0.86 | 1.00e-1 | 1.08 | 2.67e-1 | 0.80 | 1.47e-1 | 1.04 | 0.759 |
| 40 | 6.16e-3 | 1.00 | 1.68e-1 | 0.74 | 4.72e-2 | 0.98 | 1.53e-1 | 0.69 | 7.01e-2 | 1.01 | 0.683 |
| 80 | 3.07e-3 | | 1.01e-1 | | 2.39e-2 | | 9.46e-3 | | 3.47e-2 | | 0.645 |
| Problem 2 with $d = 6$ and $\ell = 2$ with analytic initial value $Y_0 = 1.4599$. | | | | | | | | | | | |
| 5 | 7.25e-2 | 1.23 | 5.65e-1 | 0.90 | 1.20 | 0.90 | 4.43e-1 | 0.80 | 3.51e-1 | 1.10 | 1.80 |
| 10 | 3.10e-2 | 1.10 | 3.02e-1 | 0.76 | 7.41e-1 | 0.49 | 2.54e-1 | 0.69 | 1.63e-1 | 0.94 | 1.55 |
| 20 | 1.45e-2 | 0.87 | 1.79e-1 | 0.63 | 5.26e-1 | 0.28 | 1.57e-1 | 0.60 | 8.51e-2 | 0.82 | 1.55 |
| 40 | 7.96e-3 | 0.35 | 1.15e-1 | 0.53 | 4.34e-1 | 0.18 | 1.04e-1 | 0.54 | 4.81e-2 | 0.69 | 1.51 |
| 80 | 6.24e-3 | | 7.96e-2 | | 3.84e-1 | | 7.15e-2 | | 3.00e-2 | | 1.49 |
| Problem 3 with $d = 25$ and $\ell = 1$ with analytic initial value $Y_0 = 11.348$. | | | | | | | | | | | |
| 5 | 2.25e-1 | 1.86 | 1.93 | 0.68 | 3.40 | - | 1.29 | 0.50 | 1.43 | 0.99 | 12.78 |
| 10 | 6.19e-2 | 0.97 | 1.21 | 0.42 | 2.54 | - | 9.15e-1 | 0.31 | 0.72 | 0.47 | 12.07 |
| 20 | 5.66e-2 | 0.57 | 9.00e-1 | 0.29 | 2.73 | - | 7.40e-1 | 0.19 | 0.52 | 0.53 | 11.87 |
| 40 | 2.48e-2 | 0.25 | 7.37e-1 | 0.069 | 2.72 | - | 6.47e-1 | 0.023 | 0.36 | 0.53 | 11.71 |
| 80 | 2.09e-2 | | 7.03e-1 | | 3.06 | | 6.37e-1 | | 0.25 | | 11.60 |

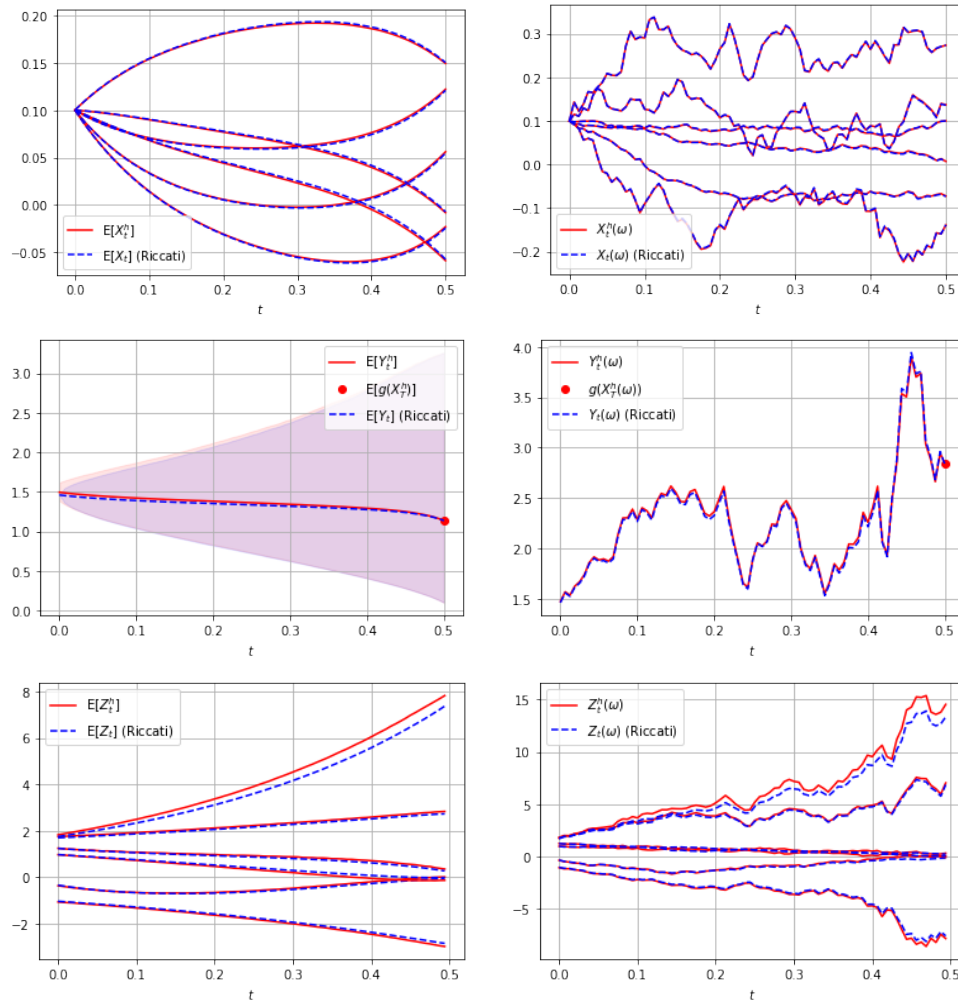


FIG. 3. Average of solutions and a single solution path compared to their analytic counterparts for the first LQ control problem from section 5.1.2, i.e., the problem with $d = 6$ and $\ell = 2$. The shaded area represents an empirical credible interval for Y , defined as the area between the 5th and the 95th percentile at each time point. We do not include credible intervals for X and Z in this figure to facilitate visualization. For X and Z , we see one realization of each of the six components.

than a similar but 100-dimensional symmetric problem). Nonsymmetric problems in the literature are [20, 40, 47, 52, 58], and the dimensions are 4, 5, 3, 4, and 2, respectively. A symmetric problem in 100 dimensions is found in [40].

Despite the challenging nature of the problem and its relatively high dimensionality, we achieve acceptable results, which are displayed in the bottom third of Table 1. It should, however, be pointed out that the error source induced by the time discretization is no longer dominating. This means that we do not see a convergence with the number of time steps for the approximation of the Z -process. All the other discretization errors decrease with the step size, but it is clear that we have other significant error sources. Figure 4 shows that visually the performance for 40 time steps is acceptable, even though some of the components of the Z -process oscillate

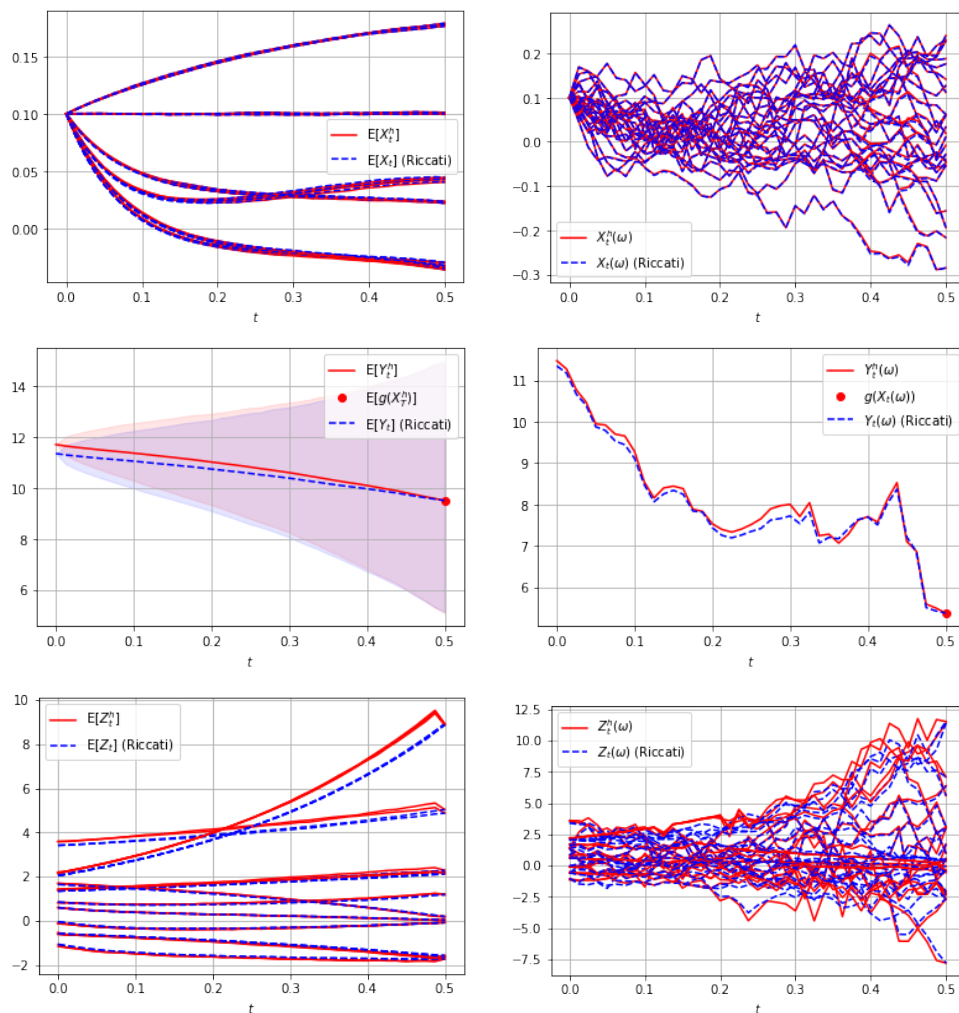


FIG. 4. Average of solutions and a single solution path compared to their analytic counterparts for the second LQ control problem from section 5.1.2, i.e., the problem with $d = 25$ and $\ell = 1$. The shaded area represents an empirical credible interval for Y , defined as the area between the 5th and the 95th percentile at each time point. We do not include credible intervals for X and Z in this figure to facilitate visualization. For X and Z , we see one realization of each of the 25 components.

close to the terminal time. The phenomena of accurate X - and Y -processes and less accuracy in some of the components of the Z -process could, at least heuristically, be explained by the mapping $\mathbb{R}^{25} \ni Z_t \mapsto u_t \in \mathbb{R}$. It is reasonable to assume that some of the components of the Z -process are more influential in the abovementioned mapping, which is what we have seen empirically in our experiments. Moreover, we have noticed that the components of the Z -process with the lowest magnitudes are less accurately approximated (relatively), which by the form of the feedback control also justifies the above reasoning.

We use the following parameters: $T = 0.5$, $d = 25$, $\ell = 1$, $A = \text{diag}([1, 2, 3, \dots, 1, 2, 3, 1])$, $B = (1, 1, 0.5, 1, 0, 0, 1, 1, 0.5, 1, 0, 0, 1, 1, 0.5, 1, 0, 0, 1, 1, 0.5, 1, 0, 0, 1)$, $C = (-0.2, -0.1, 0, 0, 0.1, 0.2, -0.2, -0.1, 0, 0, 0.1, 0.2, -0.2, -0.1, 0, 0, 0.1, 0.2, -0.2, -0.1, 0, 0, 0.1, 0.2, -0.2, -0.1, 0, 0, 0.1)$,

Downloaded 10/12/23 to 131.211.12.11 . Redistribution subject to SIAM license or copyright; see <https://epubs.siam.org/terms-privacy>

Let $x_0 \in \mathbb{R}^d$, $A, \Sigma, \in \mathbb{R}^{d \times d}$, $R_x, G \in \mathbb{S}_+^d$, $R_\alpha \in \mathbb{S}_+^\ell$, and $B \in \mathbb{R}^{d \times \ell}$ be of full rank and $C \in \mathbb{R}^d$. The state equation and cost functional of a nonlinear quadratic control problem are given by

Due to the linear dependence of the control and the quadratic cost functional, the optimal feedback control is again given by $u_t^* = -\frac{1}{2}R_u^{-1}B^\top D_x V(t, X_t)$. Similar to the above, V is the solution to the associated HJB equation. By setting $Y_t = V(t, X_t)$ and $Z_t = D_x V(t, X_t)$, we obtain the FBSDE

Particularly, we consider a three-dimensional problem with control in two dimensions, i.e., $d = 3$ and $\ell = 2$, and use the following matrices for the state:

$$\Sigma = \text{diag}([0.1, 0.1, 0.1]), \quad x_0 = (0.1, 0.1, 0.1)^\top, \quad T = 0.25.$$

$$R_x = \text{diag}([5, 1, 1]), \quad R_u = \text{diag}([1, 1]), \quad G = \text{diag}([1, 5, 1]).$$

TABLE 2

| | $\ Y_N^h - g(X_N^h)\ _{L_h^2}$ | | $ Y_0 - Y_0^h $ | | Y_0^h |
|-----|--------------------------------|------|-----------------|------|---------|
| N | Error | EOC | Error | EOC | Value |
| 5 | 2.69e-2 | 0.59 | 9.80e-3 | 1.01 | 0.2297 |
| 10 | 1.79e-2 | 0.53 | 4.85e-3 | 1.03 | 0.2241 |
| 20 | 1.24e-2 | 0.50 | 2.38e-3 | 0.99 | 0.2219 |
| 40 | 8.76e-3 | 0.49 | 1.20e-3 | 0.98 | 0.2207 |
| 80 | 6.27e-3 | | 6.07e-4 | | 0.2200 |

Acknowledgments. We very much thank Boualem Djehiche for his kindness in bringing some very useful references to our knowledge. We also thank the two referees for very useful comments that have helped us to improve the manuscript.

REFERENCES

- [1] K. ANDERSSON, *Approximate Stochastic Control Based on Deep Learning and Forward Backward Stochastic Differential Equations*, Master's thesis, Chalmers University of Technology, 2019.
- [2] F. ANTONELLI, *Backward Forward Stochastic Differential Equations*, Ph.D. thesis, Purdue University, 1993.
- [3] K. J. ÅSTRÖM, *Introduction to Stochastic Control Theory*, Courier Corporation, 2012.
- [4] C. BECK, S. BECKER, P. CHERIDITO, A. JENTZEN, AND A. NEUFELD, *Deep splitting method for parabolic PDEs*, SIAM J. Sci. Comput., 43 (2021), pp. A3135–A3154.
- [5] C. BECK, S. BECKER, P. GROHS, N. JAAFARI, AND A. JENTZEN, *Solving the Kolmogorov PDE by means of deep learning*, J. Sci. Comput., 88 (2021), pp. 1–28.
- [6] C. BECK, M. HUTZENTHALER, A. JENTZEN, AND B. KUCKUCK, *An Overview on Deep Learning-Based Approximation Methods for Partial Differential Equations*, preprint, <https://arxiv.org/abs/2012.12348>, 2020.
- [7] C. BECK, E. WEINAN, AND A. JENTZEN, *Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations*, J. Nonlinear. Sci., 29 (2019), pp. 1563–1619.
- [8] C. BENDER AND R. DENK, *A forward scheme for backward SDEs*, Stochastic Process. Appl., 117 (2007), pp. 1793–1812.
- [9] C. BENDER AND J. STEINER, *Least-squares Monte Carlo for backward SDE*, in Numerical Methods in Finance, Springer, New York, 2012, pp. 257–289.
- [10] J. BERNER, P. GROHS AND A. JENTZEN, *Analysis of the generalization error: Empirical risk minimization over deep artificial neural networks overcomes the curse of dimensionality in the numerical approximation of Black-Scholes partial differential equations*, SIAM J. Math. Data Sci., 2 (2020), pp. 631–657.
- [11] W.-J. BEYN AND R. KRUSE, *Two-sided error estimates for the stochastic theta method*, Discrete Contin. Dyn. Syst. Ser. B, 14 (2010), pp. 389–407.
- [12] J. F. BONNANS, J. GIANATTI, AND F. J. SILVA, *On the time discretization of stochastic optimal control problems: The dynamic programming approach*, ESAIM Control Optim. Calc. Var., 25 (2019), 63.
- [13] B. BOUCHARD AND N. TOUZI, *Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations*, Stochastic Process. Appl., 111 (2004), pp. 175–206.
- [14] G. BUTTAZZO AND G. DAL MASO, *γ -convergence and optimal control problems*, J. Optim. Theory Appl., 38 (1982), pp. 385–407.
- [15] Q. CHAN-WAI-NAM, J. MIKAEL, AND X. WARIN, *Machine learning for semi linear PDEs*, J. Sci. Comput., 79 (2019), pp. 1667–1712.
- [16] K. W. CHAU AND C. W. OOSTERLEE, *Stochastic grid bundling method for backward stochastic differential equations*, Int. J. Comput. Math., 96 (2019), pp. 2272–2301.
- [17] J. CHESSARI AND R. KAWAI, *Numerical Methods for Backward Stochastic Differential Equations: A Survey*, preprint, <https://arxiv.org/abs/2101.08936>, 2021.
- [18] D. CRISAN, K. MANOLARAKIS, AND N. TOUZI, *On the Monte Carlo simulation of BSDEs: An improvement on the Malliavin weights*, Stochastic Process. Appl., 120 (2010), pp. 1133–1158.
- [19] G. CYBENKO, *Approximations by superpositions of a sigmoidal function*, Math. Control Signals Systems, 2 (1989), pp. 183–192.
- [20] B. DAI, V. R. SURABHI, P. KRISHNAMURTHY, AND F. KHORRAMI, *Learning Locomotion Controllers for Walking Using Deep FBSDE*, preprint, <https://arxiv.org/abs/2107.07931>, 2021.
- [21] D. ELBRÄCHTER, P. GROHS, A. JENTZEN, AND C. SCHWAB, *DNN expression rate analysis of high-dimensional PDEs: Application to option pricing*, Constr. Approx., 55 (2022), pp. 3–71.
- [22] A. FAHIM, N. TOUZI, AND X. WARIN, *A probabilistic numerical method for fully nonlinear parabolic PDEs*, Ann. Appl. Probab., 21 (2011), pp. 1322–1364.
- [23] F. FANG AND C. W. OOSTERLEE, *A novel pricing method for European options based on Fourier-cosine series expansions*, SIAM J. Sci. Comput., 31 (2009), pp. 826–848.

- [24] M. FUJII, A. TAKAHASHI, AND M. TAKAHASHI, *Asymptotic expansion as prior knowledge in deep learning method for high dimensional BSDEs*, Asia-Pac. Financ. Mark., 26 (2019), pp. 391–408.
- [25] A. GNOATTO, C. REISINGER, AND A. PICARELLI, *Deep xVA Solver—A Neural Network Based Counterparty Credit Risk Management Framework*, SSRN 3594076, 2020.
- [26] E. GOBET AND C. LABART, *Error expansion for the discretization of backward stochastic differential equations*, Stochastic Process. Appl., 117 (2007), pp. 803–829.
- [27] E. GOBET AND C. LABART, *Solving BSDE with adaptive control variate*, SIAM J. Numer. Anal., 48 (2010), pp. 257–277.
- [28] E. GOBET, J. G. LÓPEZ-SALAS, P. TURKEDJIEV, AND C. VÁZQUEZ, *Stratified regression Monte Carlo scheme for semilinear PDEs and BSDEs with large scale parallelization on GPUs*, SIAM J. Sci. Comput., 38 (2016), pp. C652–C677.
- [29] P. GROHS, F. HORNING, A. JENTZEN, AND P. VON WURSTEMBERGER, *A Proof that Artificial Neural Networks Overcome the Curse of Dimensionality in the Numerical Approximation of Black–Scholes Partial Differential Equations*, <https://arxiv.org/abs/1809.02362>, 2018.
- [30] J. HAN, *Deep Learning Approximation for Stochastic Control Problems*, preprint, <https://arxiv.org/abs/1611.07422>, 2016.
- [31] J. HAN, A. JENTZEN, AND E. WEINAN, *Solving high-dimensional partial differential equations using deep learning*, Proc. Natl. Acad. Sci. USA, 115 (2018), pp. 8505–8510.
- [32] J. HAN AND J. LONG, *Convergence of the deep BSDE method for coupled FBSDEs*, Probab. Uncertain. Quant. Risk, 5 (2020), pp. 1–33.
- [33] P. HENRY-LABORDERE, *Deep Primal-Dual Algorithm for BSDEs: Applications of Machine Learning to CVA and IM*, SSRN 3071506, 2017.
- [34] T. HUIJSKENS, M. J. RUITER, AND C. W. OOSTERLEE, *Efficient numerical fourier methods for coupled forward-backward SDEs*, J. Comput. Appl. Math., 296 (2016), pp. 593–612.
- [35] C. HURÉ, H. PHAM, AND X. WARIN, *Some Machine Learning Schemes for High-Dimensional Nonlinear PDEs*, preprint, <https://arxiv.org/abs/1902.01599>, 2019.
- [36] C. HURÉ, H. PHAM, AND X. WARIN, *Deep backward schemes for high-dimensional nonlinear PDEs*, Math. Comp., 89 (2020), pp. 1547–1579.
- [37] M. HUTZENTHALER, A. JENTZEN, AND T. KRUSE, *On multilevel Picard numerical approximations for high-dimensional nonlinear parabolic partial differential equations and high-dimensional nonlinear backward stochastic differential equations*, J. Sci. Comput., 79 (2019), pp. 1534–1571.
- [38] M. HUTZENTHALER, A. JENTZEN, T. KRUSE, AND T. A. NGUYEN, *A proof that rectified deep neural networks overcome the curse of dimensionality in the numerical approximation of semilinear heat equations*, SN Partial Differ. Equ. Appl., 1 (2020), pp. 1–34.
- [39] A. JENTZEN, D. SALIMOVA, AND T. WELTI, *A Proof that Deep Artificial Neural Networks Overcome the Curse of Dimensionality in the Numerical Approximation of Kolmogorov Partial Differential Equations with Constant Diffusion and Nonlinear Drift Coefficients*, preprint, <https://arxiv.org/abs/1809.07321>, 2018.
- [40] S. JI, S. PENG, Y. PENG, AND X. ZHANG, *Three algorithms for solving high-dimensional fully coupled FBSDEs through deep learning*, IEEE Intell. Syst., 35 (2020), pp. 71–84.
- [41] S. JI, S. PENG, Y. PENG, AND X. ZHANG, *A Control Method for Solving High-Dimensional Hamiltonian Systems Through Deep Neural Networks*, preprint, <https://arxiv.org/abs/2111.02636>, 2021.
- [42] Y. JIANG AND J. LI, *Convergence of the Deep BSDE Method for FBSDEs with Non-Lipschitz Coefficients*, preprint, <https://arxiv.org/abs/2101.01869>, 2021.
- [43] L. KAPLLANI AND L. TENG, *Deep Learning Algorithms for Solving High Dimensional Nonlinear Backward Stochastic Differential Equations*, preprint, <https://arxiv.org/abs/2010.01319>, 2020.
- [44] D. P. KINGMA AND J. BA, *Adam: A Method for Stochastic Optimization*, preprint, <https://arxiv.org/abs/1412.6980>, 2014.
- [45] A. KOHATSU-HIGA, A. LEJAY, AND K. YASUDA, *Weak rate of convergence of the Euler–Maruyama scheme for stochastic differential equations with non-regular drift*, J. Comput. Appl. Math., 326 (2017), pp. 138–158.
- [46] R. KRUSE, *Characterization of bistability for stochastic multistep methods*, BIT, 52 (2012), pp. 109–140.
- [47] Y. LIU, Y. WANG, Z. ZHUANG, AND X. GUO, *Deep FBSDE Controller for Attitude Control of Hypersonic Aircraft*, in Proceedings of the 6th IEEE International Conference on Advanced Robotics and Mechatronics, IEEE, 2021, pp. 294–299.
- [48] J. MA, P. PROTTER, J. SAN MARTÍN, AND S. TORRES, *Numerical method for backward stochastic differential equations*, Ann. Appl. Probab., 12 (2002), pp. 302–316.

- [49] J. MA, P. PROTTER, AND J. YONG, *Solving forward-backward stochastic differential equations explicitly – A four step scheme*, Probab. Theory Related Fields, 98 (1994), pp. 339–359.
- [50] R. MIKULEVICIUS AND E. PLATE, *Rate of convergence of the Euler approximation for diffusion processes*, Math. Nachr., 151 (1991), pp. 233–239.
- [51] B. NEGYESI, K. ANDERSSON, AND C. W. OOSTERLEE, *The One Step Malliavin scheme: New Discretization of BSDEs Implemented with Deep Learning Regressions*, preprint, <https://arxiv.org/abs/2110.05421>, 2021.
- [52] M. PEREIRA, Z. WANG, I. EXARCHOS, AND E. A. THEODOROU, *Learning Deep Stochastic Optimal Control Policies Using Forward-Backward SDEs*, preprint, <https://arxiv.org/abs/1902.03986>, 2019.
- [53] M. RAISSI, *Forward-Backward Stochastic Neural Networks: Deep Learning of High-Dimensional Partial Differential Equations*, preprint, <https://arxiv.org/abs/1804.07010>, 2018.
- [54] R. T. ROCKAFELLAR AND R. J.-B. WETS, *Variational Analysis*, Grundlehren Math. Wiss. 317, Springer Science & Business, Berlin, 2009.
- [55] M. J. RUIJTER AND C. W. OOSTERLEE, *A Fourier cosine method for an efficient computation of solutions to BSDEs*, SIAM J. Sci. Comput., 37 (2015), pp. A859–A889.
- [56] M. J. RUIJTER AND C. W. OOSTERLEE, *Numerical Fourier method and second-order Taylor scheme for backward SDEs in finance*, Appl. Numer. Math., 103 (2016), pp. 1–26.
- [57] L. TENG, *A review of tree-based approaches to solve forward-backward stochastic differential equations*, J. Comput. Finance, 25 (2021), pp. 125–159.
- [58] Y. WANG AND Y.-H. NI, *Deep BSDE-ML Learning and Its Application to Model-Free Optimal Control*, preprint, <https://arxiv.org/abs/2201.01318>, 2022.
- [59] Y. YU, B. HIENTZSCH AND N. GANESAN, *Backward Deep BSDE Methods and Applications to Nonlinear Problems*, preprint, <https://arxiv.org/abs/2006.07635>, 2020.