Modern machine learning methods for high-dimensional forward-backward SDEs with applications in mathematical finance

Négyesi, B.

DOI
10.4233/uuid:bb7bee07-8690-4575-83c7-cdf4aabb4263

Publication date
2025

Document Version
Final published version

Citation (APA)
Négyesi, B. (2025). *Modern machine learning methods for high-dimensional forward-backward SDEs with applications in mathematical finance*. [Dissertation (TU Delft), Delft University of Technology]. https://doi.org/10.4233/uuid:bb7bee07-8690-4575-83c7-cdf4aabb4263

Modern machine learning methods for high-dimensional forward-backward SDEs with applications in mathematical finance

Bálint Négyesi

Doctoral dissertation
2025

TU Delft Delft University of Technology

MODERN MACHINE LEARNING METHODS FOR HIGH-DIMENSIONAL FORWARD-BACKWARD SDES WITH APPLICATIONS IN MATHEMATICAL FINANCE

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus prof. dr. ir. T.H.J.J. van der Hagen,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op vrijdag 26 september 2025 om 12:30 uur

door

Bálint NÉGYESI

Master of Science in Applied Mathematics,
Technische Universiteit Delft, Nederland,
geboren te Boedapest, Hongarije.

Dit proefschrift is goedgekeurd door de promotoren.

Samenstelling promotiecommissie:

Rector Magnificus,	voorzitter
Prof. dr. ir. C.W. Oosterlee,	Universiteit Utrecht, promotor
Prof. dr. ir. M.C. Veraar,	Technische Universiteit Delft, promotor
Onafhankelijke leden:	
Prof. dr. C. Bender,	Universität des Saarlandes, Duitsland
Dr. F. Fang,	Technische Universiteit Delft
Prof. dr. A. Gnoatto,	Università degli Studi di Verona, Italië
Prof. dr. A. Papapantoleon,	Technische Universiteit Delft
Prof. dr. N.F.F. Schweizer,	Tilburg University, Nederland
Prof. dr. A.W. van der Vaart,	Technische Universiteit Delft, reservelid

Het onderzoek voor dit proefschrift is mede gefinancierd door het Peter Paul Peterich fonds.

Keywords: forward-backward stochastic differential equations, deep BSDE, neural networks, machine learning, One Step Malliavin scheme, Greeks, Gamma approximations, discretely reflected BSDE, delta-gamma hedging, fully-coupled FBSDE, a posteriori convergence, stochastic optimal control, BCOS, Fourier cosine expansions, second-order Taylor schemes

Printed by: Ipskamp Printing

ISBN 978-94-6518-086-1

An electronic version of this dissertation is available at
http://repository.tudelft.nl/.

“almost nothing important that ever happens to you happens because you engineer it. Destiny has no beeper; destiny always leans trenchcoated out of an alley with some sort of ’psst’ that you usually can’t even hear because you’re in such a rush to or from something important you’ve tried to engineer.”

David Foster Wallace (Infinite Jest)

CONTENTS

SUMMARY

Backward stochastic differential equations (BSDE) are a fundamental tool in the mathematical modelling of financial problems. Through the famous nonlinear extensions to the Feynman-Kac formula, they do not merely provide a stochastic representation of the solution to large classes of partial differential equations such as pricing- or Hamilton-Jacobi-Bellman equations, but also include sensitivities, corresponding to derivatives of the solution, which are crucial in many financial mathematical applications. Henceforth, they simultaneously represent option pricing and hedging problems, and form a natural framework for the numerical treatment of stochastic optimal control.

The main challenge in the numerical approximation of such equations is the computation of conditional expectations over potentially high-dimensional spaces. In classical settings, where the dimensionality of the underlying randomness is moderate, many approaches have been proposed in the literature. However, for high-dimensional problems, one has to resort to Monte Carlo methods. In recent years, a new class of regression Monte Carlo methods has arisen in the literature, so called deep BSDE methods, which practically approximate the solution of BSDEs in a neural network regression Monte Carlo framework, after forming a suitable loss function motivated either by stochastic optimal control or the martingale representation theorem. These classes of methods can roughly be divided into two main categories. Forward methods, where the solution of the associated backward SDE is simultaneously optimized in a global optimization, minimizing a loss function stemming from a stochastic target problem reformulation. Alternatively, backward methods have been investigated, where the numerical resolution of the equation is decomposed into smaller sub-optimizations corresponding to a discrete set of points in a suitable time discretization. These methods enabled the numerical treatment of longstanding open challenges, such as the pricing and delta-hedging of multi-asset financial options up to $d = 100$ risk factors and beyond.

The goal of this thesis is to analyze such modern machine learning based numerical methods, and apply them in the financial mathematical context. We propose numerical extensions of these methods in high-dimensional frameworks, analyze their convergence properties in discrete time, and investigate their robustness and accuracy in practical applications such as hedging and stochastic optimal control. Our main contributions in each chapter can be summarized as follows.

In chapter 2 a novel discretization, the One Step Malliavin (OSM) scheme, is presented for decoupled forward-backward stochastic differential equations with differentiable coefficients, simultaneously solving the BSDE and its Malliavin sensitivity problem. The control process is estimated by the corresponding linear BSDE driving the trajectories of the Malliavin derivatives of the solution pair, which implies the need to provide accurate second derivatives, Γ estimates. The approximation is based on a merged formulation given by the Feynman-Kac formulae and the Malliavin chain rule. In order to allow for an efficient numerical solution of the arising semi-discrete conditional

expectations in possibly high-dimensions, it is fundamental that the chosen approach admits to differentiable estimates. Two fully-implementable schemes are considered: the BCOS method as a reference in the one-dimensional framework and neural network Monte Carlo regressions in case of high-dimensional problems, similarly to the recently emerging class of deep BSDE methods. An error analysis is carried out to show $\mathbb{L}^2$ convergence of order 1/2, under standard Lipschitz assumptions and additive noise in the forward diffusion. Numerical experiments are provided for a range of different semi-linear equations up to 50 dimensions, demonstrating that the proposed scheme yields a significant improvement in the control estimations.

Extending the OSM scheme to reflected BSDEs, chapter 3 presents a deep BSDE approach for the pricing and delta-gamma hedging of high-dimensional Bermudan options, with applications in portfolio risk management. Large portfolios of a mixture of multi-asset European and Bermudan derivatives are cast into the framework of discretely reflected BSDEs. This system is discretized by the One Step Malliavin scheme, and solved in a deep BSDE formulation, involving a Γ process, corresponding to second-order sensitivities of the associated option prices. The resulting option Deltas and Gammas are used to discretely rebalance the corresponding replicating strategies. Numerical experiments are presented on both high-dimensional basket options and large portfolios consisting of multiple options with varying early exercise rights, moneyness and volatility. These examples demonstrate the robustness and accuracy of the method up to 100 risk factors. The resulting hedging strategies significantly outperform benchmark methods both in the case of standard delta- and delta-gamma hedging.

In the last two chapters, we turn our attention to stochastic optimal control problems. First, in chapter 4, we are concerned with high-dimensional coupled forward-backward stochastic differential equations (FBSDE) approximated by the deep BSDE method of Han et al. (2018). It was shown by Han and Long (2020) that the errors induced by the deep BSDE method admit an a posteriori estimate depending on the loss function, whenever the backward equation only couples into the forward diffusion through the Y process. We generalize this result to drift coefficients that may also depend on Z, and give sufficient conditions for convergence under standard assumptions. The resulting conditions are directly verifiable for any equation. Consequently, unlike in earlier theory, our convergence analysis enables the treatment of FBSDEs stemming from stochastic optimal control problems. In particular, we provide a theoretical justification for the non-convergence of the deep BSDE method observed in recent literature, and present direct guidelines for when convergence can be guaranteed in practice. Our theoretical findings are supported by several numerical experiments in high-dimensional settings.

Chapter 5 presents a higher-order numerical method for scalar valued, coupled FBSDEs. Unlike most classical references, the forward component is not only discretized by an Euler-Maruyama approximation but also by higher-order Taylor schemes. This includes the famous Milstein scheme, providing an improved strong convergence rate of order 1; and the simplified order 2.0 weak Taylor scheme exhibiting weak convergence rate of order 2. In order to have a fully-implementable scheme in case of these higher-order Taylor approximations, which involve the derivatives of the decoupling fields, we use a Fourier cosine expansion method, the COS method, built on the known conditional characteristic function of each transition, to approximate the conditional expectations

arising from the backward component. Even though higher-order numerical approximations for the backward equation are deeply studied in the literature, to the best of our understanding, the present numerical scheme is the first which achieves strong convergence of order 1 for the whole coupled system, including the forward equation, which is often the main interest in applications such as stochastic control. Numerical experiments demonstrate the proclaimed higher-order convergence, both in case of strong and weak convergence rates, for various equations ranging from decoupled to the fully-coupled settings.

The findings of this thesis contribute to the frontiers of the numerical approximation of high-dimensional BSDEs in several different aspects of mathematical finance. From the numerical analysis point of view, the derivation and convergence of the One Step Malliavin scheme provides a novel discrete time approximation scheme including second-order option Greeks. Consequently, the convergence of its deep BSDE formulation, enables the delta-gamma hedging of large basket options with early-exercise features. The generalized convergence of the forward deep BSDE method proves sufficient conditions under which an *a posteriori* convergence bound can be guaranteed. Consequently, our results allow for the coupling of the control process in the forward diffusion, and therefore, enable the deep BSDE treatment of stochastic optimal control problems stemming from the dynamic programming principle in case of drift control. Finally, in a more classical framework, we generalize the COS method to more general second-order Taylor discretizations of the underlying forward diffusion. Doing so, we obtain numerical methods which have higher-order convergence rates both in the weak- and strong senses than reference approaches.

SAMENVATTING

Achterwaartse stochastische differentiaalvergelijkingen (BSDE) vormen een fundamenteel instrument in de wiskundige modellering van financiële vraagstukken. Door niet-lineaire uitbreidingen van de Feynman-Kac-formule bieden zij niet alleen een stochastische representatie voor oplossingen van grote klassen van partiële differentiaalvergelijkingen zoals prijsbepalings- of Hamilton-Jacobi-Bellman-vergelijkingen, maar omvatten zij ook afgeleiden die cruciaal zijn in veel financiële wiskundige toepassingen. Hierdoor vertegenwoordigen zij zowel optieprijsbepalings- als hedgingproblemen en bieden zij een kader voor de numerieke behandeling van stochastische optimale controle.

De voornaamste uitdaging bij de numerieke benadering van dergelijke vergelijkingen is de berekening van conditionele verwachtingen in mogelijk hoog-dimensionale ruimtes. In klassieke settings, waar de dimensie van de onderliggende stochastiek beperkt is, zijn er veel methoden bekend in de literatuur. Echter, voor hoog-dimensionale problemen moet men terugvallen op Monte Carlo-methoden. In recente jaren is er echter een nieuwe klasse van regressie-Monte Carlo-methoden in de literatuur verschenen, de zogenaamde deep BSDE-methoden, die de oplossing van BSDE's benaderen in een neuraal netwerk-regressiekader met Monte Carlo simulatie, na het opstellen van een geschikte verliesfunctie die gemotiveerd is door stochastische optimale regeltechniek of de martingaalrepresentatiestelling. Deze methoden kunnen grofweg in twee hoofdcategorieën worden onderverdeeld. Voorwaartse methoden, waarbij de oplossing van de bijbehorende achterwaartse SDE gelijktijdig wordt geoptimaliseerd in een globale optimalisatie, door minimalisatie van een verliesfunctie die voortkomt uit een stochastische doelprobleemherformulering. Anderzijds zijn achterwaartse methoden voorgesteld, waarbij de numerieke oplossing van de vergelijking wordt opgesplitst in suboptimalisaties die overeenkomen met een discrete verzameling punten in een geschikte tijdsdiscretisatie. Deze methoden hebben de numerieke behandeling mogelijk gemaakt van langdurig openstaande uitdagingen, zoals de prijsbepaling en delta-hedging van multi-asset financiële opties tot $d = 100$ risicofactoren en meer.

Het doel van dit proefschrift is om dergelijke moderne machine learning-gebaseerde numerieke methoden te analyseren en toe te passen binnen de financiële wiskunde. We stellen numerieke uitbreidingen van deze methoden voor in hoog-dimensionale kaders, analyseren hun convergentie-eigenschappen in discrete tijd en onderzoeken hun robuustheid en nauwkeurigheid in praktische toepassingen zoals hedging en stochastische optimale regeltechniek toepassingen. Onze belangrijkste bijdragen in elk hoofdstuk kunnen als volgt worden samengevat.

In hoofdstuk 2 wordt een nieuwe discretisatie, het One Step Malliavin-schema (OSM), gepresenteerd voor ontkoppelde voorwaarts-achterwaartse stochastische differentiaalvergelijkingen met differentieerbare coëfficiënten, waarbij de BSDE en het Malliavin-afgeleide probleem gelijktijdig worden opgelost. Het controleproces wordt geschat door de overeenkomstige lineaire BSDE die de trajecten genereert van de

Malliavin-afgeleiden van het oplossingspaar, wat impliceert dat nauwkeurige Γ-schattingen nodig zijn. De benadering is gebaseerd op een formulering gegeven door de Feynman-Kac-formules en de Malliavin-ketenregel. Om een efficiënte numerieke oplossing mogelijk te maken van de resulterende semi-discrete conditionele verwachtingen in hoge dimensies, is het fundamenteel dat de gekozen methode differentieerbare schattingen toelaat. Twee implementeerbare schema's worden beschouwd: de BCOS-methode als referentie in de ééndimensionale context en neurale netwerk Monte Carlo-regressies voor hoog-dimensionale problemen, vergelijkbaar met de recent opkomende klasse van deep BSDE-methoden. Een foutenanalyse wordt uitgevoerd om $\mathbb{L}^2$-convergentie van orde $1/2$ aan te tonen, onder standaard Lipschitz-voorwaarden en additieve ruis in de voorwaartse diffusie. Numerieke experimenten worden gepresenteerd voor verschillende semi-lineaire vergelijkingen tot 50 dimensies, waaruit blijkt dat het voorgestelde schema een significante verbetering biedt.

Door het OSM-schema uit te breiden naar gereflecteerde BSDE's, wordt in hoofdstuk 3 een deep BSDE-benadering gepresenteerd voor de prijsbepaling en delta-gamma-hedging van hoog-dimensionale Bermuda opties, met toepassingen in portefeuillerisicobeheer. Grote portefeuilles van een mix van multi-asset Europese en Bermuda derivaten worden in het kader van discreet gereflecteerde BSDE's gegoten. Dit systeem wordt gediscretiseerd met het One Step Malliavin-schema, waarbij een Γ-proces betrokken is, dat overeenkomt met de tweede-orde afgeleiden van de bijbehorende optieprijzen. Het gediscretiseerde systeem wordt efficiënt opgelost met een neurale netwerkregressie Monte Carlo-methode voor een groot aantal onderliggende activa. De resulterende optie-Deltas en -Gammas worden gebruikt om de bijbehorende replicatiestrategieën discreet te herbalanceren. Numerieke experimenten worden gepresenteerd voor zowel hoog-dimensionale opties als grote portefeuilles met meerdere opties met variërende uitoefenrechten, moneyness en volatiliteit. Deze voorbeelden demonstreren de robuustheid en nauwkeurigheid van de methode tot 100 risicofactoren. De resulterende hedgingstrategieën presteren aanzienlijk beter dan benchmarkmethoden in zowel standaard delta- als delta-gamma-hedging.

In de laatste twee hoofdstukken richten we ons op hoog-dimensionale stochastische optimale controleproblemen. Eerst behandelen we in hoofdstuk 4 hoog-dimensionale gekoppelde voorwaarts-achterwaartse stochastische differentiaalvergelijkingen (FBSDE) die worden benaderd door de deep BSDE-methode van Han et al. (2018). Han en Long (2020) toonden aan dat de fouten geïnduceerd door de deep BSDE-methode een a posteriori afschatting toelaten, afhankelijk van de verliesfunctie, wanneer de achterwaartse vergelijking alleen in de voorwaartse diffusie koppelt via het Y proces. Wij generaliseren dit resultaat naar driftcoëfficiënten die ook afhankelijk kunnen zijn van Z, en geven voldoende voorwaarden voor convergentie onder standaard aannames. De resulterende voorwaarden zijn direct verifieerbaar voor elke vergelijking. Onze theoretische bevindingen worden ondersteund door meerdere numerieke experimenten in hoog-dimensionale settings.

Hoofdstuk 5 presenteert een hogere-orde numerieke methode voor scalaire, gekoppelde FBSDE's. In tegenstelling tot de meeste klassieke benaderingen wordt de voorwaartse component niet alleen gediscretiseerd met een Euler-Maruyama-approximate maar ook met hogere-orde Taylor-schema's. Deze omvatten het beroemde Milstein-

schema, dat een verbeterde sterke convergentiesnelheid van orde 1 biedt, en het vereenvoudigde orde 2.0 zwakke Taylor-schema dat een zwakke convergentie van orde 2 vertoont. Numerieke experimenten demonstreren de hogere-orde convergentie, zowel voor sterke als zwakke convergentiesnelheden, voor verschillende vergelijkingen variërend van ontkoppelde tot volledig gekoppelde problemen.

De bevindingen van dit proefschrift dragen bij aan de vooruitgang van de numerieke benadering van hoog-dimensionale BSDE's op verschillende gebieden binnen de financiële wiskunde. Vanuit het perspectief van numerieke analyse biedt de afleiding en convergentie van het One Step Malliavin-schema een nieuw discrete-tijd-approximatieschema dat tweede-orde optie-Greeks omvat. Daardoor maakt de convergentie van de deep BSDE-formulering de delta-gamma-hedging van hoog-dimensionale opties met vroegtijdige uitoefeningskenmerken mogelijk. De gegeneraliseerde convergentie van de voorwaartse deep BSDE-methode geeft ons voldoende voorwaarden waaronder een *a posteriori* convergentiegrens kan worden gegarandeerd. Bijgevolg stellen onze resultaten de koppeling van het controleproces in de voorwaartse diffusie mogelijk, en daardoor de deep BSDE-behandeling van stochastische optimale controleproblemen die voortkomen uit het dynamische programmeringsprincipe bij driftcontrole. Ten slotte, in een meer klassiek kader, generaliseren we de COS-methode naar meer algemene tweede-orde Taylor-discretisaties van de onderliggende voorwaartse diffusie. Door dit te doen verkrijgen we numerieke methoden met een hogere-orde convergentiesnelheid, zowel in de zwakke als in de sterke zin, dan referentiebenaderingen.

1

INTRODUCTION

Spring maar in de lampen
want de vloer begint te dampen[1]

DJ Maurice & Snollebollekes (De vloer is lava)

[1] Jump into the lights
because the floor is starting to steam

DJ Maurice & Snollebollekes (The floor is lava)

This dissertation is concerned with the numerical approximation of backward stochastic differential equations (BSDE). Such equations in their standard form can be written as follows

$$Y_t = \xi + \int_t^T f(t, Y_t, Z_t)\mathrm{d}t - \int_t^T Z_t \mathrm{d}W_t, \tag{1.1}$$

where ξ is a random variable, f a random field admitting certain integrability conditions, and $W_t, t \geq 0$ is a d-dimensional Brownian motion defined over an appropriate probability space. These equations were discovered in the seminal papers of Bismut in [20, 19] in the linear case, and thoroughly studied and extended by a series of works of Etienne Pardoux, Shige Peng and their co-authors in the 1990s, we mention [127, 126, 129, 128, 131, 130] among others. Throughout the last several decades, BSDEs enjoyed great attention from several areas of the mathematics community. From the stochastic analysis perspective, much work has gone into proving sufficient well-posedness of the equations in various settings [103, 125, 124, 123, 7, 100, 97, 52, 80], regularity and density estimates [167, 96, 95, 78, 109, 108, 1] or Malliavin differentiability [46, 110, 57, 79]. From the numerical analysis point of view, the main challenge in the numerical approximation of equations as (1.1) comes down to efficient and accurate approximations of conditional expectations over potentially high-dimensional domains. Recently, the equations gained a renowned attention due to advancements in the numerical approximation of high-dimensional regression problems, which is the core subject of this work.

The connection between BSDEs and conditional expectations is already apparent from the theoretical perspective. In a similar way as stochastic differential equations (SDE) provide a non-linear extension to Itô integrals, BSDEs are a non-linear extension to the martingale representation theorem, see e.g. [89]. To this end, let us consider an $\mathscr{F}_T$ measurable, L^2 integrable random variable ξ. According to the martingale representation theorem, we have that under sufficient conditions this random variable admits the following predictable representation

$$\xi = \mathbb{E}[\xi] + \int_0^T Z_s \mathrm{d}W_s,$$

and uniqueness of the martingale integrand, and thus the representation, can be ensured by imposing specific integrability conditions [42]. Defining $Y_t := \mathbb{E}[\xi|\mathscr{F}_t]$ as the Doob martingale of the aforementioned random variable, one immediately has by standard properties of the stochastic integral and conditional expectations

$$Y_t = \mathbb{E}[\xi] + \int_0^t Z_s \mathrm{d}W_s.$$

Combining the last two expressions results in the zero driver BSDE with $f \equiv 0$

$$Y_t = \xi - \int_t^T Z_s \mathrm{d}W_s.$$

Equation (1.1) can be thought of as a non-linear generalization of this relation, due to the additional integral term.

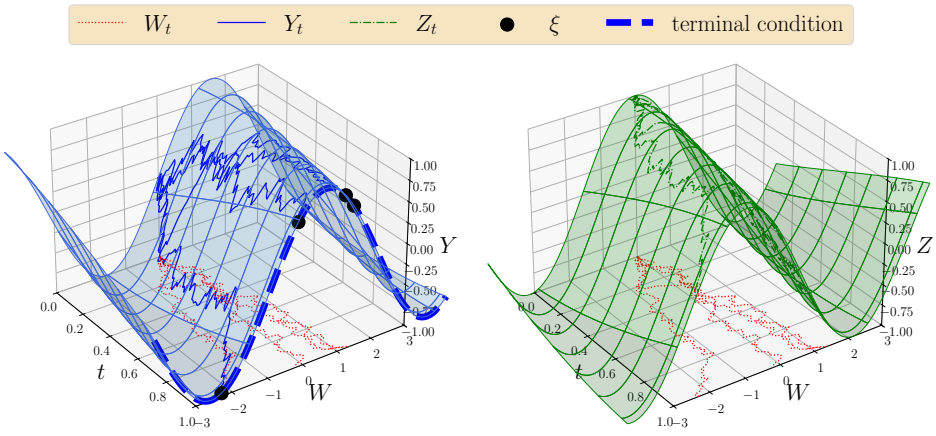

Figure 1.1: BSDE illustration, equation from [144, ex. 1]. On the left Y, on the right Z. Realizations of the Brownian motion in dotted red of the horizontal plane. Realizations of the solution pair in solid blue and dashed green, respectively. Realizations of the terminal condition marked by black dots. The Markovian mappings, mapping the diffusion process in (1.3) to ξ, Y, Z are marked by, dashed blue, and the surfaces in blue and green, respectively.

Unlike in the case of standard, forward SDEs, the solution of (1.1) is a *pair* of stochastic processes $\{(Y_t, Z_t)\}_{0\leq t\leq T}$, such that the equation is satisfied almost surely. The control process, Z, ensures that the corresponding solution pair is *adapted* with respect to the corresponding filtration, subtracting the right amount of randomness. An illustration of the stochastic processes involved in (1.1) is given in figure 1.1. The terms involved in (1.1) are generally referred to as ξ being the (random) terminal condition, f as the driver or generator.

1.1. SOURCES OF RANDOMNESS, DIFFERENT TYPES OF BSDES

Equation (1.1) is incomplete on its own, in that it does not specify the source of randomness determining the coefficients ξ, f. In the BSDE literature, there are several variants of (1.1), depending on this. In this thesis, we consider two special cases: so called *decoupled* and *coupled* systems of forward-backward stochastic differential equations (FBSDE). In these cases, the backward equation can be written in terms of *deterministic mappings* g and f

$$Y_t = g(X_T) + \int_t^T f(s, X_s, Y_s, Z_s)\mathrm{d}s - \int_t^T Z_s \mathrm{d}W_s, \tag{1.2}$$

where both the terminal condition and the driver depend on an extra stochastic process X, which – in case of FBSDEs – solves a forward SDE itself

$$X_t = x_0 + \int_0^t \mu(s, X_s, Y_s, Z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s, Y_s, Z_s)\mathrm{d}W_s. \tag{1.3}$$

In case the deterministic coefficients in (1.3) do not depend on the arguments Y, Z, the SDE can be solved independently of (1.2), hence the equations are *decoupled*. On the contrary, when μ and σ take Y, Z as arguments, (1.3)-(1.1) form a *coupled* system of equations, whose solution is a triplet of stochastic processes $\{(X_t, Y_t, Z_t)\}_{0\le t\le T}$, satisfying certain integrability assumptions. In this dissertation, both cases are considered. In chapters 2 and 3, we consider decoupled FBSDE systems, i.e. restrict our analysis to the special case of (1.6) determining the randomness in (1.2). On the contrary, we loosen this assumption in chapters 4 and 5, which are concerned with the coupled setting, i.e. the solution pair of the backward equation enters the random dynamics in (1.3).

On reflected BSDEs. Additionally, chapter 3 relies on the concept of reflected backward stochastic differential equations. Heuristically speaking, these equations generalize (1.1) by ensuring that the Y part of the solution always stays above an auxiliary lower boundary process. In mathematical formalism, this leads to the following set of equations

$$\begin{aligned} Y_t &= \xi + \int_t^T f(s, Y_s, Z_s)\mathrm{d}s - \int_t^T Z_s \mathrm{d}W_s + K_T - K_t, \\ Y_t &\ge L_t, \quad \int_0^T [Y_t - L_t]\,\mathrm{d}K_t = 0. \end{aligned} \tag{1.4}$$

In the above, the lower boundary process is denoted L_t, and the last term is the so called *Skorohod condition* [161]. The additional terms compared to (1.1) establish that the process Y stays above the lower boundary process L, and that K is the minimum force term that can achieve this condition. The solution of (1.4) is a triplet of stochastic processes $\{(Y_t, Z_t, K_t)\}_{0\le t\le T}$, such that (1.4) is satisfied almost surely, and the processes admit natural integrability and continuity conditions. For a recent account on the well-posedness of such reflected equations, we refer to [137] and the references therein.

1.2. CONNECTIONS WITH PDES

From the numerical perspective, a key motivation behind the study of BSDEs lies in their innate connections with second-order partial differential equations (PDE), generalizing the Feynman-Kac formula to non-linear settings. The Feynman-Kac formula establishes an – in some sense – equivalence relation between the solution of a second-order linear parabolic PDE and an Itô process. In fact, the solution of the following linear PDE

$$\begin{aligned} &\partial_t u(t,x) + 1/2\,\mathrm{tr}\{\sigma(t,x)\sigma^T(t,x)\mathrm{Hess}u(t,x)\} \\ &\quad + \mu^T(t,x)\nabla u(t,x) + r(t,x)u(t,x) + f(t,x) = 0, \qquad (t,x)\in[0,T]\times\mathbb{R}^d, \\ &u(T,x) = g(x), \qquad x\in\mathbb{R}^d, \end{aligned} \tag{1.5}$$

can be represented by the diffusion process

$$X_t = x_0 + \int_0^t \mu(s, X_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \tag{1.6}$$

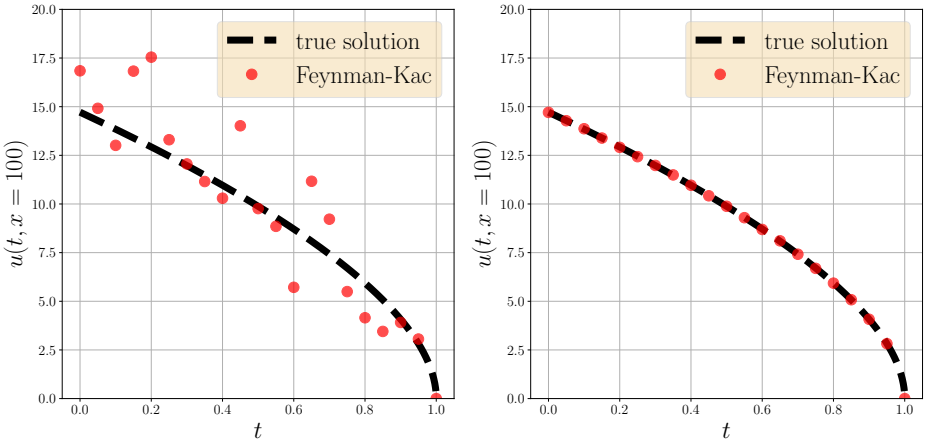

(a) small Monte Carlo sample

(b) large Monte Carlo sample

Figure 1.2: Feynman-Kac formula (1.7) illustration. The linear PDE (1.5) corresponds to a vanilla, European call option under the Black-Scholes model, admitting an analytical solution to (1.5) in closed form. Monte Carlo approximations of $u(t, x = 100)$ by (1.7). On the left, small Monte Carlo sample ($M = 2^5$); on the right, large MC sample ($M = 2^{20}$).

subject to some $\mathscr{F}_0$ measurable initial condition, due to the Feynman-Kac formula, see e.g. [121, 148],

$$u(t,x) \equiv \mathbb{E}\left[e^{\int_t^T r(s,X_s)\mathrm{d}s} g(X_T) + \int_t^T e^{\int_t^s r(\tau,X_\tau)\mathrm{d}\tau} f(s,X_s)\mathrm{d}s | X_t = x \right]. \tag{1.7}$$

Herein, the conditional expectation is taken under the probability measure under which $\{W_t\}_{0\leq t}$ is a Brownian motion. The main implication of this formula is that one, by forming an appropriate diffusion process such as (1.6), and taking expectations of its solution as in (1.7), can solve linear PDEs of the form (1.5). The representation opens up tons of different probabilistic ways for the numerical treatment of (1.5), especially in case the spatial dimensions of the domain is high (large d), when classical numerical PDE methods are intractable due to the curse of dimensionality. An illustration of a Monte Carlo type approximation enabled by the Feynman-Kac formula is given in figure 1.2.

One of the main motivations behind the study of BSDEs is given by non-linear extensions of the Feynman-Kac relations above. Given the following semi-linear extension of (1.5)

$$\begin{aligned} &\partial_t u(t,x) + 1/2\,\mathrm{tr}\{\sigma(t,x)\sigma^T(t,x)\mathrm{Hess}u(t,x)\} \\ &\quad + \mu^T(t,x)\nabla u(t,x) + f(t,x,u,\nabla_x u\sigma) = 0, \qquad (t,x) \in [0,T]\times\mathbb{R}^d, \\ &u(T,x) = g(x), \qquad x \in \mathbb{R}^d, \end{aligned} \tag{1.8}$$

one can define a BSDE, where the randomness is given by deterministic mappings of the solution of a forward SDE as in (1.6), and the corresponding Markov BSDE can be written in the form of (1.2). Then, under mild assumptions – see [161] and the references

therein –, the solution of (1.8) preserves a relation similar to (1.7), called the *nonlinear Feynman-Kac formulae*

$$Y_t = u(t, X_t), \quad Z_t = \nabla_x u(t, X_t)\sigma(t, X_t), \tag{1.9}$$

satisfied almost surely. The main implication of (1.9) is that allows for a stochastic representation of a large class of second-order partial differential equations. In particular, from a numerical analysis perspective, the nonlinear Feynman-Kac expressions open up the possibility for the numerical solution of nonlinear PDEs in a Monte Carlo fashion. Such numerical algorithms are especially relevant in high-dimensional applications, where the domain over which the PDE is formulated lies in a high-dimensional vector space.

Generalizations of (1.9) hold true even beyond the semi-linear case. In case of quasi-linear equations, allowing non-linear dependence on the solution and its gradients in the coefficient term of the second-derivative in (1.8), the associated diffusion and backward equation form a coupled FBSDE system, as in (1.3)-(1.2). We refer to theorem 4.3.1 in chapter 4 for full details. Moreover, similar stochastic representations can be given for free-boundary PDEs, established by reflected BSDEs as in (1.4) – see e.g. [46, 44, 45]. Although it falls out of the scope of this dissertation, we remark that expressions similar to (1.9) can be established even in the case of fully nonlinear second-order PDEs, see e.g. [33, 150, 138, 135, 136, 14].

In this thesis, we will consider each of the aforementioned extensions – apart from the fully nonlinear equations. Chapter 2 is concerned with the semi-linear (1.8) high-dimensional setting. In chapter 3 this is extended to the case of free boundaries. Finally, chapters 4 and 5 are written in the quasi-linear framework. In the former, we are concerned with equations over high-dimensional domains, whereas in the latter we consider higher-order convergence for equations over the real line.

1.3. APPLICATIONS IN MATHEMATICAL FINANCE

Applications, where nonlinear PDEs are solved numerically over high-dimensional domains naturally arise in mathematical finance. For instance, in the context of pricing basket options issued on many underlyings, depending on a large number of risk factors. Alternatively, another important application is optimal asset allocation, portfolio optimization, where an investor has to make optimal decisions allocating her wealth in a large number of different assets, with the objective of either maximizing some utility function, or minimizing an arbitrary cost functional. In what follows we give a soft introduction to each of these subjects, as they form the core applications of the results in this thesis.

1.3.1. OPTION PRICING AND HEDGING

One of the main challenges in mathematical finance is to price derivative contracts issued on underlyings whose prices evolve according to a random phenomenon. A financial option gives its holder the right or obligation to purchase or sell a given instrument at a pre-specified point in the future. Usually, the underlying instruments are modeled by forward stochastic differential equations of some form, the pricing of such a contract depends on a certain functional of the modeled asset prices.

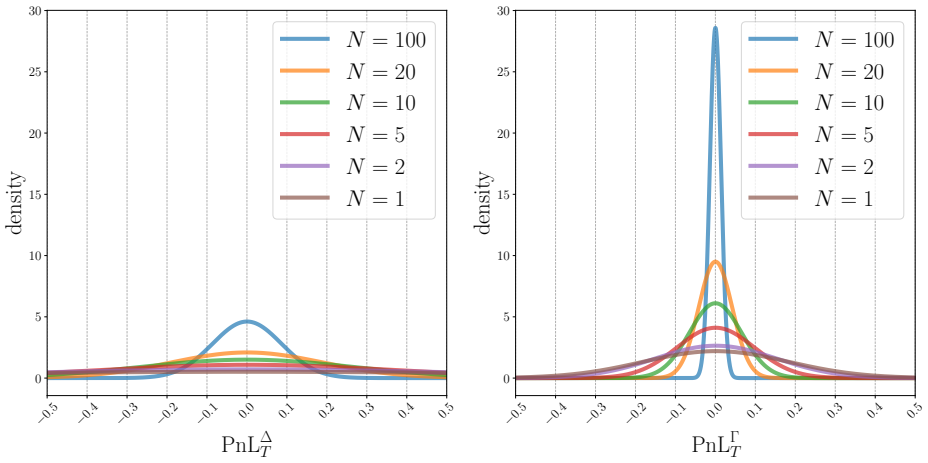

(a) Delta hedging with (1.10)

(b) Gamma hedging with (1.12)

Figure 1.3: Discrete replication illustration. Delta- vs delta-gamma hedging of a European, vanilla, Black-Scholes call option for several rebalancing frequencies N. Figure from [116], see fig. 3.1 in chapter 3.

From the hedging perspective, one's goal is to protect oneself from random movements in the underlying's stochastic dynamics. In order to achieve this, an investor can construct a delta hedging replicating portfolio consisting of the following positions

- short position in a derivative contract to be hedged;
- long positions on the underlyings assets the derivative contract was issued on;
- deposit in a bank account.

Given these positions, the value of the investor's hedging portfolio, P, evolves according to the following SDE

$$\mathrm{d}P_t = -\mathrm{d}v(t, X_t) + \sum_{i=1}^{m} \alpha_t^i \mathrm{d}S_t^i + \mathrm{d}B_t, \tag{1.10}$$

subject to the usual self-financing condition $P_0 = 0$. The variance minimizing first-order conditions lead to

$$\alpha_t^* = \underset{(\alpha_t^1,\ldots,\alpha_t^m)\in\mathbb{R}^m}{\operatorname{arginf}} \quad \mathbb{V}\mathrm{ar}\left[\mathrm{d}P_t\right] = \nabla_x v(t, X_t), \tag{1.11}$$

offsetting all first-order *sensitivities* of the replicating portfolio's value. A direct application of Itô's formula shows that in the continuously rebalanced case, the condition above provides a perfect replication of the option, i.e. $P_t \equiv 0$ for all $t \in [0, T]$ almost surely. However, this is no longer the case if the hedging weights are only updated over a discrete set of points in time, and the investor is exposed to risk through the random movements in the underlyings. These risks are even more profound for highly volatile

underlyings, whose first-order derivatives, *Deltas* change rapidly. In order to mitigate this risk in the discrete time framework, one could offset higher order sensitivities of the replicating portfolio's value, and construct a so called delta-gamma hedging portfolio, which accounts for changes in the second-derivative, *Gamma*, as well. Nonetheless, as the underlyings have constant zero Gammas, one needs to augment the replicating portfolio in (1.10) with associated *gamma hedging instruments*, whose Gammas are non-vanishing. To this end, consider K many financial securities, whose prices depend on the same underlying risk factors as above. Subsequently, the associated delta-gamma hedging portfolio's value evolves as follows

$$\mathrm{d}P_t^\Gamma = -\mathrm{d}V_t + \sum_{i=1}^{m} \alpha_t^i \mathrm{d}S_t^i + \sum_{k=1}^{K} \beta_t^k \mathrm{d}u^k(t, X_t) + \mathrm{d}B_t, \tag{1.12}$$

where u^k denotes the price of the corresponding gamma hedging instrument, and β_t^k the hedging weight with which the corresponding hedging instrument is held. It can be shown – see e.g. [64] and also chapter 3 below – that the variance minimizing first- and second-order conditions, simultaneously offsetting both the replicating portfolio's (1.12) first- and second-order derivatives with respect to the underlying risk factors result in the following linear system of equations

$$\sum_{k=1}^{K} \beta_t^k \partial_{li}^2 u^k(t, X_t) = \partial_{li}^2 v(t, X_t), \qquad li \in \mathscr{I}, \tag{1.13a}$$

$$\alpha_t^i = \partial_i v(t, X_t) - \sum_{k=1}^{K} \beta_t^k \partial_i u^k(t, X_t), \qquad 1 \le i \le m \tag{1.13b}$$

solved by the optimal hedging weights α_t^i, β_t^k. An illustration of the comparison between discretely rebalanced delta- (1.10) and delta-gamma (1.12) portfolios is given in figure 1.3.

In light of the Feynman-Kac relations in (1.7), it is natural to consider the problem outlined above in the framework of FBSDEs. Indeed, given an SDE such as (1.6) modelling the risk factors, an associated BSDE in (1.2) does not merely include the prices of option contracts issued on these underlyings, but also first-order sensitivities needed to compute the optimal Delta hedging weights on the right-hand side of (1.11). Hence, solving a BSDE is, in a sense, equivalent to delta replication in (1.10). Nonetheless, the formulae in (1.7) only include first-order *Greeks* of the associated option, and are thus insufficient in the context of delta-gamma hedging with moderated discrete time replication errors in (1.12). In order to be able to compute the right-hand sides in (1.13b), one needs an additional stochastic representation including second-order derivatives of the corresponding option prices, representing second-order Greeks, Gammas.

One of the main contributions of this thesis, is to establish that in the discrete time framework. To do so, we build on a theoretical representation formula for the Z process, originally from Pardoux and Peng [126], and later extended to the non Markovian case by El Karoui et al. in [46]. In fact, under suitable assumptions, it turns out that the solution pair in (1.1) is differentiable in the *Malliavin sense* – see [119]. Moreover, these Malliavin

derivatives $\{(D_s Y_t, D_s Z_t)\}_{0\leq s,t\leq T}$ satisfy a *linear* BSDE themselves

$$\begin{aligned} D_s Y_t = D_s\xi + \int_t^T D_s f(r, Y_r, Z_r) + \partial_y f(s, Y_s, Z_s) D_s Y_r + \partial_z f(r, Y_r, Z_r) D_s Z_r \mathrm{d}r \\ - \int_t^T (D_s Z_r^T \mathrm{d}W_r)^T, \qquad 0 \leq s \leq t \leq T, \\ D_s Y_t = 0, \quad D_s Z_t = 0, \qquad 0 \leq t < s \leq T. \end{aligned} \tag{1.14}$$

In particular, the control process Z in (1.1) admits

$$Z_t = D_t Y_t, \tag{1.15}$$

almost surely, for each $t \in [0, T]$. In light of (1.15) and the Malliavin chain rule – see lemma 2.2.1 in chapter 2 –, a simultaneous solution to the BSDEs (1.1) and (1.14), would then extend the Feynman-Kac formulae, providing not just associated option prices (Y), Deltas (Z) but also Gammas through $\{D_t Z_t\}_{0\leq t\leq T}$ for the whole time horizon. Subsequently, these solutions can be plugged into the first- and second-order conditions of delta-gamma hedging given by (1.13), and options' delta-gamma replication is simultaneously solved.

The missing ingredients in this solution is that, unlike in the sole case of (1.1), one first needs to discretize the system of equations (1.1)-(1.14), prove the convergence of the corresponding discretizations, and choose an appropriate machinery which approximates the resulting conditional expectations in a potentially high-dimensional framework. Indeed, we remark that for a large number of risk factors in (1.6), (1.14) is a vector-valued BSDE, and $\{D_s Z_s\}_{0\leq s\leq T}$ is a matrix-valued process with $d \times d$ elements, resulting in $d \times d$ many conditional expectations to be computed in a numerical setting. Finally, the resulting discrete time approximations need to be plugged into the delta-gamma hedging conditions in (1.13) in order to replicate the option price in a discrete time framework.

We carry out this program throughout chapters 2 and 3 of this thesis in the Markovian framework. Chapter 2 introduces a new discretization called the *One Step Malliavin* scheme, which simultaneously solves the pair of BSDEs (1.1)-(1.14). The convergence of this scheme is proved in theorem 2.4.1, showing optimal convergence rates under standard assumptions. Subsequently, a fully-implementable machine learning regression based approach is proposed in order to approximate the conditional expectations arising in the One Step Malliavin scheme – see also section 1.4.2. The convergence of this method is proved in theorem 2.5.2. Building upon these results, chapter 3 resolves the discretely rebalanced delta-gamma hedging problem outlined in (1.12) in the portfolio context, including (discretely) reflected BSDEs (1.4) corresponding to options with early exercise features.

1.3.2. STOCHASTIC OPTIMAL CONTROL: PORTFOLIO OPTIMIZATION

Chapters 4 and 5 are concerned with coupled FBSDEs of the form (1.3)-(1.2). A fundamental motivation for the study of such systems of equations is given by their inherent connections with stochastic control problems. In what follows, we give a brief introduction to these classes of problems, with a focus on financial mathematics.

A classical illustration of stochastic optimal control theory is the optimal portfolio allocation model of Merton [111, 112]. This model is concerned with an agent, whose aim is to distribute her wealth between a risky asset S and a risk free bond with rate r, and decide her consumption in a way that maximizes her utility over a finite time horizon T. At each point in time t, the agent faces two decisions: how much of her total wealth she spends on immediate consumption c_t, and what proportion α_t of the leftover wealth she keeps in the risky asset. These two decisions are described by the control process $\{u_t := (\alpha_t, c_t)\}_{0\le t\le T}$. When the risky asset's price is described by a geometric Brownian motion

$$\mathrm{d}S_t = \mu S_t \mathrm{d}t + \sigma S_t \mathrm{d}W_t,$$

the wealth process of the agent satisfies the following controlled SDE

$$X_t^u = x_0 + \int_0^t \left(r + \alpha_s(\mu - r)\right) X_s^u - c_s \mathrm{d}s + \int_0^t \alpha_s \sigma X_s^u \mathrm{d}W_s,$$

depending on the choices of the control parameters. The set of admissible controls is a pair of stochastic processes $\{u_t := (\alpha_t, c_t)\}_{0\le t\le T}$ which takes values in $\mathbb{R}\times\mathbb{R}_+$, describing portfolio allocation policies and consumption over the finite investment horizon T, such that $\int_0^T |\alpha_t|^2 + |c_t|^2 \mathrm{d}t < \infty$ almost surely. This condition is imposed in order to ensure the well-posedness of the controlled diffusion.

The agent's objective is to choose her control parameters in a way that maximizes her expected discounted utility at a future point in time. In particular, given two utility functions, $U_c : \mathbb{R}_+ \to \mathbb{R}$, $U_w : \mathbb{R} \to \mathbb{R}$, which describe the agent's preferences in consumption and terminal wealth, respectively, the goal is to carry out the following stochastic optimization problem

$$\max_{u=(\alpha,c)} \mathbb{E}\left[e^{-rT} U_w(X_T^u) + \int_0^T e^{-rt} U_c(c_t) \mathrm{d}t \right].$$

It turns out, when the agent's consumption and wealth preferences are described by utility functions exhibiting the *constant relative risk aversion* (CRRA) property, see e.g. [132], the problem above admits a closed-form solution

$$\alpha_t^* = \frac{\mu - r}{\sigma^2(1-\gamma)}, \quad c_t^* = b^{\frac{1}{1-\gamma}}(t) X_t^{u^*},$$

where γ corresponds to the risk aversion coefficient, and the deterministic function $b(t)$ solves an ODE [112]. The main conclusions of the Merton model is that a CRRA utility maximizing agent needs to keep a constant ratio, the *Merton fraction*, of her wealth in risky assets, independently of the current wealth, only depending on the risk premium $\mu - r$, volatility σ and her own risk aversion parameter γ. For extensions on other types of utility functions, we refer to [101] and the references therein.

The Merton model is a very special case of more general stochastic optimal control problems, where the goal is to optimize an objective functional subject to some random phenomenon that is described by a controlled stochastic differential equation. One, in

general, is interested in the solution of the following problem

$$J(0,x_0) := \inf_{u\in\mathscr{U}} \mathbb{E}\left[g(X_T^u) + \int_0^t \bar{f}(t,X_t^u,u_t)\mathrm{d}t|\mathscr{F}_t \right], \tag{1.16a}$$

$$X_t^u = x_0 + \int_0^t \bar{\mu}(s,X_s^u,u_s)\mathrm{d}s + \int_0^t \bar{\sigma}(s,X_s^u,u_s)\mathrm{d}W_s. \tag{1.16b}$$

Unlike in the case of the famous Merton fraction, such problems do not generally admit closed-form analytical expressions for the solution of the optimal control

$$u_\cdot^* := \arg\inf_{u\in\mathscr{A}} J(0,x_0).$$

From the numerical analysis perspective, there are two main classes of algorithms that tackle the approximation of the problem described in (1.16). One, is based on Bellman's dynamic programming principle, which gives rise to a *Hamilton-Jacobi-Bellman* equation, a fully nonlinear PDE, solved by the value function

$$\begin{aligned} -\partial_t v(t,x) - \inf_{u\in\mathscr{A}} \left[\mathfrak{L}^u v(t,x) - \bar{f}(t,x,u) \right] = 0, \qquad & (t,x)\in[0,T)\times\mathbb{R}^d, \\ v(T,x) = g(x), \qquad & x\in\mathbb{R}^d, \end{aligned} \tag{1.17}$$

where the infinitesimal generator associated with the diffusion (1.16b) is given by

$$\mathfrak{L}^u v := \nabla_x \bar{\mu}^T(t,x,u)v + \frac{1}{2}\mathrm{tr}\{\bar{\sigma}(t,x,u)\bar{\sigma}^T(t,x,u)\mathrm{Hess}_x v\}.$$

The terminal boundary condition associated with the HJB equation coincides with that of the terminal part of the objective functional in (1.16a). Subsequently, a numerical PDE solver can be applied to directly solve (1.17), and therefore resolve the associated stochastic optimal control problem in (1.16) by means of the value function – see e.g. [21, 82] and the references therein.

The main challenges, again, are hidden in the dimensionality of the associated, controlled forward diffusion X^u in (1.16b). In applications where X^u takes values in $\mathbb{R}^d$ – like a Merton type portfolio allocation with many risky assets – such classical PDE solvers become inapplicable due to their exponentially increasing computational complexity. A phenomenon often referred to as the curse of dimensionality. Alternatively, motivated by the nonlinear Feynman-Kac relations (1.7) establishing a connection between PDEs like (1.17), one can derive a coupled FBSDE system associated with the stochastic optimal control problem (1.16). There are two distinct ways to derive this associated coupled FBSDE system. First, when the control in the diffusion process in (1.16b) is such that $\bar{\sigma}$ *does not* take u as an argument, the corresponding control problem is called *drift control*. In that case, one can show – see e.g. [132, 156] – that the HJB equation in (1.17) turns into a semi-linear equation, similar to (1.8). Consequently, the value function and its gradient admit a stochastic representation provided by a BSDE as in (1.2), with the special property that the corresponding drift $\bar{\mu}$, diffusion $\bar{\sigma}$ and driver $\bar{f}$ only depend on (t,x,z) and not on y. On the contrary, in case $\bar{\sigma}$ also takes u as an argument, one is faced with a *diffusion control* problem, where the associated HJB equation is fully nonlinear.

In that case, one can rely on the stochastic maximum principle (SMP) – see e.g. [156] –, which gives rise to a vector-valued, fully-coupled coupled FBSDE system as in (1.3)-(1.2).

In both cases, given a robust and accurate approximation method that numerically resolves an associated BSDE, the stochastic optimal control problem can be solved by an equivalent coupled FBSDE system. In this thesis, we contribute in two ways to the literature of the numerical methods built on these types of stochastic representations. First, in chapter 4 we are concerned with the numerical solution of coupled FBSDE systems with a high-dimensional controlled diffusion process. Using the deep BSDE method originally proposed by [43, 68] – see also section 3.4 below –, we extend the convergence result of Han and Long [69] to the case of general drift coefficients μ in (1.3), i.e. drift terms which can depend on the Z process. We derive an *a posteriori* error bound, which ensures the convergence of the method up to the convergence of the terminal *loss function*. Moreover, in chapter 5, we investigate an alternative to PDE methods directly applied on the HJB equation (1.17) in the low-dimensional setting, i.e. for scalar valued controlled diffusions in (1.16b). The COS method – see also section 1.4.3 below – applied in chapter 5 provides a higher-order temporal discretization of the associated FBSDE system, in the fully-coupled framework. We show that the resulting approximations admit higher-order convergence rates than related reference methods, in the step size of a discrete time partition, both in the sense of *strong* and *weak* convergence rates.

1.4. MODERN NUMERICAL METHODS

Solving equations of the type (1.1) usually requires numerical methods approximating the solution pair (Y, Z) given an appropriate sampling mechanism dealing with the random data (ξ, f). Such numerical methods consist of an appropriate time discretization which is a sequence of backward, recursively nested conditional expectations going from terminal time T to $t = 0$. Thereafter, these conditional expectations are approximated numerically, by a suitable approach. In what follows, we discuss further details of such numerical methods.

1.4.1. TIME DISCRETIZATION

A naive idea to tackle (1.1) in a numerical framework would be to simulate the random data (ξ, f), and perform backward Euler stepping from $t = T$ to $t = 0$, similar to the numerical approximation of ordinary differential equations, with terminal conditions. The reason why this is not suitable in the context of BSDEs is due to the violation of adaptivity conditions imposed on the solution pair (Y_t, Z_t). In order to provide adaptive, discrete time resolution of the stochastic processes in (1.1), one needs compute conditional expectations instead. In fact, taking conditional expectations of (1.1) and the same equa-

tion multiplied by a Brownian increment $(W_t - W_s)$ with $t \geq s$ gives the following

$$Y_s = \mathbb{E}\left[Y_t + \int_s^t f(r, Y_r, Z_r)\mathrm{d}r | \mathscr{F}_s \right],$$
$$0 \equiv \mathbb{E}\left[(W_t - W_s) Y_s | \mathscr{F}_s\right] = \mathbb{E}\left[(W_t - W_s) Y_s | \mathscr{F}_s\right] + \mathbb{E}\left[(W_t - W_s) \int_s^t f(r, Y_r, Z_r)\mathrm{d}r | \mathscr{F}_s \right]$$
$$- \mathbb{E}\left[\int_s^t Z_r \mathrm{d}r | \mathscr{F}_s \right],$$

due to Itô's isometry. Discretizing the continuous integral in the above, and iterating over a discrete time partition $\pi(N) \coloneqq \{t_n : t_0 = 0, t_N = T, t_{n-1} < t_n < t_{n+1} \text{ for } n = 1, \dots, N-1\}$ yields the following abstract time discretization, approximating the solution pair at each point in the discrete time grid

$$\begin{aligned} Y^\pi_{t_N} &= \xi, \\ Z^\pi_{t_n} &= \frac{1}{t_{n+1} - t_n} \mathbb{E}\left[(W_{t_{n+1}} - W_{t_n}) Y^\pi_{t_{n+1}} | \mathscr{F}_s \right], \\ Y^\pi_{t_n} &= (t_{n+1} - t_n) f(t_n, Y^\pi_{t_n}, Z^\pi_{t_n}) + \mathbb{E}\left[Y^\pi_{t_{n+1}} | \mathscr{F}_s \right], \end{aligned} \tag{1.18}$$

for $n = N-1, \dots, 0$. The conditional expectations ensure that the resulting sequence of discrete time approximations $\{(Y^\pi_{t_n}, Z^\pi_{t_n})\}_{n=0,\dots,N}$ are adapted with respect the to filtration. On the other hand, they induce a major computational challenge in terms of the numerical approximations needed to estimate the right-hand sides of (1.18).

Over the past several decades, many numerical methods have been considered to approximate conditional expectations of the form above, in the context of FBSDEs. Nonetheless, whenever the randomness in the data (ξ, f) is determined by some high-dimensional stochastic process, all these methods suffer from the *curse of dimensionality*, since their corresponding computational complexities scale exponentially in the number of spatial dimensions. In what follows, we discuss a recently emerging class of numerical methods that are aimed to tackle this computational drawback for high-dimensional problems.

1.4.2. DEEP BSDE: MODERN REGRESSION MONTE CARLO

In this thesis, we primarily consider two different classes of algorithms for the numerical approximation of the conditional expectations appearing in (1.18). In high-dimensional frameworks, we use so called *deep BSDE* methods, which, as explained below, can be thought of as generalized Monte Carlo regression approaches with neural network parametrizations. In what follows, we give a brief introduction to the main variants of this method used throghout the thesis. Additionally, for low-dimensional forward diffusions, we also consider applications of the COS method, which is a Fourier cosine expansion method approximating conditional expectations by appropriate Fourier transforms, given a characteristic function available in closed for. The introduction of the latter we postpone to section 1.4.3, and in what follows we focus on the deep BSDE type algorithms.

A straightforward, brute force way to compute conditional expectations such as the ones in (1.18), would be to perform *inner Monte Carlo* sampling at each state of the state

process. However, such an approach is computationally very expensive, even more so in the high-dimensional framework. Alternatively, a classical methodology to tackle the approximation of the associated conditional expectation is *least-squares Monte Carlo* (LSMC) regression. This method was first proposed by Longstaff and Schwartz in [98] in the context of American option pricing, and has since been extensively studied in the framework of discrete time approximations of BSDEs – without the sake of completeness, we mention [24, 63, 62]. The LSMC method is built on the fundamental property of Markovianity. Namely, given a BSDE (1.2) that has the Markov property, both the Y and Z processes turn out to be deterministic mappings of the underlying state process X, and $Y_t = u(t, X_t)$ and $Z_t = v(t, X_t)$ for some deterministic mappings of time and space $u : [0,T] \times \mathbb{R}^d \to \mathbb{R}^q$, $v : [0,T] \times \mathbb{R}^d \to \mathbb{R}^{q\times d}$ – see e.g. [46] and (1.7). In particular, the deterministic mapping $u(t,\cdot)$ solves the following minimization problems

$$u(t,\cdot) \in \underset{h:\mathbb{R}^d\to\mathbb{R}^q,\,\text{measurable}}{\operatorname{arginf}} \mathbb{E}\left[\left|g(X_T) + \int_t^T f(s, X_s, Y_s, Z_s)\mathrm{d}s - h(X_t)\right|^2\right], \tag{1.19}$$

and similarly in case of v. Consequently, the solutions of the corresponding conditional expectations are reduced to an L^2 projection for each $\mathscr{F}_t$ at every time t. Equation (1.19) is, of course, not directly implementable. Indeed, a standard LSMC programme would typically consist of the following *approximations* that enable to compute the associated conditional expectations:

- time discretization replacing the continuous integrals in (1.19) by an appropriate discrete sum, as e.g. in (1.18);
- parametrization truncating the infinite dimensional space of all measurable mappings into a finite subspace of basis functions $\{\phi_k(t,\cdot)\}_{k=0,\dots,K-1}$;
- Monte Carlo approximations replace the true expectation $\mathbb{E}$ by the empirical mean of a discrete sample of M independent realizations of the data (ξ, f);
- backward recursion from T to $t = 0$.

Combining these approximations, one can derive a least-squares regression problem associated with the minimization in (1.19), which yields a closed-form expression for the approximated solution

$$u(t,\cdot) \approx \sum_{k=0}^{K-1} \widehat{\alpha}_k \phi_k(t,\cdot), \tag{1.20}$$

where the regression weights solve the standard OLS formula $\widehat{\alpha} = (\widehat{\alpha}_0, \dots, \widehat{\alpha}_{K-1}) \coloneqq (\Phi^T\Phi)^{-1}\Phi^Y\Psi$, with Φ and Ψ corresponding to the regression matrix and labels, respectively.

In terms of the numerical approximation of high-dimensional BSDEs, the main challenge in the approach outlined above is to find an appropriate parametrization of finite basis functions with sufficient representation power, such that high-dimensional nonlinear phenomena can accurately be modeled in a finite Monte Carlo sample. In recent

years, a new class of methods has been proposed by the scientific computing and numerical analysis communities which tackle this challenge by means of neural networks. Doing so, one generalizes standard OLS regressions to nonlinear parametrizations with higher expressive power. This comes with the price to pay that one loses the analytical expression giving the true minimizers $\widehat{\alpha}$ of the regression problem. Instead, one typically needs to form an L^2 loss function associated with the projection problem in (1.19). This loss function is subsequently minimized by an iterative method, such as stochastic gradient descent. The methods which follow these ideas are often referred to as *deep BSDE methods*. In what follows, after a general introduction to deep learning, we explain the main ideas behind such *modern regression Monte Carlo methods*.

NEURAL NETWORKS

Neural networks can be thought of as nonlinear generalizations to the parametrizations of linear combinations of some basis functions. In order to achieve nonlinearity, one takes an affine combination as in (1.20), and passes it through some nonlinear *activation function* a. These two operations together form a *hidden layer* of a neural network. Taking a sequence of such compositions, *deep neural networks* can be formulated as a hierarchical sequence of hidden layers

$$\Psi(x|\Theta) \coloneqq a_{\text{out}} \circ A_{L+1}(\cdot|\theta_{L+1}) \circ a \circ A_L(\cdot|\theta_L) \circ a \circ \cdots \circ a \circ A_1(\cdot|\theta_1) \circ x.$$

Out of the standard activation functions, in this thesis we use rectified linear units (relu) and the hyperbolic tangent. The parameters in the affine combinations A^l, $l = 1, \dots, L$ are called *weights* and *biases*, and finding optimal values such that the corresponding parametrization best describes some phenomenon – such as the L^2 projection in (1.19) –, constitutes the *training* of the neural network. Finding the optimal parameter space is often done in an iterative optimization method that aims to minimize a *loss function* via stochastic gradient descent (SGD). For more details, we refer to [65].

The claim that deep neural networks are suitable function approximators in high-dimensional frameworks is often justified by their *universal approximation property*. Informally, this characteristic, established by universal approximation theorems (UAT), ensures that large classes of functions can be approximated by (deep) neural networks with arbitrary accuracy. We refer to theorem 2.5.1 in chapter 2 for a precise statement of the UAT due to Hornik et al. [71]. We remark that finding minimal conditions for UAT type characteristics is a subject of active research, for an overview see [134].

FORWARD DEEP BSDE METHOD.

The first deep BSDE method was proposed in the groundbreaking papers of [43, 68]. Their idea can roughly be summarized as follows. In case of a decoupled, Markovian FBSDE such as the system (1.6)-(1.2), one can derive the following auxiliary BSDE whose solution pair coincides with that of (1.2) by merging the Feynman-Kac relations in (1.7) with (1.2)

$$\begin{aligned} u(t, X_t) = g(X_T) &+ \int_t^T f(s, X_s, u(s, X_s), \nabla_x u(s, X_s)\sigma(s, X_s))\mathrm{d}s \\ &- \int_t^T \nabla_x u(s, X_s)\sigma(s, X_s)\mathrm{d}W_s, \quad 0 \le t \le T. \end{aligned}$$

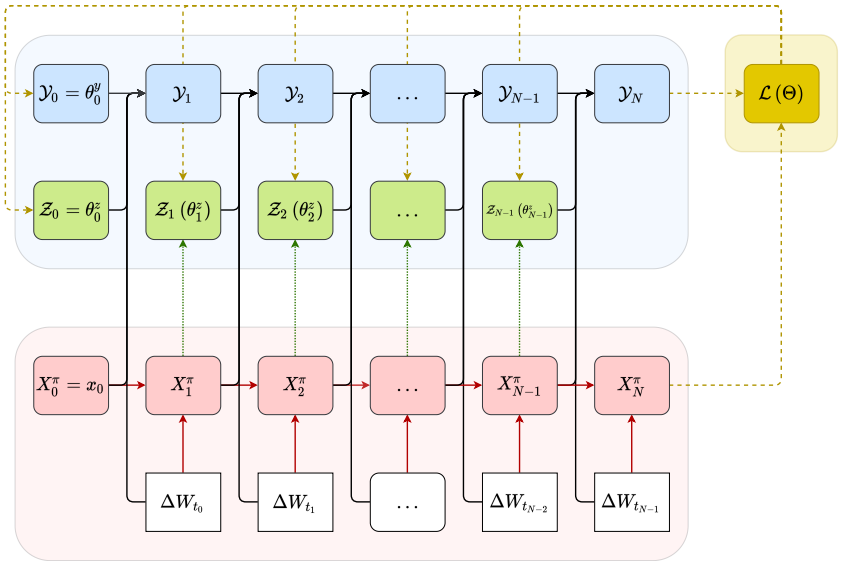

Figure 1.4: Forward deep BSDE method illustration. SDE (1.3), BSDE (1.2) and loss function (1.25) minimization in red, blue and yellow, respectively. Figure taken from [118, fig. 4.2].

In the above, $\{X_t\}_{0\le t\le T}$ is the solution of (1.6). In particular, the equation above can be rewritten into the following forward SDE controlled by some deterministic mapping u of the state process

$$Y_t^u = u(0,X_0) + \int_0^t f(s,X_s,Y_s^u,\nabla_x u(s,X_s)\sigma(s,X_s))\mathrm{d}s - \int_0^t \nabla_x u(s,X_s)\sigma(s,X_s)\mathrm{d}W_s, \tag{1.21}$$

for each $0 \le t \le T$. When the control mapping u is given by the true solution of the associated nonlinear PDE as in (1.8), the solution of (1.21) coincides with that of (1.2). Even though, in general, one does not know the solution u, the observation above motivates to cast this problem into the framework of stochastic optimal control (1.16). To this end, treating the random variable $u(0,X_0)$ and stochastic process $\nabla_x u(t,X_t)\sigma(t,X_t)$ as control parameters, one can formally define the following stochastic optimal control problem as in (1.16)

$$\inf_{\alpha=(\varphi,\zeta)} \mathbb{E}\left[|g(X_T) - Y_T^\alpha|^2\right], \tag{1.22a}$$

$$Y_t^\alpha = \varphi + \int_0^t f(s,X_s,Y_s^\alpha,\zeta_s)\mathrm{d}s - \int_0^t \zeta_s \mathrm{d}W_s. \tag{1.22b}$$

By well-posedness of (1.2), and thus (1.21), $(Y_0,\{Z_t\}_{0\le t\le T})$ is a minimizer of the objective functional in (1.22a). Consequently, the solution of the *stochastic target problem* in (1.22) – see e.g. [153] –, also solves the BSDE (1.2), and by uniqueness of the latter, this solution is also unique.

The continuous time equation (1.22b) can be approximated in a discrete time framework in two steps. First, if the continuous dynamics of the forward SDE described by (1.6) is not analytically resolvable, one replaces X in (1.22) with appropriate discrete time approximations X^π such as an Euler-Maruyama scheme [90]. For alternative SDE discretizations we refer to chapter 5. Subsequently, the continuous time integrals in (1.21) are discretized by a forward Euler-Maruyama scheme performed over a discrete time partition $\pi(N) \coloneqq \{t_n : t_0 = 0, t_{n-1} < t_n < t_{n+1}, n = 1,\dots,N-1, t_N = T\}$. This yields the following discrete time approximations for (1.22b)

$$\begin{aligned} Y_0^{\alpha,\pi} &= \varphi, \\ Y_{t_{n+1}}^{\alpha,\pi} &= Y_{t_n}^{\alpha,\pi} + f(t_n, X_{t_n}^\pi, Y_{t_n}^{\alpha,\pi}, \zeta_{t_n})(t_{n+1} - t_n) + \zeta_{t_n}(W_{t_{n+1}} - W_{t_n}), \quad n = 0,\dots,N-1. \end{aligned} \tag{1.23}$$

Finally, by virtue of the Feynman-Kac formulae in (1.7), one can replace the stochastic control variables in (1.23) by the *deterministic* mappings of the state process in (1.6) using the relations $\varphi = u(0, X_0)$ and $\zeta_{t_n} = \nabla_x u(t_n, X_{t_n})\sigma(t_n, X_{t_n})$. To make the model fully-implementable, these *unknown* deterministic mappings are parameterized by neural networks $\varphi(\cdot|\theta_0^y)$ and $\zeta(\cdot|\theta_n^z)$, depending on some parameter sets θ_0^y and θ_n^z, respectively, for each point in the discrete time partition $t_n, n = 0,\dots,N-1$. Given the collection of all neural networks involved, one subsequently has a parametric model depending on the collection of parameters $(\theta_0^y, \theta_0^z, \dots, \theta_{N-1}^z)$ and the corresponding discrete time approximations of the controlled diffusion process in (1.22b) by

$$\begin{aligned} \mathscr{Y}_0^{u,\pi}(\theta_0^y) &= \varphi(X_0^\pi|\theta_0^y), \\ \mathscr{Y}_{t_{n+1}}^{u,\pi}(\Theta_n) &= \mathscr{Y}_{t_n}^{u,\pi}(\Theta_{n-1}) - f(t_n, X_{t_n}^\pi, \mathscr{Y}_{t_n}^{u,\pi}(\Theta_{n-1}), \zeta(X_{t_n}^\pi|\theta_n^z))(t_{n+1} - t_n) \\ &\qquad + \zeta(X_{t_n}^\pi|\theta_n^z)(W_{t_{n+1}} - W_{t_n}), \quad n = 0,\dots,N-1, \end{aligned} \tag{1.24}$$

with $\Theta_n \coloneqq (\theta_0^y, \theta_0^z, \dots, \theta_{n-1}^z, \theta_n^z)$ for $n = 0,\dots N-1$ and $\Theta_{-1} = \theta_0^y$. The complete parameter set of the model is therefore $\Theta \coloneqq \Theta_{N-1}$. In order to find an optimal parameter set, one can impose the natural discrete time version of the L^2 loss functional in (1.22a), which leads to the following *loss function*

$$\mathscr{L}(\Theta) \coloneqq \mathbb{E}\left[|\mathscr{Y}_{t_N}^{u,\pi}(\Theta) - g(X_{t_N}^\pi)|^2\right]. \tag{1.25}$$

A suitable discrete time parametric model, approximating the continuous time solutions of (1.22), minimizes this loss function, leading to optimal parameters $\bar{\Theta} \in \arg\inf_\Theta \mathscr{L}(\Theta)$. Up to the convergence of the discrete time approximations X^π and $Y^{\alpha,\pi}$ in (1.23), the optimal discrete time approximations $\mathscr{Y}_{t_n}^{u,\pi}(\bar{\Theta}_{n-1})$ and $\zeta(X_{t_n}^\pi|\bar{\theta}_n^z)$ then approximate the solution pair (Y_{t_n}, Z_{t_n}) of the BSDE (1.2) for each point in the discrete time partition.

The key idea behind the forward deep BSDE is to replace the direct approximation of the conditional expectations in (1.18) by a discretized version of the stochastic target problem in (1.22). Thereafter, by passing batches of realizations of the Brownian motion, one can train the corresponding machine learning model by a suitable iterative optimization method – such as stochastic gradient descent – in order to minimize the loss function (1.25). These numerical approximations can be efficiently carried out in high-dimensional settings, i.e. when X is an $\mathbb{R}^d$ valued stochastic process for a large d. The method has been extended in many different directions, e.g. 2BSDEs relating to

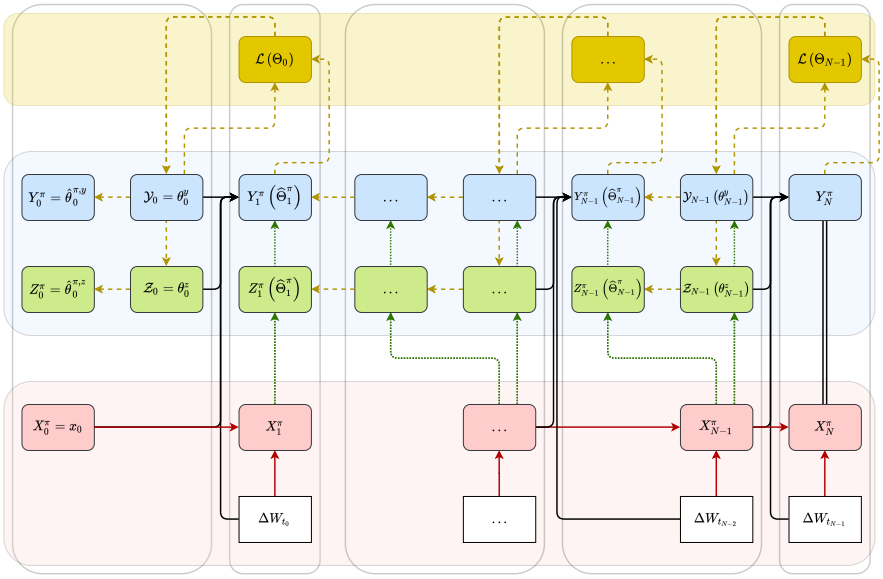

Figure 1.5: Backward deep BSDE method illustration. SDE (1.3), BSDE (1.2) and loss function (1.25) minimization in red, blue and yellow, respectively. Figure taken from [118, fig. 4.3].

fully nonlinear PDEs [11, 10], coupled FBSDE systems [69], FBSDEs with jumps [3, 5], mean-field equations [143, 67], operator valued equations [120], and using several different approaches for the parametrizations of the stochastic target problem – see [141, 84, 86, 157, 158, 4, 151, 94, 51, 155, 29]. For a survey, we refer to [34], and in the context of financial mathematical applications we highlight [91, 93, 61, 26].

BACKWARD DEEP BSDE METHODS

A different, backward approach first appeared in the literature in the paper of Huré et al. in [77]. Here, instead of one simultaneously trained forward optimization, the discrete time approximations of Y and Z are decomposed into N many separate optimization problems, locally in time. Each sub-optimization problem minimizes an associated L^2 loss function, that is related to the L^2 projection in (1.19). In particular, by the merged formulation of the Feynman-Kac relations and the BSDE, (1.21) leads to the following equation between two adjacent time steps $t_n < t_{n+1}$ in a discrete time partition $\pi(N)$

$$u(t_n, X_{t_n}) = u(t_{n+1}, X_{t_{n+1}}) + \int_{t_n}^{t_{n+1}} f(s, X_s, u(s, X_s), \nabla_x u(s, X_s))\mathrm{d}s - \int_{t_n}^{t_{n+1}} \nabla_x u(s, X_s)\sigma(s, X_s)\mathrm{d}W_s.$$

A suitable discretization – see (1.23) –, followed by parametrizations $\varphi(\cdot|\theta_n^y) : \mathbb{R}^d \to \mathbb{R}$ and $\zeta(\cdot|\theta_n^z) : \mathbb{R}^d \to \mathbb{R}^d$ of the unknown deterministic mappings $u(s,\cdot)$ and $\nabla_x u(t_n,\cdot)\sigma(t_n,\cdot)$, respectively, similar to (1.24), subsequently leads to the following parametric approxi-

mations

$$u(t_{n+1}, X^\pi_{t_{n+1}}) \approx \varphi(X^\pi_{t_n}|\theta^y_n) - f(t_n, X^\pi_{t_n}, \varphi(X^\pi_{t_n}|\theta^y_n), \zeta(X^\pi_{t_n}|\theta^z_n))(t_{n+1} - t_n) \\ + \zeta(X^\pi_{t_n}|\theta^z_n)(W_{t_{n+1}} - W_{t_n}) \\ =: F(t_n, X^\pi_{t_n}, \varphi(X^\pi_{t_n}|\theta^y_n), \zeta(X^\pi_{t_n}|\theta^z_n), t_{n+1}, W_{t_{n+1}} - W_{t_n}), \tag{1.26}$$

for each $n = 0, \dots, N-1$.

As the deterministic mapping corresponding to the terminal condition of the BSDE (1.2) is known by $u(t_N, X^\pi_{t_N}) \equiv g(X^\pi_{t_N})$, as with the backward iteration in (1.18), we set the approximations at $t = T$ to $\mathscr{Y}^\pi_{t_N} := g(X^\pi_{t_n})$. Then, one can gather sufficient approximations of the parameters $\theta^y_{N-1}, \theta^z_{N-1}$, such that (1.26) holds in an appropriate L^2 sense. In order to do so, [77] proposed the following loss function, measuring the mean-squared distance between the left- and right-hand sides of (1.26) for $n = N-1$

$$\mathscr{L}_n(\theta^y_n, \theta^z_n) = \mathbb{E}\Big[\big|\mathscr{Y}^\pi_{t_{n+1}} - \varphi(X^\pi_{t_n}|\theta^y_n) \\ + f(t_n, X^\pi_{t_n}, \varphi(X^\pi_{t_n}|\theta^y_n), \zeta(X^\pi_{t_n}|\theta^z_n))(t_{n+1} - t_n) \\ - \zeta(X^\pi_{t_n}|\theta^z_n)(W_{t_N} - W_{t_n})\big|^2\Big]. \tag{1.27}$$

A suitable minimizer $(\bar\theta^y_{N-1}, \bar\theta^z_{N-1}) \in \arg\inf_{\theta^y, \theta^z} \mathscr{L}_{N-1}(\theta^y, \theta^z)$ then approximates the relation (1.26), and thus setting $(\mathscr{Y}^\pi_{t_{N-1}}, \mathscr{Z}^\pi_{t_{N-1}}) := (\varphi(X^\pi_{t_{N-1}}|\bar\theta^y_{N-1}), \zeta(X^\pi_{t_{N-1}}|\bar\theta^z_{N-1}))$ approximates the solution pair $(Y_{t_{N-1}}, Z_{t_{N-1}}) = (u(t_{N-1}, X_{t_{N-1}}), \nabla_x u(t_{N-1}, X_{t_{N-1}}))$ of the BSDE (1.2). The scheme is subsequently made fully implementable in a backward recursion $n = N-1, \dots, 0$, where a separate SGD optimization is carried out at each point n, depending on the previously collected discrete time approximations $\mathscr{Y}^\pi_{t_{n+1}}$. In fact, as shown by [77, thm. 4.1, steps 3-4] and also in steps 1-2 in theorem 2.5.2 of chapter 2, given a *rich enough* space of parametric functions the minimizer of the loss function (1.27) provides simultaneous approximations of the conditional expectations in (1.18) for each $n = 0, \dots, N-1$. An illustration of the backward dynamic programming deep BSDE method outlined above is given in figure 1.5. Without the sake of completeness, we mention related extensions [58, 133, 28].

COMPARISON OF DEEP BSDE METHODS

The main differences between the forward and backward type methods can be summarized as follows. In case of the approach of [43, 68], the loss function is imposed on all model parameters simultaneously, meaning that the corresponding stochastic gradient descent optimization approximates the discrete time approximations of all time steps globally. On the other hand, the backward formulation decomposes the approximation over the time interval $[0, T]$ into N separate optimization problems, each having its own loss function and corresponding to approximations only at the given time step. This allows for simpler gradient descent steps, smoother loss functions, and more manageable memory constraints.

In terms of theoretical convergence guarantees, Han and Long in [69] showed that the forward deep BSDE method admits an *a posteriori* error bound, depending on the value of the discretized loss functional in (1.25). Their result holds in the even more general case of coupled FBSDEs corresponding to (1.3), where μ and σ are allowed to take Y

as an argument. This result was extended in [87] to the case of non Lipschitz coefficients and in [143] in the framework of mean-field equations. Huré et al. [77] show a somewhat stronger *a priori* result for the convergence of the backward deep BSDE method. In fact, their analysis ensures that the sequence of optimizations corresponding to (1.27) yields discrete time approximations preserving the optimal L^2 convergence rate of order $1/2$ of the backward Euler scheme in (1.18), with an additional *regression bias* term associated with the representation error of the chosen function bases, i.e. expressivity of the neural networks included in the model. Up to a universal approximation type argument, this gives a convergence guarantee of the backward deep BSDE method, at least in the case of converging SGD approximations.

In this thesis, we provide contributions to the literature of both methods. The deep BSDE method approximating the One Step Malliavin scheme in chapter 2 is inspired by the deep backward dynamic programming type method of Huré et al. [77]. On the other hand, chapter 4 is concerned with the convergence of the forward deep BSDE method for coupled FBSDEs, as in [68]. Therein, we extend the a posteriori convergence result of Han and Long [69] to more generally coupled FBSDEs.

1.4.3. COS METHOD FOR BSDES

The second class of numerical methods for the approximation of conditional expectations, that we analyze in this thesis is built of Fourier cosine expansions. The COS method [48] is a Fourier method designed to approximate Markovian conditional expectations of the form in (1.18). It is built on the key insight that even in case the density function corresponding to the solution of (1.6) is not available, its Fourier transform, the characteristic function, often is – at least in case of affine diffusion models. Subsequently, a Markovian conditional expectation of the form

$$y(t,x) = \mathbb{E}[v(s,X_s)|X_t = x] = \int_{\mathbb{R}} v(s,\xi)\, p_{s|t}(\xi|x)\mathrm{d}\xi, \quad s \geq t,$$

can be approximated by the following sequence of approximations. First, truncating the infinite domain to a finite interval $[a,b] \subset \mathbb{R}$ leads to

$$y_1(t,x) = \int_a^b v(t,\xi)\, p(\xi|x)\mathrm{d}\xi.$$

Replacing the conditional density function with Fourier cosine expansion $p(\xi|x) = {\sum_{k=0}}'^{\infty} \mathscr{P}_k(s|t)\cos(k\pi(\xi-a)/(b-a))$, and truncating the infinite series to a finite, K number of terms results in the subsequent approximation

$$y_2(t,x) = {\sum_{k=0}}'^{K} \mathscr{P}_k(s|t) \int_a^b v(s,\xi)\cos\left(k\pi\frac{\xi-a}{b-a}\right)\mathrm{d}\xi \equiv \frac{b-a}{2}{\sum_{k=0}}'^{K} \mathscr{P}_k(s|t)\mathscr{V}_k(s),$$

where $\mathscr{V}_k(s)$, $k = 0,\ldots,$ are the Fourier cosine expansion coefficients of the deterministic mapping $v(s,\cdot)$. Thereafter, one exploits the relation between the density expansion

coefficients and the Fourier transform of its density to write

$$\begin{aligned}\mathscr{P}_k(s|t) &:= \frac{2}{b-a}\int_a^b p(\xi|x)\cos\left(k\pi\frac{\xi-a}{b-a}\right)\mathrm{d}\xi \\ &= \frac{2}{b-a}\,\mathrm{Re}\left\{e^{-ik\pi a/(b-a)}\int_a^b e^{ik\pi x/(b-a)}p(\xi|x)\mathrm{d}\xi\right\}.\end{aligned}$$

In the final approximation of the COS method, the last term above is replaced by the characteristic function $\phi(u|x)$ of $X_s|X_t = x$, leading to the COS formula

$$y_3(t,x) = \sum_{k=0}^{K-1}{}' \mathscr{V}_k(s)\,\mathrm{Re}\left\{\phi\left(\frac{k\pi}{b-a}|x\right)e^{-k\pi a/(b-a)}\right\}. \tag{1.28}$$

The COS formula in (1.28) has been applied to many problems related to financial mathematics, such as option pricing [49, 48, 50, 159], stochastic control [147] or valuation of electricity storage contracts [22]. Most relevantly to the subject of this thesis, the method has been applied to the numerical solution of scalar valued FBSDE system, first in [144] for the case of decoupled FBSDEs with constant drift and diffusion coefficients in (1.6), later extended to more general SDE coefficients in [145], higher dimensional diffusion processes in [139] and finally coupled equations in [76].

This thesis has two contributions to the COS literature. In section 2.5.1 of chapter 2, we propose a COS method for the solution of the One Step Malliavin scheme in the one-dimensional framework. Finally, chapter 5 proposes a higher-order numerical method for fully-coupled FBSDEs (1.3)-(1.2) built on COS approximations.

1.5. CONTRIBUTIONS AND OVERVIEW

The dissertation is structured in a way that each chapter reads as a standalone piece. Therefore, we may repeat some common notation in each chapter separately. The dissertation is structured as follows.

In chapter 2, we present the *One Step Malliavin scheme (OSM)*, a novel discretization that simultaneously solves Markovian BSDEs (1.1) and the BSDE of the solution pair's Malliavin derivatives (1.14). The scheme includes Γ estimates, making it suitable to the application of delta-gamma hedging. The OSM scheme is made fully-implementable by two different approaches: BCOS method in the one-dimensional framework, and a backward deep BSDE method [77] as in (1.27) for high-dimensional equations. In theorem 2.4.1, we show the convergence of the discrete time approximations with optimal order 1/2 rate, under standard assumptions. Theorem 2.5.2 establishes the convergence of the associated deep BSDE approximations, up to a universal approximation type argument. Numerical experiments showcase the accuracy of the method for several semilinear equations up to 50 dimensions.

Chapter 3 applies the OSM scheme in the context of delta-gamma hedging in (1.13). We extend the OSM scheme to discretely reflected BSDEs, and use that for the pricing and delta-gamma hedging of portfolios of high-dimensional Bermudan options. The resulting option Deltas and Gammas are used to discretely rebalance the corresponding replicating strategies. Numerical experiments are presented on options with varying

early exercise rights, moneyness and volatility. These examples demonstrate the robustness and accuracy of the method up to 100 risk factors.

In chapter 4, we are concerned with high-dimensional coupled FBSDE systems approximated by the forward deep BSDE method in (1.22). We extend that a posteriori convergence guarantee of Han and Long [69] to drift coefficients in (1.3), which may take Z as an argument. Theorem 4.3.3 establishes sufficient conditions for an a posteriori convergence guarantee. As a result, in contrast to earlier theories, our convergence analysis facilitates the treatment of FBSDEs arising from stochastic optimal control problems as in (1.16). Our theoretical insights are further validated through multiple numerical experiments in high-dimensional settings.

Chapter 5 presents a higher-order numerical method for scalar valued, coupled forward-backward stochastic differential equations. This method is built on COS approximations (1.28). In contrast to classical reference methods, the forward diffusion is discretized by general second-order Taylor schemes, including *Milstein* and *simplified 2.0 weak Taylor* approximations, implying higher order strong and weak convergence rates, respectively. Numerical experiments confirm the stated higher-order convergence for both strong and weak convergence rates across various equations, ranging from decoupled to fully coupled settings.

Finally, we summarize the findings of this thesis in chapter 6 and point out several open research directions worth to discover.

2

THE ONE STEP MALLIAVIN SCHEME: NEW DISCRETIZATION OF BSDES IMPLEMENTED WITH DEEP LEARNING REGRESSIONS

És egy pár üveg sörtől majd mindenki elfelejti,
hogy az lett amit így pont nem akart itt senki.
És nézi az ördög és hátrébb nézi az Isten,
és végül a szemük előtt táncol majd minden.[1]

Kispál és a Borz (Disznók tánca)

This chapter is based on the article [114]: B. Negyesi, K. Andersson, and C. W. Oosterlee. "The One Step Malliavin scheme: new discretization of BSDEs implemented with deep learning regressions". In: *IMA Journal of Numerical Analysis* 44.6 (2024), pp. 3595–3647. DOI: 10.1093/imanum/drad092.

[1] And a few bottles of beer will make everyone forget
that it has ended being something no one wanted.
And there watches the devil and God watches from behind,
and in the end everything will dance in front of their eyes.

Kispál és a Borz (Dance of the pigs)

2.1. INTRODUCTION

In this chapter, we are concerned with the numerical solution of a system of forward-backward stochastic differential equations (FBSDE) where the randomness in the backward equation (BSDE) is driven by a forward stochastic differential equation (SDE). These systems are written in the general form

$$X_t = x_0 + \int_0^t \mu(s, X_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\mathrm{d}W_s, \tag{2.1a}$$

$$Y_t = g(X_T) + \int_t^T f(s, X_s, Y_s, Z_s)\mathrm{d}s - \int_t^T \left(Z_s \mathrm{d}W_s\right)^T, \tag{2.1b}$$

where $\{W_t\}_{0\leq t\leq T}$ is a d-dimensional Brownian motion and $\mu : [0,T]\times\mathbb{R}^{d\times 1} \to \mathbb{R}^{d\times 1}$, $\sigma : [0,T]\times\mathbb{R}^{d\times 1} \to \mathbb{R}^{d\times d}$, $g : \mathbb{R}^{d\times 1} \to \mathbb{R}^{q\times 1}$ and $f : [0,T]\times\mathbb{R}^{d\times 1}\times\mathbb{R}^{q\times 1}\times\mathbb{R}^{q\times d} \to \mathbb{R}^{q\times 1}$ are all deterministic mappings of time and space, with some fixed $T > 0$. Adhering to the stochastic control terminology, we often refer to Z as the *control process*. We shall work under the standard well-posedness assumptions of Pardoux and Peng [126], which require Lipschitz continuity of the corresponding coefficients in order to ensure the existence of a unique solution pair $\{(Y_t, Z_t)\}_{0\leq t\leq T}$ adapted to the augmented natural filtration. The main motivation to study FBSDE systems lies in their connection with parabolic, second-order partial differential equations (PDE), generalizing the well-known Feynman-Kac relations to non-linear settings. Indeed, considering the semi-linear, parabolic terminal problem

$$\begin{aligned} \partial_t u(t,x) + \frac{1}{2}\mathrm{Tr}\left\{\sigma\sigma^T(t,x)\mathrm{Hess}_x u(t,x)\right\} + \nabla_x u(t,x)\mu(t,x) + f(t,x,u,\nabla_x u(t,x)\sigma(t,x)) = 0 \\ u(T,x) = g(x), \end{aligned} \tag{2.2}$$

the Markov solution to (2.1) coincides with the solution of (2.2) in an almost sure sense, provided by the *non-linear Feynman-Kac* relations

$$Y_t = u(t, X_t), \qquad Z_t = \nabla_x u(t, X_t)\sigma(t, X_t). \tag{2.3}$$

Consequently, the BSDE formulation provides a stochastic representation to the simultaneous solution of a parabolic problem and its gradient, which is an advantageous feature for several applications in stochastic control and finance, where sensitivities play a fundamental role. These relations can be extended to *viscosity solutions* in case (2.2) does not admit to a classical solution – see [126]. Moreover, it is known – see [126, 46, 73, 110] – that under suitable regularity assumptions the solution pair of the backward equation is differentiable in the Malliavin sense [119], and the Malliavin derivatives $\{(D_s Y_t, D_s Z_t)\}_{0\leq s,t\leq T}$ satisfy a linear BSDE themselves, where the Z process admits to a continuous modification provided by $Z_t = D_t Y_t$.

From a numerical standpoint, the main challenge in solving BSDEs stems from the approximation of conditional expectations. Indeed, a discretization of the backward equation in (2.1b) yields a sequence of recursively nested conditional expectations at each point in the discretized time window. Over the years, several methods have been

proposed to tackle the solution of the FBSDE system using: PDE methods in [102]; forward Picard iterations in [13]; quantization techniques in [9]; chaos expansion formulas in [25]; Fourier cosine expansions in [144, 145] and regression Monte Carlo approaches in [63, 24, 16]. These methods have shown great results in low-dimensional settings, however, the majority of them suffers from the curse of dimensionality, meaning that their computational complexity scales exponentially in the number of dimensions. Although, regression Monte Carlo methods have been successfully proven to overcome this burden, they are difficult to apply beyond $d = 10$ dimensions due to the necessity of a finite regression basis. The primary challenge in the numerical solution of BSDEs is related to the approximation of the Z process. In particular, the standard backward Euler discretization results in a conditional expectation estimate of Z which scales inverse proportionally with the step size of the time discretization – see [24]. This phenomenon poses a significant amount of difficulty in least-squares Monte Carlo frameworks, as the corresponding regression targets have diverging conditional variances in the continuous limit.

Recently, the field has received renewed attention due to the pioneering paper of Han et al. [68], in which they reformulate the backward discretization in a forward fashion, parametrize the control process of the solution by deep neural networks and train the resulting sequence of networks in a global optimization given by the terminal condition of (2.1b). Their method has enjoyed various modifications and extensions, see, e.g., [53, 11]. In particular, Huré et al. in [77] proposed an alternative where the optimization of the sequence of neural networks is done in a backward recursive manner, similarly to classical regression Monte Carlo approaches. We refer to the class of these deep learning based formulations as *Deep BSDE* methods, which have shown remarkable empirical results in solving high-dimensional problems. Note, however, that the approach of [68] solely captures the deterministic mapping connecting the forward diffusion in (2.1) to the solution pair of the BSDE at $t = 0$. Even though the extension of [77] gives such approximations at future time steps, the accuracy of both methods degrades significantly in the Z part of the solution. The total approximation errors of such Deep BSDE methods have been investigated in [69, 77, 58]. The results in [69] provide a *posteriori estimate* driven by the error in the terminal condition, whereas the analyses in [77, 58] show that due to the universal approximation theorem (UAT) of deep neural networks, the total approximation error of neural network parametrizations is consistent with the discretization in terms of regression biases.

The main motivation behind the present chapter roots in the observations above. In order to provide more accurate solutions for the Z process, we exploit the aforementioned relation between the Malliavin derivative of Y and the control process by solving the linear BSDE driving the trajectories of DY. Hence, we are faced with the solution of one scalar-valued BSDE and one d-dimensional BSDE at each point in time. This raises the need for a new discrete scheme, which we call the *One Step Malliavin (OSM)* scheme. The discretization of the linear BSDE of the Malliavin derivatives is based on a merged formulation of the Feynman-Kac formulae in theorem 2.3 and the chain rule formula of Malliavin calculus [119]. As we shall see, the resulting discrete time approximation of the Z process possesses the same order of conditional variance as the ones of the Y process, making the scheme significantly more attractive in a regression Monte Carlo

framework compared to classical Euler discretizations. On the other hand, our formulation carries an extra layer of difficulty, in that we are forced to approximate the *"the Z of the Z, i.e. Γ processes"* [62, Pg.1184] in the Malliavin BSDE which are, in light of theorem 2.3, related to the Hessian matrix of the solution of the corresponding parabolic problem (2.2). In this regard, our setting shares similarities with *second-order backward SDEs (2BSDEs)* [33] and fully non-linear problems [47]. We analyze the discrete time approximation errors and show that under certain assumptions the new scheme has the same $\mathbb{L}^2$ convergence rate of order $1/2$ as the backward Euler scheme of BSDEs [24].

Two fully-implementable approaches are investigated to solve the resulting discretization. First, we provide an extension to the BSDE-COS (BCOS) method [144] and approximate solutions to one-dimensional problems by Fourier cosine expansions. Ultimately, the presence of Γ estimates induces d^2 many additional conditional expectations to be approximated at each point in time, which makes the OSM scheme less tractable for classical Monte Carlo parametrizations when d is large. Thereafter, inspired by the encouraging results of Deep BSDE methods in case of high-dimensional equations, we propose a neural network least-squares Monte Carlo approach similar to the one of [77], where the Y, Z and Γ processes are parametrized by fully-connected, feedforward deep neural networks. Subsequently, parameters of these networks are optimized in a recursive fashion, backwards over time, where at each time step two distinct gradient descent optimizations are performed, minimizing losses corresponding to the aforementioned discretization. Motivated by the UAT property of neural networks in Sobolev spaces, similarly to [77], we consider two variants of the latter approach: one in which the Γ process is parametrized by a matrix-valued deep neural network; and one in which the Γ process is approximated as the Jacobian of the parametrization of the Z process, inspired by theorem 2.3. The total approximation error is investigated similarly to [58, 77] and shown to be consistent with the discretization under the assumption of perfectly converging gradient descent iterations. We demonstrate the accuracy and robustness of our problem formulation with numerical experiments. In particular, using BCOS as a benchmark method for one-dimensional problems, we empirically assess the regression errors induced by gradient descent. We provide examples up to $d = 50$ dimensions.

The rest of the chapter is organized as follows. In section 2.2 we provide the necessary theoretical foundations, followed by section 2.3 where the new discrete scheme is formulated. In section 2.4 a discrete time approximation error analysis is given, bounding the total discretization error of the proposed scheme. Section 2.5 is concerned with the implementation of the discretization scheme, giving two fully-implementable approaches for the arising conditional expectations. First, the BCOS method [144] is extended in case of one-dimensional problems, then a Deep BSDE [68, 77] approach is formulated for high-dimensional equations. A complete regression error analysis is provided, building on the universal approximation properties of neural networks. Our analysis is concluded by numerical experiments presented in 2.6, which confirm the theoretical results and showcase great accuracy over a wide range of different problems.

2.2. BACKWARD STOCHASTIC DIFFERENTIAL EQUATIONS AND MALLIAVIN CALCULUS

In the following section we introduce the notions of BSDEs and Malliavin calculus used throughout the chapter.

2.2.1. PRELIMINARIES

Let us fix $0 \leq T < \infty$ and $d, q, n, k \in \mathbb{N}^+$. We are concerned with a filtered probability space $(\Omega, \mathscr{F}, \mathbb{P}, \{\mathscr{F}\}_{0\leq t\leq T})$, where $\mathscr{F} = \mathscr{F}_T$ and $\{\mathscr{F}\}_{0\leq t\leq T}$ is the natural filtration generated by a d-dimensional Brownian motion $\{W_t\}_{0\leq t\leq T}$ augmented by $\mathbb{P}$-null sets of Ω. In what follows, all equalities concerning $\mathscr{F}_t$-measurable random variables are meant in the $\mathbb{P}$-a.s. sense and all expectations – unless otherwise stated – are meant under $\mathbb{P}$. Throughout the whole chapter we rely on the following notations

- $|x| \coloneqq \mathrm{Tr}\{x^T x\}$ for the Frobenius norm of any $x \in \mathbb{R}^{q\times d}$. In case of scalar and vector inputs this coincides with the standard Euclidean norm. Additionally, we put $\langle x|y\rangle$ for the Euclidean inner product of $x, y \in \mathbb{R}^d$.
- $\mathbb{S}^p(\mathbb{R}^{q\times d})$ for the space of continuous and progressively measurable stochastic processes $Y : \Omega \times [0,T] \to \mathbb{R}^{q\times d}$ such that $\mathbb{E}\left[\sup_{0\leq t\leq T} |Y|^p\right] < \infty$.
- $\mathbb{H}^p(\mathbb{R}^{q\times d})$ for the space of progressively measurable stochastic processes $Z : \Omega \times [0,T] \to \mathbb{R}^{q\times d}$ such that $\mathbb{E}\left[\left(\int_0^T |Z_t|^2 \mathrm{d}t\right)^{p/2}\right] < \infty$.
- $\mathbb{L}^p_{\mathscr{F}_t}(\mathbb{R}^{q\times d})$ for the space of $\mathscr{F}_t$-measurable random variables $\xi : \Omega \to \mathbb{R}^{q\times d}$ such that $\mathbb{E}\left[|\xi|^p\right] < \infty$.
- $L^2([0,T];\mathbb{R}^q)$ for the Hilbert space of deterministic functions $h : [0,T] \to \mathbb{R}^q$ such that $\int_0^T |h(t)|^2 \mathrm{d}t < \infty$. Additionally, we denote its inner product by $\langle h|g\rangle_{L^2} \coloneqq \int_0^T \langle h(t)|g(t)\rangle \mathrm{d}t$.
- $\nabla_x f \coloneqq \left(\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_d}\right)$ for the gradient of a scalar-valued, multivariate function $(t,x,y,z) \mapsto f(t,x,y,z)$ with respect to $x \in \mathbb{R}^d$, defined as a row vector, and analogously for $\nabla_y f, \nabla_z f$. Similarly, we denote the Jacobian matrix of a vector-valued function $\psi : \mathbb{R}^d \to \mathbb{R}^q$ by $\nabla_x \psi \in \mathbb{R}^{q\times d}$. For notational convenience, we set the Jacobian matrix of row and column vector-valued functions in the same fashion.
- $C^k_b(\mathbb{R}^d;\mathbb{R}^q), C^k_p(\mathbb{R}^d;\mathbb{R}^q)$ for the set of k-times continuously differentiable functions $\varphi : \mathbb{R}^d \to \mathbb{R}^q$ such that all partial derivatives up to order k are bounded or have polynomial growth, respectively.
- $\mathbb{E}_n[\Phi] \coloneqq \mathbb{E}\left[\Phi|\mathscr{F}_{t_n}\right]$ for conditional expectations with respect to the natural filtration, given a time partition $0 = t_0 < t_1 < \cdots < t_N = T$. We occasionally use the notation $\mathbb{E}^x_n[\Phi] \coloneqq \mathbb{E}\left[\Phi|X_{t_n} = x\right]$ when the filtration is generated by a Markov process X.
- $\mathbf{1}_{q,d}, \mathbf{0}_{q,d}$ for $\mathbb{R}^{q\times d}$ matrices full of ones and zeros, respectively.

By slight abuse of notation we put $\mathbb{S}^p(\mathbb{R}) \coloneqq \mathbb{S}^p(\mathbb{R}^{1\times 1})$, $\mathbb{H}^p(\mathbb{R}^d) \coloneqq \mathbb{H}^p(\mathbb{R}^{1\times d})$, $\mathbf{1}_d \coloneqq \mathbf{1}_{1,d}$ and $\mathbf{0}_d \coloneqq \mathbf{0}_{1\times d}$.

We recall the most important notions of Malliavin differentiability and refer to [119] for a more detailed account on the subject. Consider the space of random processes $W(h) \coloneqq \int_0^T h(t)\mathrm{d}W_t$ with $h \in L^2([0,T];\mathbb{R}^{1\times d})$. Let us now define the subspace $\mathscr{R} \subseteq \mathbb{L}^2_{\mathscr{F}_T}$ of smooth, scalar-valued random variables which are of the form $\Phi = \varphi(W(h_1),\dots,W(h_n))$ with some $\varphi \in C_p^\infty(\mathbb{R}^n;\mathbb{R})$. The Malliavin derivative of Φ is then defined as the $\mathbb{R}^{1\times d}$-valued stochastic process $D_s\Phi \coloneqq \sum_{i=1}^n \partial_i\varphi(W(h_1),\dots,W(h_n))h_i(s)$. The derivative operator can be extended to the closure of $\mathscr{R}$ with respect to the norm

$$\|\Phi\|_{\mathbb{D}^{1,p}} \coloneqq \left(\mathbb{E}\left[|\Phi|^p + \left(\int_0^T |D_s\Phi|^2\mathrm{d}s\right)^{p/2}\right]\right)^{1/p},$$

see [119, Prop.1.2.1]. We denote this closure as the space of Malliavin differentiable, $\mathbb{R}$-valued random variables by $\mathbb{D}^{1,p}(\mathbb{R})$. For the space of vector-valued $\Phi = (\Phi_1,\dots,\Phi_q)$ Malliavin differentiable random variables, we put $\Phi \in \mathbb{D}^{1,p}(\mathbb{R}^q)$ when $\Phi_i \in \mathbb{D}^{1,p}(\mathbb{R})$ for each $i = 1,\dots,q$. The Malliavin derivative $D_s\Phi \in \mathbb{R}^{q\times d}$ is then the matrix-valued stochastic process whose i'th row is $D_s\Phi_i$. The final result which extends the chain rule of elementary calculus to the Malliavin differentiation operator is fundamental for the present chapter, essentially enabling the formulation of the upcoming discrete scheme.

Lemma 2.2.1 (Malliavin chain rule lemma)
Let $\psi \in C_b^1(\mathbb{R}^d;\mathbb{R}^q)$ and fix $p \geq 1$. Consider $F \in \mathbb{D}^{1,p}(\mathbb{R}^d)$. Then $\psi(F) \in \mathbb{D}^{1,p}(\mathbb{R}^q)$, furthermore for each $0 \leq s \leq T$

$$D_s\psi(F) = \nabla_x\psi(F)D_sF.$$

The lemma can be relaxed to the case where ψ is only Lipschitz continuous – see [119, Prop.1.2.4].

2.2.2. BACKWARD STOCHASTIC DIFFERENTIAL EQUATIONS

We first provide the necessary theoretical foundations for the well-posedness of the underlying FBSDE system in (2.1) guaranteeing the existence of a unique solution triple. Given the stronger assumptions later required for their Malliavin differentiability, we restrict the presentation to standard Lipschitz assumptions. For a more general exposure we refer to [31] and the references therein.

It is well-known – see, e.g., [89] – that the SDE in (2.1a) admits to a unique strong solution $\{X_t\}_{0\leq t\leq T} \in \mathbb{S}^p(\mathbb{R}^{d\times 1})$ whenever $x_0 \in \mathbb{L}^p_{\mathscr{F}_0}(\mathbb{R}^{d\times 1})$ and μ,σ are Lipschitz continuous in the spatial variable, i.e.

$$\big|\mu(t,x_1)-\mu(t,x_2)\big| + |\sigma(t,x_1)-\sigma(t,x_2)| \leq L_{\mu,\sigma}|x_1-x_2|$$

for all $t \in [0,T]$, $x_1,x_2 \in \mathbb{R}^{d\times 1}$, with some $L_{\mu,\sigma} > 0$. Additionally, the solution $\{X_t\}_{0\leq t\leq T}$ satisfies the following estimates for all $p \geq 1$

$$\mathbb{E}\left[\sup_{0\leq t\leq T}|X_t|^p\right] \leq C_p, \qquad \mathbb{E}\left[|X_t - X_s|^p\right] \leq C_p|t-s|^{p/2}, \tag{2.4}$$

with constant C_p only depending on p, T, d. In case of the Arithmetic Brownian Motion (ABM) with constant μ and σ, (2.1a) admits to the unique solution $X_t = x_0 + \mu t + \sigma W_t$. In particular, the Malliavin chain rule formula in lemma 2.2.1 implies that $D_s X_t = \mathbb{1}_{s\leq t}\sigma$.

The well-posedness of the backward equation in (2.1b) is guaranteed – see, e.g., [46] – by the Lipschitz continuity of the driver, on top of the polynomial growth of the terminal condition

$$\begin{aligned}\left|f(t,x,y_1,z_1) - f(t,x,y_2,z_2)\right| &\leq L_{f,g}\left(\left|y_1 - y_2\right| + |z_1 - z_2|\right),\\ \left|f(t,x,y,z)\right| + \left|g(x)\right| &\leq L_{f,g}\left(1 + |x|^p\right),\end{aligned}$$

for any $t \in [0,T]$, $y_1, y_2 \in \mathbb{R}^q$, $z_1, z_2 \in \mathbb{R}^{q\times d}$, with some $L_{f,g} > 0$ and $p \geq 2$. These conditions, combined with the ones for the SDEs above, imply the existence of a unique solution pair $Y \in \mathbb{S}^p(\mathbb{R}^q)$, $Z \in \mathbb{H}^p(\mathbb{R}^{q\times d})$ satisfying (2.1b). Let us now fix $q = 1$ and restrict the further analysis to scalar-valued backward equations. Thereafter, under the aforementioned conditions, the FBSDE system in (2.1) admits to a unique solution triple $\{(X_t, Y_t, Z_t)\}_{0\leq t\leq T} \in \mathbb{S}^p(\mathbb{R}^{d\times 1}) \times \mathbb{S}^p(\mathbb{R}) \times \mathbb{H}^p(\mathbb{R}^{1\times d})$.

2.2.3. MALLIAVIN DIFFERENTIABLE FBSDE SYSTEMS

This chapter is focused on a special class of FBSDE systems such that the solution triple $\{(X_t, Y_t, Z_t)\}_{0\leq t\leq T}$ is differentiable in the Malliavin sense. The Malliavin differentiability of the forward equation is guaranteed by the following theorem due to Nualart in [119, Thm.2.2.1].

Lemma 2.2.2 (Malliavin differentiability of SDEs, [119])
Let $x_0 \in \mathbb{L}^p_{\mathscr{F}_0}(\mathbb{R}^{d\times 1})$, $\mu \in C_b^{0,1}([0,T]\times\mathbb{R}^{d\times 1};\mathbb{R}^{d\times 1})$, $\sigma \in C_b^{0,1}([0,T]\times\mathbb{R}^{d\times 1};\mathbb{R}^{d\times d})$ and $\mu(t,0)$, $\sigma(t,0)$ be uniformly bounded for all $0 \leq t \leq T$. Put $\{X_t\}_{0\leq t\leq T}$ for the unique solution of (2.1a). Then for all $t \in [0,T]$, $X_t \in \mathbb{D}^{1,p}(\mathbb{R}^{d\times 1})$ and there exists a continuous modification of its Malliavin derivative $\{D_s X_t\}_{0\leq s,t\leq T} \in \mathbb{S}^p(\mathbb{R}^{d\times d})$ which satisfies the linear SDE

$$D_s X_t = \mathbb{1}_{s\leq t}\left\{\sigma(s,X_s) + \int_s^t \nabla_x\mu(r,X_r)D_s X_r\,\mathrm{d}r + \int_s^t \nabla_x\sigma(r,X_r)D_s X_r\,\mathrm{d}W_r.\right\}, \tag{2.5}$$

where $\nabla_x\sigma$ denotes a $\mathbb{R}^{d\times d\times d}$-valued tensor with $[\nabla_x\sigma]_{ijk} = \partial_k[\sigma]_{ij}$. Furthermore, there exists a constant C_p, only depending on p, T, d, such that

$$\sup_{s\in[0,T]}\mathbb{E}\left[\sup_{t\in[s,T]}|D_s X_t|^p\right] \leq C_p, \quad \mathbb{E}\left[|D_s X_r - D_s X_t|^p\right] \leq C_p|r-t|^{p/2}, \quad \forall r,t \geq s. \tag{2.6}$$

The main implication of the proposition above is that under relatively mild assumptions on the bounded continuous differentiability of the coefficients in (2.1a), the Malliavin derivative of the solution satisfies a linear SDE, where the random coefficients depend on the solution of the SDE itself. Intriguingly, a similar assertion can be made about the solution pair of the backward equation in (2.1b), which – on top of establishing their Malliavin differentiability – also creates a connection between the Malliavin derivative DY and the control process. This is stated by the following theorem originally from Pardoux and Peng in [126], which we state under the loosened conditions of El Karoui et al. [46, Prop.5.9].

Theorem 2.2.1 (Malliavin differentiability of BSDEs, [46])
Let the coefficients of (2.1a) satisfy the conditions of lemma 2.2.2 and assume $f \in C_b^{0,1,1,1}([0,T]\times\mathbb{R}^{d\times 1},\mathbb{R},\mathbb{R}^{1\times d};\mathbb{R})$, $g\in C_b^1(\mathbb{R}^{d\times 1};\mathbb{R})$. *Fix* $p\geq 2$. *Put* $\{(Y_t,Z_t)\}_{0\leq t\leq T}$ *for the unique solution pair of (2.1b). Then for all* $t\in[0,T]$ $Y_t\in\mathbb{D}^{1,2}(\mathbb{R})$, $Z_t\in\mathbb{D}^{1,2}(\mathbb{R}^{1\times d})$ *and there exist modifications of their Malliavin derivatives* $\{D_sY_t\}_{0\leq s,t\leq T}\in\mathbb{S}^p(\mathbb{R}^{1\times d})$, $\{D_sZ_t\}_{0\leq,s,t\leq T}\in\mathbb{H}^p(\mathbb{R}^{d\times d})$ *which satisfy the following linear BSDE*

$$\begin{aligned}
D_sY_t &= \nabla_x g(X_T)D_sX_T \\
&\quad + \int_t^T \nabla_x f(r,X_r,Y_r,Z_r)D_sX_r + \nabla_y f(r,X_r,Y_r,Z_r)D_sY_r \\
&\qquad + \nabla_z f(r,X_r,Y_r,Z_r)D_sZ_r\,\mathrm{d}r \\
&\quad - \int_t^T \left((D_sZ_r)^T\mathrm{d}W_r\right)^T, \quad 0\leq s\leq t\leq T, \\
D_sY_t &= \mathbf{0}_d, \quad D_sZ_t = \mathbf{0}_{d,d}, \quad 0\leq t<s\leq T.
\end{aligned} \tag{2.7}$$

Furthermore, there exists a continuous modification of the control process such that $Z_t = D_tY_t$ *almost surely for all* $0\leq t\leq T$.

We emphasize the linearity of (2.7) and remark that the corresponding random coefficients of the linear equation depend on the solution of (2.1). Henceforth, in light of lemma 2.2.2 and theorem 2.2.1, we define $\{D_sX_t\}_{0\leq s,t\leq T}$ and $\{D_sY_t\}_{0\leq s,t\leq T}, \{D_sZ_t\}_{0\leq s,t\leq T}$ as the versions of the corresponding Malliavin derivatives satisfying (2.5) and (2.7), respectively. For the rest of the chapter, in order to ease the presentation, we introduce the notations $\mathbf{X}_t := (X_t,Y_t,Z_t)$, $\mathbf{D}_s\mathbf{X}_t := (D_sX_t,D_sY_t,D_sZ_t)$ and $f^D(t,\mathbf{X}_t,\mathbf{D}_s\mathbf{X}_t) := \nabla_x f(t,\mathbf{X}_t)D_sX_t + \nabla_y f(t,\mathbf{X}_t)D_sY_t + \nabla_z f(t,\mathbf{X}_t)D_sZ_t$ for all $0\leq s,t\leq T$.

Path regularity and Hölder continuity. For $\{X_t\}_{0\leq t\leq T}\in\mathbb{S}^p(\mathbb{R}^{d\times 1})$ we have that the solution of the forward SDE is a continuous $\mathbb{R}^{d\times 1}$-valued random process which is bounded in the supremum norm. Similar statements can be made about its Malliavin derivative $\{D_sX_t\}_{0\leq s,t\leq T}$. In particular, the Hölder regularity estimates in (2.4) and (2.6) ensure that the corresponding processes are not just continuous but also have a modification admitting to α-Hölder continuous trajectories of order $\alpha\in(0,1/2)$ provided by the Kolmogorov-Chentsov theorem – see, e.g., [89]. Since the $1/2$-Hölder regularity of (Y,Z) plays a crucial role in the convergence analysis of the discrete scheme – see theorem 2.4.1 in particular –, we elaborate on the conditions under which the continuous parts of the solutions to (2.1b) and (2.7) admit to similar estimates. Indeed, one can show that if the solutions $(Y,Z)\in\mathbb{S}^p(\mathbb{R})\times\mathbb{H}^p(\mathbb{R}^{d\times 1})$ of (2.1b) satisfy the condition $\sup_{0\leq t\leq T}\mathbb{E}\left[|Z_t|^p\right]<\infty$ then there exists a constant C_p such that

$$\mathbb{E}\left[|Y_t-Y_s|^p\right]\leq C_p|t-s|^{p/2},$$

see [73, Corollary 2.7]. In particular, the Y process admits to an α-Hölder continuous modification of order $\alpha\in(0,1/2-1/p)$. Under the conditions of theorem 2.2.1, this is naturally guaranteed, and for $p=2$ it implies the *mean-squared continuity* of the Y process. Moreover, the Z process admits to a continuous modification solving (2.7), which

guarantees $Z \in \mathbb{S}^p(\mathbb{R}^{1\times d})$ and, in particular, boundedness in the supremum norm. Under stronger assumptions one can also establish a similar path regularity result of the control process. Imkeller and Dos Reis in [78, Thm.5.5] show that with additional conditions, essentially requiring second-order bounded differentiability of the corresponding coefficients μ, σ, f and g, the following also holds for all $p \geq 2$

$$\mathbb{E}\left[|Z_t - Z_s|^p\right] \leq C_p |t-s|^{p/2}. \tag{2.8}$$

Hu et al. prove a similar result in [73, Thm.2.6] under slightly different assumptions in the general non Markovian framework. We omit the explicit presentation of the necessary conditions for (2.8) to hold, nevertheless emphasize that assumption 2.4.1 of the convergence analysis in sec. 2.4 ensures the path regularity of the Z process and in particular implies mean-squared continuous trajectories.

2.3. THE DISCRETE SCHEME

In the following section the proposed discretization scheme is introduced. The objective of the discretization is to simultaneously solve the pair of FBSDE systems given by (2.1) and the FBSDE system of its Malliavin derivatives provided by lemma 2.2.2 and theorem 2.2.1. Therefore, we are concerned with the solution to the following pair of FBSDE systems

$$X_t = x_0 + \int_0^t \mu(r, X_r)\mathrm{d}r + \int_0^t \sigma(r, X_r)\mathrm{d}W_r, \tag{2.9a}$$

$$Y_t = g(X_T) + \int_t^T f(r, \mathbf{X}_r)\mathrm{d}r - \int_t^T Z_r \mathrm{d}W_r, \tag{2.9b}$$

$$D_s X_t = \mathbb{1}_{s\leq t}\left[\sigma(s, X_s) + \int_s^t \nabla_x \mu(r, X_r) D_s X_r \mathrm{d}r + \int_s^t \nabla_x \sigma(r, X_r) D_s X_r \mathrm{d}W_r\right], \tag{2.9c}$$

$$D_s Y_t = \mathbb{1}_{s\leq t}\left[\nabla_x g(X_T) D_s X_T + \int_t^T f^D(r, \mathbf{X}_r, \mathbf{D}_s\mathbf{X}_r)\mathrm{d}r - \int_t^T \left((D_s Z_r)^T \mathrm{d}W_r\right)^T\right]. \tag{2.9d}$$

The solution is a pair of triples of stochastic processes $\{(X_t, Y_t, Z_t)\}_{0\leq t\leq T}$ and $\{(D_s X_t, D_s Y_t, D_s Z_t)\}_{0\leq s,t\leq T}$ such that (2.9) holds $\mathbb{P}$ almost surely. Consider a discrete time partition $\pi^N := \{t_0, \ldots, t_N\}$ with $0 = t_0 < t_1 < \cdots < t_N = T$ and set $\Delta W_n := W_{t_{n+1}} - W_{t_n}$, $\Delta t_n := t_{n+1} - t_n$, $|\pi| := \max_{0\leq n\leq N-1} t_{n+1} - t_n$. We denote the discrete time approximations by $\mathbf{X}_n^\pi := (X_n^\pi, Y_n^\pi, Z_n^\pi)$ and $\mathbf{D}_n\mathbf{X}_m^\pi := (D_n X_m^\pi, D_n Y_m^\pi, D_n Z_m^\pi)$ for each $0 \leq n, m \leq N$.

The forward component in (2.9a) is approximated by the classical Euler-Maruyama scheme, i.e.,

$$X_0^\pi := x_0, \quad X_{n+1}^\pi := X_n^\pi + \mu(t_n, X_n^\pi)\Delta t_n + \sigma(t_n, X_n^\pi)\Delta W_n^\pi, \tag{2.10}$$

for each $n = 0, \ldots, N-1$. It is well-known – see, e.g., [90] – that under standard Lipschitz assumptions on the drift and diffusion coefficients, these estimates admit to

$$\limsup_{|\pi|\to 0} \frac{1}{|\pi|}\mathbb{E}\left[\left|X_{t_n} - X_n^\pi\right|^2\right] < \infty. \tag{2.11}$$

Classically, the backward component in (2.9b) is approximated in two steps. In order to meet the necessary adaptivity requirements of the solution pair (Y,Z), one takes appropriate conditional expectations of (2.9b) and the same equation multiplied with the Brownian increment ΔW_n^T. Using standard properties of stochastic integrals, Itô's isometry and a *theta-discretization* of the remaining time integrals with parameters $\vartheta_y,\vartheta_z>0$ subsequently give – see, e.g., [144]

$$\begin{aligned} Y_N^\pi &= g(X_N^\pi), \quad Z_N^\pi = \nabla_x g(X_N^\pi)\sigma(t_N,X_N^\pi),\\ Z_n^\pi &= -\frac{1-\vartheta_z}{\vartheta_z}\mathbb{E}_n\left[Z_{n+1}^\pi\right] + \frac{1}{\Delta t_n \vartheta_z}\mathbb{E}_n\left[\Delta W_n^T Y_{n+1}^\pi\right] \\ &\qquad + \frac{1-\vartheta_z}{\vartheta_z}\mathbb{E}_n\left[\Delta W_n^T f(t_{n+1},\mathbf{X}_{n+1}^\pi)\right], \end{aligned} \tag{2.12a}$$

$$Y_n^\pi = \Delta t_n \vartheta_y f(t_n,X_n^\pi,Y_n^\pi,Z_n^\pi) + \mathbb{E}_n\left[Y_{n+1}^\pi\right] + \Delta t_n(1-\vartheta_y)\mathbb{E}_n\left[f(t_{n+1},\mathbf{X}_{n+1}^\pi)\right]. \tag{2.12b}$$

In case $\vartheta_y=\vartheta_z=1$, this scheme is called the standard *Euler scheme for BSDEs*.

2.3.1. THE OSM SCHEME

The novelty of the hereby proposed discretization is that on top of solving (2.9b), we also solve the linear BSDE in (2.9d) driving the Malliavin derivatives of the solution pair. Exploiting the relation between DY and Z established by theorem 2.2.1, we set the control estimates according to the discrete time approximations of the Malliavin BSDE. As in the case of the forward component itself, the Malliavin derivative in (2.9c) is approximated by an Euler-Maruyama discretization, giving estimates

$$D_n X_m^\pi := \begin{cases} \mathbb{1}_{m=n}\sigma(t_n,X_n^\pi), & 0\le m\le n\le N,\\ D_n X_{m-1}^\pi + \nabla_x\mu(t_{m-1},X_{m-1}^\pi)D_n X_{m-1}^\pi \Delta t_{m-1} & \\ \qquad + \nabla_x\sigma(t_{m-1},X_{m-1}^\pi)D_n X_{m-1}^\pi \Delta W_{m-1}, & 0\le n<m\le N. \end{cases} \tag{2.13}$$

Unlike in the case of X_n^π, the convergence of these approximations is not straightforward due to the fact that the initial condition $D_nX_n^\pi=\sigma(t_n,X_n^\pi)$ already depends on the discrete approximation X_n^π provided by (2.10). Nonetheless, as we shall soon see, our discretization of the linear BSDE in (2.9d) only relies on the approximations $D_nX_{n+1}^\pi$ for each $n=0,\dots,N-1$. This is a significant relaxation of the convergence criterion, as it can be shown that under relatively mild assumptions on the coefficients in (2.9a), $D_nX_{n+1}^\pi$ defined by (2.13) inherits the convergence rate of (2.11) – see appendix 2.A for details.

The discretization of the backward component in (2.9d) is done as follows. For any $n=0,\dots,N-1$

$$D_{t_n}Y_{t_n} = D_{t_n}Y_{t_{n+1}} + \int_{t_n}^{t_{n+1}} f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r)\mathrm{d}r - \int_{t_n}^{t_{n+1}}\left((D_{t_n}Z_r)^T\mathrm{d}W_r\right)^T, \tag{2.14}$$

subject to the terminal condition. Multiplying this equation with ΔW_n from the left, Itô's isometry implies

$$\begin{aligned} \mathbb{E}_n\left[\int_{t_n}^{t_{n+1}} D_{t_n}Z_r\mathrm{d}r\right] &= \mathbb{E}_n\left[\Delta W_n\left(D_{t_n}Y_{t_{n+1}} + \int_{t_n}^{t_{n+1}} f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r)\mathrm{d}r\right)\right],\\ D_{t_n}Y_{t_n} &= \mathbb{E}_n\left[D_{t_n}Y_{t_{n+1}} + \int_{t_n}^{t_{n+1}} f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r)\mathrm{d}r\right]. \end{aligned} \tag{2.15}$$

In order to avoid implicitness on Y, we approximate the continuous time integrals with the left- and right rectangle rules, respectively, and obtain discrete time approximations

$$\begin{aligned} D_n Z_n^\pi &= \frac{1}{\Delta t_n}\mathbb{E}_n\left[\Delta W_n\left(D_n Y_{n+1}^\pi + \Delta t_n f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1}^\pi)\right)\right], \\ D_n Y_n^\pi &= \mathbb{E}_n\left[D_n Y_{n+1}^\pi + \Delta t_n f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1}^\pi)\right]. \end{aligned} \tag{2.16}$$

At this point, to make the scheme viable, one relies on estimates $D_n Y_m^\pi, D_n Z_m^\pi$ on top of the Euler-Maruyama approximations of DX given by (2.13). This is done by a merged formulation of the Feynman-Kac formulae in (2.3) and the Malliavin chain rule in lemma 2.2.1. Indeed, given the Markov nature of the FBSDE system, the solutions of (2.9b) can be written as $Y_t = y(t, X_t), Z_t = z(t, X_t)$ for some sufficiently smooth deterministic functions $y : [0,T] \times \mathbb{R}^{d\times 1} \to \mathbb{R}$, $z : [0,T] \times \mathbb{R}^{d\times 1} \to \mathbb{R}^{1\times d}$. Moreover, the Malliavin chain rule implies that

$$D_{t_n} Y_r = \nabla_x y(r, X_r) D_{t_n} X_r, \qquad D_{t_n} Z_r = \nabla_x z(r, X_r) D_{t_n} X_r =: \gamma(r, X_r) D_{t_n} X_r,$$

for some deterministic functions $y : [0,T] \times \mathbb{R}^{d\times 1} \to \mathbb{R}$ and $z : [0,T] \times \mathbb{R}^{d\times 1} \to \mathbb{R}^{1\times d}$, where we defined $\gamma : [0,T] \times \mathbb{R}^{d\times 1} \to \mathbb{R}^{d\times d}$ as the Jacobian matrix of $z(r, X_r)$, and similarly $\Gamma_t := \gamma(t, X_t)$. Furthermore, due to the Feynman-Kac relations we also have $z(r, X_r) = \nabla_x y(r, X_r)\sigma(r, X_r)$ and therefore

$$D_{t_n} Y_r = z(r, X_r)\sigma^{-1}(r, X_r) D_{t_n} X_r, \qquad D_{t_n} Z_r = \gamma(r, X_r) D_{t_n} X_r. \tag{2.17}$$

Motivated by these relations, we approximate the discretized Malliavin derivatives in (2.16) according to

$$D_n Y_m^\pi := Z_m^\pi \sigma^{-1}(t_m, X_m^\pi) D_n X_m^\pi, \qquad D_n Z_m^\pi := \Gamma_m^\pi D_n X_m^\pi, \qquad 0 \le n, m \le N. \tag{2.18}$$

Henceforth, the discrete approximations of the Y process driven by (2.9b) are given in an identical fashion to (2.12b) with $\vartheta_y \in [0,1]$ as a free parameter of the discretization. Moreover, in order to be able to control the $\mathbb{L}^2$ projection error of $D_n Z_m^\pi$ with discrete Grönwall estimates – see Step 1 of theorem 2.4.1 in particular –, we make the $\nabla_z f$ part of f^D implicit in $D_n Z_n^\pi$, and introduce the notation $\mathbf{D}_n \mathbf{X}_{n+1,n}^\pi := \left(D_n X_{n+1}^\pi, D_n Y_{n+1}^\pi, D_n Z_n^\pi\right)$. Subject to the terminal conditions in (2.9b) and (2.9d), on top of the Malliavin chain rule estimates in (2.18), this leads to the following discrete scheme, which we shall call the *One Step Malliavin* (OSM) scheme

$$Y_N^\pi = g(X_N^\pi), \quad Z_N^\pi = \nabla_x g(X_N^\pi)\sigma(t_N, X_N^\pi), \quad \Gamma_N^\pi = [\nabla_x(\nabla_x g\sigma)](t_N, X_N^\pi),$$

$$\Gamma_n^\pi \sigma(t_n, X_n^\pi) = D_n Z_n^\pi = \frac{1}{\Delta t_n}\mathbb{E}_n\left[\Delta W_n\left(D_n Y_{n+1}^\pi + \Delta t_n f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1,n}^\pi)\right)\right], \tag{2.19a}$$

$$Z_n^\pi = \mathbb{E}_n\left[D_n Y_{n+1}^\pi + \Delta t_n f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1,n}^\pi)\right], \tag{2.19b}$$

$$Y_n^\pi = \vartheta_y \Delta t_n f(t_n, X_n^\pi, Y_n^\pi, Z_n^\pi) + \mathbb{E}_n\left[Y_{n+1}^\pi + (1-\vartheta_y)\Delta t_n f(t_{n+1}, \mathbf{X}_{n+1}^\pi)\right]. \tag{2.19c}$$

The scheme is made fully implementable by an appropriate parametrization to approximate the arising conditional expectations.

Remark 2.3.1 (Comparison of discretizations)
There are two key differences between the standard Euler discretization in (2.12) and the OSM scheme in (2.19). First, unlike in the former, the OSM scheme's solution is a triple of discrete random processes, including an additional layer of Γ estimates. Moreover, it can be seen that the estimate in (2.19b) exhibits a better conditional variance than that of (2.12a). In case of the standard Euler discretization, the Z process is approximated through Itô's isometry and the corresponding discrete time approximations include a $1/\Delta t_n$ factor – second term in (2.12a) – which leads to a quadratically exploding conditional variance of the resulting estimates. Several variance reduction techniques have been proposed to mitigate this problem – we mention [62, 2]. On the other hand, within the OSM scheme, the Z process is approximated by the continuous solution of the Malliavin BSDE in (2.9d) and therefore it carries the same conditional variance behavior as the Y estimate. In case of a fully-implementable regression Monte Carlo setting, this explains why the OSM scheme may provide more accurate control approximations.

Alternative formulations. Equation (2.19) is not the first approach to the BSDE problem building on theorem 2.2.1. Turkedjiev in [154] proposed a discrete time approximation scheme, where the Z process is estimated by an integration by parts formula stemming from Malliavin calculus and discovered in [105, Thm.3.1]. Hu et al. in [73] proposed an explicit scheme in the case of non Markovian BSDEs, where the control process is estimated using a representation formula implied by the linearity of the Malliavin BSDE (2.9d) – see [46, Prop.5.5]. Briand and Labart in [25] offer a different approach to BSDEs, where building on chaos expansion formulas, the Z process is taken as the Malliavin derivative of Y given by theorem 2.2.1. The difference between these formulations and (2.19) is mostly twofold. The OSM scheme is concerned with solving the entire pair of FBSDE systems (2.9) and not just the backward component in (2.9b). This means that unlike in [154, 25, 73], discrete time approximations give Γ estimates as well. Additionally, one important difference in the OSM scheme compared to the approaches [154, 73] is that the conditional expectations in (2.19) project $\mathscr{F}_{t_{n+1}}$-measurable random variables onto $\mathscr{F}_{t_n}$, whereas in the case of those works the arguments of the conditional expectations are $\mathscr{F}_T$-measurable. An important implication of this difference is that – unlike [154, 73] –, in order to simulate the arguments of the arising conditional expectations in (2.19), one does not rely on discrete time approximations of the Malliavin derivatives $D_n X^\pi_m$ over the whole time window ($n \le m \le N$) but only in between adjacent time steps $D_n X^\pi_{n+1}$. This is an advantage from the convergence analysis perspective whenever one does not have analytical access to the trajectories of $\{D_s X_t\}_{0\le s,t\le T}$. In fact, ensuring the convergence of the Euler-Maruyama scheme for the Malliavin derivative in (2.13) for any $n \le m \le N$ is known to be non-trivial, see [73, Remark 5.1]. On the other hand, as shown in appendix 2.A, under suitable regularity assumptions, $D_n X^\pi_{n+1}$ converges in the $\mathbb{L}^2$-sense with a rate of $1/2$, which renders the convergence of the discrete time approximations of the OSM scheme possible.

2.4. DISCRETIZATION ERROR ANALYSIS

Having introduced the discrete scheme simultaneously solving the FBSDE system itself and the FBSDE system of its solutions' Malliavin derivatives, we investigate the errors induced by the discretization of the continuous processes in (2.19). It is known – see [24] – that the $\mathbb{L}^2$ discretization errors of the backward Euler scheme in (2.12) admit to

$$\max_{0\leq n\leq N}\mathbb{E}\left[\left|Y_{t_n}-Y_n^\pi\right|^2\right]+\mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}}\left|Z_r-Z_n^\pi\right|^2\mathrm{d}r\right] \leq C\left(\mathbb{E}\left[\left|g(X_T)-g(X_N^\pi)\right|^2\right]+\varepsilon^Z(|\pi|)+|\pi|\right), \quad (2.20)$$

where $\varepsilon^Z(|\pi|) := \mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}}\left|Z_r-\bar{Z}_n^{n+1}\right|^2\mathrm{d}r\right]$ with $\bar{Z}_n^{n+1} := 1/\Delta t_n\mathbb{E}_n\left[\int_{t_n}^{t_{n+1}}Z_r\mathrm{d}r\right]$ according to [160]. The purpose of the following section is to show a similar result for the proposed OSM scheme and prove that it is *consistent* in the $\mathbb{L}^2$-sense, i.e. the discrete time approximations errors converge to zero as the mesh size of the time partition $|\pi|$ vanishes. In particular, we shall see that under standard Lipschitz assumptions on the driver f of the BSDE (2.9b) and the driver f^D of the linear Malliavin BSDE (2.9d), and additive noise in the forward diffusion, the convergence is of order $\mathcal{O}(|\pi|^{1/2})$.

Assumption 2.4.1
The following assumptions are in place.

($\mathbf{A}^{\mu,\sigma}$) *SDE*

- ($\mathbf{A}_1^{\mu,\sigma}$) *the forward equation has constant drift and diffusion coefficients (Arithmetic Brownian motion);*
- ($\mathbf{A}_2^{\mu,\sigma}$) *the forward SDE has a uniformly elliptic diffusion coefficient, i.e. for any $\zeta\in\mathbb{R}^{1\times d}$ there exists a $\beta>0$ such that $\zeta\sigma\sigma^T\zeta^T>\beta|\zeta|^2$* [2]*;*

($\mathbf{A}^{f,g}$) *BSDE*

- ($\mathbf{A}_1^{f,g}$) *$g\in C_b^{2+\alpha}(\mathbb{R})$ with some $\alpha>0$, furthermore g is also bounded;*
- ($\mathbf{A}_2^{f,g}$) *$f\in C_b^{0,2,2,2}(\mathbb{R})$;*
- ($\mathbf{A}_3^{f,g}$) *f and its partial derivatives $\nabla_x f,\nabla_y f,\nabla_z f$ are all $1/2$-Hölder continuous in time.*

The conditions above are not minimal – see also sec. 2.4.2. Nevertheless, for the sake of the present analysis they are sufficient. In particular, since bounded continuous differentiability implies Lipschitz continuity due to the mean-value theorem, by theorem 2.2.1 we have that under assumption 2.4.1 the FBSDE (2.9a)–(2.9b) is Malliavin differentiable, and the Malliavin derivatives of its solutions satisfy the FBSDE (2.9c)–(2.9d). Additionally, due to [40, Thm. 2.1], we can also exploit the following useful result from the theory of parabolic PDEs.

[2] We remark that this condition is equivalent to $A=\sigma\sigma^T$ being a positive definite matrix.

Lemma 2.4.1 ([40])
Under assumption 2.4.1 the parabolic PDE in (2.2) admits to a unique solution $u \in C_b^{1,2}(\mathbb{R})$.

Thereafter, provided by lemma 2.4.1, one can use the merged formulation of the Malliavin chain rule lemma 2.2.1 and the non-linear Feynman-Kac relations given by (2.17), in order to get the explicit formulas for the solutions of (2.9d) depending only on time and the state variable. We remark that in our setting $\sigma \in \mathbb{R}^{d\times d}$, the existence of the inverse is guaranteed by the uniform ellipticity condition set on σ in assumption 2.4.1. In case the Brownian motion and the forward diffusion have different dimensions, similar statements can be made about right inverses – see [154]. Another important implication of the estimate above is that assumption 2.4.1, through lemma 2.4.1, also implies that the driver of the Malliavin BSDE f^D is Lipschitz continuous in its spatial arguments within the bounded domain. Indeed, the mean-value theorem for $f \in C_b^{0,2,2,2}(\mathbb{R})$ implies that f and all its first-order derivatives in (x,y,z) are Lipschitz continuous, consequently for any uniformly bounded argument (DX, DY, DZ) the following holds

$$\begin{aligned}
&\left|f(t_1,\mathbf{x}_1)-f(t_2,\mathbf{x}_2)\right| \leq L_f\left(|t_1-t_2|^{1/2}+|x_1-x_2|+\left|y_1-y_2\right|+|z_1-z_2|\right),\\
&|\xi_1|,\left|\eta_1\right|,|\zeta_1| \leq L_{f^D}:\\
&\left|f^D(t_1,\mathbf{x}_1,\boldsymbol{\xi}_1)-f^D(t_2,\mathbf{x}_2,\boldsymbol{\xi}_2)\right| \leq L_{f^D}\Big(|t_1-t_2|^{1/2}+|x_1-x_2|+\left|y_1-y_2\right|+|z_1-z_2|\\
&\qquad\qquad +|\xi_1-\xi_2|+\left|\eta_1-\eta_2\right|+|\zeta_1-\zeta_2|\Big),
\end{aligned} \tag{2.21}$$

with $\mathbf{x}_i = \left(x_i, y_i, z_i\right)$, $\boldsymbol{\xi}_i := \left(\xi_i, \eta_i, \zeta_i\right)$, $i = 1,2$; for all $t_i \in [0,T]$, $x_i \in \mathbb{R}^{d\times 1}$, $y_i \in \mathbb{R}$, $z_i, \eta_i \in \mathbb{R}^{1\times d}$ and $\xi_i, \zeta_i \in \mathbb{R}^{d\times d}$, where $L_f, L_{f^D} > 0$. Here we also used the assumption of Hölder continuity established by $(\mathbf{A}_3^{f,g})$.

Given the usual time partition, it is clear that the discrete approximations (2.19) are deterministic functions of X_n^π and thereupon we put $Y_n^\pi =: y^\pi(t_n, X_n^\pi) =: y_n^\pi(X_n^\pi)$, $Z_n^\pi =: z^\pi(t_n, X_n^\pi) =: z_n^\pi(X_n^\pi)$, $\Gamma_n^\pi =: \gamma^\pi(t_n, X_n^\pi) =: \gamma_n^\pi(X_n^\pi)$. In light of (2.18), we use the approximations

$$D_n Y_{n+1}^\pi = Z_{n+1}^\pi \sigma^{-1}(t_{n+1}, X_{n+1}^\pi) D_n X_{n+1}^\pi, \qquad D_n Z_n^\pi = \Gamma_n^\pi D_n X_n^\pi. \tag{2.22}$$

We introduce the short-hand notations $\Delta X_n^\pi := X_{t_n} - X_n^\pi$, $\Delta Y_n^\pi = Y_{t_n} - Y_n^\pi, \Delta Z_n^\pi = Z_{t_n} - Z_n^\pi$, $\Delta D_n X_{n+1}^\pi := D_{t_n} X_{t_{n+1}} - D_n X_{n+1}^\pi$, $\Delta D_n Y_{n+1}^\pi := D_{t_n} Y_{t_{n+1}} - D_n Y_{n+1}^\pi$ and $\Delta\Gamma_n^\pi := \Gamma_{t_n} - \Gamma_n^\pi$. Under the conditions of assumption 2.4.1, provided by lemma 2.2.2 and theorem 2.2.1, we have that the processes (X, Y, Z, DX, DY) are all mean-squared continuous in time, i.e. there exists a general constant C such that for all $s, t, r \in [0,T]$

$$\begin{aligned}
&\mathbb{E}\left[|X_t - X_r|^2\right] \leq C|t-r|, \quad \mathbb{E}\left[|Y_t - Y_r|^2\right] \leq C|t-r|, \quad \mathbb{E}\left[|Z_t - Z_r|^2\right] \leq C|t-r|,\\
&\mathbb{E}\left[|D_s Y_t - D_s Y_r|^2\right] \leq C|t-r|, \quad \mathbb{E}\left[|D_s X_t - D_s X_r|^2\right] \leq C|t-r|, \quad \forall r, t \geq s.
\end{aligned} \tag{2.23}$$

Finally, we use

$$\overline{DZ}_n^{n+1} := \frac{1}{\Delta t_n}\mathbb{E}_n\left[\int_{t_n}^{t_{n+1}} D_{t_n} Z_r \mathrm{d}r\right] \tag{2.24}$$

for the $\mathbb{L}^2$-projection of the corresponding Malliavin derivative with respect to the $\mathscr{F}_{t_n}$ σ-algebra, with which we can define the $\mathbb{L}^2(\mathbb{R}^{d\times d})$-regularity of DZ as follows

$$\varepsilon^{DZ}(|\pi|) := \sum_{n=0}^{N-1} \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right]. \tag{2.25}$$

Under the condition of constant diffusion coefficients in assumption 2.4.1, we have that $D_{t_n} Z_r = D_{t_m} Z_r = \Gamma_r \sigma$ for any $t_n, t_m < r$. Thereafter, exploiting the fact that due to assumption 2.4.1 the terminal condition of the Malliavin BSDE (2.9d) is also Lipschitz continuous, one can apply [160, Thm.3.1] and get

$$\limsup_{|\pi|\to 0} \frac{1}{|\pi|} \varepsilon^{DZ}(|\pi|) < \infty. \tag{2.26}$$

2.4.1. DISCRETE-TIME APPROXIMATION ERROR

The main goal of this section is to give an upper bound for the discrete time approximation errors defined by

$$\mathscr{E}^{\pi}(|\pi|) := \max_{0\le n\le N} \mathbb{E}\left[\left|\Delta Y_n^{\pi}\right|^2\right] + \max_{0\le n\le N} \mathbb{E}\left[\left|\Delta Z_n^{\pi}\right|^2\right] + \mathbb{E}\left[\sum_{n=0}^{N-1} \int_{t_n}^{t_{n+1}} \left|\left(\Gamma_r - \Gamma_n^{\pi}\right)\sigma\right|^2 \mathrm{d}r\right]$$
$$\le C|\pi|. \tag{2.27}$$

This is established by the following theorem.

Theorem 2.4.1 (Consistency of the OSM scheme)
Under assumption 2.4.1, the scheme defined by (2.19) for any $\vartheta_y \in [0,1]$ has $\mathbb{L}^2$-convergence of order 1/2, i.e.

$$\limsup_{|\pi|\to 0} \frac{1}{|\pi|} \mathscr{E}^{\pi}(|\pi|) < \infty. \tag{2.28}$$

Proof. Throughout the proof C denotes a constant independent of the time partition, whose value may vary from line to line. We proceed in steps and prove estimates for each component of the discretization error.

Step 1. *Estimate for DZ.* First, we establish an estimate for the corresponding discretization error of the DZ-component with respect to the $\mathbb{L}^2$-projection $\overline{DZ}_n^{n+1}$. Let us fix $n = 0,\dots,N-1$. Comparing (2.15) with (2.24), we find

$$\Delta t_n \overline{DZ}_n^{n+1} = \mathbb{E}_n\left[\Delta W_n D_{t_n} Y_{t_{n+1}}\right] + \mathbb{E}_n\left[\Delta W_n \int_{t_n}^{t_{n+1}} f^D(r, \mathbf{X}_r, \mathbf{D}_{t_n}\mathbf{X}_r)\mathrm{d}r\right].$$

Combining this with the definition of the discrete scheme ((2.19a)) gives

$$\begin{aligned}
&\Delta t_n (\overline{DZ}_n^{n+1} - D_n Z_n^{\pi}) \\
&\quad = \mathbb{E}_n\left[\Delta W_n \left(\Delta D_n Y_{n+1}^{\pi} - \mathbb{E}_n\left[\Delta D_n Y_{n+1}^{\pi}\right]\right)\right] \\
&\qquad + \mathbb{E}_n\left[\Delta W_n \left(\int_{t_n}^{t_{n+1}} f^D(r, \mathbf{X}_r, \mathbf{D}_{t_n}\mathbf{X}_r) - f^D(t_{n+1}, \mathbf{X}_{n+1}^{\pi}, \mathbf{D}_n \mathbf{X}_{n+1,n}^{\pi})\mathrm{d}r\right)\right],
\end{aligned}$$

using the tower property of conditional expectations. In Frobenius norm, the conditional $\mathbb{L}^2(\mathbb{R}^d)$ Cauchy-Schwarz inequality subsequently implies

$$\begin{aligned}
&\Delta t_n \left| \overline{DZ}_n^{n+1} - D_n Z_n^\pi \right| \\
&\quad \le (d\Delta t_n)^{1/2} \left(\mathbb{E}_n \left[\left| \Delta D_n Y_{n+1}^\pi - \mathbb{E}_n \left[\Delta D_n Y_{n+1}^\pi \right] \right|^2 \right] \right)^{1/2} \\
&\qquad + (d\Delta t_n)^{1/2} \left(\mathbb{E}_n \left[\left| \int_{t_n}^{t_{n+1}} f^D(r, \mathbf{X}_r, \mathbf{D}_{t_n} \mathbf{X}_r) - f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1,n}^\pi) \mathrm{d}r \right|^2 \right] \right)^{1/2},
\end{aligned}$$

by the independence of Brownian increments. Hence, due to the $L^2([0,T];\mathbb{R}^d)$ Cauchy-Schwarz inequality, we gather

$$\begin{aligned}
&\Delta t_n \left| \overline{DZ}_n^{n+1} - D_n Z_n^\pi \right| \\
&\quad \le (d\Delta t_n)^{1/2} \left(\mathbb{E}_n \left[\left| \Delta D_n Y_{n+1}^\pi - \mathbb{E}_n \left[\Delta D_n Y_{n+1}^\pi \right] \right|^2 \right] \right)^{1/2} \\
&\qquad + d^{1/2} \Delta t_n \left(\mathbb{E}_n \left[\int_{t_n}^{t_{n+1}} \left| f^D(r, \mathbf{X}_r, \mathbf{D}_{t_n} \mathbf{X}_r) - f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1,n}^\pi) \right|^2 \mathrm{d}r \right] \right)^{1/2}.
\end{aligned}$$

Using the inequality $a, b \in \mathbb{R} : (a+b)^2 \le 2(a^2+b^2)$ we collect the following $\mathbb{L}^2(\mathbb{R}^{d\times d})$ upper bound

$$\begin{aligned}
&\Delta t_n \mathbb{E} \left[\left| \overline{DZ}_n^{n+1} - D_n Z_n^\pi \right|^2 \right] \\
&\quad \le 2d \left(\mathbb{E} \left[\left| \Delta D_n Y_{n+1}^\pi \right|^2 \right] - \mathbb{E} \left[\left| \mathbb{E}_n \left[\Delta D_n Y_{n+1}^\pi \right] \right|^2 \right] \right) \\
&\qquad + 2d\Delta t_n \mathbb{E} \left[\int_{t_n}^{t_{n+1}} \left| f^D(r, \mathbf{X}_r, \mathbf{D}_{t_n} \mathbf{X}_r) - f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1,n}^\pi) \right|^2 \mathrm{d}r \right].
\end{aligned} \tag{2.29}$$

According to (2.21), the uniform boundedness of $\mathbf{D}_{t_n} \mathbf{X}_r$ implies that f^D is Lipschitz continuous in all its spatial arguments and 1/2-Hölder continuous in time, with a universal constant L_{f^D}. This, combined with the mean-squared continuities of the $X, Y, Z, D_{t_n} X$ and $D_{t_n} Y$ in (2.23), implies

$$\begin{aligned}
&\Delta t_n \mathbb{E} \left[\left| \overline{DZ}_n^{n+1} - D_n Z_n^\pi \right|^2 \right] \\
&\quad \le 2d \left(\mathbb{E} \left[\left| \Delta D_n Y_{n+1}^\pi \right|^2 \right] - \mathbb{E} \left[\left| \mathbb{E}_n \left[\Delta D_n Y_{n+1}^\pi \right] \right|^2 \right] \right) \\
&\qquad + 14 d L_{f^D}^2 \Delta t_n \Big\{ C\Delta t_n^2 + 2\Delta t_n \left(\mathbb{E} \left[\left| \Delta X_{n+1}^\pi \right|^2 \right] + \mathbb{E} \left[\left| \Delta Y_{n+1}^\pi \right|^2 \right] + \mathbb{E} \left[\left| \Delta Z_{n+1}^\pi \right|^2 \right] \right) \\
&\qquad\qquad + 2\Delta t_n \left(\mathbb{E} \left[\left| \Delta D_n X_{n+1}^\pi \right|^2 \right] + \mathbb{E} \left[\left| \Delta D_n Y_{n+1}^\pi \right|^2 \right] \right) \\
&\qquad\qquad + \mathbb{E} \left[\int_{t_n}^{t_{n+1}} \left| D_{t_n} Z_r - D_n Z_n^\pi \right|^2 \mathrm{d}r \right] \Big\},
\end{aligned} \tag{2.30}$$

where we again used $(a+b)^2 \leq 2(a^2+b^2)$ for $a,b \in \mathbb{R}$. By the definition of $\overline{DZ}_n^{n+1}$ in (2.24), the last term can be split as follows

$$\mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - D_n Z_n^\pi\right|^2 \mathrm{d}r\right] = \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] + \Delta t_n \mathbb{E}\left[\left|\overline{DZ}_n^{n+1} - D_n Z_n^\pi\right|^2\right]. \tag{2.31}$$

Plugging this back in (2.30) yields

$$\begin{aligned}\Delta t_n \mathbb{E}\left[\left|\overline{DZ}_n^{n+1} - D_n Z_n^\pi\right|^2\right] \leq\ & 2d\left(\mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right] - \mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right]\right) \\ & + 14dL_{f^D}^2 \Delta t_n \Big\{ C\Delta t_n^2 + 2\Delta t_n \mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] \\ & \quad + 2\Delta t_n \left(\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\right) \\ & \quad + 2\Delta t_n \left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right]\right) \\ & \quad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] \\ & \quad + \Delta t_n \mathbb{E}\left[\left|\overline{DZ}_n^{n+1} - D_n Z_n^\pi\right|^2\right] \Big\}.\end{aligned}$$

For sufficiently small time steps satisfying $14dL_{f^D}^2 \Delta t_n \leq 1/2$, we can therefore gather the estimate

$$\begin{aligned}&\Delta t_n \mathbb{E}\left[\left|\overline{DZ}_n^{n+1} - D_n Z_n^\pi\right|^2\right] \\ &\quad \leq 4d\left\{\mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right] - \mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right]\right\} \\ &\qquad + 28dL_{f^D}^2 \Delta t_n \Big\{ C\Delta t_n^2 + 2\Delta t_n \mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] \\ &\qquad\quad + 2\Delta t_n \left(\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\right) \\ &\qquad\quad + 2\Delta t_n \left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right]\right) \\ &\qquad\quad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] \Big\}.\end{aligned} \tag{2.32}$$

Step 2. *Estimate for Z*. With the above result in hand, we give an estimate for the control process. Under assumption 2.4.1, provided by theorem 2.2.1, we identify the control process Z by its continuous modification given by DY and establish pointwise estimates. Indeed, from (2.15) and the definition of the discrete scheme in (2.19b), it follows

$$\begin{aligned}\Delta Z_n^\pi =\ & \mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right] \\ & + \mathbb{E}_n\left[\int_{t_n}^{t_{n+1}} f^D(r, \mathbf{X}_r, \mathbf{D}_{t_n}\mathbf{X}_r) - f^D(t_{n+1}, \mathbf{X}_{n+1}^\pi, \mathbf{D}_n \mathbf{X}_{n+1,n}^\pi)\mathrm{d}r\right].\end{aligned} \tag{2.33}$$

Applying the Young-inequality of the form $(a+b)^2 \leq (1+\rho\Delta t_n)a^2 + (1+\frac{1}{\rho\Delta t_n})b^2$ with any $\rho > 0$; using the Jensen- and $L^2([0,T];\mathbb{R}^d)$ Cauchy-Schwarz inequalities gives

$$\begin{aligned}\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right] &\leq (1+\rho\Delta t_n)\mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right] \\ &\quad + \frac{1}{\rho}(1+\rho\Delta t_n)\mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r) - f^D(t_{n+1},\mathbf{X}_{n+1}^\pi,\mathbf{D}_n\mathbf{X}_{n+1,n}^\pi)\right|^2 \mathrm{d}r\right].\end{aligned} \tag{2.34}$$

Exploiting the Lipschitz- and Hölder continuity of f^D in (2.21) and using the mean-squared continuities of $X, Y, Z, D_{t_n}X$ and $D_{t_n}Y$ in (2.23), we subsequently gather

$$\begin{aligned}&\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right] \\ &\quad\leq (1+\rho\Delta t_n)\mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right] \\ &\qquad + \frac{7L_{f^D}^2}{\rho}(1+\rho\Delta t_n)\Big\{C\Delta t_n^2 \\ &\qquad\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\right) \\ &\qquad\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right]\right) \\ &\qquad\qquad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n}Z_r - D_n Z_n^\pi\right|^2 \mathrm{d}r\right]\Big\}.\end{aligned} \tag{2.35}$$

Splitting the last term according to (2.31), substituting the upper bound (2.32) and choosing $\rho^* := 28dL_{f^D}^2$ then yields

$$\begin{aligned}&\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right] \\ &\quad\leq (1+\rho^*\Delta t_n)\mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right] \\ &\qquad + \frac{1+\rho^*\Delta t_n}{2}\Big\{C\Delta t_n^2 \\ &\qquad\qquad + (1+28dL_{f^D}^2\Delta t_n)\Delta t_n\Big(\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] \\ &\qquad\qquad\qquad + \mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\Big) \\ &\qquad\qquad + (1+28dL_{f^D}^2\Delta t_n)\Delta t_n\left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right]\right) \\ &\qquad\qquad + \frac{1+28dL_{f^D}^2\Delta t_n}{2}\mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right]\Big\},\end{aligned} \tag{2.36}$$

for any sufficiently small $\Delta t_n < 1$. At this point, we can make use of the fact that due to $(\mathbf{A}_1^{\mu,\sigma})$ in assumption 2.4.1 $X_n^\pi = \sigma W_{t_n} = X_{t_n}$ and $D_n X_{n+1}^\pi = \sigma \equiv D_{t_n}X_{t_{n+1}}$, which in particular implies $X_{t_n} - X_n^\pi \equiv 0, D_{t_n}X_{t_{n+1}} - D_n X_{n+1}^\pi \equiv 0$ and

$$\Delta D_n Y_{n+1}^\pi = \Delta Z_{n+1}^\pi, \qquad D_{t_n}Z_{t_n} - D_n Z_n^\pi = \Delta\Gamma_n^\pi\sigma, \tag{2.37}$$

in light of (2.22). Plugging these estimates back in (2.36) subsequently gives

$$\begin{aligned}\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right] &\leq (1+C_z\Delta t_n)\mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right] \\ &\quad + C_z\left\{\Delta t_n^2+\Delta t_n\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right]+\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r-\overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right]\right\}.\end{aligned} \tag{2.38}$$

Step 3. *Estimate for Y.* Given f's Lipschitz continuity in (x,y,z) and 1/2-Hölder continuity in t by (2.21), the mean-squared continuities of X,Y and Z in (2.23); through subsequent applications of the Young-, Jensen- and Cauchy-Schwarz inequalities analogously to the previous steps, we derive the following inequality from the dynamics of Y in (2.9b) and the discrete scheme in (2.19c)

$$\begin{aligned}\mathbb{E}\left[\left|\Delta Y_n^\pi\right|^2\right] &\leq (1+\beta\Delta t_n)\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] \\ &\quad + \frac{8L_f^2}{\beta}(1+\beta\Delta t_n)\Big\{C\Delta t_n^2+2\vartheta_y^2\Delta t_n\left(\mathbb{E}\left[\left|\Delta Y_n^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right]\right) \\ &\qquad + 2(1-\vartheta_y)^2\Delta t_n\left(\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\right)\Big\},\end{aligned} \tag{2.39}$$

with any $\beta>0$.

Step 4. *Combined estimate for Y and Z.* Combining the estimates in (2.38) and (2.39) gives

$$\begin{aligned}&\left(1-\frac{16L_f^2(1+\beta)\vartheta_y^2}{\beta}\Delta t_n\right)\left(\mathbb{E}\left[\left|\Delta Y_n^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right]\right) \\ &\qquad \leq (1+C_y\Delta t_n)\left(\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\right) \\ &\qquad\quad + C\left\{\Delta t_n^2+\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r-\overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right]\right\},\end{aligned}$$

with $C_y=\beta+\frac{16L_f^2(1+\beta)}{\beta}(1-\vartheta_y)^2+C_z$. Then, for any given $\beta>0$ and sufficiently small time step admitting to $\frac{16L_f^2(1+\beta)\vartheta_y^2}{\beta}\Delta t_n<1$, we derive

$$\begin{aligned}&\mathbb{E}\left[\left|\Delta Y_n^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right] \leq (1+C\Delta t_n)\left(\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right]\right) \\ &\qquad + C\left\{\Delta t_n^2+\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r-\overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right]\right\}.\end{aligned}$$

Thereupon, the discrete Grönwall lemma implies that

$$\begin{aligned}&\max_{0\leq n\leq N}\mathbb{E}\left[\left|\Delta Y_n^\pi\right|^2\right]+\max_{0\leq n\leq N}\mathbb{E}\left[\left|\Delta Z_n^\pi\right|^2\right] \\ &\qquad \leq C\Big\{\mathbb{E}\left[\left|g(X_T)-g(X_N^\pi)\right|^2\right] \\ &\qquad\quad + \mathbb{E}\left[\left|\nabla_x g(X_T)\sigma(t_N,X_T)-\nabla_x g(X_{t_N}^\pi)\sigma(t_N,X_N^\pi)\right|^2\right] \\ &\qquad\quad + \varepsilon^{DZ}(|\pi|)+|\pi|\Big\},\end{aligned} \tag{2.40}$$

where we also used the definition in (2.25). The proclaimed estimate for the (Y,Z) part then follows from the observation that under assumption 2.4.1 the terminal conditions of both the BSDE in (2.9b) and the Malliavin BSDE in (2.9d) are analytically observed; and the fact that, according to (2.26), $\varepsilon^{DZ}(|\pi|)$ is also $\mathcal{O}(|\pi|)$.

Step 5. *Final estimate for* Γ. It remains to show the consistency of the Γ estimate. From (2.29) and (2.31), we get

$$\begin{aligned}
&\mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - D_n Z_n^\pi\right|^2 \mathrm{d}r\right] \\
&\qquad \le \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] + 2d\left\{\mathbb{E}\left[\left|\Delta D_n Y_{n+1}^\pi\right|^2\right] - \mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right]\right\} \\
&\qquad\quad + 2d\Delta t_n \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r) - f^D(t_{n+1},\mathbf{X}_{n+1}^\pi,\mathbf{D}_n\mathbf{X}_{n+1,n}^\pi)\right|^2 \mathrm{d}r\right].
\end{aligned}$$

Summation from $n = 0,\dots,N-1$ thus gives

$$\begin{aligned}
&\mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - D_n Z_n^\pi\right|^2 \mathrm{d}r\right] \\
&\qquad \le \mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] + 2d\mathbb{E}\left[\left|\Delta D_{N-1} Y_N^\pi\right|^2\right] \\
&\qquad\quad + 2d\sum_{n=1}^{N-1}\left\{\mathbb{E}\left[\left|\Delta D_{n-1} Y_n^\pi\right|^2\right] - \mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right]\right\} \\
&\qquad\quad + 2d\sum_{n=0}^{N-1}\Delta t_n \mathbb{E}\Big[\int_{t_n}^{t_{n+1}} \left|f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r) - f^D(t_{n+1},\mathbf{X}_{n+1}^\pi,\mathbf{D}_n\mathbf{X}_{n+1,n}^\pi)\right|^2 \mathrm{d}r\Big],
\end{aligned}$$

where we changed the summation index for the first part of the third term. Using the relations in (2.37) implied by assumption 2.4.1, we can upper bound the summation term by the estimate (2.34)

$$\begin{aligned}
&\mathbb{E}\Big[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - D_n Z_n^\pi\right|^2 \mathrm{d}r\Big] \\
&\qquad \le \mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}} \left|D_{t_n} Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] + 2d\mathbb{E}\left[\left|\Delta D_{N-1} Y_N^\pi\right|^2\right] \\
&\qquad\quad + 2d\varrho\sum_{n=1}^{N-1}\Delta t_n \mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n Y_{n+1}^\pi\right]\right|^2\right] \\
&\qquad\quad + 2d\sum_{n=0}^{N-1}(1/\varrho + 2\Delta t_n)\mathbb{E}\Big[\int_{t_n}^{t_{n+1}} \left|f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r) - f^D(t_{n+1},\mathbf{X}_{n+1}^\pi,\mathbf{D}_n\mathbf{X}_{n+1,n}^\pi)\right|^2 \mathrm{d}r\Big],
\end{aligned}$$

for any $\varrho > 0$. Similar steps as in (2.35) subsequently give

$$\begin{aligned}
&\mathbb{E}\Big[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}}\big|D_{t_n}Z_r - D_nZ_n^\pi\big|^2\mathrm{d}r\Big] \\
&\quad\le \mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right] + 2d\mathbb{E}\left[\left|\Delta D_{N-1}Y_N^\pi\right|^2\right] \\
&\qquad + 2d\varrho\sum_{n=1}^{N-1}\Delta t_n\mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_nY_{n+1}^\pi\right]\right|^2\right] \\
&\qquad + 14L_{f_D}^2 d\sum_{n=0}^{N-1}(1/\varrho + 2\Delta t_n)\Big\{C\Delta t_n^2 + 2\Delta t_n\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] + 2\Delta t_n\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad + 2\Delta t_n\mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right] + 2\Delta t_n\mathbb{E}\left[\left|\Delta D_nX_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad + 2\Delta t_n\mathbb{E}\left[\left|\Delta D_nY_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\big|D_{t_n}Z_r - D_nZ_n^\pi\big|^2\mathrm{d}r\right]\Big\}.
\end{aligned}$$

By choosing $\varrho^* = 56L_{f_D}^2 d$, we have that for any sufficiently small $|\pi|$ satisfying $28L_{f_D}^2 d|\pi| < 1/4$

$$\begin{aligned}
&\mathbb{E}\Big[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}}\big|D_{t_n}Z_r - D_nZ_n^\pi\big|^2\mathrm{d}r\Big] \\
&\quad\le 2\mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right] + 4d\mathbb{E}\left[\left|\Delta D_{N-1}Y_N^\pi\right|^2\right] \\
&\qquad + 4d\varrho^*\sum_{n=1}^{N-1}\Delta t_n\mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_nY_{n+1}^\pi\right]\right|^2\right] \\
&\qquad + \sum_{n=0}^{N-1}(1/2 + 56L_{f_D}^2 d\Delta t_n)\Big\{C\Delta t_n^2 + 2\Delta t_n\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] + 2\Delta t_n\mathbb{E}\left[\left|\Delta Y_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad + 2\Delta t_n\Big(\mathbb{E}\left[\left|\Delta Z_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad + \mathbb{E}\left[\left|\Delta D_nX_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta D_nY_{n+1}^\pi\right|^2\right]\Big)\Big\}.
\end{aligned}$$

Once again applying the relations in (2.37), Jensen's inequality, the convergence of the $\mathbb{L}^2$-regularity of DZ in (2.26) and the estimate (2.40) proven in the previous step, now shows the proclaimed convergence of the Γ estimates.

This concludes the proof.

□

The final result in (2.28) expresses that the $\mathbb{L}^2$ convergence rate of the discrete time approximations induced by (2.19) is of order $\mathcal{O}(|\pi|^{1/2})$ under the conditions imposed in assumption 2.4.1. Comparing the convergence bound of theorem 2.4.1 to that of the classical backward Euler discretization in (2.20), three observations need to be made.

First, in contrast to the backward Euler discretization, the OSM scheme admits to a bound where the Z process is controlled by the maximum error over the discrete time steps – see (2.27). This is due to the fact that under the OSM formulation, theorem 2.2.1 guarantees a continuous version of the control process bounded in the supremum norm, and thus allows for pointwise estimates. Additionally, we see that even though the hereby proposed discretization solves a *larger problem* by incorporating Γ estimates, it exhibits the same, optimal rate of convergence well-known for the classical backward Euler discretization of BSDEs in (2.20). At last, unlike in the aforementioned case, our final estimate does not include the strong discretization errors of the terminal conditions of the BSDEs (2.9b) and (2.9d). This is merely due to the fact that under assumption 2.4.1 we assumed constant diffusion coefficients, which led to the corresponding terms canceling in (2.40). Similarly, we exploited that under our conditions the Malliavin BSDE's terminal condition is Lipschitz continuous, leading to an $\mathcal{O}(|\pi|^{1/2})$ convergence of the $\mathbb{L}^2$-regularity of DZ according to (2.26). In case of irregular terminal conditions and non-analytical forward diffusions, it is expected that the corresponding terms would also contribute to the final estimate.

2.4.2. ASSUMPTIONS REVISITED

In order to conclude the discussion on the discrete time approximation errors, we elaborate on the conditions set in assumption 2.4.1. Key aspects of their relevance are highlighted and potential ways to generalize the results are pointed out in order to encourage further research.

Not surprisingly, compared to classical discretizations excluding the Malliavin components, necessarily stricter conditions need to be posed in order to ensure Malliavin differentiability of the original FBSDE system in (2.9a) – (2.9b). The differentiability requirements on the coefficients f and g in ($\mathbf{A}_1^{f,g}$)–($\mathbf{A}_2^{f,g}$) are inherently linked to the Malliavin differentiability of the FBSDE in (2.9). However, the Malliavin differentiability of the solution pair holds under significantly milder assumptions. We refer to [110] for a recent account on the subject, where it is shown that first-order continuous differentiability, with not necessarily bounded $\nabla_x g, \nabla_x f$ is sufficient.

The reason why we nonetheless decided to restrict the assumptions to second-order bounded differentiability is mostly related to lemma 2.4.1 and the Lipschitz continuity of f^D in (2.21). Although the Lipschitz continuity of $\nabla_x f, \nabla_y f, \nabla_z f$ are all guaranteed by the $C_b^{0,2,2,2}$ assumption, the same cannot be said about the Malliavin derivative arguments $\mathbf{D}_s\mathbf{X}_t$ of f^D. More precisely, in order to have Lipschitz continuity in all spatial arguments, one – on top of the boundedness of the partial derivatives of f – also needs to have the uniform boundedness of all the Malliavin derivatives (DX, DY, DZ). Due to the Malliavin chain rule estimates in (2.17), under the assumption of constant diffusion coefficients in ($\mathbf{A}_1^{\mu,\sigma}$), the uniform boundedness of the Malliavin derivatives is implied by the twice bounded differentiability of the solution of the parabolic problem in (2.2). This is guaranteed by lemma 2.4.1, requiring the conditions in ($\mathbf{A}_1^{f,g}$)–($\mathbf{A}_2^{f,g}$) to be satisfied. In case the uniform boundedness of (DY, DZ) is not readily available, one can truncate the corresponding arguments of f^D similarly to [31], and discretize the truncated Malliavin problem accordingly. Thereafter, the total discrete time approximation error can be decomposed into a truncation and discretization component, which guarantee conver-

gence for an appropriately chosen, adaptive truncation range.

Throughout the analysis, we also often relied on the assumption that the underlying forward diffusion admits to constant drift and diffusion coefficients due to ($\mathbf{A}_1^{\mu,\sigma}$). In particular, this assumption allowed us to neglect the contribution of error terms such as $\mathbb{E}\left[\left|X_{t_n} - X_n^\pi\right|^2\right]$ and $\mathbb{E}\left[\left|D_{t_n}X_{t_{n+1}} - D_n X_{n+1}^\pi\right|^2\right]$ – see, e.g., (2.37). However, it is well-known that the strong convergence of Euler-Maruyama approximations is of order 1/2 – see (2.11) –, carrying the same order of convergence as the rest of the terms in our estimates. The convergence of the Malliavin derivative $D_n X_{n+1}^\pi$ with respect to an Euler-Maruyama discretization in (2.13) is more troublesome. In fact, as highlighted by related works in the literature – see [73, Remark 5.1] –, it is difficult to guarantee the convergence of $D_s X^\pi$ over the whole time horizon. It is important to highlight that the OSM scheme in (2.19) does not require approximations of the corresponding Malliavin derivative over the whole time window but only in between adjacent time steps $D_n X_{n+1}^\pi$. This is a major relieve in terms of convergence as one can easily show that within this one time stepping (OSM) scheme, $D_n X_{n+1}^\pi$ inherits the convergence properties of the forward diffusion under mild assumptions – see appendix 2.A.

The main difficulty with respect to general forward diffusions is related to the Malliavin chain rule approximations given by (2.17). In fact, when $D_n X_{n+1}^\pi \neq D_{t_n} X_{t_{n+1}}$ one needs to deal with product terms such as

$$\begin{aligned} D_{t_n} Y_{t_{n+1}} - D_n Y_{n+1}^\pi = &\left[Z_{t_{n+1}}\sigma^{-1}(t_{n+1}, X_{t_{n+1}}) - Z_{n+1}^\pi \sigma^{-1}(t_{n+1}, X_{n+1}^\pi)\right] D_{t_n} X_{t_{n+1}} \\ &+ Z_{n+1}^\pi \sigma^{-1}(t_{n+1}, X_{n+1}^\pi)\left[D_{t_n} X_{t_{n+1}} - D_n X_{n+1}^\pi\right]. \end{aligned} \tag{2.41}$$

These pose a significant amount of difficulty when one – unlike in the case of ($\mathbf{A}_1^{\mu,\sigma}$) – does not have the uniform boundedness of σ^{-1} and $\{D_s X_t\}_{0\leq s,t\leq T}$. Additionally, in order to ensure the boundedness of the discrete estimates Z_{n+1}^π, a certain truncation procedure would be required, further complicating the analysis. Therefore, we decided to restrict the assumptions to constant diffusion coefficients and to leave the general case for future research.

Remark 2.4.1 (Non-constant drift and Girsanov's theorem)
We remark that the assumption of a constant drift coefficient is mostly a matter convenience. Indeed, with a straightforward change of measure argument via the Girsanov theorem, one can merge the corresponding non-constant drift contribution onto the driver of the BSDE and – as long as the drift itself satisfies the continuously bounded differentiable assumptions posed on $\nabla_x f$ *– the same analysis holds.*

2.5. FULLY IMPLEMENTABLE SCHEMES WITH DIFFERENTIABLE FUNCTION APPROXIMATORS AND NEURAL NETWORKS

Having established a convergence result for the discrete time approximation's error induced by (2.19), we now turn to fully-implementable schemes where the appearing conditional expectations are numerically approximated by a certain machinery. In other words, we are concerned with the following modification of the discrete scheme in (2.19)

$$\widehat{Y}_N^\pi = g(X_N^\pi), \quad \widehat{Z}_N^\pi = \nabla_x g(X_N^\pi)\sigma(t_N, X_N^\pi), \quad \widehat{\Gamma}_N^\pi = [\nabla_x(\nabla_x g\sigma)](t_N, X_N^\pi),$$

$$\check{\Gamma}_n^\pi \sigma(t_n, X_n^\pi) = D_n \check{Z}_n^\pi = \frac{1}{\Delta t_n}\mathbb{E}_n\left[\Delta W_n \left(D_n \widehat{Y}_{n+1}^\pi + \Delta t_n f^D(t_{n+1}, \widehat{\mathbf{X}}_{n+1}^\pi, \mathbf{D}_n \check{\mathbf{X}}_{n+1,n}^\pi)\right)\right], \tag{2.42a}$$

$$\widehat{\Gamma}_n^\pi \leftarrow \mathscr{P}(\check{\Gamma}_n^\pi),$$

$$\check{Z}_n^\pi = \mathbb{E}_n\left[D_n \widehat{Y}_{n+1}^\pi + \Delta t_n f^D(t_{n+1}, \widehat{\mathbf{X}}_{n+1}^\pi, \mathbf{D}_n \widehat{\mathbf{X}}_{n+1,n}^\pi)\right], \tag{2.42b}$$

$$\widehat{Z}_n^\pi \leftarrow \mathscr{P}(\check{Z}_n^\pi),$$

$$\check{Y}_n^\pi = \vartheta_y \Delta t_n f(t_n, X_n^\pi, \check{Y}_n^\pi, \widehat{Z}_n^\pi) + \mathbb{E}_n\left[\widehat{Y}_{n+1}^\pi + (1-\vartheta_y)\Delta t_n f(t_{n+1}, \widehat{\mathbf{X}}_{n+1}^\pi)\right], \tag{2.42c}$$

$$\widehat{Y}_n^\pi \leftarrow \mathscr{P}(\check{Y}_n^\pi),$$

with the notations $\widehat{\mathbf{X}}_{n+1}^\pi := \left(\widehat{X}_{n+1}^\pi, \widehat{Y}_{n+1}^\pi, \widehat{Z}_{n+1}^\pi\right)$, $\mathbf{D}_n \check{\mathbf{X}}_{n+1,n}^\pi := (D_n X_{n+1}^\pi, D_n \widehat{Y}_{n+1}^\pi, D_n \check{Z}_n^\pi)$ and $\mathbf{D}_n \widehat{\mathbf{X}}_{n+1,n}^\pi := \left(D_n X_{n+1}^\pi, D_n \widehat{Y}_{n+1}^\pi, D_n \widehat{Z}_n^\pi\right)$, where $D_n \widehat{Y}_{n+1}^\pi := \widehat{Z}_{n+1}^\pi \sigma^{-1}(t_{n+1}, X_{n+1}^\pi) D_n X_{n+1}^\pi$ and $D_n \widehat{Z}_n^\pi := \widehat{\Gamma}_n^\pi D_n X_n^\pi$ – similarly as in (2.18). The final approximations are denoted by $(\widehat{Y}_n^\pi, \widehat{Z}_n^\pi, \widehat{\Gamma}_n^\pi)$ and $\mathscr{P}$ denotes a *machinery* which, given approximations at future time steps, estimates the *true* conditional expectations $(\check{Y}_n^\pi, \check{Z}_n^\pi, \check{\Gamma}_n^\pi)$. It is worth to notice that (2.42b) is explicit, whereas (2.42a) and (2.42c) are both implicit when $\vartheta_y > 0$. Due to the Markov feature of the corresponding problem, we can write all estimates as deterministic functions of the state process $\check{Y}_n^\pi =: \check{y}_n^\pi(X_n^\pi)$, $\check{Z}_n^\pi =: \check{z}_n^\pi(X_n^\pi)$, $\check{\Gamma}_n^\pi =: \check{\gamma}_n^\pi(X_n^\pi)$ and $\widehat{Y}_n^\pi =: \widehat{y}_n^\pi(X_n^\pi)$, $\widehat{Z}_n^\pi =: \widehat{z}_n^\pi(X_n^\pi)$, $\widehat{\Gamma}_n^\pi =: \widehat{\gamma}_n^\pi(X_n^\pi)$ at each time instance.

In the literature there exist several techniques to numerically approximate conditional expectations, see, e.g., [9, 25, 24]. In what follows, we investigate two specific approaches in the context of the OSM scheme. We first give an extension to the BCOS method [144] which shall later be used as a benchmark method for one-dimensional problems. Our main approximation tool is based on a least-squares Monte Carlo formulation similar to those of the Deep BSDE methods [68, 77], where the functions parametrizing the solution triple are fully-connected, feedforward neural networks. Due to the universal approximation properties of neural networks in Sobolev spaces, this will allow us to distinguish between two variants. In the first one, the Γ process is parametrized by a matrix-valued neural network whose parameters are optimized in a stochastic gradient descent iteration. In the second, this parametrization is circumvented and, in light of (2.3), the Γ estimates are directly calculated as the Jacobian of the Z process. However, such directly linked estimates induce an additional source of error, which shall be addressed in theorem 2.5.2, where we give an error bound for the complete approximation error of the fully-implementable OSM scheme, given the cumulative regression errors of neural network regressions, similarly to the ones proven in [69, 77].

2.5.1. THE BCOS METHOD

We recall the most fundamental notions of the BCOS method [144]. In order to keep the presentation concise, for the sake of this section we restrict ourselves to the one-dimensional case. BCOS is an extension of the COS method [48] to the setting of FB-

2

SDE systems, whose main idea is to recover the probability densities of certain random variables given that their characteristic function is available. The key idea of the BCOS method can be summarized as follows. In general, for a Markov problem, conditional expectations are of the form

$$I(x) := \mathbb{E}\left[v(t_{n+1}, X^{\pi}_{n+1}) | X^{\pi}_{n} = x\right] = \int_{\mathbb{R}} v(t_{n+1}, \rho) p(\rho|x) \mathrm{d}\rho,$$

where $p(\rho|x)$ is the conditional transition density function from state (t,x) to state (t_{n+1}, ρ). Assuming that the integrand above decays in the infinite limit, one can truncate the integration range to a sufficiently wide finite domain $[a,b]$. Thereafter, the Fourier cosine expansion of the deterministic mapping $v(t_{n+1}, \cdot) : [a,b] \to \mathbb{R}$ reads as[3]

$$v(t_{n+1}, \rho) = \sum_{k=0}^{\infty}{}' \mathcal{V}_k(t_{n+1}) \cos\Big(k\pi \frac{\rho - a}{b-a}\Big),$$

where the series coefficients are given by $\mathcal{V}_k(t_{n+1}) := \frac{2}{b-a}\int_a^b v(t_{n+1}, \rho)\cos\big(k\pi\frac{\rho-a}{b-a}\big)\mathrm{d}\rho$. Plugging these estimates back in the conditional expectation, with an additional truncation of the Fourier expansion to a finite number of K coefficients, gives the approximation [48]

$$I(x) \approx \widehat{I}(x) := \sum_{k=0}^{K-1}{}' \mathcal{V}_k(t_{n+1}) \mathrm{Re}\{\Phi(k|x)\}, \tag{2.43}$$

where $\Phi(k|x) := \phi(\frac{k\pi}{b-a}|x) e^{ik\pi\frac{x-a}{b-a}}$ and $\phi(u|x)$ is the conditional characteristic function of the Markov transition. In case the underlying Markov process is an Euler-Maruyama approximation of the solution to a forward SDE, the conditional characteristic function is given by $\phi(u|x) = \exp\big(iu\mu(t_n,x)\Delta t_n - \frac{1}{2}u^2\sigma^2(t_n,x)\Delta t_n\big)$. Using an integration by parts argument – see [144, Appendix A.1] and appendix 2.B – similar results can be constructed for conditional expectations of the forms

$$\begin{aligned} J(x) &:= \mathbb{E}^x_n\left[v(t_{n+1}, X^{\pi}_{n+1})\Delta W_n\right] \qquad (2.44)\\ &\approx \widehat{J}(x) := \Delta t_n \sigma(t_n,x) \sum_{k=0}^{K-1}{}' -\frac{k\pi}{b-a}\mathcal{V}_k(t_{n+1})\mathrm{Im}\{\Phi(k|x)\},\\ K(x) &:= \mathbb{E}^x_n\left[v(t_{n+1}, X^{\pi}_{n+1})(\Delta W_n)^2\right] \qquad (2.45)\\ &\approx \widehat{K}(x) := \Delta t_n \sum_{k=0}^{K-1}{}' \mathcal{V}_k(t_{n+1})\mathrm{Re}\{\Phi(k|x)\}\\ &\qquad - \Delta t_n^2 \sigma^2(t_n,x) \sum_{k=0}^{K-1}\left(\frac{k\pi}{b-a}\right)^2 \mathcal{V}_k(t_{n+1})\mathrm{Re}\{\Phi(k|x)\}. \end{aligned}$$

Built on these approximations, the BCOS method goes as follows. One first needs to recover the coefficients of the terminal conditions either analytically or via Discrete Cosine Transforms (DCT). These coefficients are plugged into conditional expectations of

[3] We adhere to the standard notation where $\sum_{k=0}^{K-1}{}' a_k := a_0/2 + \sum_{k=1}^{K-1} a_k$, i.e. the first element is multiplied by 1/2.

the form (2.43), (2.44) and (2.45), providing estimates for the solutions at t_{N-1}. In order to make the scheme fully-implementable, one also relies on a machinery which recovers these coefficients while going to time step n, from time step $n+1$ in a backward recursive algorithm. This step can either be done by Fast Fourier Transforms (FFT) [144] when the coefficients of the SDE are constant, or with DCT when they are not [145]. When one is faced with an implicit conditional expectation ($\vartheta_y > 0$) Picard iterations are performed, which – under Lipschitz assumptions and sufficiently small time steps – converge exponentially fast to the unique fixed point solution.

In particular, the BCOS approximations for (2.42) read as follows – for a more detailed derivation, see appendix 2.C

$$
\begin{aligned}
\widehat{y}_N^\pi(x) &= g(x), \quad \widehat{z}_N^\pi(x) = \partial_x g(x)\sigma(T,x), \quad \widehat{\gamma}_N^\pi(x) = \partial_x\left(\partial_x g\sigma\right)(T,x),\\
\widehat{\gamma}_n^\pi(x)\sigma(t_n,x) &= \sum_{k=0}^{K-1}{}' \widehat{\mathscr{D}\mathscr{Z}}_k(t_{n+1})\cos\Big(k\pi\frac{x-a}{b-a}\Big),\\
\widehat{z}_n^\pi(x) &= \sigma(t_n,x)(1+\partial_x\mu(t_n,x)\Delta t_n)\sum_{k=0}^{K-1}{}' \mathscr{W}_k(t_{n+1})\,\mathrm{Re}\{\Phi(k|x)\}\\
&\quad - \sigma^2(t_n,x)\partial_x\sigma(t_n,x)\Delta t_n\sum_{k=0}^{K-1}{}' \frac{k\pi}{b-a}\mathscr{W}_k(t_{n+1})\,\mathrm{Im}\{\Phi(k|x)\}\\
&\quad + \Delta t_n\widehat{\gamma}_n^\pi(x)\sigma(t_n,x)\sum_{k=0}^{K-1}{}' \mathscr{F}_k^z(t_{n+1})\,\mathrm{Re}\{\Phi(k|x)\},\\
\widehat{y}_n^\pi(x) &= \sum_{k=0}^{K-1}{}' \widehat{\mathscr{Y}}_k(t_n)\cos\Big(k\pi\frac{x-a}{b-a}\Big),
\end{aligned}
$$

where we defined

$$
\begin{aligned}
h_{n+1}^\pi(X_{n+1}^\pi) &:= \widehat{y}_{n+1}^\pi(X_{n+1}^\pi) + (1-\vartheta_y)\Delta t_n f(t_{n+1}, X_{n+1}^\pi, \widehat{y}_{n+1}^\pi(X_{n+1}^\pi), \widehat{z}_{n+1}^\pi(X_{n+1}^\pi)),\\
w_{n+1}^\pi(X_{n+1}^\pi) &:= \left(1+\partial_y f(t_{n+1}, \widehat{\mathbf{X}}_{n+1}^\pi)\right)\widehat{z}_{n+1}^\pi(X_{n+1}^\pi)\sigma^{-1}(t_{n+1}, X_{n+1}^\pi) + \Delta t_n\partial_x f(t_n, \widehat{\mathbf{X}}_{n+1}^\pi)
\end{aligned}
\tag{2.47}
$$

for the explicit parts of the discrete approximations (2.42c) and (2.42b), respectively. The coefficients

$$
\begin{aligned}
\mathscr{W}_k(t_{n+1}) &:= \frac{2}{b-a}\int_a^b w_{n+1}^\pi(\rho)\cos\Big(k\pi\frac{\rho-a}{b-a}\Big)\mathrm{d}\rho,\\
\mathscr{H}_k(t_{n+1}) &:= \frac{2}{b-a}\int_a^b h_{n+1}^\pi(\rho)\cos\Big(k\pi\frac{\rho-a}{b-a}\Big)\mathrm{d}\rho,\\
\mathscr{F}_k^z(t_{n+1}) &:= \frac{2}{b-a}\int_a^b \partial_z f(t_{n+1},\rho)\cos\Big(k\pi\frac{\rho-a}{b-a}\Big)\mathrm{d}\rho
\end{aligned}
$$

are approximated by their DCT counterparts $\widehat{\mathscr{W}_k}(t_{n+1}), \widehat{\mathscr{H}_k}(t_{n+1}), \widehat{\mathscr{F}_k^z}(t_{n+1})$, respectively. $\widehat{\mathscr{D}\mathscr{Z}}_k(t_{n+1})$ is recovered with DCT on the approximations $\mathbb{E}_n^x\left[\Delta t_n^{-1}\Delta W_n w_{n+1}^\pi(X_{n+1}^\pi)D_n X_{n+1}^\pi\right] / \left(1-\mathbb{E}_n^x\left[\Delta W_n\partial_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)\right]\right)$. Thereafter, the BCOS formulas in (2.43), (2.44) and (2.45), together with the Euler-Maruyama estimates (2.13), imply the estimates for Γ and Z. The Z estimates are plugged into the approximation of the Y process in (2.42c). The coefficients $\widehat{\mathscr{Y}_k}(t_n)$ are recovered from the estimates

$y_n^{P,\pi}(x) = \vartheta_y \Delta t_n f(t_n, x, y_n^{P-1,\pi}(x), \widehat{z}_n^{\pi}(x)) + \mathbb{E}_n^x\left[h_{n+1}^{\pi}\right]$ after a sufficient number of P Picard iterations are taken. This completes the BCOS algorithm for the OSM scheme.

For a detailed account on the contributions of the corresponding truncation and approximation errors of the BCOS method, we refer to [144, 145, 48] and the references therein. Although the method can be extended to higher-dimensional diffusion processes, it suffers from the curse of dimensionality through the inevitable spatial discretization required in the Fourier frequency domain.

2.5.2. NEURAL NETWORKS

In recent years, neural networks have shown excellent empirical results when deployed in a regression Monte Carlo framework for BSDEs [68, 77, 53]. In what follows, we are concerned with the class of feedforward, fully-connected deep neural networks, particularly in the context of approximating high-dimensional conditional expectations. This family of functions $\Psi(\cdot|\Theta) : \mathbb{R}^{d\times 1} \to \mathbb{R}^{q\times d}$ can be described as a hierarchical sequence of compositions

$$\Psi(x|\Theta) := a_{\text{out}} \circ A_{L+1}(\cdot|\theta_{L+1}) \circ a \circ A_L(\cdot|\theta_L) \circ a \circ \cdots \circ a \circ A_1(\cdot|\theta_1) \circ x.$$

The affine transformations $A_l, l = 1,\dots,L$ are called *hidden layers* and are of the form $A_l(y|\theta_l := (W_{l-1}^l, b_l)) := W_{l-1}^l y + b_l$, with $W_{l-1}^l \in \mathbb{R}^{S_l\times S_{l-1}}$ being a matrix of *weights* and $b_l \in \mathbb{R}^{S_l\times 1}$, $S_{l-1}, S_l \in \mathbb{N}$ the *biases*. Furthermore, $a : \mathbb{R} \to \mathbb{R}$ describes a non-linear *activation* function, which is applied element-wise on the output of each affine transformation. The size S_l denotes how many *neurons* are contained in the given layer. The *output layer* is defined by $A_{L+1}(y|\theta_{L+1} := (W_L^{L+1}, b_{L+1})) := W_L^{L+1} y + b_{L+1}$ with $W_L^{L+1} \in \mathbb{R}^{q\times d\times S_L}, b_{L+1} \in \mathbb{R}^{q\times d}$. The complete parameter space of such an architecture is therefore given by $\Theta := (\theta_1,\dots,\theta_{L+1}) \in \mathbb{R}^{q\times d\times(S_L+1)+\sum_{l=1}^{L} S_{l-1}\times S_l+S_l}$. Widely common choices for the non-linearity include: Rectified Linear Units (ReLU), sigmoid and the hyperbolic tangent activations. The optimal parameter space Θ^* is usually approximated by first formulating a *loss function* which measures an abstract distance from the desired behavior, and then iteratively minimizing this loss through a *stochastic gradient descent* (SGD) type algorithm. For more details, we refer to [65].

The use of deep learning is often motivated by the so-called *Universal Approximation Theorems (UAT)* which establish that neural networks can approximate a wide class of functions with arbitrary accuracy. The first version of the UAT property was proven by Cybenko in [37]. However, as in the applications of this chapter derivative approximations play an important role, we present the following extension of Hornik et al. [72], which extends the UAT property to *Sobolev spaces*. In what follows, we use the common notations for $W^{k,p}(U) := \{f \in L^p(U) : \|f\|_{W^{k,p}} := (\sum_{|\alpha|\le k} \int_U |D^\alpha f|^p \mathrm{d}\lambda)^{1/p} < \infty\}$ for Sobolev spaces, where α denotes a multi-index, D^α is the differentiation operator in the weak sense and λ is the Lebesgue measure. In particular, we use $H^k(U) := W^{k,2}(U)$. Then the UAT in Sobolev spaces can be stated as follows – for a proof see [72, Corollary 3.6].

Theorem 2.5.1 (Universal Approximation Theorem in Sobolev Spaces, [72])
Let $a : \mathbb{R} \to \mathbb{R}$ be an ℓ-finite activation function, i.e. $a \in C^l(\mathbb{R})$ and $\int_{\mathbb{R}} |D^\ell a| < \infty$. Let $U \subseteq \mathbb{R}^{d\times 1}$ be a compact subset. Denote the class of single hidden layer neural networks by $\Sigma(a) := \{\psi : \mathbb{R}^{d\times 1} \to \mathbb{R}^{1\times q} : \psi(x|\Theta = (W_0^1, b_1, W_1^2, b_2)) = W_1^2 a(W_0^1 x + b_1) + b_2, W_0^1 \in$

$\mathbb{R}^{S_1\times d}, b_1 \in \mathbb{R}^{S_1\times 1}, W_1^2 \in \mathbb{R}^{1\times q\times S_1}, b_2 \in \mathbb{R}^{1\times q}, S_1 \in \mathbb{N}\}$. *Then* $\Sigma(a)$ *is dense in* $W^{m,p}(U)$ *for each* $0 \le m \le \ell$, *i.e. for any* $\epsilon > 0$ *and* $f \in W^{m,p}(U)$ *there exists a* $\psi \in \Sigma(a)$ *such that* $\|\psi - f\|_{W^{m,p}} < \epsilon$.

In particular, we have that for any $\ell = 1$*-finite activation* a, $f \in H^1(U)$ *and* $\epsilon > 0$, *there exists a* $\psi \in \Sigma(a)$ *such that*

$$\int_U |\psi - f|^2 \mathrm{d}\lambda + \int_U |\nabla_x \psi - Df|^2 \mathrm{d}\lambda < \epsilon.$$

The main implication of the UAT property is that given a compact domain on $\mathbb{R}^d$ and an appropriate activation function, one can approximate any Sobolev function by shallow neural networks[4] with arbitrary accuracy. It is worth to highlight that in the context of a regression Monte Carlo application, this does not establish an implementable *regression bias* due to the lack of bounds on the width of the hidden layer. We remark that the above version is not a state of the art result and refer to [134] for a classical survey on the subject.

Layer Normalization. Normalization is a standard tool to enhance the convergence of stochastic gradient descent like algorithms [65]. In standard examples [68], this is usually done by a so-called *batch normalization* technique. However, as we shall see, in our setting batch normalization is computationally intensive as it ruins batch independence and implies quadratic dependence of the Jacobian tensor on the chosen batch size. Hence, we instead deploy *layer normalization* [8], where normalization takes place across the output activations of a given hidden layer. Therefore, the final network architecture considered in section 2.6 is described by the sequence of compositions

$$\Psi(x|\bar{\Theta}) := a^{\text{out}} \circ A^{L+1}(\cdot|\theta^{L+1}) \circ a \circ A^L(\cdot|\theta^L) \circ \bar{n} \circ a \circ \cdots \circ \bar{n} \circ A^1(\cdot|\theta^1) \circ x, \tag{2.48}$$

with $\bar{n}(\cdot|\beta_l)$ and $\bar{\Theta} := (\Theta, \beta_1, \ldots, \beta_{L-1})$, where β_l denotes the l'th normalization layer's parameters – see [8].

2.5.3. A DEEP BSDE APPROACH

In what follows, we formulate a Deep BSDE approach similar to [77], which scales well in high-dimensional settings and tackles the fully-implementable scheme (2.42) in a neural network least-squares Monte Carlo framework. The main difference between our approach and that of [77] is that, unlike in the discretization problem (2.12), we solve the d-dimensional linear BSDE of the Malliavin derivatives in (2.9d) – on top of the scalar BSDE (2.9b). We separate the solutions of these two BSDEs and perform two distinct neural network regressions at each time step. We distinguish between two approaches. The first involves an additional layer of parametrization in which the matrix-valued Γ process is approximated by an $\mathbb{R}^{d\times d}$-valued neural network. In the second, we take advantage of neural networks being dense function approximators in Sobolev spaces provided by theorem 2.5.1, circumvent parametrizing the Γ process and instead obtain it as the direct derivative of the Z process via automatic differentiation – in a way very similar to the second scheme (DBDP2) of [77]. In doing so, we require a so-called Jacobian training where the loss is dependent of the derivative of the neural network involved.

[4] It is clear that the above statement generalizes to deep neural networks containing multiple hidden layers.

2

In order to motivate the merged problem formulation, notice that by assumption 2.4.1 on the coefficients of the BSDE, the arguments of the conditional expectations in (2.42) are all $\mathbb{L}^2$-integrable random variables. Consequently, (2.42b), combined with the martingale representation theorem, implies the existence of a unique random process $D_n\widetilde{Z}_r$ such that

$$D_n\widehat{Y}^\pi_{n+1} + \Delta t_n f^D(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}, \mathbf{D}_n\breve{\mathbf{X}}^\pi_{n+1,n}) = \breve{Z}^\pi_n + \int_{t_n}^{t_{n+1}} \left((D_n\widetilde{Z}_r)^T \mathrm{d}W_r\right)^T. \tag{2.49}$$

Itô's isometry implies that the $\mathbb{L}^2$-projection of $D_n\widetilde{Z}_r$ coincides with $D_n\breve{Z}^\pi_n$ in (2.42a)

$$D_n\breve{Z}^\pi_n = \frac{1}{\Delta t_n}\mathbb{E}_n\left[\int_{t_n}^{t_{n+1}} D_n\widetilde{Z}_r \mathrm{d}r\right]. \tag{2.50}$$

Thereupon, $\breve{Z}^\pi_n + ((D_n\breve{Z}^\pi_n)^T\Delta W_n)^T$ is not just the best $\mathbb{L}^2$-projection of the left-hand side of (2.49) but also of the arguments of the conditional expectations on the right-hand side of (2.42a). Hence, it simultaneously solves the discretization problems (2.42a) and (2.42b).

Motivated by these observations, the Deep BSDE approach then goes as follows – the complete algorithm is collected in algorithm 1. We set $\widehat{Y}^\pi_N = g(X^\pi_N)$, $\widehat{Z}^\pi_N = \nabla_x g(X^\pi_N)\sigma(T, X^\pi_N)$ and $\widehat{\Gamma}^\pi_N = \nabla_x(\nabla_x g\sigma)(T, X^\pi_N)$. Thereafter, each time step's Y, Z and Γ are parametrized by three independent fully-connected feedforward neural networks $\varphi(\cdot|\theta^y) : \mathbb{R}^{d\times 1} \to \mathbb{R}$, $\psi(\cdot|\theta^z) : \mathbb{R}^{d\times 1} \to \mathbb{R}^{1\times d}$ and $\chi(\cdot|\theta^\gamma) : \mathbb{R}^{d\times 1} \to \mathbb{R}^{d\times d}$ of the type (2.48). The parameter sets $(\theta^z, \theta^\gamma)$ and θ^y are trained in two separate regressions. First, in light of (2.49), we define the loss function of the regression problem corresponding to (2.42a)–(2.42b) by

$$\begin{aligned}\mathscr{L}^{z,\gamma}_n(\theta^z, \theta^\gamma) := \mathbb{E}\Big[\Big|&(1 + \Delta t_n \nabla_y f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}))D_n\widehat{Y}^\pi_{n+1} \\ &+ \Delta t_n \nabla_x f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}) D_n X^\pi_{n+1} \\ &- \psi(X^\pi_n|\theta^z) + \Delta t_n \nabla_z f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1})\chi(X^\pi_n|\theta^\gamma)\sigma(t_n, X^\pi_n) \\ &- ((\chi(X^\pi_n|\theta^\gamma)\sigma(t_n, X^\pi_n))^T\Delta W_n)^T\Big|^2\Big],\end{aligned} \tag{2.51}$$

where we approximate $D_n Z^\pi_n$ by $\chi(X^\pi_n|\theta^\gamma)D_n X^\pi_n$, according to the Malliavin chain rule. We gather an approximation of the minimal parameter set $(\theta^{z,*}_n, \theta^{\gamma,*}_n) \in \arg\min_{(\theta^z,\theta^\gamma)} \mathscr{L}^z_n(\theta^z, \theta^\gamma)$ after minimizing an empirically observed version of the loss function through a stochastic gradient descent optimization, resulting in approximations $\widehat{\theta}^z_n$ and $\widehat{\theta}^\gamma_n$ – see alg. 1. The final approximations are given by $\widehat{Z}^\pi_n := \psi(X^\pi_n|\widehat{\theta}^z_n)$ and $\widehat{\Gamma}^\pi_n := \chi(X^\pi_n|\widehat{\theta}^\gamma_n)$.

Similarly to the second scheme in [77], an alternative formulation can be given which avoids parametrizing the Γ process, and instead approximates it as the direct derivative of the Z process provided by the Malliavin chain rule lemma 2.2.1. Eventually, this implies the direct connection $\chi(X^\pi_n|\theta^\gamma) \equiv \nabla_x\psi(X^\pi_n|\theta^z)$, with which the corresponding loss

function becomes

$$\begin{aligned}\mathscr{L}_n^{z,\nabla z}(\theta^z) := \mathbb{E}\Big[\Big|&(1+\Delta t_n \nabla_y f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi))D_n\widehat{Y}_{n+1}^\pi \\ &+\Delta t_n \nabla_x f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)D_n X_{n+1}^\pi \\ &-\psi(X_n^\pi|\theta^z)+\Delta t_n \nabla_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)\nabla_x\psi(X_n^\pi|\theta^z)\sigma(t_n,X_n^\pi) \\ &-((\nabla_x\psi(X_n^\pi|\theta^z)\sigma(t_n,X_n^\pi))^T\Delta W_n)^T\Big|^2\Big],\end{aligned} \tag{2.52}$$

where we exploited the relation between the Γ and Z processes, provided by the Malliavin chain rule, and set $D_n\widehat{Z}_n^\pi = \nabla_x\widehat{z}_n^\pi(X_n^\pi)D_nX_n^\pi$. The SGD approximation of the optimal parameter space $\theta_n^{z,*} \in \operatorname{argmin}_{\theta^z}\mathscr{L}_n^{z,\nabla z}(\theta^z)$ is denoted by $\widehat{\theta}_n^z$, and the final approximations are of the form $\widehat{Z}_n^\pi := \psi(X_n^\pi|\widehat{\theta}_n^z)$ and $\widehat{\Gamma}_n^\pi := \nabla_x\psi(X_n^\pi|\widehat{\theta}_n^z)$.

Subsequently, these approximations are plugged into the regression problem of (2.42c). This step, apart from the additional theta-discretization, is identical to that of [77] and the loss function reads as

$$\begin{aligned}\mathscr{L}_n^y(\theta^y) := \mathbb{E}\Big[\Big|&\widehat{Y}_{n+1}^\pi+(1-\vartheta_y)\Delta t_n f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi) \\ &-\varphi(X_n^\pi|\theta^y)+\vartheta_y\Delta t_n f(t_n,X_n^\pi,\varphi(X_n^\pi|\theta^y),\widehat{Z}_n^\pi)-\widehat{Z}_n^\pi\Delta W_n\Big|^2\Big].\end{aligned} \tag{2.53}$$

The stochastic gradient descent approximation of the optimal parameter space $\theta_n^{y,*} \in \operatorname{argmin}_{\theta^y}\mathscr{L}_n^y(\theta^y)$ is denoted by $\widehat{\theta}_n^y$ and the final approximation is given by $\widehat{Y}_n^\pi := \varphi(X_n^\pi|\widehat{\theta}_n^y)$. At last, motivated by the continuity of the processes $\{(Y_t,Z_t)\}_{0\le t\le T}$ in the Malliavin framework, we initialize the parameters of the next time step's parametrizations according to

$$\theta^z = \widehat{\theta}_n^z, \qquad \theta^\gamma = \widehat{\theta}_n^\gamma, \qquad \theta^y = \widehat{\theta}_n^y. \tag{2.54}$$

Such a *transfer learning* trick guarantees a good initialization of the SGD iterations for $\check{Y}_{n-1}^\pi, \check{Z}_{n-1}^\pi, \check{\Gamma}_{n-1}^\pi$, simplifying the learning problem and reducing the number of iteration steps required for convergence. For an empirical assessment on the efficiency of this transfer learning trick, we refer to [32, Sec.5.3].

Dimensionality, linearity and vector-Jacobian products. The main reason why no numerical scheme has been proposed to solve the Malliavin BSDE in (2.9d) is related to dimensionality. Since the Γ process is an $\mathbb{R}^{d\times d}$-valued process, its computational complexity in a least-squares Monte Carlo method has a quadratic dependence on the number of dimensions d. Indeed, a least-squares Monte Carlo approach for the BSDE (2.1b) essentially comes down to the approximation of $d+1$-many conditional expectations. If, in addition, one would also like to solve the Malliavin BSDE (2.9d) this leads to d^2 additional conditional expectations to be approximated, induced by the Γ process. This observation justifies the use of deep neural network parametrizations which enable good scalability in high-dimensions. Moreover, notice that the training of the loss function (2.52) through an SGD optimization requires differentiating the loss with respect to the parameters θ^z in each step. With the loss already depending on the Jacobian of the mapping $\psi(\cdot|\theta^z)$, this in particular implies that in each SGD step one needs to

calculate the Hessian of a vector-valued mapping ψ with respect to the parameters θ^z. Consequently, for high-dimensional problems the training of (2.52) becomes excessively intensive from a computational point of view. Nonetheless, what makes the Deep BSDE approach corresponding to (2.52) efficiently implementable is the linearity of the Malliavin BSDE (2.9d). In fact, due to linearity, one can circumvent explicitly calculating the Jacobian matrix of Z as it suffices to calculate the vector-Jacobian product

$$\nabla_z f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}) \nabla_x \psi(X^\pi_n|\theta^z) = \nabla_x \left\langle v \middle| \psi(X^\pi_n|\theta^z) \right\rangle, \qquad v := \nabla_z f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}),$$

which boils down to computing a gradient instead. This mitigates the computational costs of minimizing the automatic differentiated loss function in (2.52) in an SGD iteration.

2.5.4. REGRESSION ERROR ANALYSIS

In order to conclude the discussion on fully-implementable schemes for (2.42), we extend the discretization error results established by theorem 2.4.1, so that it incorporates the approximation errors of the arising conditional expectations. Even though we focus on the Deep BSDE approach, our arguments naturally extend to the BCOS estimates. We consider shallow neural networks, with S_1-many hidden neurons and a hyperbolic tangent activation. While distinguishing between the parametrized and automatic differentiated Γ variants – see (2.51) and (2.52), respectively –, we rely on the following subclass of shallow neural networks introduced in theorem 2.5.1

$$\Sigma_{C^2_b}(\tanh) := \left\{ \psi(x|\theta^z(S_1)) \in \Sigma(\tanh) : \sum_{i=1}^{d} \sum_{j=1}^{S_1} \left| [W^2_1(S_1)]_{1,i,j} \right| + \left| [W^1_0(S_1)]_{j,i} \right| \leq \Upsilon(S_1) \right\},$$

for some dominating sequence $\Upsilon : \mathbb{N}_+ \to \mathbb{R}$. Then, due to the boundedness of the hyperbolic tangent function and its first two derivatives, the following upper bounds are in place for any $\psi(\cdot|\theta^z) \in \Sigma_{C^2_b}(\tanh)$

$$\begin{aligned} &\sup_{x \in \mathbb{R}^{d\times 1}} \left| \psi(x|\theta^z) \right| \leq \Upsilon(S_1), \qquad \sup_{x \in \mathbb{R}^{d\times 1}} \left| \nabla_x \psi(x|\theta^z) \right| \leq \Upsilon^2(S_1), \\ &\sup_{x \in \mathbb{R}^{d\times 1}} \left| \mathrm{Hess}_x \psi(x|\theta^z) \right| \leq \Upsilon^3(S_1). \end{aligned} \tag{2.55}$$

In light of theorem 2.5.1, the hyperbolic tangent function is $\ell = 1$-finite. Subsequently the family of shallow networks of the form (2.48) is dense in $H^1(U)$ for any compact subset $U \subset \mathbb{R}^{d\times 1}$.

The final approximations are denoted by $\widehat{Y}^\pi_n := \widehat{y}^\pi_n(X^\pi_n) =: \varphi(X^\pi_n|\widehat{\theta}^y_n)$, $\widehat{Z}^\pi_n := \widehat{z}^\pi_n(X^\pi_n) =: \psi(X^\pi_n|\widehat{\theta}^z_n)$ and $\widehat{\Gamma}^\pi_n := \widehat{\gamma}^\pi_n(X^\pi_n) =: \chi(X^\pi_n|\widehat{\theta}^\gamma_n)$. We introduce the notations $\Delta\widecheck{Y}^\pi_n := Y_{t_n} - \widecheck{Y}^\pi_n$, $\Delta\widecheck{Z}^\pi_n := Z_{t_n} - \widecheck{Z}^\pi_n$, $\Delta\widecheck{\Gamma}^\pi_n := \Gamma_{t_n} - \widecheck{\Gamma}^\pi_n$, and $\Delta\widehat{Y}^\pi_n := Y_{t_n} - \widehat{Y}^\pi_n$, $\Delta\widehat{Z}^\pi_n := Z_{t_n} - \widehat{Z}^\pi_n$, $\Delta\widehat{\Gamma}^\pi_n := \Gamma_{t_n} - \widehat{\Gamma}^\pi_n$. In light of the UAT property in theorem 2.5.1, we define the *regression biases*

$$\begin{aligned} \epsilon^y_n &:= \inf_{\theta^y} \mathbb{E}\left[\left| \widecheck{y}^\pi_n(X^\pi_n) - \varphi(X^\pi_n|\theta^y) \right|^2 \right], \\ \epsilon^z_n &:= \inf_{\theta^z} \mathbb{E}\left[\left| \widecheck{z}^\pi_n(X^\pi_n) - \psi(X^\pi_n|\theta^z) \right|^2 \right], \epsilon^\gamma_n := \inf_{\theta^\gamma} \mathbb{E}\left[\left| \left(\widecheck{\gamma}^\pi_n(X^\pi_n) - \chi(X^\pi_n|\theta^\gamma)\right) \sigma(t_n, X^\pi_n) \right|^2 \right], \\ \epsilon^{z,\nabla z}_n &:= \inf_{\theta^z} \mathbb{E}\left[\left| \widecheck{z}^\pi_n(X^\pi_n) - \psi(X^\pi_n|\theta^z) \right|^2 + \Delta t_n \left| \left(\nabla_x \widecheck{z}^\pi_n(X^\pi_n) - \nabla_x \psi(X^\pi_n|\theta^z)\right) \sigma(t_n, X^\pi_n) \right|^2 \right]. \end{aligned} \tag{2.56}$$

The goal is to establish an upper bound for the total approximation error defined by

$$\widehat{\mathscr{E}}^{\pi}(|\pi|) := \max_n \mathbb{E}\left[\left|\Delta \widehat{Y}_n^{\pi}\right|^2\right] + \max_n \mathbb{E}\left[\left|\Delta \widehat{Z}_n^{\pi}\right|^2\right] + \mathbb{E}\left[\sum_{n=0}^{N-1}\int_{t_n}^{t_{n+1}} \left|\Gamma_r - \widehat{\Gamma}_n^{\pi}\right|^2 \mathrm{d}r\right],$$

depending on not just the discretization but also the regression errors arising from the approximations of the conditional expectations in (2.42).

Theorem 2.5.2
Let the conditions of assumption 2.4.1 be in place. Assume the time partition satisfies $N\Delta t_n \geq c$ for each $0 \leq n \leq N-1$, with some constant c. Then, for sufficiently small $|\pi|$, the total approximation error of the OSM scheme defined by the loss function (2.51) admits to

$$\widehat{\mathscr{E}}^{\pi}(|\pi|) \leq C\left(|\pi| + N\sum_{n=0}^{N-1}\{\epsilon_n^y + \epsilon_n^z\} + \sum_{n=0}^{N-1}\epsilon_n^{\gamma}\right). \tag{2.57}$$

Furthermore, in case the Γ process is taken as the direct derivative of the Z process, as in (2.52), the total error can be bounded by

$$\widehat{\mathscr{E}}^{\pi}(|\pi|) \leq C\left(|\pi| + N\sum_{n=0}^{N-1}\{\epsilon_n^y + \epsilon_n^{z,\nabla z}\} + \frac{\Upsilon^6(S_1)}{N}\right), \tag{2.58}$$

where C is a constant independent of the time partition π^N.

Proof. Throughout the proof C denotes a constant independent of the time partition, whose value may vary from line to line. We only highlight arguments which significantly differ from the ones of theorem 2.4.1.

Step 1. *Regression errors induced by the loss functions.* Using (2.49), the relation (2.50) and the total law of probability, the loss function in (2.51) can be rewritten as follows

$$\begin{aligned}
&\mathscr{L}_n^{z,\gamma}(\theta^z,\theta^{\gamma}) \\
&\quad= \mathbb{E}\left[\left|\check{Z}_n^{\pi} - \psi(X_n^{\pi}|\theta^z) + \Delta t_n \nabla_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^{\pi})\left(\chi(X_n^{\pi}|\theta^{\gamma}) - \check{\Gamma}_n^{\pi}\right)\sigma(t_n,X_n^{\pi})\right|^2\right] \\
&\qquad+ \Delta t_n \mathbb{E}\left[\left|\left(\check{\Gamma}_n^{\pi} - \chi(X_n^{\pi}|\theta^{\gamma})\right)\sigma(t_n,X_n^{\pi})\right|^2\right] + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^{\pi}\right|^2\mathrm{d}r\right] \\
&\qquad+ 2\Delta t_n \mathbb{E}\Big[\nabla_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^{\pi})\left(\chi(X_n^{\pi}|\theta^{\gamma}) - \check{\Gamma}_n^{\pi}\right)\sigma(t_n,X_n^{\pi}) \\
&\qquad\qquad\qquad\qquad\times \int_{t_n}^{t_{n+1}}(D_n\widetilde{Z}_r - \chi(X_n^{\pi}|\theta^{\gamma})\sigma(t_n,X_n^{\pi}))^T\mathrm{d}W_r\Big] \\
&\quad=: \widetilde{\mathscr{L}}_n^{z,\gamma}(\theta^z,\theta^{\gamma}) + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^{\pi}\right|^2\mathrm{d}r\right] + \widetilde{I}_n^{\gamma}(\theta^{\gamma}).
\end{aligned} \tag{2.59}$$

The inequality $(a+b)^2 \leq (1+\varrho_1\Delta t_n)a^2 + (1+1/(\varrho_1\Delta t_n))b^2$, on top of the bounded differ-

entiability of f provided by assumption 2.4.1, implies

$$\widetilde{\mathscr{L}}_n^{z,\gamma}(\theta^z,\theta^\gamma) \le (1+\varrho_1\Delta t_n)\mathbb{E}\left[\left|\check{Z}_n^\pi - \psi(X_n^\pi|\theta^z)\right|^2\right] + \left[\frac{L_{\nabla f}^2}{\rho_1}(1+\varrho_1\Delta t_n)+1\right]\Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi - \chi(X_n^\pi|\theta^\gamma)\right)\sigma(t_n,X_n^\pi)\right|^2\right]. \tag{2.60}$$

By the inequality $(a+b)^2 \ge (1-\varrho_2\Delta t_n)a^2 + (1-1/(\varrho_2\Delta t_n))b^2$, the following also holds

$$\widetilde{\mathscr{L}}_n^{z,\gamma}(\theta^z,\theta^\gamma) \ge (1-\varrho_2\Delta t_n)\mathbb{E}\left[\left|\check{Z}_n^\pi - \psi(X_n^\pi|\theta^z)\right|^2\right] + \left(1+\frac{L_{\nabla f}^2}{\varrho_2}(\varrho_2\Delta t_n-1)\right)\Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi - \chi(X_n^\pi|\theta^\gamma)\right)\sigma(t_n,X_n^\pi)\right|^2\right]. \tag{2.61}$$

The Cauchy-Schwarz inequality, (2.50) and the ϵ-Young inequality $ab \le a^2/(2\varepsilon)+\varepsilon b^2/2$, with $\varepsilon = 4L_{\nabla f}^2$, yield

$$\left|\widetilde{I}_n^\gamma(\theta^\gamma)\right| \le (1/4+4L_{\nabla f}^2\Delta t_n)\Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi - \chi(X_n^\pi|\theta^\gamma)\right)\sigma(t_n,X_n^\pi)\right|^2\right] + 4L_{\nabla f}^2\Delta t_n\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^\pi\right|^2 \mathrm{d}r\right]. \tag{2.62}$$

Implied by (2.59), minimizing $\mathscr{L}_n^{z,\gamma}(\theta^z,\theta^\gamma)$ is equivalent to minimizing $\bar{\mathscr{L}}_n^{z,\gamma} := \widetilde{\mathscr{L}}_n^{z,\gamma}(\theta^z,\theta^\gamma) + \widetilde{I}_n^\gamma(\theta^\gamma)$. Assuming that $(\widehat{\theta}_n^z,\widehat{\theta}_n^\gamma)$ is a perfect approximation – see remark 2.5.1 – of the minimal parameter space $(\theta_n^{z,*},\theta_n^{\gamma,*}) \in \operatorname{argmin}_{\theta^z,\theta^\gamma}\mathscr{L}_n^{z,\gamma}(\theta^z,\theta^\gamma)$, we have $\bar{\mathscr{L}}_n^{z,\gamma}(\widehat{\theta}_n^z,\widehat{\theta}_n^\gamma) \le \bar{\mathscr{L}}_n^{z,\gamma}(\theta^z,\theta^\gamma)$ for any (θ^z,θ^γ). Combined with (2.60), (2.61), the triangle inequality and (2.62), this implies

$$\begin{aligned}
&(1-\varrho_2\Delta t_n)\mathbb{E}\left[\left|\check{Z}_n^\pi - \widehat{Z}_n^\pi\right|^2\right] + \left(3/4 - L_{\nabla f}^2/\varrho_2 - 3L_{\nabla f}^2\Delta t_n\right)\Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi - \widehat{\Gamma}_n^\pi\right)\sigma(t_n,X_n^\pi)\right|^2\right] \\
&\quad\le (1+\varrho_1\Delta t_n)\mathbb{E}\left[\left|\check{Z}_n^\pi - \psi(X_n^\pi|\theta^z)\right|^2\right] \\
&\qquad + \left[\frac{L_{\nabla f}^2}{\varrho_1}(1+\varrho_1\Delta t_n) + 5/4 + 4L_{\nabla f}^2\Delta t_n\right]\Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi - \chi(X_n^\pi|\theta^\gamma)\right)\sigma(t_n,X_n^\pi)\right|^2\right] \\
&\qquad + 8L_{\nabla f}^2\Delta t_n\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^\pi\right|^2\mathrm{d}r\right],
\end{aligned} \tag{2.63}$$

for any (θ^z,θ^γ), $\varrho_1,\varrho_2>0$. In particular, choosing $\varrho_2^* := 8L_{\nabla f}^2$, for any sufficiently small Δt_n such that $3L_{\nabla f}^2\Delta t_n < 1/8$ and $\varrho_2^*\Delta t_n \le 1/2$, we derive

$$\mathbb{E}\left[\left|\check{Z}_n^\pi - \widehat{Z}_n^\pi\right|^2\right] + \Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi - \widehat{\Gamma}_n^\pi\right)\sigma(t_n,X_n^\pi)\right|^2\right] \le C(\epsilon_n^z + \Delta t_n\epsilon_n^\gamma) + 16L_{\nabla f}^2\Delta t_n\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^\pi\right|^2\mathrm{d}r\right], \tag{2.64}$$

recalling the definitions in (2.56). Through analogous steps to [77, Thm. 4.1, step 3-4], a similar estimate can be established for the loss function (2.53), ultimately giving

$$\mathbb{E}\left[\left|\check{Y}_n^\pi - \widehat{Y}_n^\pi\right|^2\right] \leq C \inf_{\theta^y} \mathbb{E}\left[\left|\check{Y}_n^\pi - \varphi(X_n^\pi|\theta^y)\right|^2\right] =: C\epsilon_n^y. \tag{2.65}$$

Step 2. $\mathbb{L}^2$-*regularity of* $D_n\widetilde{Z}_r$. In what follows, we will need an estimate controlling the so-called $\mathbb{L}^2$-regularity of the stochastic integrand $D_n\widetilde{Z}_r$, corresponding to the last term in (2.64). This term admits to the following bound

$$\begin{aligned}\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^\pi\right|^2 \mathrm{d}r\right] \leq\ & 3\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_{t_n}Z_r\right|^2 \mathrm{d}r\right] \\ &+ 3\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2 \mathrm{d}r\right] \\ &+ 3\Delta t_n \mathbb{E}\left[\left|\overline{DZ}_n^{n+1} - D_n\check{Z}_n^\pi\right|^2\right] =: 3(R_1+R_2+R_3).\end{aligned} \tag{2.66}$$

The second term of the right-hand side corresponds to the $\mathbb{L}^2$-regularity of DZ given by (2.25). For the first term, notice that by Itô's isometry, (2.49) and (2.14), we have

$$R_1 = \mathbb{E}\left[\left|\Delta\check{Z}_n^\pi - \Delta D_n\widehat{Y}_{n+1}^\pi + \int_{t_n}^{t_{n+1}} f^D(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi, \mathbf{D}_n\check{\mathbf{X}}_{n+1,n}^\pi) - f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r)\mathrm{d}r\right|^2\right].$$

(2.42b) implies an identity similar to (2.33). Then, by the law of total probability, the $L^2([0,T];\mathbb{R}^d)$ Cauchy-Schwarz and Jensen inequalities, it follows that

$$\begin{aligned}R_1 \leq\ & 2\mathbb{E}\left[\left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2 - \left|\mathbb{E}_n\left[\Delta D_n\widehat{Y}_{n+1}^\pi\right]\right|^2\right] \\ &+ 4\Delta t_n\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|f^D(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi,\mathbf{D}_n\check{\mathbf{X}}_{n+1,n}^\pi) - f^D(r,\mathbf{X}_r,\mathbf{D}_{t_n}\mathbf{X}_r)\right|^2\mathrm{d}r\right].\end{aligned} \tag{2.67}$$

Notice that the second term above implicitly depends on R_3. Similarly to step 1 in theorem 2.4.1 – see (2.30) in particular –, we also gather the following estimate

$$\begin{aligned}R_3 :=\ & \Delta t_n\mathbb{E}\left[\left|\overline{DZ}_n^{n+1} - D_n\check{Z}_n^\pi\right|^2\right] \\ \leq\ & 4d\left(\mathbb{E}\left[\left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right] - \mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n\widehat{Y}_{n+1}^\pi\right]\right|^2\right]\right) \\ &+ 28dL_{f^D}^2\Delta t_n\Big\{C\Delta t_n^2 + 2\Delta t_n\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] \\ &\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta\widehat{Y}_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta\widehat{Z}_{n+1}^\pi\right|^2\right]\right) \\ &\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right]\right) \\ &\qquad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right]\Big\},\end{aligned} \tag{2.68}$$

for any sufficiently small time step satisfying $14dL_{f^D}^2\Delta t_n \le 1/2$. Plugging the combined estimate resulting from (2.67) and (2.68) into (2.66), subsequently gives

$$\begin{aligned}
&\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^\pi\right|^2\mathrm{d}r\right] \\
&\qquad\le 3(2+8d)\left(\mathbb{E}\left[\left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right]-\mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n\widehat{Y}_{n+1}^\pi\right]\right|^2\right]\right) \\
&\qquad\quad+84L_{f^D}^2(1+2d)\Delta t_n\Big\{C\Delta t_n^2+2\Delta t_n\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad\qquad+2\Delta t_n\left(\mathbb{E}\left[\left|\Delta\widehat{Y}_{n+1}^\pi\right|^2\right]+\mathbb{E}\left[\left|\Delta\widehat{Z}_{n+1}^\pi\right|^2\right]\right) \\
&\qquad\qquad\qquad+2\Delta t_n\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2\right] \\
&\qquad\qquad\qquad+2\Delta t_n\mathbb{E}\left[\left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right]\Big\} \\
&\qquad\quad+C\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r-\overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right],
\end{aligned}\tag{2.69}$$

establishing an upper bound for the $\mathbb{L}^2$-regularity of $D_n\widetilde{Z}_r$.

Step 3. *Approximation error bound in the parametrized case.* The total approximation errors can be decomposed into discretization and regression errors as follows

$$\begin{aligned}
\Delta t_n\mathbb{E}\left[\left|\overline{DZ}_n^{n+1}-D_n\widehat{Z}_n^\pi\right|^2\right] &\le 2\Delta t_n\mathbb{E}\left[\left|\overline{DZ}_n^{n+1}-D_n\check{Z}_n^\pi\right|^2\right] \\
&\quad+2\Delta t_n\mathbb{E}\left[\left|D_n\check{Z}_n^\pi-D_n\widehat{Z}_n^\pi\right|^2\right],
\end{aligned}\tag{2.70}$$

$$\begin{aligned}
(1-\beta_z\Delta t_n)\mathbb{E}\left[\left|\Delta\widehat{Z}_n^\pi\right|^2\right] &\le \mathbb{E}\left[\left|\Delta\check{Z}_n^\pi\right|^2\right]+\frac{1}{\beta_z\Delta t_n}\mathbb{E}\left[\left|\check{Z}_n^\pi-\widehat{Z}_n^\pi\right|^2\right], \\
(1-\beta_y\Delta t_n)\mathbb{E}\left[\left|\Delta\widehat{Y}_n^\pi\right|^2\right] &\le \mathbb{E}\left[\left|\Delta\check{Y}_n^\pi\right|^2\right]+\frac{1}{\beta_y\Delta t_n}\mathbb{E}\left[\left|\check{Y}_n^\pi-\widehat{Y}_n^\pi\right|^2\right],
\end{aligned}\tag{2.71}$$

for any $\beta_z,\beta_y>0$. Combined with (2.64), (2.71) leads to the following estimate

$$\begin{aligned}
&(1-\beta_z\Delta t_n)\mathbb{E}\left[\left|\Delta\widehat{Z}_n^\pi\right|^2\right] \\
&\qquad\le \mathbb{E}\left[\left|\Delta\check{Z}_n^\pi\right|^2\right]+\frac{C}{\beta_z\Delta t_n}(\epsilon_n^z+\Delta t_n\epsilon_n^\gamma)+\frac{16L_{\nabla f}^2}{\beta_z}\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r-D_n\check{Z}_n^\pi\right|^2\mathrm{d}r\right],
\end{aligned}$$

for any $\beta_z > 0$. Similar arguments as in step 2 in theorem 2.4.1 subsequently give

$$\begin{aligned}
&(1-\beta_z\Delta t_n)\mathbb{E}\left[\left|\Delta\widehat{Z}_n^\pi\right|^2\right] \\
&\quad\le (1+\varrho\Delta t_n)\mathbb{E}\left[\left|\mathbb{E}_n\left[\Delta D_n\widehat{Y}_{n+1}^\pi\right]\right|^2\right] \\
&\qquad + \frac{7L_{f^D}^2}{\varrho}(1+\varrho\Delta t_n)\Big\{C\Delta t_n^2 \\
&\qquad\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2 + \left|\Delta\widehat{Y}_{n+1}^\pi\right|^2 + \left|\Delta\widehat{Z}_{n+1}^\pi\right|^2\right]\right) \\
&\qquad\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2 + \left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right]\right) \\
&\qquad\qquad + \Delta t_n\mathbb{E}\left[\left|D_n\check{Z}_n^\pi - \overline{DZ}_n^{n+1}\right|^2\right] \\
&\qquad\qquad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right]\Big\} \\
&\qquad + \frac{C}{\beta_z\Delta t_n}(\epsilon_n^z + \Delta t_n\epsilon_n^\gamma) + \frac{16L_{\nabla f}^2}{\beta_z}\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r - D_n\check{Z}_n^\pi\right|^2\mathrm{d}r\right],
\end{aligned}$$

for any $\varrho > 0$. Plugging in the estimates established by (2.68) and (2.69), choosing $\varrho^* = 56dL_{f^D}^2$ and $\beta_z^* = 96(2+8d)L_{\nabla f}^2$, we derive

$$\begin{aligned}
&(1-\beta_z^*\Delta t_n)\mathbb{E}\left[\left|\Delta\widehat{Z}_n^\pi\right|^2\right] \le (1+C_z\Delta t_n)\mathbb{E}\left[\left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right] \\
&\qquad + C_z\Big\{C\Delta t_n^2 + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta X_{n+1}^\pi\right|^2 + \left|\Delta\widehat{Y}_{n+1}^\pi\right|^2 + \left|\Delta\widehat{Z}_{n+1}^\pi\right|^2\right]\right) \\
&\qquad\qquad + 2\Delta t_n\left(\mathbb{E}\left[\left|\Delta D_n X_{n+1}^\pi\right|^2 + \left|\Delta D_n\widehat{Y}_{n+1}^\pi\right|^2\right]\right) \\
&\qquad\qquad + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right] \\
&\qquad\qquad + \frac{\epsilon_n^z + \Delta t_n\epsilon_n^\gamma}{\beta_z^*\Delta t_n}\Big\}.
\end{aligned} \tag{2.72}$$

By analogous computations for Y, similar arguments as in theorem 2.4.1 imply

$$\begin{aligned}
&(1-\beta^*\Delta t_n)\left(\mathbb{E}\left[\left|\Delta\widehat{Y}_n^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta\widehat{Z}_n^\pi\right|^2\right]\right) \\
&\qquad\le (1+C\Delta t_n)\left(\mathbb{E}\left[\left|\Delta\widehat{Y}_{n+1}^\pi\right|^2\right] + \mathbb{E}\left[\left|\Delta\widehat{Z}_{n+1}^\pi\right|^2\right]\right) \\
&\qquad\quad + C\Big\{\Delta t_n^2 + \mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_{t_n}Z_r - \overline{DZ}_n^{n+1}\right|^2\mathrm{d}r\right] \\
&\qquad\qquad + \frac{\epsilon_n^y + \epsilon_n^z + \Delta t_n\epsilon_n^\gamma}{\Delta t_n}\Big\},
\end{aligned} \tag{2.73}$$

with some $\beta^* > 0$, depending on both β_z^*, β_y^*. Thereafter, for any sufficiently small time step admitting to $\beta^*\Delta t_n < 1$, an application of the discrete Grönwall lemma implies the total approximation error of Y and Z in (2.57).

The Γ estimate then follows in a similar manner to step 5 in theorem 2.4.1 using (2.70), (2.68), (2.64), (2.72); and observing that $(1+C\Delta t_n)/(1-\beta^*\Delta t_n)-1$ is $\mathcal{O}(|\pi|)$ given $\beta^*\Delta t_n<1$. This completes the total approximation error of (2.57).

2

Step 4. *Derivative representation error of Z and Γ.* In order to prove (2.58), we need to establish an error estimate bounding the difference between the spatial derivative of (2.42b) and the target of (2.42a). Notice that under the conditions of assumption 2.4.1 and (2.55), the arguments of the conditional expectations are all C_b^2 in x. Then, formal differentiation of (2.42b) with the Leibniz rule and the integration-by-parts formula in (2.80) applied on (2.42a), gives

$$\begin{aligned}\left(\nabla_x \check{z}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)\right)\sigma = \Delta t_n[\left(\widehat{\gamma}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)\right)\sigma]^T\mathbb{E}_n\left[\nabla_x\nabla_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)\right]\sigma\\ +\Delta t_n\mathbb{E}_n\left[\nabla_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)\right]\nabla_x\widehat{\gamma}_n^\pi(X_n^\pi)\sigma^2.\end{aligned}$$

By the bounded differentiability conditions in $(\mathbf{A}_2^{f,g})$, we have that

$$\begin{aligned}\mathbb{E}\left[\left|\left(\nabla_x \check{z}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)\right)\sigma\right|^2\right] \le 2\Delta t_n^2 L^2_{\nabla^2 f}|\sigma|^2\mathbb{E}\left[\left|\left(\widehat{\gamma}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)\right)\sigma\right|^2\right]\\ +2\Delta t_n^2 L^2_{\nabla f}|\sigma|^4\mathbb{E}\left[\left|\nabla_x\widehat{\gamma}_n^\pi(X_n^\pi)\right|^2\right].\end{aligned}$$

Splitting the first term according to $\widehat{\gamma}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)=\widehat{\gamma}_n^\pi(X_n^\pi)-\nabla_x\check{z}_n^\pi(X_n^\pi)+\nabla_x\check{z}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)$, using the direct estimate $\widehat{\gamma}_n^\pi(X_n^\pi)\equiv\nabla_x\widehat{z}_n^\pi(X_n^\pi)$ implied by (2.52), and recalling the bounds in (2.55), subsequently yields

$$\mathbb{E}\left[\left|\left(\nabla_x \check{z}_n^\pi(X_n^\pi)-\check{\gamma}_n^\pi(X_n^\pi)\right)\sigma\right|^2\right] \le C\left(\mathbb{E}\left[\left|\left(\nabla_x \check{z}_n^\pi(X_n^\pi)-\nabla_x\widehat{z}_n^\pi\right)\sigma\right|^2\right]+\Upsilon^6(S_1)\right),$$

for small enough time steps admitting to $4\Delta t_n^2 L^2_{\nabla^2 f}|\sigma|^2<1$. Combining this estimate with the upper bound (2.63), recalling the definition of $\epsilon_n^{z,\nabla z}$ in (2.56), we gather

$$\begin{aligned}\mathbb{E}\left[\left|\check{Z}_n^\pi-\widehat{Z}_n^\pi\right|^2\right]+\Delta t_n\mathbb{E}\left[\left|\left(\check{\Gamma}_n^\pi-\nabla_x\widehat{Z}_n^\pi\right)\sigma(t_n,X_n^\pi)\right|^2\right] \qquad (2.74)\\ \le C\left(\epsilon_n^{z,\nabla z}+\Delta t_n^3\Upsilon^6(S_1)\right)+16L^2_{\nabla f}\Delta t_n\mathbb{E}\left[\int_{t_n}^{t_{n+1}}\left|D_n\widetilde{Z}_r-D_n\check{Z}_n^\pi\right|^2\mathrm{d}r\right],\end{aligned}$$

for small enough time steps $\Delta t_n<1$ and diverging $\Upsilon(S_1)$. The total approximation error estimate in (2.58) then follows in a similar manner, combining (2.74) with (2.65), (2.73) and the discrete Grönwall lemma, as in the previous step.

This completes the proof.

□

Theorem 2.5.2 establishes the convergence of the Deep BSDE approach to (2.42), given the UAT property of neural networks provided by theorem 2.5.1. The first terms in both (2.57) and (2.58) correspond to the discrete time approximation errors in theorem 2.4.1. The second terms correspond to the approximations of the neural network

regression Monte Carlo approach. Provided by theorem 2.5.1, the corresponding regression biases defined by (2.56) can be made arbitrarily small with the choice of shallow neural networks already. In exchange to avoid the parametrization in the automatic differentiation approach in (2.52), one needs to restrict the parametrization to the case of $\Sigma_{C_b^2}(\tanh)$ neural networks and subsequently deal with an additional error term in (2.58), which depends on the increasing sequence $\Upsilon(S_1)$, controlling the magnitude of the parameters. If this dominating sequence is such that $\Upsilon^6(S_1)/N \to 0$ while $S_1, N \to \infty$ this ensures the existence of neural networks $\varphi(\cdot|\theta^y), \psi(\cdot|\theta^z) \in \Sigma_{C_b^2}(\tanh)$ such that the total approximation error converges. We shall, however, notice that the claim above guarantees nothing more, and in fact does not guarantee the convergence of the final approximations including regression errors, which we highlight in the remark below.

Remark 2.5.1 (Limitations of theorem 2.5.2)
In the proof of theorem 2.5.2 we neglected the presence of three additional error terms. These are the following.

1. *First, the definitions in* (2.56) *only express the regression biases due to the choice of a finite number of parameters. The actual regression errors also incorporate the approximation error of the optimal parameter space* $\widehat{\theta}_n^z$ *and induce a term* $\mathbb{E}\left[\left|\varphi(X_n^\pi|\theta_n^{y,*}) - \varphi(X_n^\pi|\widehat{\theta}_n^y)\right|^2\right]$*, which stems from the fact that unlike in a linear regression method – see, e.g., [16] –, one does not have a closed-form expression for the true minimizers* $(\theta_n^{z,*}, \theta_n^{\gamma,*}), \theta_n^{y,*}$ *but can only gather an approximation of them with a stochastic gradient descent (SGD) optimization. The present understanding of this term is poor, mainly due to the non-convexity of the corresponding target function – see [83] and the references therein. Currently, there exists no theoretical guarantee which would ensure the convergence of SGD iterations in the FBSDE context. Furthermore, the second term at the right-hand side of* (2.57) *(respectively,* (2.58)*) implies that, in order to preserve the convergence of the total approximation error* $\widehat{\mathscr{E}}^\pi(|\pi|)$*, one needs* $\epsilon_n^y + \epsilon_n^z$ $(\epsilon_n^y + \epsilon_n^{z,\nabla z})$ *to be at least* $\mathcal{O}(N^{-2})$ *for each* $n = 0, \dots, N-1$*. In case of the regression biases defined by* (2.56)*, this can be achieved by the UAT property in theorem 2.5.1. Establishing a similar theoretical guarantee for the regression* errors *stemming from SGD approximations is currently not possible due to the aforementioned reasons. Nonetheless, in figure 2.3, we provide empirical evidence which suggests that this condition may indeed be satisfied in practice, encouraging further research in this direction.*

2. *The second term arises due to the fact that in practice one can only calculate an empirical version of the expectations in* $\mathscr{L}_n^y, \mathscr{L}_n^{z,\gamma}, \mathscr{L}_n^{z,\nabla z}$*. This induces a Monte Carlo simulation error of finitely many samples. However, as we shall see in the upcoming numerical section, thanks to the soft memory limitation of a single SGD step, one can pass so many realizations of the underlying Brownian motion throughout the optimization cycle that the magnitude of the corresponding error term becomes negligible compared to other sources of error.*

3. *The final observation that needs to be highlighted is the compactness assumption on the domain in theorem 2.5.1. This error term can be dealt with in a similar fashion*

to [77, Remark 4.2], where a localization argument is constructed in such a way that – under suitable truncation ranges – convergence is ensured.

2

2.6. NUMERICAL EXPERIMENTS

In order to show the accuracy and robustness of the proposed scheme, we present results of numerical experiments on three different types of problems. We distinguish between the two Deep BSDE approaches for the OSM scheme, based on whether the Γ process is parametrized with an $\mathbb{R}^{d\times d}$-valued neural network – see (2.51) –, or it is obtained as the direct Jacobian of the parametrization of the Z process via automatic differentiation – as in (2.52). We label these variants by (P) and (D), respectively. As a reference method, we compare the results of the OSM scheme to the first scheme (DBDP1) of Huré et al. [77], which corresponds to the Euler discretization of (2.12) when $\vartheta_y = \vartheta_z = 1$. In accordance with their findings, we found the parametrized version (DBDP1) more robust than the automatic differentiated one (DBDP2) in high-dimensional settings.

Each BSDE is discretized with N equidistant time intervals, giving $\Delta t_n = T/N$ for all $n = 0,\dots,N-1$. For the implicit ϑ_y parameter of the discretization in (2.19), we choose values $\vartheta_y \in \{0, 1/2, 1\}$. In all upcoming examples we use fully-connected, feedforward neural networks of $L = 2$ hidden layers with $S_l = 100 + d$ neurons in each layer. In line with theorem 2.5.2, a hyperbolic tangent activation is deployed, yielding continuously differentiable parametrizations. Layer normalization [8] is applied in between the hidden layers. For the stochastic gradient descent iterations, we use the Adam optimizer with the adaptive learning rate strategy of [32] – see $\eta(i)$ in algorithm 1. The optimization is done as follows: in each backward recursion we allow $I = 2^{15}$ SGD iterations for the $N-1$'th time step. Thereafter, we make use of the transfer learning initialization given by (2.54), and reduce the number of iterations to $I = 2^{11}$ for all preceding time steps. In each iteration step, the optimization receives a new, independent sample of the underlying forward diffusion with $B = 2^{10}$ sample paths, meaning that in total the iteration processes 2^{25} and 2^{21} many realizations of the Brownian motion at time step $n = N-1$ and $n < N-1$, respectively. In order to speed up normalization, neural network trainings were carried out with single floating point precision. For the implementation of the BCOS method, we choose $K = 2^9$ Fourier coefficients, $P = 5$ Picard iterations and truncate the infinite integrals to a finite interval of $[a,b] = [x_0 + \kappa_\mu - L\sqrt{\kappa_\sigma}, x_0 + \kappa_\mu + L\sqrt{\kappa_\sigma}]$, where $\kappa_\mu = \mu(0,x_0)\,T$, $\kappa_\sigma = \sigma(0,x_0)\,T$. As in [145], we fix $L = 10$.

The OSM method has been implemented in TensorFlow 2. In order to exploit static graph efficiency, all core methods are decorated with `tf.function` decorators. The library used in this chapter will be publicly accessible under github. All experiments below were run on a DELL Alienware Aurora R10 machine, equipped with an AMD Ryzen 9 3950X CPU (16 cores, 64Mb cache, 4.7GHz) and an Nvidia GeForce RTX 3090 GPU (24Gb). In order to assess the inherent stochasticity of both the regression Monte Carlo method and the SGD iterations, we run each experiment 5 times and report on the mean and standard deviations of the resulting independent approximations. $\mathbb{L}^2$-errors are estimated over an independent sample of size $M = 2^{10}$ produced by the same machinery

as the one used for the simulations. Hence, the final error estimates are calculated as

$$\widehat{\mathbb{E}}\left[\left|\Delta\widehat{Y}_n^\pi\right|^2\right] = \frac{1}{M}\sum_{m=1}^{M}\left|\Delta\widehat{Y}_n^\pi(m)\right|^2, \quad \widehat{\mathbb{E}}\left[\left|\Delta\widehat{Z}_n^\pi\right|^2\right] = \frac{1}{M}\sum_{m=1}^{M}\left|\Delta\widehat{Z}_n^\pi(m)\right|^2,$$
$$\widehat{\mathbb{E}}\left[\left|\Delta\widehat{\Gamma}_n^\pi\right|^2\right] = \frac{1}{M}\sum_{m=1}^{M}\left|\Delta\widehat{\Gamma}_n^\pi(m)\right|^2$$

where $\Delta Y_n^\pi(m)$ corresponds to the m'th path of test sample, and similarly for other error measures.

2.6.1. EXAMPLE 1: REACTION-DIFFUSION WITH DIMINISHING CONTROL

The first, *reaction-diffusion* type equation is taken from [62, Example 2]. Such equations are common in financial applications. The coefficients of the BSDE (2.1) are as follows

$$\mu = \mathbf{0}_d, \qquad \sigma = I_d,$$
$$f(t,x,y,z) = \frac{\omega(t,\lambda x)}{[1+\omega(t,\lambda x)]^2}\left[\lambda^2 d(y-\gamma) - 1 - \frac{\lambda^2}{2}d\right], \qquad g(x) = \gamma + \frac{\omega(T,\lambda x)}{1+\omega(T,\lambda x)}, \tag{2.75}$$

where $\omega(t,x) = \exp\left(t+\sum_{i=1}^d x_i\right)$. These parameters satisfy assumption 2.4.1. The driver is independent of Z and f^D does not depend on the Y process. Consequently, the solutions of (2.9b) and (2.9d) can be separated into two disjoint problems. The analytical solutions are given by

$$X_t = W_t, \quad y(t,x) = \frac{\omega(t,\lambda x)}{1+\omega(t,\lambda x)}, \quad z(t,x) = \lambda\frac{\omega(t,\lambda x)}{(1+\omega(t+\lambda x))^2}\mathbf{1}_d,$$
$$\gamma(t,x) = \lambda^2\frac{\omega(t,\lambda x)(1-\omega(t,\lambda x))}{(1+\omega(t,\lambda x))^3}\mathbf{1}_{d,d}.$$

We choose $T = 0.5$, $\gamma = 0.6, \lambda = 1$ and fix $x_0 = \mathbf{1}_d$. We consider $d \in \{1, 10\}$ with $\vartheta_y \in \{0, 1\}$.

In figure 2.1, the convergence of the two fully-implementable schemes is assessed. Figure 2.1a depicts the convergence for $d = 1$. The BCOS estimates, drawn with lines, show the same order of convergence as in theorem 2.4.1, confirming the theoretical findings of the discretization error analysis. The Deep BSDE approximations, depicted with scattered error bars, exhibit higher error figures, showcasing the presence of an additional regression component. Nevertheless, the complete approximation error of the corresponding regression estimates admit to the same order of convergence as in theorem 2.5.2. The Γ approximations corresponding to the parametrized (P) and automatic differentiated (D) cases, demonstrate the difference between the bounds in (2.57) and (2.58). Indeed, we observe an extra error stemming from the bounded differentiability component of the neural networks – see (2.55). The convergence of the regression approximations flattens out for the finest time partition $N = 100$ – see the regression error of Y in particular – at a level of $\sim \mathcal{O}(10^{-7})$, indicating the presence of a regression bias induced by the restriction on a finite number of parameters. In figure 2.1b, the same dynamics are depicted for $d = 10$, where we observe the same order of convergence, in accordance with theorem 2.5.2. Note that the regression estimates of the Z process

Table 2.1: Example 1 in (2.75), $d = 10$, $N = 100$. Summary of Deep BSDE estimates. Mean-squared errors are calculated over an independent sample of $M = 2^{10}$ realizations of the underlying Brownian motion. Means and standard deviations (in parentheses) obtained over 5 independent runs of the algorithm. Best estimates within one standard deviation highlighted in gray. Γ estimates from Huré et al. in [77] are obtained via automatic differentiation.

	OSM($\vartheta_y = 0$)		OSM($\vartheta_y = 1$)		Huré et al. (2020)
	(P)	(D)	(P)	(D)	
$\lvert\Delta\widehat{Y}_0^\pi\rvert/\lvert Y_0\rvert$	3e−4 (3e−4)	3e−4 (2e−4)	6e−4 (2e−4)	2e−4 (2e−4)	1.1e−3 (4e−4)
$\lvert\Delta\widehat{Z}_0^\pi\rvert/\lvert Z_0\rvert$	7e−3 (3e−3)	8e−3 (2e−3)	9e−3 (2e−3)	9e−3 (5e−3)	9e−3 (2e−3)
$\lvert\Delta\widehat{\Gamma}_0^\pi\rvert$	1.2e−2 (3e−3)	8e−3 (3e−3)	9e−3 (1e−3)	8e−3 (2e−3)	9.9e+2 (8e+1)
$\max_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{Y}_n^\pi\rvert^2]$	2.4e−5 (5e−6)	2.4e−5 (7e−6)	2.7e−5 (8e−6)	2.1e−5 (4e−6)	2.9e−5 (6e−6)
$\max_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{Z}_n^\pi\rvert^2]$	1.3e−4 (2e−5)	9e−5 (1e−5)	1.1e−4 (2e−5)	1.0e−4 (3e−5)	7.4e−4 (9e−5)
$\sum_{n=0}^{N-1} \Delta t_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{\Gamma}_n^\pi\rvert^2]$	8e−4 (2e−4)	5.0e−4 (7e−5)	8e−4 (2e−4)	5e−4 (1e−4)	5.0e+3 (8e+2)
runtime (s)	1.20e+3 (1e+1)	1.44e+3 (2e+1)	1.19e+3 (1e+1)	1.43e+3 (5e+1)	5.7e+2 (3e+1)

converge until, and including, the finest time partition $N = 100$ in case of the OSM discretization. On the other hand, with the approach of Huré et al. [77] the decay stops at $N = 50$, indicating the impact of diverging conditional variances, as anticipated in remark 2.3.1. Table 2.1 contains the means and standard deviations of a collection of error measures with respect to 5 independent runs of the same regression Monte Carlo method. It can be seen that – regardless of the value of ϑ_y – the OSM scheme yields an order of magnitude improvement in the approximation of the Z process, while showing identical error figures in the Y process. Errors under the automatic differentiated case (D) with (2.52) are slightly better than in the parametrized approach (P). The Γ approximations show comparable accuracies. The total runtime of the OSM regressions is approximately double of that of [77], which is intuitively explained by the fact that (2.42) solves two BSDEs at each point in time. Execution times under the automatic differentiated variant are slightly higher than in the parameterized case, confirming the extra computational complexity of Jacobian training in (2.52). The neural network regression Monte Carlo method yields sharp, robust estimates with small standard deviations over independent runs of the algorithm, in particular corresponding the Z process.

2.6.2. EXAMPLE 2: HAMILTON-JACOBI-BELLMAN WITH LQG CONTROL

The Hamilton-Jacobi-Bellman (HJB) equation is a non-linear PDE derived from Bellman's dynamic programming principle, whose solution is the *value function* of a corresponding *stochastic control* problem. In what follows, we consider the linear-quadratic-Gaussian (LQG) control, which describes a linear system driven by additive noise [68]. The FBSDE system (2.1), associated with the HJB equation has the following coefficients

$$\mu = \mathbf{0}_d, \qquad \sigma = \sqrt{2} I_d, \qquad f(t,x,y,z) = |z|^2, \qquad g(x) = x^T A x + v^T x + c, \tag{2.76}$$

where $A \in \mathbb{R}^{d\times d}, v \in \mathbb{R}^{d\times 1}, c \in \mathbb{R}$. Unlike in [68], the hereby considered terminal condition is a quadratic mapping of space. This choice is made so that we have access to *semi-analytical*, pathwise reference solutions $\{(Y_t, Z_t, \Gamma_t)\}_{0\le t\le T}$. Indeed, considering the associated parabolic problem (2.2), it is straightforward to show that the solution is given

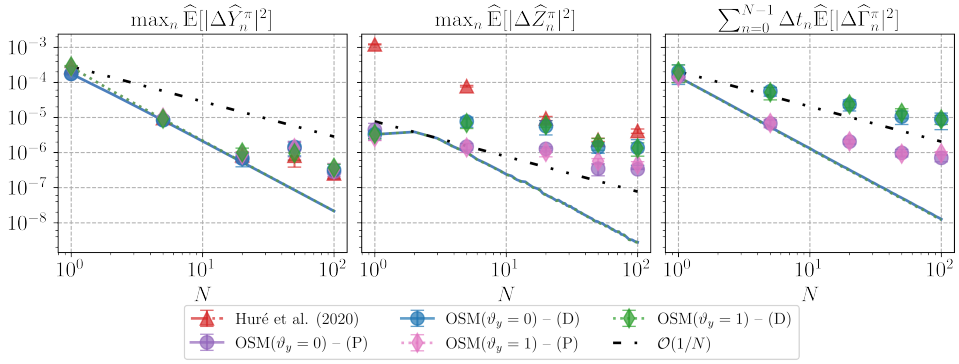

(a) BCOS and Deep BSDE, $d = 1$. From left to right: maximum mean-squared approximation errors of Y and Z; average mean-squared approximation error of Γ. Lines correspond to BCOS estimates, scattered error bars to the means and standard deviations of 5 independent neural network regressions.

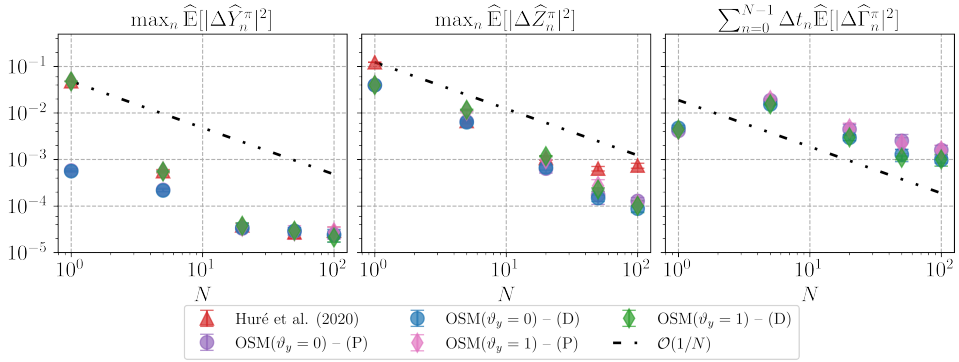

(b) Deep BSDE, $d = 10$. From left to right: maximum mean-squared approximation errors of Y and Z; average mean-squared approximation error of Γ. Means and standard deviations are calculated over 5 independent runs of the algorithm.

Figure 2.1: Example 1 in (2.75). Convergence of approximation errors. Mean-squared errors are calculated over an independent sample of $M = 2^{10}$ realizations of the underlying Brownian motion.

by

$$X_t = \sigma W_t, \quad y(t,x) = x^T P(t) x + Q^T(t) x + R(t),$$
$$z(t,x) = \sigma\left(\left[P(t) + P^T(t)\right] x + Q(t)\right), \quad \gamma(t,x) = \sigma\left[P(t) + P^T(t)\right],$$

where the purely time dependent functions $P : [0,T] \to \mathbb{R}^{d\times d}, Q : [0,T] \to \mathbb{R}^{d\times 1}, R : [0,T] \to \mathbb{R}$ satisfy the following set of Riccati type ordinary differential equations (ODE)

$$\begin{aligned} \dot{P}(t) - \left[P(t) + P^T(t)\right]^2 = 0, \quad \dot{Q}(t) - 2\left[P(t) + P^T(t)\right] Q(t) = 0, \\ \dot{R}(t) + \mathrm{Tr}\left\{P(t) + P^T(t)\right\} - |Q(t)|^2 = 0, \\ P(T) = A, \quad Q(T) = \nu, \quad R(T) = c, \end{aligned} \tag{2.77}$$

with $\dot{P} = \mathrm{d}P/\mathrm{d}t$, $\dot{Q} = \mathrm{d}Q/\mathrm{d}t$ and $\dot{R} = \mathrm{d}R/\mathrm{d}t$. The reference solution is then obtained by integrating (2.77) over a refined time grid of $N_{\mathrm{ODE}} = 10^4$ intervals.[5] We take $A = I_d, \nu = \mathbf{0}_d, c = 0$, $T = 0.5$ and fix $x_0 = \mathbf{1}_d$. An interesting feature of the FBSDE system defined by (2.76) is that the driver is independent of Y meaning that the Malliavin BSDE in (2.9d) can be solved separately from the backward equation. Consequently, the discrete time approximations of Z and Γ in (2.19) do not depend on ϑ_y. Moreover, the driver is quadratically growing in Z, in particular, it is only Lipschitz continuous over compact domains. Nevertheless, we include this problem to show promising results beyond assumption 2.4.1. We pick $\vartheta_y = 1/2$ and investigate the solution in $d \in \{1, 50\}$.

In figure 2.2 the regression errors of the Deep BSDE approach are assessed in $d = 1$. The true regression targets in (2.42) are benchmarked according to BCOS. In fact, at time step n, the corresponding cosine expansion coefficients are recovered by means of DCT, given neural network approximations $\widehat{Y}^{\pi}_{n+1}, \widehat{Z}^{\pi}_{n+1}, \widehat{\Gamma}^{\pi}_{n+1}$. These coefficients are subsequently plugged in (2.46) to gather BCOS estimates. For large enough Fourier domains and sufficiently many Picard iterations, the COS error becomes negligible compared to the discretization component and the resulting estimates approximate the true regression labels $\check{Y}^{\pi}_n, \check{Z}^{\pi}_n, \check{\Gamma}^{\pi}_n$. Hence, they can then be used to assess the regression errors induced by the Monte Carlo method. Figure 2.2a depicts these regression errors over time for $N = 100$. As it can be seen, the model of Huré et al. [77] and the OSM scheme result in similar regression error components for the Y process. However, the regression errors of the Z process are three orders of magnitude worse in case of the reference method [77], and in fact, dominate the total approximation error at $n = N - 1$. In contrast, the OSM estimates – middle plot of figure 2.2a – exhibit the same order of regression error as for the Y process. This demonstrates the advantageous conditional variance behavior of the corresponding OSM estimates, as pointed out in remark 2.3.1. The regression errors of the Γ process show comparable figures. The cumulative regression errors, corresponding to the second term in theorem 2.5.2, are collected in figure 2.3b. In case of the model in [77], the cumulative regression error of the Z process blows up as the mesh size $|\pi| = T/N$ decreases. On the contrary, the cumulative regression errors in all processes (Y, Z, Γ) are at a constant level of $\mathcal{O}(10^{-5})$ for the OSM scheme. In light of remark 2.5.1, this indicates that the chosen, finite network architecture incorporates a regression bias which cannot be further reduced. In our experiments, we found that it is difficult to decrease this component by changing the number of hidden layers L or neurons per hidden layer S_l. Assessing this phenomenon requires a better understanding of both narrow UAT estimates and the convergence of SGD iterations.

In figure 2.3 the $d = 50$ dimensional case is depicted. In order to have dimension independent scales, *relative* mean-squared errors are reported. Figure 2.3a collects the relative approximation error over the discretized time window when $N = 100$. Compared to [77], the OSM estimates yield a significant improvement in each part of the solution triple. In particular, the approximation errors of the Z process are three orders of magnitude better with both the parametrized (P) and automatic differentiated (D) approaches. In case of the Γ process, two observations can be made. First, the corresponding curve demonstrates that naive automatic differentiation of the Z approximations in [77] does not provide reliable Γs. Moreover, it can be seen that the parametrized version (P) of the

[5] This is done using scipy.integrate.odeint.

Deep BSDE approach given by (2.51) provides an order of magnitude better average Γ errors. The convergence of the total approximation errors is depicted in figure 2.3b. The neural network regression estimates converge for both the parametrized (P) and the automatic differentiated (D) loss functions until $N = 50$, when the regression bias becomes apparent. Additionally, the convergence of the Γ approximations is significantly better in the parametrized case, suggesting that for such a quadratically scaling driver the last term of (2.58) is a driving error component.

In table 2.2 means and standard deviations of a collection of error measures are gathered, with respect to 5 independent runs of the same regression Monte Carlo method, for both $d = 1$ and $d = 50$. The numbers are in line with the observations above. In particular, we highlight that the error terms corresponding to the Z and Γ approximations are four orders of magnitude better than in case of the reference method [77]. The parametrized version (P) of the Deep BSDE shows consistently better convergence. The total runtime of the neural network regression Monte Carlo approach is moderately increased between $d = 1$ and $d = 50$. In fact, the average execution time of a single SGD step for the parametrized (P) case in (2.51) increases from 2.8e−3(4e−4) to 3.3e−3(4e−4) seconds in between $d = 1$ and $d = 50$. The same numbers for the automatic differentiated formulation (D) in (2.52) are 3.8(4e−4) and 4.4e−3(5e−4) seconds. These figures demonstrate the aforementioned methods' scalability for high-dimensional FBSDE systems. Finally, we point out that the OSM estimates are robust over independent runs of the algorithm as showcased by the small standard deviations in table 2.2.

2.6.3. EXAMPLE 3: SPACE-DEPENDENT DIFFUSION COEFFICIENTS

Our final example is taken from [113, 145] and it is meant to demonstrate that the conditions in assumption 2.4.1 can be substantially relaxed. The FBSDE system (2.1) is defined by the following coefficients

$$\begin{aligned}
\mu_i(t,x) &= \frac{(1+x_i^2)}{\left(2+x_i^2\right)^3}, \quad \sigma_{ij}(t,x) = \frac{1+x_i x_j}{2+x_i x_j}\delta_{ij}, \\
f(t,x,y,z) &= \frac{1}{\lambda(t+\tau)}\exp\left(-\frac{x^T x}{\lambda(t+\tau)}\right)\left[4\sum_{i=1}^{d}\frac{x_i^2(1+x_i^2)}{(2+x_i^2)^3}\right. \\
&\qquad \left. + \sum_{i=1}^{d}\frac{(1+x_i^2)^2}{(2+x_i^2)^2}\left(1-2\frac{x_i^2}{\lambda(t+\tau)}\right) - \sum_{i=1}^{d}\frac{x_i^2}{t+\tau}\right] \\
&\quad + \sqrt{\frac{1+y^2+\exp\left(-\frac{2x^T x}{\lambda(t+\tau)}\right)}{1+2y^2}}\sum_{i=1}^{d}\frac{z_i x_i}{(2+x_i^2)^2}, \qquad g(x) = \exp\left(-\frac{x^T x}{\lambda(T+\tau)}\right).
\end{aligned} \tag{2.78}$$

The analytical solutions are given by

$$y(t,x) = \exp\left(-\frac{x^T x}{\lambda(t+\tau)}\right), \quad z_j(t,x) = -\frac{1+x_j^2}{2+x_j^2}\frac{2\exp\left(-\frac{x^T x}{\lambda(t+\tau)}\right)}{\lambda(t+\tau)}x_j, \quad \gamma_{ij}(t,x) = \partial_j z_i(t,x).$$

We use $T = 10, \lambda = 10, \tau = 1$, $d = 1$ and fix $x_0 = \mathbf{1}_d$. Notice that μ and σ are both C_b^2. In conjecture with appendix 2.A, this implies that the Euler-Maruyama schemes in (2.10)

2

Table 2.2: Example 2 in (2.76). Summary of Deep BSDE estimates. Mean-squared errors are calculated over an independent sample of $M = 2^{10}$ realizations of the underlying Brownian motion. Means and standard deviations (in parentheses) obtained over 5 independent runs of the algorithm. Best estimates within one standard deviation highlighted in gray. Γ estimates from Huré et al. in [77] are obtained via automatic differentiation.

(a) $d = 1, N = 100$.

	OSM($\vartheta_y = 1/2$) (P)	(D)	Huré et al. (2020)
$\lvert\Delta\widehat{Y}_0^\pi\rvert/\lvert Y_0\rvert$	1.1e−3 (5e−4)	2e−3 (1e−3)	1.5e−3 (3e−4)
$\lvert\Delta\widehat{Z}_0^\pi\rvert/\lvert Z_0\rvert$	1.3e−4 (9e−5)	8e−5 (9e−5)	1e−3 (1e−3)
$\lvert\Delta\widehat{\Gamma}_0^\pi\rvert/\lvert\Gamma_0\rvert$	1.0e−4 (5e−5)	2e−4 (1e−4)	1.05 (7e−2)
$\max_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{Y}_n^\pi\rvert^2]$	8e−6 (2e−6)	8e−6 (3e−6)	1.1e−4 (1e−5)
$\max_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{Z}_n^\pi\rvert^2]$	8e−7 (3e−7)	1.4e−6 (6e−7)	6.4e−3 (3e−4)
$\sum_{n=0}^{N-1}\Delta t_n\widehat{\mathbb{E}}[\lvert\Delta\widehat{\Gamma}_n^\pi\rvert^2]$	8e−7 (4e−7)	2.8e−6 (9e−7)	5.5e−3 (7e−4)
runtime (s)	1.18e+3 (4e+1)	1.41e+3 (3e+1)	5.7e+2 (4e+1)

(b) $d = 50, N = 100$.

	OSM($\vartheta_y = 1/2$) (P)	(D)	Huré et al. (2020)
$\lvert\Delta\widehat{Y}_0^\pi\rvert/\lvert Y_0\rvert$	8e−4 (5e−4)	1e−3 (1e−3)	1.7e−1 (8e−2)
$\lvert\Delta\widehat{Z}_0^\pi\rvert/\lvert Z_0\rvert$	5.0e−3 (5e−4)	1.4e−2 (3e−3)	2.8e−1 (7e−2)
$\lvert\Delta\widehat{\Gamma}_0^\pi\rvert/\lvert\Gamma_0\rvert$	3.1e−2 (2e−3)	4.9e−2 (7e−3)	3.5 (1e−1)
$\max_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{Y}_n^\pi\rvert^2]$	2.7 (1e−1)	2.5 (3e−1)	7e+1 (4e+1)
$\max_n \widehat{\mathbb{E}}[\lvert\Delta\widehat{Z}_n^\pi\rvert^2]$	3.4e−2 (1e−3)	3.1e−2 (3e−3)	2.8e+2 (1e+1)
$\sum_{n=0}^{N-1}\Delta t_n\widehat{\mathbb{E}}[\lvert\Delta\widehat{\Gamma}_n^\pi\rvert^2]$	4.1e−4 (6e−5)	3.3e−3 (2e−4)	2.9 (2e−1)
runtime (s)	1.36e+3 (1e+1)	1.62e+3 (4e+1)	6.16e+2 (1e+1)

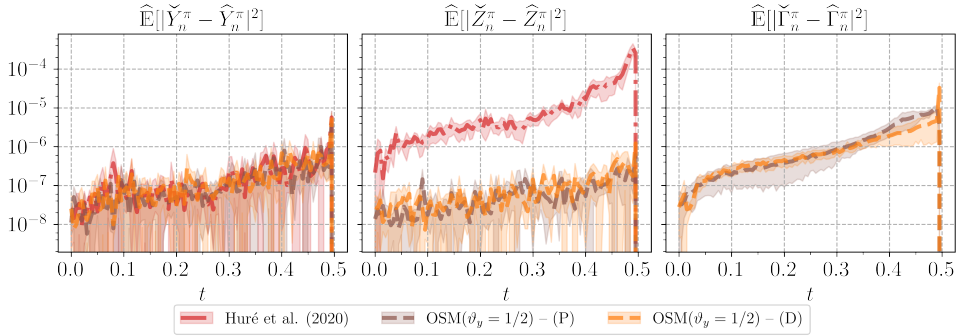

(a) Regression errors over time, $d = 1$, $N = 100$. From left to right: mean-squared regression errors of the Y, Z and Γ approximations over the discrete time window.

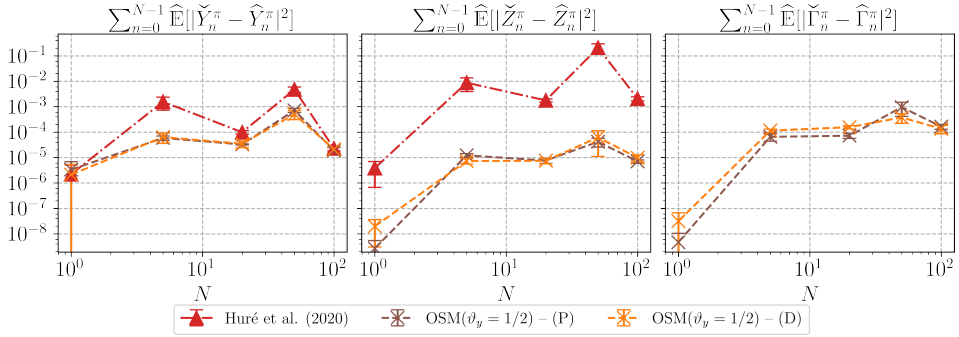

(b) Convergence of cumulative regression errors, $d = 1$. From left to right: cumulative regression errors of the Y, Z and Γ approximations over the number of time steps N.

Figure 2.2: Example 2 in (2.76). Neural network regression errors in $d = 1$. The true regression targets of (2.42) are identified by BCOS estimates. Mean-squared errors are calculated over an independent sample of $M = 2^{10}$ realizations of the underlying Brownian motion. Means and standard deviations are obtained over 5 independent runs of the algorithm.

and (2.13) have an $\mathbb{L}^2$ convergence rate of order 1/2. Additionally, by Itô's formula, the unique solution of the SDE is given by the closed form expression [113]

$$X_t = \Lambda(x_0 + \arctan(x_0) + W_t), \tag{2.79}$$

where $\Lambda : \mathbb{R} \to \mathbb{R}$ is defined implicitly $\Lambda(r) + \arctan(r) := r$ for any $r \in \mathbb{R}$, and applied element-wise. It is straightforward to check that $\Lambda \in C_b^1(\mathbb{R};\mathbb{R})$, in particular $\Lambda'(r) = \frac{1+\Lambda^2(r)}{2+\Lambda^2(r)}$ implying that Λ is a bijective. In light of the Malliavin chain rule formula in lemma 2.2.1, we then also have

$$D_s X_t = \frac{1 + \Lambda^2(x + \arctan(x) + W_t)}{2 + \Lambda^2(x + \arctan(x) + W_t)} \mathbb{1}_{s \le t}.$$

We assess the convergence of the Euler-Maruyama estimates in (2.10)–(2.13) by solving the non-linear equation in (2.79) for each realization of the Brownian motion.[6] The re-

[6] This is done by `scipy.optimize.root`'s `df-sane` algorithm which deploys the method in [92].

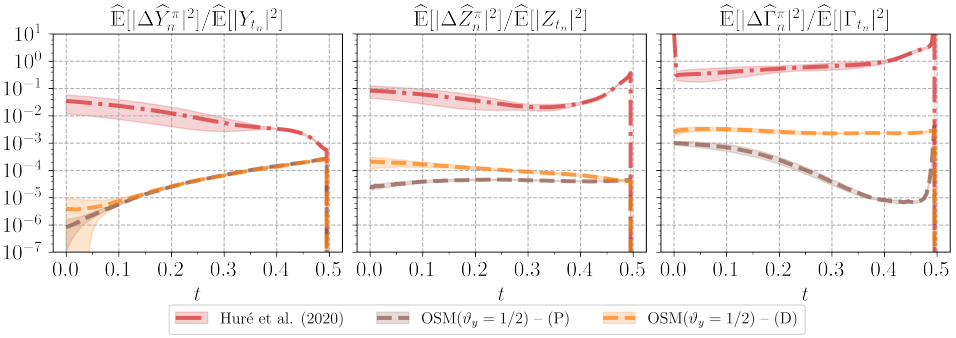

(a) Relative approximation errors over time, $d = 50$, $N = 100$. From left to right: relative mean-squared approximation errors of Y, Z and Γ over the discrete time window.

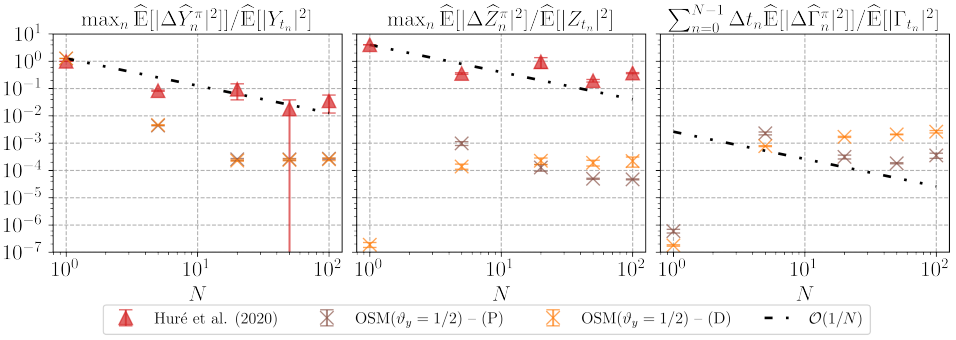

(b) Convergence of relative approximation errors, $d = 50$. From left to right: maximum relative mean-squared error of the Y, Z approximations; average relative mean-squared error of the Γ approximations.

Figure 2.3: Example 2 in (2.76). $d = 50$. Relative approximation errors. Mean-squared errors are calculated over an independent sample of $M = 2^{10}$ realizations of the underlying Brownian motion. Means and standard deviations are obtained over 5 independent runs of the algorithm. Γ estimates from Huré et al. in [77] are obtained via automatic differentiation.

sults of the numerical simulations in $d = 1$ are given in figure 2.4 for the parametrized Deep BSDE case and $\vartheta_y = 0, 1/2, 1$. We see that, in line with appendix 2.A, $D_n X_{n+1}^\pi$ inherits the convergence rate of X_n^π. The convergence rates of $(\widehat{Y}_n^\pi, \widehat{Z}_n^\pi, \widehat{\Gamma}_n^\pi)$ are of the same order as in theorem 2.5.2. The BCOS estimates and the Deep BSDE approach exhibit coinciding error figures until a magnitude of $\mathcal{O}(10^{-6})$ is reached, when the regression bias becomes apparent. Similar convergence behavior is observed in high-dimensions. The results suggest that the convergence of the OSM scheme can be extended to the non-additive noise case.

2.7. CONCLUSION

In this chapter we introduced the One Step Malliavin (OSM) scheme, a new discretization for Malliavin differentiable FBSDE systems where the control process is estimated by solving the linear BSDE driving the Malliavin derivatives of the solution pair. The

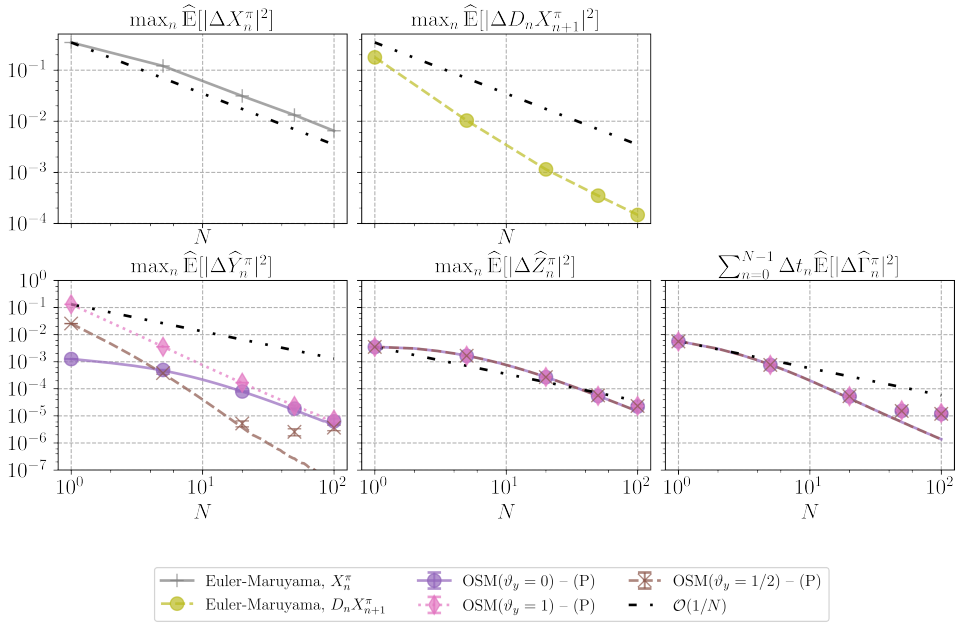

Figure 2.4: Example 3 in (2.78). Convergence of approximation errors for $d = 1$. From left to right, top to bottom: maximum mean-squared errors of Euler-Maruyama approximations of X and DX; maximum mean-squared approximation errors of Y and Z; average mean-squared approximation error of Γ. Lines correspond to BCOS estimates, scattered error bars to the means and standard deviations of 5 independent neural network regressions. The mean errors are obtained over an independent sample of $M = 2^{10}$ trajectories of the underlying Brownian motion.

main contributions can be summarized as follows. The discretization in (2.19) includes Γ estimates, linked to the Hessian matrix of the associated parabolic problem. In theorem 2.4.1 we have shown that under standard Lipschitz assumptions and additive noise in the forward diffusion, the aforementioned discrete time approximations admit to an $\mathbb{L}^2$ convergence of order 1/2. We gave two fully-implementable schemes. In case of one-dimensional problems, we extended the BCOS method [144], and gathered approximations via Fourier cosine expansions in (2.46). For high-dimensional equations, similarly to recent Deep BSDE methods [68, 77], we formulated a neural network regression Monte Carlo approach, where the corresponding processes of the solution triple are parametrized by fully-connected, feedforward neural networks. We carried out a complete regression error analysis in theorem 2.5.2 and showed that the neural network parametrizations are consistent with the discretization, in terms of regression biases controlled by the universal approximation property. We supported our theoretical findings by numerical experiments and demonstrated the accuracy and robustness of the proposed approaches for a range of high-dimensional problems. Using BCOS estimates as benchmarks for one-dimensional equations, we empirically assessed the regression errors induced by stochastic gradient descent. Our findings with the Deep BSDE approach showcase accurate approximations for each process in (2.42), and in particular

exhibit significantly improved approximations of the Z process for heavily control dependent equations.

2.A. CONVERGENCE OF $D_n X^\pi_{n+1}$

We show the convergence of $D_n X^\pi_{n+1}$ estimates of the Euler-Maruyama discretization (2.13) under the assumptions

($\widetilde{\mathbf{A}}_1^{\sigma,\mu}$) σ is uniformly bounded;

($\widetilde{\mathbf{A}}_2^{\sigma,\mu}$) $\mu \in C_b^{0,1}(\mathbb{R}^{d\times 1};\mathbb{R})$, $\sigma \in C_b^{0,1}(\mathbb{R}^{d\times 1};\mathbb{R}^{d\times d})$. In particular both of them are Lipschitz continuous in x.

From the estimation (2.13) and the linear SDE of the Malliavin derivative in (2.9c) – using the inequality $(a+b+c)^2 \le 3(a^2+b^2+c^2)$, on top of the $L^2([0,T];\mathbb{R}^{d\times d})$ Cauchy-Schwarz inequality and Itô's isometry –, it follows

$$\begin{aligned}
&\mathbb{E}\left[\left|D_{t_n}X_{t_{n+1}} - D_n X^\pi_{n+1}\right|^2\right] \\
&\qquad \le 3\mathbb{E}\left[\left|\sigma(t_n,X_{t_n}) - \sigma(t_n,X^\pi_n)\right|^2\right] \\
&\qquad\quad + 3\Delta t_n \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|\nabla_x\mu(r,X_r)D_{t_n}X_r - \nabla_x\mu(t_n,X^\pi_n)\sigma(t_n,X^\pi_n)\right|^2 \mathrm{d}r\right] \\
&\qquad\quad + 3\mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left|\nabla_x\sigma(r,X_r)D_{t_n}X_r - \nabla_x\sigma(t_n,X^\pi_n)\sigma(t_n,X^\pi_n)\right|^2 \mathrm{d}r\right].
\end{aligned}$$

Bounded continuous differentiability in assumption ($\widetilde{\mathbf{A}}_2^{\sigma,\mu}$), in particular, implies Lipschitz continuity. Furthermore, by the uniform boundedness of the diffusion coefficient and the mean-squared continuity of $D_{t_n}X$ in (2.6), we gather

$$\mathbb{E}\left[\left|D_{t_n}X_{t_{n+1}} - D_n X^\pi_{n+1}\right|^2\right] \le 3L_\sigma^2 \mathbb{E}\left[\left|X_{t_n} - X^\pi_n\right|^2\right] + C\Delta t_n,$$

for any $\Delta t_n < 1$. Then, due to the discretization error of the Euler-Maruyama estimates given by (2.11), we conclude $\limsup_{|\pi|\to 0} \frac{1}{|\pi|}\mathbb{E}\left[\left|D_{t_n}X_{t_{n+1}} - D_n X^\pi_{n+1}\right|^2\right] < \infty$.

2.B. INTEGRATION BY PARTS FORMULAS

For the formula in (2.44) we refer to [144, A.1]. In order to prove (2.45), let $v : [0,T]\times\mathbb{R} \to \mathbb{R}$ and consider

$$\mathbb{E}_n^x\left[v(t_{n+1},X^\pi_{n+1}(\Delta W_n))\Delta W_n^2\right] = \mathbb{E}_n^x\left[\frac{1}{\sqrt{2\pi\Delta t_n}}\int_{\mathbb{R}} v(t_{n+1},X^\pi_{n+1}(\nu))\nu^2 e^{-\frac{1}{2\Delta t_n}\nu^2}\mathrm{d}\nu\right],$$

with the Euler-Maruyama approximations $X^\pi_{n+1}(\Delta W_n) = x + \mu(t_n,x)\Delta t_n + \sigma(t_n,x)\Delta W_n$. For a sufficiently smooth v, integration by parts implies

$$\begin{aligned}\mathbb{E}^x_n\Big[\frac{1}{\sqrt{2\pi\Delta t_n}}\int_{\mathbb{R}} v(t_{n+1},X^\pi_{n+1}(\nu))\nu^2 e^{-\frac{1}{2\Delta t_n}\nu^2}\mathrm{d}\nu\Big]\\ = \frac{1}{\sqrt{2\pi\Delta t_n}}\mathbb{E}^x_n\Big[-\Delta t_n\Big[\nu v(t_{n+1},X^\pi_{n+1}(\nu))e^{-\nu^2/(2\Delta t_n)}\Big]^{+\infty}_{-\infty}\\ + \Delta t_n\int_{\mathbb{R}} v(t_{n+1},X^\pi_{n+1}(\nu))e^{-\frac{1}{2\Delta t_n}\nu^2}\mathrm{d}\nu\\ + \Delta t_n\sigma(t_n,x)\int_{\mathbb{R}}\partial_x v(t_{n+1},X^\pi_{n+1}(\nu))\nu e^{-\frac{1}{2\Delta t_n}\nu^2}\mathrm{d}\nu\Big].\end{aligned}$$

For a v with sufficient radial decay, we therefore conclude that

$$\mathbb{E}^x_n\left[v(t_{n+1},X^\pi_{n+1})\Delta W_n^2\right] = \Delta t_n\mathbb{E}^x_n\left[v(t_{n+1},X^\pi_{n+1})\right] + \Delta t_n^2\sigma^2(t_n,x)\mathbb{E}^x_n\left[\partial^2_{xx}v(t_{n+1},X^\pi_{n+1})\right],$$

by the estimate in (2.44).

Thereupon, given a cosine expansion approximation $v(t_{n+1},\rho) \approx \sum'^{K-1}_{k=0}\mathcal{V}_k(t_{n+1})\cos\big(k\pi\frac{\rho-a}{b-a}\big)$, the corresponding spatial derivative approximations are given by $\partial_x v(t_{n+1},\rho) \approx \sum'^{K-1}_{k=0} -\mathcal{V}_k(t_{n+1})\frac{k\pi}{b-a}\sin\big(k\pi\frac{\rho-a}{b-a}\big)$, $\partial^2_{xx}v(t_{n+1},\rho) \approx \sum'^{K-1}_{k=0} -\mathcal{V}_k(t_{n+1})\Big(\frac{k\pi}{b-a}\Big)^2\cos\big(k\pi\frac{\rho-a}{b-a}\big)$. Then (2.44)–(2.45) follow from the expressions $\mathbb{E}^x_n\Big[\sin\Big(k\pi\frac{X^\pi_{n+1}-a}{b-a}\Big)\Big] = \mathrm{Im}\{\Phi(k|x)\}$, $\mathbb{E}^x_n\Big[\cos\Big(k\pi\frac{X^\pi_{n+1}-a}{b-a}\Big)\Big] = \mathrm{Re}\{\Phi(k|x)\}$, where $\Phi(k|x)$ is defined as in section 2.5.1.

Multi-dimensional extensions. In case the underlying forward process is an $\mathbb{R}^{d\times 1}$-dimensional Brownian motion, the following extension can be given. Let $v : [0,T]\times\mathbb{R}^{d\times 1} \to \mathbb{R}$ be a scalar-valued. Then reasoning similar to [144, A.1] shows that $\mathbb{E}_n\left[(\Delta W_n)_{i1}\, v(t_{n+1},X^\pi_{n+1})\right] = \sum_{k=1}^d \Delta t_n\mathbb{E}_n\left[\partial_k v(t_{n+1},X^\pi_{n+1})\right]\big(\sigma(t_n,X^\pi_n)\big)_{ki}$. In matrix notation

$$\big(\mathbb{E}_n\left[\Delta W_n v(t_{n+1},X^\pi_{n+1})\right]\big)^T = \Delta t_n\mathbb{E}_n\left[\nabla_x v(t_{n+1},X^\pi_{n+1})\right]\sigma(t_n,X^\pi_n).$$

Alternatively, for a vector-valued mapping $\psi : [0,T]\times\mathbb{R}^{d\times 1} \to \mathbb{R}^{1\times d}$, similar arguments give the following, component-wise formula $\mathbb{E}_n\Big[(\Delta W_n)_{i1}\big(\psi(t_{n+1},X_{n+1})\big)_{1j}\Big] = \sum_{k=1}^d \Delta t_n\mathbb{E}_n\Big[\partial_k\big(\psi(t_{n+1},X^\pi_{n+1})\big)_{1j}\Big]\big(\sigma(t_n,X^\pi_n)\big)_{ki}$. In matrix notation

$$\big(\mathbb{E}_n\left[\Delta W_n\psi(t_{n+1},X^\pi_{n+1})\right]\big)^T = \Delta t_n\mathbb{E}_n\left[\nabla_x\psi(t_{n+1},X^\pi_{n+1})\right]\sigma(t_n,X^\pi_n), \tag{2.80}$$

where $\nabla_x\psi$ is the Jacobian matrix of ψ.

2.C. BCOS ESTIMATES

Let us fix $d = 1$. The BCOS approximations of the OSM scheme in (2.46) can be derived as follows. Using the definition in (2.47) and the Euler-Maruyama estimates in (2.13), the

Γ estimates in (2.42a) can be written according to

$$\begin{aligned} D_n\breve{Z}_n^\pi = \breve{\gamma}_n^\pi(x)\sigma(t_n,x) = {} & \frac{1}{\Delta t_n}\sigma(t_n,x)\left(1+\Delta t_n\partial_x\mu(t_n,x)\right)\mathbb{E}_n^x\left[\Delta W_n\, w_{n+1}^\pi(\widehat{\mathbf{X}}_{n+1}^\pi)\right] \\ & + \frac{1}{\Delta t_n}\sigma(t_n,x)\partial_x\sigma(t_n,x)\mathbb{E}_n^x\left[\Delta W_n^2\, w_{n+1}^\pi(\widehat{\mathbf{X}}_{n+1}^\pi)\right] \\ & + \mathbb{E}_n^x\left[\Delta W_n\partial_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)\right]\breve{\gamma}_n^\pi(x)\sigma(t_n,x). \end{aligned}$$

A cosine expansion approximation for $w_{n+1}^\pi(\widehat{\mathbf{X}}_{n+1}^\pi)$ and $\partial_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^\pi)$ can be obtained by means of DCT, yielding approximations $\{\widehat{\mathcal{W}_k}(t_{n+1})\}_{k=0,\dots,K-1}$, $\{\widehat{\mathscr{F}_k^z}(t_{n+1})\}_{k=0,\dots,K-1}$ respectively. Consequently, plugging these approximations combined with the integration by parts formulas in (2.44)–(2.45), in the estimate above yields

$$\begin{aligned} \widehat{\gamma}_n^\pi(x)\sigma(t_n,x) = {} & -\sigma^2(t_n,x)(1+\partial_x\mu(t_n,x)\Delta t_n)\sum_{k=0}^{K-1}{}'\frac{k\pi}{b-a}\widehat{\mathcal{W}_k}(t_{n+1})\,\mathrm{Im}\{\Phi(k|x)\} \\ & + \sigma(t_n,x)\partial_x\sigma(t_n,x)\sum_{k=0}^{K-1}{}'\widehat{\mathcal{W}_k}(t_{n+1})\,\mathrm{Re}\{\Phi(k|x)\} \\ & - \Delta t_n\sigma^3(t_n,x)\partial_x\sigma(t_n,x)\sum_{k=0}^{K-1}{}'\left(\frac{k\pi}{b-a}\right)^2\widehat{\mathcal{W}_k}(t_{n+1})\,\mathrm{Re}\{\Phi(k|x)\} \\ & - \widehat{\gamma}_n^\pi(x)\Delta t_n\sigma^2(t_n,x)\sum_{k=0}^{K-1}{}'\frac{k\pi}{b-a}\widehat{\mathscr{F}_k^z}(t_{n+1})\,\mathrm{Im}\{\Phi(k|x)\}. \end{aligned}$$

The approximation $D_n\widehat{Z}_n^\pi = \widehat{\Gamma}_n^\pi\sigma(t_n,X_n^\pi)$ subsequently follows. The coefficients $\{\widehat{\mathscr{D}\mathscr{Z}}_k(t_{n+1})\}_{k=0,\dots,K-1}$ are calculated by DCT and subsequently plugged into the approximations of the Z process, which follows analogously using the formulas in (2.43)–(2.44). The approximation of the Y process in (2.42c) is identical to [144] and therefore omitted.

Algorithm 1 One-Step Malliavin Algorithm (OSM)

Require: $\pi(N), \vartheta_y \in [0,1]$ ▷ discretization parameters

Require: $B \in \mathbb{N}^+$, $I \in \mathbb{N}$, $\eta : \mathbb{N} \to \mathbb{R}$ ▷ training parameters

return $\{(\widehat{Y}^\pi_n, \widehat{Z}^\pi_n, \widehat{\Gamma}^\pi_n)\}_{n=0,\dots,N}$ ▷ discrete time approximations over π

$\widehat{Y}^\pi_N \leftarrow g(X^\pi_N)$, $\widehat{Z}^\pi_N \leftarrow \nabla_x g(X^\pi_N)\sigma(t_N, X^\pi_N)$, $\widehat{\Gamma}^\pi_N \leftarrow \nabla_x(\nabla_x g\sigma)(t_N, X^\pi_N)$ ▷ collect terminal condition

$\varphi(\cdot|\theta^y) : \mathbb{R}^{d\times 1} \to \mathbb{R}$, $\psi(\cdot|\theta^z) : \mathbb{R}^{d\times 1} \to \mathbb{R}^{1\times d}$, $\chi(\cdot|\theta^\gamma) : \mathbb{R}^{d\times 1} \to \mathbb{R}^{d\times d}$ ▷ neural network parametrizations

for $n = N-1, \dots, 0$ **do**

  if $n = N-1$ **then**

    $\theta^{z,(0)}, \theta^{y,(0)}$ ▷ initialize parameter sets, according to [59]

  else

    $\theta^{z,(0)} \leftarrow \widehat{\theta}^z_{n+1}, \quad \theta^{y,(0)} \leftarrow \widehat{\theta}^y_{n+1}$ ▷ transfer learning initialization

  end if

  Solve (2.42b)–(2.42a)**.**

  for $i = 0, \dots, I-1$ **do**

    $\{\{X^\pi_m(b)\}_{0\le m\le N}\}_{b=1}^B$ ▷ Euler-Maruyama simulations by (2.10)

    $\{D_n X^\pi_{n+1}(b)\}_{b=1}^B$ ▷ Euler-Maruyama approximations by (2.13)

    calculate empirical loss of (2.51) or (2.52)

$$\begin{aligned}\widehat{\mathscr{L}}^{z,\gamma}_n(\theta^{z,(i)}, \theta^{\gamma,(i)}) = \frac{1}{B}\sum_{b=1}^B |&(1 + \Delta t_n \nabla_y f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}(b)))D_n \widehat{Y}^\pi_{n+1}(b) \\ &+ \Delta t_n \nabla_x f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}(b)) D_n X^\pi_{n+1}(b) - \psi(X^\pi_n(b)|\theta^{z,(i)}) \\ &+ \Delta t_n \nabla_z f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}(b))\chi(X^\pi_n(b)|\theta^{\gamma,(i)})\sigma(t_n, X^\pi_n) \\ &- ((\chi(X^\pi_n(b)|\theta^{\gamma,(i)})\sigma(t_n, X^\pi_n(b)))^T \Delta W_n(b))^T|^2\end{aligned}$$

    $(\theta^{z,(i+1)}, \theta^{\gamma,(i+1)}) \leftarrow (\theta^{z,(i)}, \theta^{\gamma,(i)}) - \eta(i)\nabla_{(\theta^z,\theta^\gamma)}\widehat{\mathscr{L}}^z_n(\theta^{z,(i)}, \theta^{\gamma,(i)})$ ▷ SGD update

  end for

  $\widehat{\theta}^z_n \leftarrow \theta^{z,(I)}, \widehat{\theta}^\gamma_n \leftarrow \theta^{\gamma,(I)}$ ▷ collect optimal parameter estimations

  $\widehat{z}^\pi_n(\cdot) \leftarrow \psi(\cdot|\widehat{\theta}^z_n), \quad \widehat{\gamma}^\pi_n(\cdot) \leftarrow \chi(\cdot|\widehat{\theta}^\gamma_n)$ ▷ collect approximations $\widehat{Z}^\pi_n, \widehat{\Gamma}^\pi_n$

  Solve (2.42c)**.**

  for $i = 0, \dots, I-1$ **do**

    $\{\{X^\pi_m(b)\}_{0\le m\le N}\}_{b=1}^B$ ▷ Euler-Maruyama simulations by (2.10)

    calculate empirical loss of (2.53)

$$\begin{aligned}\widehat{\mathscr{L}}^y_n(\theta^{y,(i)}) = \frac{1}{B}\sum_{b=1}^B |&\widehat{Y}^\pi_{n+1}(b) + (1-\vartheta_y)\Delta t_n f(t_{n+1}, \widehat{\mathbf{X}}^\pi_{n+1}(b)) - \varphi(X^\pi_n(b)|\theta^{y,(i)}) \\ &+ \vartheta_y \Delta t_n f(t_n, X^\pi_n(b), \varphi(X^\pi_n(b)|\theta^{y,(i)}), \widehat{Z}^\pi_n(b)) - \widehat{Z}^\pi_n(b)\Delta W_n(b)|^2\end{aligned}$$

    $\theta^{y,(i+1)} \leftarrow \theta^{y,(i)} - \eta(i)\nabla_\theta \widehat{\mathscr{L}}^y_n(\theta^{y,(i)})$ ▷ stochastic gradient descent step

  end for

  $\widehat{\theta}^y_n \leftarrow \theta^{y,(I)}$ ▷ collect optimal parameter estimations

  $\widehat{y}^\pi_n(\cdot) \leftarrow \varphi(\cdot|\widehat{\theta}^y_n)$ ▷ collect approximations $\widehat{Y}^\pi_n$

end for

3

A Deep BSDE approach for the simultaneous pricing and delta-gamma hedging of large portfolios of high-dimensional Bermudan options

Drink de tranen op je hand
Zwijg ervan
Erf de ogen van je kind
Kijk erdoor[1]

Spinvis (Kom terug)

This chapter is based on the preprint [116]: B. Negyesi and C. W. Oosterlee. *A deep BSDE approach for the simultaneous pricing and delta-gamma hedging of large portfolios consisting of high-dimensional multi-asset Bermudan options.* 2025. DOI: 10.48550/arXiv.2502.11706

[1] Drink the tears from your hand
Keep quiet about it
Inherit your child's eyes
Look through them

Spinvis (Come back)

3.1. INTRODUCTION

In this chapter we are concerned with the challenging hedging problem of large portfolios of high-dimensional options with early-exercise features. Consider an investor who is in possession of a portfolio consisting of J many options whose values depend on a set of common risk factors $\{X_t\}_{0\leq t\leq T}$ which can be decomposed into a set of tradeable underlying assets $\{S_t = (S_t^1,\dots,S_t^m)\}_{0\leq t\leq T}$ and some other non-tradeable component $\{\nu_t = (\nu_t^1,\dots,\nu_t^{d-m})\}_{0\leq t\leq T}$, and they together form an Itô process $X_t := (S_t,\nu_t)$ solving the following stochastic differential equation (SDE)

$$X_t = X_0 + \int_0^t \mu(s,X_s)\mathrm{d}s + \int_0^t \sigma(s,X_s)\mathrm{d}W_s, \tag{3.1}$$

where $\mu : [0,T]\times\mathbb{R}^d \to \mathbb{R}^d$, $\sigma : [0,T]\times\mathbb{R}^d \to \mathbb{R}^{d\times d}$, and $\{W_t\}_{0\leq t\leq T}$ a d-dimensional Brownian motion in an appropriate probability space.

Fixing some time horizon $0 < T < \infty$, let $\mathscr{R}^j \subseteq [0,T]$ denote the set of early exercise opportunities for each $j = 1,\dots,J$. We are concerned with the Markovian framework, i.e. options whose prices are deterministic mappings of the underlying risk factors at each point in time, and we put $v^j : [0,T]\times\mathbb{R}^d \to \mathbb{R}$, $j = 1,\dots,J$ for each of these functions. The investor's objective is to insure her positions against random movements in the underlyings and she achieves this through a hedging strategy. In particular, she constructs a delta hedging replicating portfolio consisting of a short position in all options; long position in all underlyings and a deposit in a bank account

$$\mathrm{d}P_t^\Delta = -\sum_{j=1}^J \mathrm{d}v^j(t,X_t) + \sum_{i=1}^m \alpha_t^i(\mathrm{d}S_t^i + q_t^i S_t^i \mathrm{d}t) + \mathrm{d}B_t, \quad P_0^\Delta = 0, \tag{3.2}$$

where we allow each underlying S_t^i to pay off dividends continuously with a rate q_t^i at time t. It is well-known that the variance-minimizing first-order conditions $\partial P_t^\Delta/\partial S_t^i = 0$ result in the optimal hedging weights

$$\alpha_t^i = \sum_{j=1}^J \frac{\partial v^j}{\partial S_t^i}(t,X_t), \quad i = 1,\dots,m. \tag{3.3}$$

In particular, (3.2) and (3.3) together with Itô's lemma imply that given a machinery which yields simultaneous option prices and deltas, the investor can perfectly offset her exposure in the underlyings, at least in the continuous, complete market setting.

The discussion above motivates to cast our problem into the framework of (decoupled) forward-backward stochastic differential equations (FBSDE). In fact, it is classically known, see e.g. [126, 161], that in absence of early exercise rights $\mathscr{R}^j = \{0,T\}$, the jth option is intimately related to the following standard (Markovian) BSDE

$$Y_t^j = g^j(X_T) + \int_t^T f^j(s,X_s,Y_s^j,Z_s^j)\mathrm{d}s - \int_t^T Z_s^j \mathrm{d}W_s, \tag{3.4}$$

where $g^j : \mathbb{R}^d \to \mathbb{R}$ denotes the payoff and $f^j : [0,T]\times\mathbb{R}^d\times\mathbb{R}\times\mathbb{R}^{1\times d} \to \mathbb{R}$ the driver. Namely, non-linear extensions to the Feynman-Kac relations establish the relations

$$Y_t^j = v^j(t,X_t), \quad Z_t^j = \nabla_x v^j(t,X_t)\sigma(t,X_t), \tag{3.5}$$

in an almost sure sense. Comparing (3.3) with (3.5), one can conclude that solving the BSDE associated to the option is, in a sense, equivalent with the task of (delta-)hedging. Similar relations hold in case of early-exercise rights $\mathscr{R}^j \setminus \{0, T\} \neq \emptyset$, see section 3.2 below. Over the last three decades a vast literature has been developed dealing with the numerical resolution of different types of BSDEs, see e.g. [24, 13, 25, 102, 63] and the references therein.

However, whenever the aforementioned portfolio (3.2) is high-dimensional, i.e. d or J is large, one either has to deal with a high-dimensional BSDE (3.4), or a large number of equations simultaneously, potentially both. This makes classical numerical methods intractable in the context of this work, as they all suffer from the curse of dimensionality. In recent years, initiated by the pioneering paper [68], a rapidly growing research line has been developed by the numerical analysis community, where BSDEs of the type (3.4) are approximated in regression Monte Carlo frameworks using deep neural networks to parameterize the (Markovian) solution pair of (3.4). Without the sake of completeness we mention [77, 58, 12, 28, 32]. These methods have shown remarkable empirical results tackling the numerical solution of (3.4), and by now some results are also known about their convergence properties up to universal approximation type, see [58, 69, 114]. Moreover, they have successfully been applied in the context of hedging single options, see [32, 12, 60].

However, all these aforementioned methods solely focus on solving the hedging problem explained by (3.2). Nonetheless, whenever rebalancing is only done over a finite set of dates in time, delta hedging does not achieve a perfect replication and due to the discrete time approximations, the corresponding portfolio entails risk. In particular, in a volatile economic climate, corresponding to large volatilities in the diffusion component of (3.1), the deltas on the right hand side of (3.3) change rapidly and the corresponding discrete replication error of (3.2) also grows accordingly. In order to mitigate the effect of fluctuating deltas, one can impose additional second-order constraints on top of (3.3), which effectively set the second order sensitivities, Gammas, of the accordingly constructed replicating portfolio to zero, in terms of the underlying risk factors. Doing so, one encounters two additional challenges. First, as assets themselves have vanishing gammas, in order to be able to formulate the corresponding second order conditions, one needs to augment the replicating portfolio with *gamma hedging instruments*, whose prices and Greeks are available at all points in time. Second, the resulting second order conditions involve appropriate second order sensitivities of the underlying options that are meant to be hedged. This implies additional modelling error, as in order to effectively carry out the gamma hedging strategy, the investor does not merely have to efficiently model the underlying options' prices and deltas, but also their *gammas*, even in the high-dimensional setting, consisting of many risk factors. For details, we refer to section 3.3 below.

The main objective of the present chapter is to develop a deep BSDE methodology which efficiently tackles the aforementioned challenges in the high-dimensional portfolio framework. In fact, motivated by the ideas in [114], one can derive an additional vector-valued, linear BSDE related to (3.4), whose solution pair involves a *matrix-valued* process corresponding to second-order sensitivities of the underlying option. Consequently, solving this additional BSDE together with (3.4) naturally extends the Feynman-

Kac relations in (3.5), and results in a *triple* of stochastic processes, which coincide with option prices, deltas and gammas, respectively. Given robust and efficient numerical approximations of this stochastic triple, one can mitigate the additional modelling errors arising in delta-gamma hedging, and assess the accumulating discrete replication errors in the sole delta-hedging framework.

Our main contributions are as follows. We propose a deep BSDE methodology for the portfolio hedging problem outlined above. Therefore, as a side result, we first extend the application of deep BSDE methods to the context of delta-hedging large portfolios, instead of only single options as in [32, 12, 60], in the complete Bermudan setting. This is done by casting the method of [77] into the vector-valued BSDE framework, and thus simultaneously solving all J options' pricing and delta-hedging problems. Thereafter, and most importantly, in order to reduce the discrete replication error of the delta-hedging portfolio, we propose a Gamma hedging strategy, which on top of the first-order conditions in (3.3), also aims to offset second-order terms in the portfolio's value by imposing second-order conditions, depending on second-order sensitivities, i.e. *Gammas* of the option's value. Such hedging strategies result in less frequent rebalancing and more accurate replication. However, by doing so, the investor exposes herself to additional model risk, as the Gammas must accurately be approximated in the numerical setting. In order to address this, we use recent results on the One Step Malliavin scheme from chapter 2, first proposed in [114] for standard BSDEs and later extended to discretely reflected equations, and provide a robust and accurate, fully-implementable deep BSDE method for the simultaneous delta-gamma-hedging of large portfolios. We demonstrate that this novel approach may provide a significant improvement to standard delta-hedging strategies in the discrete time framework, whenever the underlyings' deltas are highly volatile, resulting in less frequent rebalancing, and sharper Profit-and-Loss (PnL) distributions.

The chapter is organized as follows. Section 3.2 gives a short summary of the necessary theoretical basis by establishing the connections between Bermudan options and discretely reflected FBSDEs. In section 3.3, we present the delta-gamma hedging strategies and their corresponding first- and second-order conditions. Thereafter, we apply the discretizations in [114] in the vector-valued portfolio framework, and explain how the resulting approximations are applicable in a portfolio hedging context. In section 3.4 we explain how the previous works from [77, 114] extend to vector-valued equations and can be used to approximate the collection of BSDEs corresponding to problem (3.2) in a Deep BSDE approach. Finally, we demonstrate the accuracy and robustness of these strategies by numerical experiments performed on high-dimensional portfolios in section 3.5.

3.2. BERMUDAN OPTIONS AS DISCRETELY REFLECTED FBSDES

We fix $0 \leq T < \infty$ and let $J, j, d, m, k \in \mathbb{N}_+$. Throughout the chapter we are working on a filtered probability space $(\Omega, \mathscr{F}, \mathbb{P}, \{\mathscr{F}_t\}_{t\in[0,T]})$ with $\mathscr{F} = \mathscr{F}_T$, where $\mathbb{F}$ is the natural filtration generated by a d-dimensional Brownian motion $\{W_t\}_{t\in[0,T]}$, augmented by the usual $\mathbb{P}$-null sets. In what follows, all equalities concerning $\mathscr{F}_t$ measureable random variables are meant in the $\mathbb{P}$ almost sure sense, and all expectations are taken under $\mathbb{P}$, unless otherwise stated. As usual, we put $|x| := \mathrm{tr}\{x^\top x\}$, for any $x \in \mathbb{R}^{j\times d}$, and remark

that this coincides with the Euclidean norm in case of scalars and vectors. We define $\mathbb{H}^p(\mathbb{R}^{j\times d})$ to be the space of $\mathbb{R}^{j\times d}$ valued, progressively measurable stochastic processes such $Z \in \mathbb{H}^p(\mathbb{R}^{j\times d})$: $\mathbb{E}[(\int_0^T |Z_t|^2 \mathrm{d}t)^{p/2}] < \infty$. Similarly, $\mathbb{S}^p(\mathbb{R}^{j\times d}) \subset \mathbb{H}^p(\mathbb{R}^{j\times d})$, for which in addition $Y \in \mathbb{S}^p(\mathbb{R}^{j\times d})$ is also continuous and admits to $\mathbb{E}[\sup_{t\in[0,T]} |Y_t|^p] < \infty$. In what follows, for any multivariate $f : [0,T]\times\mathbb{R}^d \to \mathbb{R}^j$ function we set $\nabla_x f$ to be the Jacobian matrix taking values in $\mathbb{R}^{j\times d}$. In particular, for a scalar valued function $v : [0,T]\times\mathbb{R}^d \to \mathbb{R}$, we use $\partial_i v$ to denote the i'th partial derivative in space, and $\partial^2_{ji} v$ for the corresponding element in the Hessian matrix. Given a time partition $\mathcal{N} := \{0 = t_0 < t_1 < \cdots < t_N = T\}$ we set $\mathbb{E}_n[\cdot] := \mathbb{E}[\cdot|\mathscr{F}_{t_n}]$.

With the above notation at hand, we can formulate discretely reflected BSDEs. Heuristically speaking, a reflected BSDE is a generalization of (3.4) such that the solution is also forced to stay above a (Markovian) lower boundary process. The forcing is referred to as *reflection*. Discretely reflected BSDEs are special cases of reflected BSDEs, see e.g. [44, 161], where reflection can only occur over a finite set of points in time $\mathscr{R}^j = \{r^j_i, i = 0,\dots,R^j | r^j_0 = 0, r^j_{R^j} = T\}$. The solution to a discretely reflected BSDE indexed by j is a pair of stochastic processes $(Y^j, Z^j) \in \mathbb{S}^2(\mathbb{R}) \times \mathbb{H}^2(\mathbb{R}^{1\times d_j})$ such that

$$\begin{aligned} Y^j_T &= \widetilde{Y}^j_T := g^j(X_T), \\ \widetilde{Y}^j_t &= Y^j_{\underline{r}^j_t} + \int_t^{\overline{r}^j_t} f^j(s, X_s, \widetilde{Y}^j_s, Z^j_s)\mathrm{d}s - \int_t^{\overline{r}^j_t} Z^j_s \mathrm{d}W_s, \quad t \in [0,T], \\ Y^j_t &:= \widetilde{Y}^j_t + \mathbf{1}_{t\in\mathscr{R}^j\setminus\{0,T\}} \mathbf{1}_{l^j(X_t) > \widetilde{Y}^j_t} [l^j(X_t) - \widetilde{Y}^j_t] =: \mathfrak{R}^j_y(t, X_t, \widetilde{Y}^j_t), \end{aligned} \tag{3.6}$$

where $\underline{r}^j_t := \sup\{r \in \mathscr{R}^j : r \le t\}$, $\overline{r}^j_t := \inf\{r \in \mathscr{R}^j : r > t\}$ and $g^j, l^j : \mathbb{R}^{d_j} \to \mathbb{R}$ are deterministic mappings corresponding to the terminal condition and (Markovian) lower boundary process, respectively. Equation (3.6) together with (3.1) forms a discretely reflected FBSDE system which, under appropriate assumptions, admits a unique solution triple, see e.g. [44]. Reflected BSDEs are inherently related to second-order, semi-linear parabolic, free-boundary PDEs. In fact, under suitable assumptions, they retain a Feynman-Kac type relation similar to (3.5) in the standard case, see [44]. From the financial mathematics perspective, the main relevance of such equations is that they are a natural model for optimal stopping problems, such as Bermudan options, where $\mathscr{R}^j$ corresponds to the set of early-exercise dates, $g^j \equiv l^j$ to the instantaneous payoff, $\widetilde{Y}^j_t, Y^j_t, Z^j_t$ to the continuation value, option price and option delta at time t, respectively. Notice that (3.6) includes the standard, Markovian BSDE framework associated with European options by letting $\mathscr{R}^j = \{0, T\}$. Furthermore, from the numerical point of view, a suitable discretization of (3.6) is an approximation of American options in the asymptotic $R^j \to \infty$. Hence, in what follows we refer to all of the options in (3.2) as solutions to a discretely reflected FBSDE.

Similar to (3.1), we define $g := (g^1;\dots;g^J) : \mathbb{R}^d \to \mathbb{R}^J$, $f := (f^1;\dots;f^J) : [0,T]\times\mathbb{R}^d\times\mathbb{R}^J\times\mathbb{R}^{J\times d} \to \mathbb{R}^J$, $\widetilde{Y}_t := (\widetilde{Y}^1_t;\dots;\widetilde{Y}^J_t)$, $Y_t := (Y^1_t;\dots;Y^J_t)$ and $Z_t := (Z^1_t;\dots;Z^J_t)$ by the row-wise concatenated solutions of (3.6) for each $j = 1,\dots,J$. Therefore, $(Y,Z) \in \mathbb{S}^2(\mathbb{R}^J)\times\mathbb{H}^2(\mathbb{R}^{J\times d})$ satisfy a *collection* of discretely reflected BSDEs. It is important to notice that even though in this formulation we cast the problem of J options into the framework of vector-valued, discretely reflected BSDEs, this indeed is of mere formal convenience, and the corresponding system is only a collection not a system, i.e. the solution pair (Y^i, Z^i) does

not enter the dynamics of (Y^j, Z^j) for $i \neq j$. This is of fundamental importance, since the well-posedness of multi-dimensional reflected BSDEs remains to be a challenging open problem due to the lack of comparison principles in the vector-valued setting, see [161] and the references therein. Nevertheless, in the context of our work, there is no cross-dependence between equations of the type (3.6), and the resulting collections can safely be treated without the aforementioned theoretical obstacles. Henceforth the system of discretely reflected BSDEs simultaneously representing all options in (3.2) reads as follows

$$\begin{aligned} Y_T &= \widetilde{Y}_T := (g^1(X_T), \dots, g^J(X_T)), \\ \widetilde{Y}_t &= (Y^1_{\overline{r}^1_t}; \dots; Y^J_{\overline{r}^J_t}) + \Big(\int_t^{\overline{r}^1_t} f^1(s, X_s, \widetilde{Y}^1_s, Z^1_s) \mathrm{d}s; \dots; \int_t^{\overline{r}^J_t} f^J(s, X_s, \widetilde{Y}^J_s, Z^J_s) \mathrm{d}s \Big) \\ &\qquad - \Big(\int_t^{\overline{r}^1_t} Z^1_s \mathrm{d}W_s; \dots; \int_t^{\overline{r}^J_t} Z^J_s \mathrm{d}W_s \Big), \quad t \in [\underline{r}^j_t, \overline{r}^j_t), \\ Y_t &:= (\mathfrak{R}^1_y(t, X_t, \widetilde{Y}^1_t; \dots; \mathfrak{R}^J_y(t, X_t, \widetilde{Y}^J_t) =: \mathfrak{R}(t, X_t, \widetilde{Y}_t). \end{aligned} \tag{3.7}$$

The standard Euler discretization of (3.6) is done in two steps. First, in case no analytical solution is available, one needs to approximate the forward diffusion in (3.1) via suitable discrete time approximations, e.g. an Euler-Maruyama scheme such as below

$$X^\pi_0 = x_0, \quad X^\pi_{n+1} = X^\pi_n + \mu(t_n, X^\pi_n)\Delta t_n + \sigma(t_n, X^\pi_n)\Delta W_n, \quad \text{for} \quad n = 0, \dots, N-1. \tag{3.8}$$

Thereafter, one gathers discrete time approximations to (3.6) through a backward recursion of conditional expectations, over a discrete time partition $0 = t_0 < t_1 < \cdots < t_N = T$, starting from $n = N-1, \dots, 0$

$$\begin{aligned} Y^{j,\pi}_N &= \widetilde{Y}^{j,\pi}_N = g^j(X^\pi_N), \quad Z^{j,\pi}_n = 1/\Delta t_n \mathbb{E}_n[Y^{j,\pi}_{n+1} \Delta W^\top_n], \\ \widetilde{Y}^{j,\pi}_n &= \Delta t_n f^j(t_n, X^{j,\pi}_n, \widetilde{Y}^{j,\pi}_n, Z^{j,\pi}_n) + \mathbb{E}_n[Y^{j,\pi}_{n+1}], \quad Y^{j,\pi}_n = \mathfrak{R}^j_y(t_n, X^{j,\pi}_n, \widetilde{Y}^{j,\pi}_n), \end{aligned} \tag{3.9}$$

where $\Delta t_n = t_{n+1} - t_n$, $\Delta W_n = W_{t_{n+1}} - W_{t_n}$. Given an appropriate machinery which approximates the conditional expectations above, one subsequently gathers numerical approximations of the solution pair of (3.6). As for the corresponding discrete time approximations errors, and the convergence of the Euler scheme (3.9) we refer to [106, 24] and the references therein.

Given the discrete time approximations in (3.9) for each option in the portfolio $j = 1, \dots, J$, one can gather discrete time approximations for the delta-hedging portfolio described by (3.2). In fact, combining the Feynman-Kac relations (3.5) with the first-order conditions in (3.3), yields the following discrete time approximations for the vector of delta-hedging weights

$$\alpha^\pi_n = \big(\sum_{j=1}^J Z^{j,\pi}_n \big) \sigma^{-1}(t_n, X^\pi_n). \tag{3.10}$$

Plugging (3.10) into (3.2) gives a discretely rebalanced approximation of the self-financing replicating portfolio. In the above, and for the rest of the chapter, we assume a constant risk-free rate of r.

3.3. DELTA-GAMMA-HEDGING THROUGH ONE STEP MALLIAVIN SCHEMES

In order to improve the replication accuracy for a fixed number of rebalancing dates, one needs to offset higher order sensitivities of the associated portfolio. In the following section we extend the delta-hedging strategy of (3.2) to the case where the second-order terms are also offset, involving second-order Greeks, Gammas, of each option. First, we formulate the general delta-gamma-hedging strategies, thereafter we present a method to deal with the additional model error induced by the presence of Gammas. The latter is done by building on the discrete time approximation schemes presented in [114], where on top of (3.6) an additional, linear BSDE is solved at each point in time, corresponding to the Malliavin derivatives of the solution pair in (3.6), involving a stochastic version of the Gamma process, similarly to (3.5) – see also chapter 2.

3.3.1. DELTA-GAMMA HEDGING

In a discrete time setting, an investor cannot perfectly hedge her exposure in the underlyings solely by offsetting the first order terms in (3.2). However, assets themselves have vanishing Gammas making them unsuitable for the purpose of Gamma hedging. This motivates to expand the hedging portfolio with a set of Gamma-hedging instruments issued on the same underlyings, whose Gammas are not equal to zero. Henceforth, the augmented delta-gamma-hedging portfolio consists of the additional long positions in $k = 1,\dots,K$ gamma-hedging instruments with weights β_t^k: $+\beta_t^k u^k(t,X_t)$, where K is a constant to be fixed, and $u^k(t,X_t)$ denotes the price of the kth gamma-hedging instrument at time t. Note that each gamma-hedging instrument is allowed (but not required) to depend on all risk factors in the portfolio. The value of the portfolio described above evolves according to the SDE below

$$\mathrm{d}P_t^\Gamma = -\sum_{j=1}^{J} \mathrm{d}v^j(t,X_t) + \sum_{i=1}^{m} \alpha_t^i(\mathrm{d}S_t^i + q_t^i S_t^i \mathrm{d}t) + \sum_{k=1}^{K} \beta_t^k \mathrm{d}u^k(t,X_t) + \mathrm{d}B_t, \quad P_0^\Gamma = 0 \tag{3.11}$$

The first- and second-order conditions require $\partial P_t^\Gamma / \partial S_t^i(t,X_t) = 0$ and $\partial^2 P_t^\Gamma / (\partial S_t^l \partial S_t^i) = 0$ for a pair of $1 \le i,l \le m$, implying that the optimal hedging weights solve the following linear system (3.3)

$$\sum_{k=1}^{K} \beta_t^k \partial_{li}^2 u^k(t,X_t) = \sum_{j=1}^{J} \partial_{li}^2 v^j(t,X_t), \qquad il \in \mathscr{I}, \tag{3.12a}$$

$$\alpha_t^i = \sum_{j=1}^{J} \partial_i v^j(t,X_t) - \sum_{k=1}^{K} \beta_t^k \partial_i u^k(t,X_t), \qquad 1 \le i \le d. \tag{3.12b}$$

Note that (3.12a) is a $|\mathscr{I}| \times K$ sized linear system whose solution, at each point in time, is a vector in $\mathbb{R}^K$. While optimizing her hedging weights according to the contraints established by (3.12), the investor has two degrees of freedom:

- the index set $\mathscr{I}$ in (3.12a), with which she can decide which elements in the corresponding Γ matrices she would like to offset;

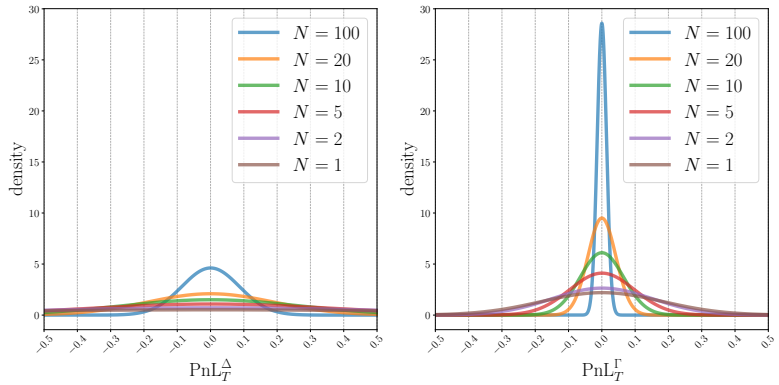

(a) delta hedging (b) delta-gamma hedging

Figure 3.1: Comparison between delta- and delta-gamma hedging, given the number of discrete rebalancing dates N. Portfolio consisting of a single European vanilla call with maturity $T = 1$, strike $K = 100$, and $r = 0$, $\sigma = 0.25$. European vanilla put used as Gamma hedging instrument with maturity $2T$, and the same strike.

- the number (and type) of gamma-hedging instruments K.

Given these choices the resulting linear system in (3.12a) may be under- or over-determined. The choice of $\mathscr{I}$ allows for the freedom to offset gammas and cross-gammas of particularly volatile assets only that may have more severe effects on the options values v^j in the portfolio.

Risk measures. Due to the finite number of rebalancing dates the resulting portfolio is not riskless in neither (3.2) nor (3.11). In order to assess the quality of a hedging strategy, one can assess the distribution of the relative Profit-and-Loss (PnL) which is an $\mathscr{F}_t$ measurable random variable defined by

$$\mathrm{PnL}_t^{\Delta} := \frac{e^{-rt}P_t^{\Delta}}{\sum_j^J v^j(0,X_0)}, \qquad \mathrm{PnL}_t^{\Gamma} := \frac{e^{-rt}P_t^{\Gamma}}{\sum_{j=1}^J v^j(0,X_0)}, \qquad \text{for all } 0 \le t \le T. \tag{3.13}$$

In particular, some of the most common risk measures include Value-at-Risk (VaR), Expected Shortfall (ES) and semivariance, which are defined below

$$\mathrm{VaR}_\alpha(\mathrm{PnL}) := \inf\{x : \mathbb{P}[\mathrm{PnL} < x] > \alpha\}, \quad \mathrm{ES}_\alpha(\mathrm{PnL}) := \mathbb{E}[\mathrm{PnL}|\mathrm{PnL} < \mathrm{VaR}_\alpha],$$
$$\mathrm{SVar}_-(\mathrm{PnL}) := \mathbb{E}[(\mathrm{PnL} - \mathbb{E}[\mathrm{PnL}])^2|\mathrm{PnL} < \mathbb{E}[\mathrm{PnL}]].$$

As an illustration of the accuracy of delta-gamma compared to delta hedging in the discrete time rebalancing framework, figure 3.1 depicts the distribution of the PnL at maturity corresponding to strategies (3.2) and (3.11), when the corresponding portfolio consists of a single, European vanilla call. In fig. 3.1 the distribution of the PnL is collected for several equidistant rebalancing dates. As can be seen, by including second-order sensitivities, one gains roughly an order of magnitude accuracy in the variance of the PnL distribution.

3.3.2. ONE STEP MALLIAVIN SCHEMES

As illustrated by figure 3.1, equation (3.11) brings an improvement to standard delta hedging in the discrete time framework by offsetting second-order terms in the series expansion of the corresponding portfolio value. However, unlike in the case of a vanilla Black-Scholes option, this comes with an additional modelling challenge. In exchange for the additional discrete replication accuracy, (3.12) exposes an investor to additional model risk, namely, accurate approximations of the options' Gammas, which is necessary in order to be able to compute the gamma-hedging weights β_t^k according to the second-order constraints – see right-hand side of (3.12a). Discretely reflected BSDEs and the Feynman-Kac relations (3.5) only provide access to the Deltas and not to Gamma. In order to address this gap in the FBSDE context, we exploit ideas presented in [114] and use the One Step Malliavin scheme to collect the necessary approximations of second-order Greeks.

In particular, it turns out that under sufficient conditions, see [23, Proposition 5.1] that the Z part of the solution pair of the (discretely) reflected BSDE (3.6) can be represented as the solution to a linear, vector-valued BSDE, corresponding to the Malliavin derivatives $\{(D_t\widetilde{Y}_u, D_tZ_u)\}_{0\le t\le r\le T}$

$$\begin{aligned}
\widetilde{Z}_t &:= D_t\widetilde{Y}_u^j = D_t\widetilde{Y}_u^j + \mathbf{1}_{u\in\mathscr{R}^j\setminus\{0,T\}}\mathbf{1}_{l^j(X_u)>\widetilde{Y}_u^j} \\
&\quad + \int_{\underline{r}_t^j}^{\overline{r}_t^j} \big[\nabla_x f(s,X_s,\widetilde{Y}_s^j,Z_s^j)D_tX_s + \nabla_y f(s,X_s,\widetilde{Y}_s^j,Z_s^j)D_t\widetilde{Y}_s^j \\
&\qquad + \nabla_z f(s,X_s,\widetilde{Y}_s^j,Z_s^j)D_tZ_s^j\big]\mathrm{d}s \\
&\quad - \int_{\underline{r}_t^j}^{\overline{r}_t^j} \big(D_tZ_s^j\big)^T\mathrm{d}W_s, \quad t\in[0,T], u\in[\underline{r}_t^j,\overline{r}_t^j), \\
Z_t^j &= D_t\widetilde{Y}_t^j + \mathbb{1}_{t\in\mathscr{R}^j}\mathbb{1}_{l^j(X_t)>\widetilde{Y}_t^j}\big[D_tl^j(X_t) - D_t\widetilde{Y}_t^j\big] =: \mathfrak{R}_z^j(t,X_t,\widetilde{Y}_t^j,D_t\widetilde{Y}_t^j).
\end{aligned} \tag{3.14}$$

Most importantly, the representation in (3.14) does not only provide a way to compute the Deltas in form of a linear BSDE, but through the DZ process and the Malliavin chain rule, the solution of (3.14) also includes second-order derivatives of the j'th option price, i.e. Gammas. In light of the Malliavin chain rule – see e.g. [119] – the DZ process coincides with the formal derivative of the Z process, and one has

$$D_tZ_t\sigma^{-1}(t,X_t) = \gamma(t,X_t) = \nabla_x z(t,X_t), \tag{3.15}$$

where $Z_t = z(t,X_t)$ is the Markovian mapping of the Z process given by the Feynman-Kac formula (3.5). Subsequently, this leads to the following expression recovering the Hessian matrix of the value function

$$\mathrm{Hess}\, v^j(t,X_t) = (\sigma^{-1}(t,X_t))^T(\gamma(t,X_t) - \nabla v^j(t,X_t)\nabla\sigma(t,X_t)). \tag{3.16}$$

The One Step Malliavin (OSM) scheme is a discretization which simultaneously solves the *pair of BSDEs* (3.6) and (3.14). Henceforth, it involves a sequence of discrete time approximations, which – in case of appropriate approximations of the resulting nested, backward recursion of conditional expectations – gives approximations for the stochastic counterparts of options' prices, Deltas and also Gammas, i.e. all components needed

in the discrete time rebalancing of the delta-gamma hedging portfolio in (3.11). In particular, given the simultaneous discrete time solution of (3.6) and (3.14) one can immediately solve the corresponding gamma hedging constraints in (3.12).

In order to be able to formulate the corresponding deep BSDE approximations, we briefly outline the time discretization of the OSM scheme. The forward diffusion is approximated according to an Euler-Maruyama scheme as in (3.8). Thereafter, the Malliavin derivative of the forward diffusion is again approximated by an Euler-Maruyama scheme leading to the following discrete time approximations

$$\begin{aligned} D_n X_n^\pi &= \sigma(t_n, X_n^\pi), \\ D_n X_{n+1}^\pi &= D_n X_n^\pi + \Delta t_n \nabla_x \mu(t_n, X_n^\pi) D_n X_n^\pi + \sum_{k=1}^{d} \nabla_x \sigma^k(t_n, X_n^\pi) D_n X_n^\pi \Delta W_n^k, \end{aligned} \tag{3.17}$$

where σ^k is the kth column of the diffusion matrix, and ΔW_n^k is the kth element of the Brownian increment vector – see e.g. [154]. Then a backward recursive sequence of conditional expectations can be formulated for the discrete approximations of the backward equations (3.6), (3.14) as follows

$$\begin{aligned} Y_N^{j,\pi} &= \widetilde{Y}_N^{j,\pi} = g^j(X_N^\pi), \quad Z_N^{j,\pi} = \widetilde{Z}_N^{j,\pi} = \nabla g(X_N^\pi)\sigma(T, X_N^\pi), \\ \Gamma_n^{j,\pi}\sigma(t_n, X_n^\pi) &:= D_n Z_n^{j,\pi} = \frac{1}{\Delta t_n}\mathbb{E}\big[\Delta W_n \big(D_n Y_{n+1}^{j,\pi} + \Delta t_n f^{D,j}(t_{n+1}, \mathbf{X}_{n+1}^{j,\pi}, \mathbf{D}_n \mathbf{X}_{n+1,n}^{j,\pi})\big)\big] \\ \widetilde{Z}_n^{j,\pi} &= D_n \widetilde{Y}_n^{j,\pi} = \mathbb{E}\big[D_n Y_{n+1}^{j,\pi} + \Delta t_n f^{D,j}(t_{n+1}, \mathbf{X}_{n+1}^{j,\pi}, \mathbf{D}_n \mathbf{X}_{n+1,n}^{j,\pi}\big], \\ \widetilde{Y}_n^{j,\pi} &= \vartheta_y \Delta t_n f^j(t_n, X_n^\pi, \widetilde{Y}_n^{j,\pi}, Z_n^{j,\pi}) \\ &\quad + \mathbb{E}_n[Y_{n+1}^{j,\pi} + (1-\vartheta_y)\Delta t_n f^j(t_{n+1}, X_{n+1}^\pi, \widetilde{Y}_{n+1}^{j,\pi}, Z_{n+1}^{j,\pi})], \\ Z_n^{j,\pi} &= \mathfrak{R}_z^j(t_n, X_n^\pi, \widetilde{Y}_n^{j,\pi}, D_n \widetilde{Y}_n^{j,\pi}), \quad Y_n^{j,\pi} = \mathfrak{R}_y^j(t_n, X_n^\pi, \widetilde{Y}_n^{j,\pi}), \end{aligned} \tag{3.18}$$

with the following approximation motivated by the Malliavin chain rule and the Feynman-Kac formula

$$D_n Y_{n+1}^{j,\pi} = Z_{n+1}^{j,\pi} \sigma^{-1}(t_{n+1}, X_{n+1}^\pi) D_n X_{n+1}^\pi.$$

Similarly to [114], in the above we use the short hand notations $\mathbf{X}_n^{j,\pi} := (X_n^\pi, \widetilde{Y}_n^{j,\pi}, Z_n^{j,\pi})$, $\mathbf{D}_n \mathbf{X}_{n+1,n}^{j,\pi} = (D_n X_{n+1}^\pi, D_n \widetilde{Y}_{n+1}^{j,\pi}, D_n Z_n^{j,\pi})$ and

$$f^{D,j}(t, \mathbf{X}_n^{j,\pi}, \mathbf{D}_n \mathbf{X}_{n+1,n}^{j,\pi}) = \nabla_x f(t, \mathbf{X}_n^{j,\pi}) D_n X_{n+1}^\pi + \nabla_y f(t, \mathbf{X}_n^{j,\pi}) D_n \widetilde{Y}_{n+1}^{j,\pi} + \nabla_z f(t, \mathbf{X}_n^{j,\pi}) D_n Z_n^\pi.$$

In particular, combining the Feynman-Kac relations in (3.5) with the Malliavin chain rule, one gathers the following one on one correspondence between the stochastic processes in (3.18) and the prices and Greeks of the options in the gamma-hedging portfolio in (3.12). Indeed, for each option j, $\widetilde{Y}^j, Y^j$ describe the continuation value and price; $\widetilde{Z}^j$ – derivative of the continuation value; Z^j the Delta and Γ the second-order Greeks. In order to recover the option Deltas and Gammas from these processes, one combines (3.5)

with (3.16), to obtain

$$\begin{aligned}\text{Delta}^{j,\pi}(t_n, X_n^\pi) &= Z_n^{j,\pi}\sigma^{-1}(t_n, X_n^\pi),\\ \text{Gamma}^{j,\pi}(t_n, X_n^\pi) &= (\sigma^{-1}(t, X_t))^T(\gamma(t, X_t) - \nabla v^j(t, X_t)\nabla\sigma(t, X_t)).\end{aligned}$$

As implied by these expressions above, the One Step Malliavin scheme for discretely reflected BSDEs in (3.18) provides discrete time approximations for the simultaneous pricing and delta-gamma hedging for each option j in the portfolio (3.11). In fact, given appropriate approximations of the conditional expectations in (3.18), one can quantify the right-hand sides of (3.12) corresponding to first- and second-order Greeks of all associated options. Then the solution of the linear system in (3.12) provides the evolution of the corresponding discretely rebalanced replicating portfolio as follows

$$P_0^{\Gamma,\pi} = 0, \qquad P_n^{\Gamma,\pi} = -\sum_{j=0}^{J} Y_n^{j,\pi} + \sum_{i=1}^{m} \alpha_n^{i,\pi}\left(S_n^{i,\pi} + q_{t_n}^i S_n^{i,\pi}\Delta t_n\right) + \sum_{k=1}^{K} \beta_n^{k,\pi} u^k(t_n, X_n^\pi) + B_n^\pi,$$

for $n = 0, \dots, N$. In the above, $B_n^\pi = e^{r\Delta t_n}B_{n-1}^\pi - \sum_{i=1}^m S_n^{i,\pi}(\alpha_n^{i,\pi} - \alpha_{n-1}^{i,\pi}) - \sum_{k=1}^K(\beta_n^{k,\pi} - \beta_{n-1}^{k,\pi})u^k(t_n, X_n^\pi)$, when $n \geq 1$, as the portfolio is self-financing. However, in order to make the scheme implementable in a high-dimensional framework, i.e. whenever the number of underlying risk factors collected in X, or the number of options in the portfolio J is large, one needs to have a method which accurately approximates the conditional expectations in (3.18). This is discussed in the following section.

3.4. DEEP BSDE APPROXIMATIONS ON THE PORTFOLIO LEVEL

In order to compute the conditional expectations in (3.9) and (3.18), accurately and robustly in a high-dimensional framework when d in (3.1) or J in (3.2)-(3.11) are large, we present a methodology based on deep neural network Monte Carlo regressions which is capable of dealing with such high-dimensional problems. Our formulation is a variant of the backward deep BSDE method. In what follows we extend this backward deep BSDE methodology in [114] – and chapter 2 of this thesis – to the delta (3.3) and delta-gamma hedging (3.12) framework of the portfolio problem in (3.2) and (3.11), i.e. the simultaneous approximation of all option prices, Deltas and Greeks $j = 1, \dots, J$. In the rest of the chapter, $\pi := \{0 = t_0 < t_1 < \cdots < t_N = T\}$ denotes a partition of the finite time interval $[0, T]$. Without loss of generality, we assume $\cup_{j=1}^J \mathscr{R}^j \subseteq \pi$, i.e. that the potential early exercise dates are included in the discrete time partition. We put $X_{t_n}^\pi = X_n^\pi$ for all discrete time time approximations over the time partition, and similarly for other processes. We define the notations $\Delta t_n := t_{n+1} - t_n$, $\Delta W_n := W_{t_{n+1}} - W_{t_n}$.

3.4.1. DEEP BSDE APPROXIMATIONS FOR THE OSM SCHEME

In order to address the second-order terms appearing in (3.12), we present the deep BSDE methodology built on the OSM scheme (3.18), providing sufficient gamma estimates that can be used in the context of delta-gamma hedging. To this end, let us put $\varphi(\cdot|\theta^y) : \mathbb{R}^d \to \mathbb{R}^J$, $\psi(\cdot|\theta^z) : \mathbb{R}^d \to \mathbb{R}^{J\times d}$ and $\chi(\cdot|\theta^\gamma) : \mathbb{R}^d \to \mathbb{R}^{J\times d\times d}$ for neural networks depending on some parameter sets $\theta^y, \theta^z, \theta^\gamma$. These parametrizations are supposed to

Algorithm 2 Deep BSDE approximations with OSM schemes

Require: $\pi(N), \varrho(R) \subset [0,T], I \in \mathbb{N}_+, \eta : \mathbb{N} \to \mathbb{R}$
Require: $\varphi(\cdot|\theta^y) : \mathbb{R}^d \to \mathbb{R}^J$, $\psi(\cdot|\theta^z) : \mathbb{R}^d \to \mathbb{R}^{J\times d}$, $\chi(\cdot|\theta^\gamma) : \mathbb{R}^d \to \mathbb{R}^{J\times d\times d}$ ▷ neural networks
Ensure: $\varrho(R) \subseteq \pi(N)$ ▷ all discrete reflection dates are contained in the discretization
  $\widehat{Y}_N^\pi \leftarrow g(X_N^\pi)$, $\widehat{Z}_N^\pi \leftarrow \nabla_x g(X_N^\pi)\sigma(T, X_N^\pi)$ ▷ collect terminal conditions of the BSDEs (3.6), (3.14)
  for $n = N-1, \dots, 0$ **do**
    if $n = N-1$ **then**
      $(\theta_n^{y,(0)}, \theta_n^{z,(0)}, \theta_n^{\gamma,(0)}) \leftarrow$ random initialization
    else
      $(\theta_n^{y,(0)}, \theta_n^{z,(0)}, \theta_n^{\gamma,(0)}) \leftarrow (\widehat{\theta}_{n+1}^y, \widehat{\theta}_{n+1}^z, \widehat{\theta}_{n+1}^\gamma)$ ▷ transfer learning in (3.26)
    end if
    for $i = 0, \dots, I-1$ **do**
      $\{X_n^\pi\}_{0\le n\le N}$ ▷ Euler-Maruyama approximation of (3.1)
      $\{D_n X_n^\pi, D_n X_{n+1}^\pi\}$ ▷ Malliavin derivative approximations in (3.17)
      $\widehat{\mathscr{L}}_n^z(\theta_n^{z,(i)}, \theta_n^{\gamma,(i)})$ ▷ empirical version of (3.19)
      $(\theta_n^{z,(i+1)}, \theta_n^{\gamma,(i+1)}) \leftarrow (\theta_n^{z,(i)}, \theta_n^{\gamma,(i)}) - \eta(i)\nabla\widehat{\mathscr{L}}_n^z(\theta_n^{z,(i)}, \theta_n^{\gamma,(i)})$ ▷ SGD step
    end for
    $(\widehat{\theta}_n^z, \widehat{\theta}_n^\gamma) \leftarrow (\theta_n^{z,(i+1)}, \theta_n^{\gamma,(i+1)})$
    $(\hat{\widetilde{Z}}_n^\pi, \widehat{\Gamma}_n^\pi\sigma(t_n, X_n^\pi)) \leftarrow (\psi(X_n^\pi|\widehat{\theta}_n^z), \chi(X_n^\pi|\widehat{\theta}_n^\gamma)\sigma(t_n, X_n^\pi))$ ▷ approximations $\widetilde{Z}, \Gamma$
    for $i = 0, \dots, I-1$ **do**
      $\widehat{\mathscr{L}}_n^y(\theta^y)$ ▷ empirical version of (3.22)
      $\theta_n^{y,(i+1)} \leftarrow \theta_n^{y,(i)} - \eta(i)\nabla\widehat{\mathscr{L}}_n^y(\theta_n^{y,(i)})$ ▷ SGD step
    end for
    $\widehat{\theta}_n^y \leftarrow \theta_n^{y,(i+1)}$
    $\hat{\widetilde{Y}}_n^\pi \leftarrow \varphi(X_n^\pi|\widehat{\theta}_n^y)$ ▷ approximation continuation value
    $\widehat{Z}_n^\pi \leftarrow \big(\mathfrak{R}_z^1(t_n, X_n^\pi, \hat{\widetilde{Y}}_n^{1,\pi}\hat{\widetilde{Z}}_n^{1,\pi}), \dots, \mathfrak{R}_z^J(t_n, X_n^\pi, \hat{\widetilde{Y}}_n^{J,\pi}\hat{\widetilde{Z}}_n^{J,\pi})\big)$ ▷ reflection Delta (3.20)
    $\widehat{Y}_n^\pi \leftarrow \big(\mathfrak{R}_y^1(t_n, X_n^\pi, \hat{\widetilde{Y}}_n^{1,\pi}), \dots, \mathfrak{R}_y^J(t_n, X_n^\pi, \hat{\widetilde{Y}}_n^{J,\pi})\big)$ ▷ reflection price process (3.21)
  end for

approximate the conditional expectations corresponding to the left hand sides of (3.18), for each option $j = 1, \dots, J$ in the portfolios (3.2), (3.11). In particular, we emphasize that the output of the *pricing* network φ is a vector containing the continuation value of each option in the portfolio; the output *delta* network ψ contains each option's delta, concatenated row by row; whereas the output of the *gamma* network χ is a tensor including each option's Gamma matrix.

In order to find appropriate parameter sets, such that φ, ψ, χ accurately approximate the conditional expectations for $\widetilde{Y}, \widetilde{Z}, DZ$ in (3.14), respectively, we define the following

L^2 loss functions motivated by the martingale representation theorem – see [114] –

$$\begin{aligned}\mathscr{L}_n^z(\theta^z,\theta^\gamma) := \frac{1}{J}\sum_{j=1}^{J}\mathbb{E}\Big[|D_nY_{n+1}^{j,\pi} + \Delta t_n \nabla_x f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^{j,\pi})D_nX_{n+1}^{\pi} \\ + \Delta t_n \nabla_y f(t_{n+1},\widehat{X}_{n+1}^{j,\pi})D_n\widetilde{Y}_{n+1}^{j,\pi} \\ + \Delta t_n \nabla_z f(t_{n+1},\widehat{\mathbf{X}}_{n+1}^{j,\pi})\chi^j(X_n^\pi|\theta^\gamma)\sigma(t_n,X_n^\pi) \\ - \psi^j(X_n^\pi|\theta^z) - \big(\chi^j(X_n^\pi|\theta^\gamma)\sigma(t_n,X_n^\pi)\big)^T\Delta W_n|^2\Big],\end{aligned} \tag{3.19}$$

where ψ^j and χ^j are the j'th row, jth element of the first axis of the output of the corresponding neural network. One can think of the loss function in (3.19) as a mean-squared error in the Frobenius matrix norm for the row-wise concatenated system consisting of (3.6) for each $j = 1,\dots,J$ – corresponding to the collection in (3.7). This way, by minimizing (3.19), one gathers simultaneous deltas and gammas for each option, without having to run separate optimization problems for each $j = 1,\dots,J$. A suitable minimizer of $\mathscr{L}_n^z$ denoted by $(\widehat{\theta}_n^z,\widehat{\theta}_n^\gamma)$ can be obtained by a Stochastic Gradient Descent (SGD) type optimization. Subsequently, the resulting parametrizations approximate the conditional expectation in (3.18) for each $j = 1,\dots,J$ in the following manner

$$\widetilde{Z}_n^{j,\pi} \approx \psi^j(X_n^\pi|\widehat{\theta}^z), \quad \Gamma_n^{j,\pi} \approx \chi^j(X_n^\pi|\widehat{\theta}^\gamma).$$

Thereafter, one can use the reflection operators defined in (3.6), (3.14), to get an approximation for the option Delta, and not just the derivative of the continuation value. In fact, given an approximation of the continuation value $\widetilde{Y}_n^{j,\pi}$, the reflection in the Z process reads as follows

$$\begin{aligned}Z_n^{j,\pi} &= \mathfrak{R}_z^j(t_n,X_n^\pi,\varphi^j(X_n^\pi|\theta^y),\psi^j(X_n^\pi|\theta^z)) \\ &= \psi^j(X_n^\pi|\theta^z) + \mathbf{1}_{t_n\in\mathscr{R}^j\setminus\{0,T\}}\mathbf{1}_{\varphi^j(X_n^\pi|\theta^y)>g^j(X_n^\pi)}\Big[\nabla_x g^j(X_n^\pi)\sigma(t_n,X_n^\pi) - \psi^j(X_n^\pi|\theta^z)\Big].\end{aligned} \tag{3.20}$$

Similarly, given the same approximation for the continuation value, the corresponding option price can be approximated by combining the continuation value with the discrete early exercise strategy implied by the reflection taking place in (3.6). The resulting approximations read as follows

$$\begin{aligned}Y_n^{j,\pi} \approx \mathbf{R}_y^j(t_n,X_n^\pi,\varphi^j(X_n^\pi|\theta^y)) = \varphi^j(X_n|\theta^y) \\ + \mathbf{1}_{t_n\in\mathscr{R}^j\setminus\{0,T\}}\mathbf{1}_{\varphi^j(X_n^\pi|\theta^y)>g^j(X_n^\pi)}\Big[g^j(X_n^\pi) - \varphi^j(X_n^\pi|\theta^y)\Big].\end{aligned} \tag{3.21}$$

The approximations in (3.20) and (3.21) can easily be vectorized over $j = 1,\dots,J$, so that they act simultaneously on the whole row-wise concatenated system of discretely reflected BSDEs in (3.7). Then, combining (3.20) with the reflection in the continuation value defined in (3.21), one can subsequently formulate the following loss function approximating the last conditional expectation in the discrete time recursion in (3.18). In

particular, we define the following loss function for the Y part of (3.6)

$$\mathscr{L}_n^y(\theta^y) := \mathbb{E}\Big[\big|\widehat{Y}_{n+1}^\pi + (1-\vartheta_y)\Delta t_n f(t_{n+1}, \mathbf{X}_{n+1}^\pi) + \vartheta_y \Delta t_n f(t_n, X_n^\pi, \varphi(X_n^\pi|\theta^y), \widehat{Z}_n^\pi) \\ - \varphi(X_n^\pi|\theta^y) - \widehat{Z}_n^\pi \Delta W_n\big|^2\Big]. \tag{3.22}$$

We emphasize that the approximation above, and its corresponding conditional expectation in (3.18), is implicit in the continuation value, not just through the $\vartheta_y > 0$ parameter, but also through the reflection occuring in (3.20). In fact, the approach in (3.19), (3.22) approximates the BSDEs in (3.14) and (3.6), respectively, meaning that the resulting approximations φ, ψ correspond to the continuation value and its gradient, respectively. In order to then approximate the option prices, and their derivatives, one needs to approximate the reflection associated with early exercising in the discrete time framework. This is achieved by (3.20) and (3.21) accordingly.

The deep BSDE method outlined above explains all necessary steps one needs to take at a given time step t_n locally. Thereafter, the method given by the loss functions (3.19) and (3.22), is made fully implementable by a backward recursion starting at terminal time, executed as follows. First, given a suitable discretization of the forward SDE and its Malliavin derivative, e.g. an Euler-Maruyama scheme such as (3.8) and (3.17), one needs to collect the terminal condition $\widehat{Y}_N^\pi = \hat{\widehat{Y}}_N^\pi = (g^1(X_N^\pi), \dots, g^J(X_N^\pi))$ and $Z_N^\pi = \tilde{Z}_N^\pi = (\nabla_x g^1(X_N^\pi), \dots, \nabla_x g^J(X_N^\pi))\sigma(T, X_N^\pi)$ as in (3.18). Then, in a backward recursion going from $n = N-1$ to 0, one parametrizes the solution pair of the Malliavin BSDE in (3.14) at time step n according to $\psi(\cdot|\theta^z), \chi(\cdot|\theta^\gamma)$. These parametrizations are optimized according to the loss function $\mathscr{L}_n^z(\theta^z, \theta^\gamma)$ defined in (3.19). Through a suitable minimization procedure such as stochastic gradient descent, one then subsequently gathers appropriate approximations of the optimal parameter set $(\widehat{\theta}_n^z, \widehat{\theta}_n^\gamma) \in \arg\min_{\theta^z, \theta^\gamma} \mathscr{L}_n^z(\theta^z, \theta^\gamma)$. Setting

$$\widehat{\Gamma}_n^\pi = \chi(X_n^\pi|\widehat{\theta}_n^\gamma), \quad \hat{\widehat{Z}}_n^\pi = \psi(X_n^\pi|\widehat{\theta}_n^z) \tag{3.23}$$

provides approximations for the first and second conditional expectations in (3.18). Combining these approximations with the discrete reflections given by (3.20) and (3.21) then gives loss function $\mathscr{L}_n^y(\theta^y)$ at time step t_n, that is meant to measure the approximation quality in the continuation values in the discretely reflected BSDE of (3.6). A second stochastic gradient descent optimization then estimates the optimal parameter set $\widehat{\theta}_n^y \in \arg\min_{\theta^y} \mathscr{L}_n^y(\theta^y)$, giving approximations for the continuation value at time t_n, defined by the last conditional expectation in (3.18)

$$\hat{\widehat{Y}}_n^\pi = \varphi(X_n^\pi|\widehat{\theta}_n^y). \tag{3.24}$$

Finally, combining the approximations of the continuation value $\hat{\widehat{Y}}_n^\pi$ and its gradient $\hat{\widehat{Z}}_n^\pi$ with the early exercise decision given by the vectorized expressions in (3.21) and (3.20) for all $j = 1, \dots, J$, gives the final discrete time approximations for all processes in the pair of backward SDEs in (3.6) and (3.14)

$$\widehat{Z}_n^\pi = \mathfrak{R}_z(t_n, X_n^\pi, \hat{\widehat{Y}}_n^\pi, \hat{\widehat{Z}}_n^\pi), \quad \widehat{Y}_n^\pi = \mathfrak{R}_y(t_n, X_n^\pi, \hat{\widehat{Y}}_n^\pi). \tag{3.25}$$

The approximations given by (3.23), (3.24), (3.25) complete the processing of time step t_n. The initialization of a stochastic gradient descent optimization has a substantial impact on the speed of convergence and also on the final accuracy of the resulting approximations. Argued by continuity in time of the stochastic processes $\widetilde{Y}, \widetilde{Z}$ in (3.6) and (3.6) we therefore initialize the parameter sets of the loss functions $\mathscr{L}^z_{n-1}$ and $\mathscr{L}^y_{n-1}$ according to the transfer learning trick

$$(\theta^z, \theta^\gamma) \leftarrow (\widehat{\theta}^z_n, \widehat{\theta}^\gamma_n), \quad \theta^y \leftarrow \widehat{\theta}^y_n, \qquad \text{for each } n = N-1, \dots, 1. \tag{3.26}$$

With these initial parameter guesses, all points $n = N-2, \dots, 0$ in time are more efficiently optimized. One then completes the algorithm by carrying out the same procedure in a backward iteration terminating at $n = 0$. The complete fully-implementable backward deep BSDE algorithm is collected in algorithm 2.

The backward deep BSDE method of Huré et al. [77]. When only the Deltas are of interest, alternative deep BSDE schemes are applicable in the high-dimensional context, see [32, 77]. The RDBDP scheme [77, sec. 3.3] is concerned with the numerical approximation of variational inequalities, which can be considered as a continuous asymptotics for (3.6) in the case $R \to \infty$. Even though, their scheme is only given for a single ($J = 1$) continuously reflected BSDE, in what follows, we naturally extend this to the collection of discretely reflected equations given in (3.7). Let us put $\varphi(\cdot|\theta^y) : \mathbb{R}^d \to \mathbb{R}^J$ and $\psi(\cdot|\theta^z) : \mathbb{R}^d \to \mathbb{R}^{J\times d}$ for two feedforward, fully-connected neural networks, depending on some potentially non-disjoint parameter sets. In fact, in [32], $\psi(\cdot|\theta^z) = \nabla_x \varphi(\cdot|\theta^y)\sigma(t_n, \cdot)$, with $\theta^z \equiv \theta^y$. These neural networks are parametrizations of the Markovian conditional expectations of the collection of conditional expectations in (3.9) for every $j = 1, \dots, J$ in (3.7). Using these parametrizations, we define the L^2 loss function

$$\begin{aligned}\mathscr{L}^{\text{Euler}}_n(\theta^y, \theta^z) := \mathbb{E}\big[\big| \widehat{Y}^\pi_{n+1} \\ - \varphi(X^\pi_n|\theta^y) - \Delta t_n f(t_n, X^\pi_n, \varphi(X^\pi_n|\theta^y), \psi(X^\pi_n|\theta^z)) + \psi(X^\pi_n|\theta^z)\Delta W_n\big|^2\big],\end{aligned} \tag{3.27}$$

which is a function of the total parameter set $\Theta := (\theta^y, \theta^z)$. Herein, we denote the row-wise concatenated solutions of each discretely reflected BSDE as in (3.7). As shown in [77, thm. 4.1, 4.4], a minimizer $(\widehat{\theta}^y, \widehat{\theta}^z)$ of the loss function defined by (3.27) provides a simultaneous approximation of the conditional expectations $\widetilde{Y}^{j,\pi}_n$ and $Z^{j,\pi}_n$ in (3.9) given by the following expressions $\hat{\widetilde{Y}}^{j,\pi}_n \approx \varphi^j(X^\pi_n|\theta^y)$, $\widehat{Z}^{j,\pi}_n \approx \psi^j(X^\pi_n|\widehat{\theta}^z)$, where φ^j, ψ^j denote the j'th row of the output layer of each neural network. In fact, by collecting each discretely reflected BSDE into the system (3.7), the loss function (3.27) takes a mean-squared loss of all equations in (3.9), for every $j = 1, \dots, J$, at the same time. The scheme is made fully-implementable by a similar backward recursion as for the OSM scheme. The main differences between the OSM approximation in algorithm 2 and those of [77, 32] is that the One Step Malliavin scheme does not only provide option prices and Deltas, like [77], but also second-order Greeks, Gammas, throughout the entire spacetime by the process DZ. This makes the deep BSDE approximations for the OSM scheme suitable for delta-gamma hedging in (3.12). Additionally, as shown in [114], the OSM scheme provides more accurate approximations for the Z process in regression Monte Carlo frameworks when the time step size Δt_n is small or the volatility is high. As we show below –

see fig. 3.5, 3.3b in particular –, this in fact results in more accurate Deltas, leading to better delta replication with the use of the OSM scheme.

3.4.2. DELTA HEDGING WITH OSM

Given the deep BSDE approximations in algorithm 2, one can subsequently solve the delta hedging problem of our investor, which comes down to the discrete time approximation of the hedging weights α in (3.2) where t_n is in a set of finite rebalancing dates $\mathscr{S}$. Without loss of generality, we assume that the deep BSDE approximation of the corresponding collection of discretely reflected BSDEs in (3.7) is given over a time partition that includes all rebalancing dates, i.e. $\mathscr{S} \subseteq \pi(N)$. Therefore, combining the first-order condition of the delta hedging weights given by (3.3) with the Feynman-Kac formula in (3.5), one immediately gathers the following approximations for all option Deltas in the portfolio

$$\begin{aligned}\widehat{\text{Delta}}_n^{\pi} = (\widehat{\text{Delta}}_n^{1,\pi},\dots,\widehat{\text{Delta}}_n^{J,\pi}) &= (\widehat{Z}_n^{1,\pi}\sigma^{-1}(t_n,X_n^{\pi});\dots;\widehat{Z}_n^{J,\pi}\sigma^{-1}(t_n,X_n^{\pi})) \\ &= \widehat{Z}_n^{\pi}\sigma^{-1}(t_n,X_n^{\pi}),\end{aligned} \tag{3.28}$$

for all $t_n \in \pi(N)$, and in particular in $\mathscr{S}$. We remark that $\widehat{\text{Delta}}_n^{\pi} \in \mathbb{R}^{J\times d}$ as it is the row-wise collection of each contract's Delta in the portfolio. Plugging this into the first-order conditions given by (3.3), one subsequently gathers the discrete time approximations of the optimal hedging weights at each rebalancing date $t_n \in \mathscr{S}$

$$\widehat{\alpha}_n^{\pi} = \sum_{j=1}^{J} \widehat{\text{Delta}}_n^{j,\pi}. \tag{3.29}$$

Given (3.29), on top of the deep BSDE approximations for the option prices in the portfolio $\widehat{Y}_n^{j,\pi}$, the corresponding discretely rebalanced delta hedged replicating portfolio's value evolves according to the following recursion

$$\begin{aligned}&\widehat{P}_0^{\Delta,\pi} = 0, \quad \widehat{P}_n^{\Delta,\pi} = -\sum_{j=1}^{J}\widehat{Y}_n^{j,\pi} + \sum_{i=1}^{m}\widehat{\alpha}_n^{i,\pi}\left(S_n^{i,\pi} + q_{t_n}^i S_n^{i,\pi}\Delta t_n\right) + B_n^{\pi}, \qquad \text{for } n=0,\dots,N,\\ &B_0^{\Delta,\pi} = 0, \qquad B_n^{\pi} = e^{r\Delta t_n}B_{n-1}^{\pi} - \sum_{i=1}^{m} S_n^{i,\pi}(\widehat{\alpha}_n^{i,\pi} - \widehat{\alpha}_{n-1}^{i,\pi}), \qquad \text{for } n=1,\dots,N.\end{aligned}$$

Notice that all necessary ingredients for the discrete time rebalancing of the hedging portfolio are included in the OSM approximations of the discretely reflected BSDE (3.6), i.e. option prices and Deltas.

3.4.3. DELTA-GAMMA HEDGING WITH OSM

The main motivation behind the One Step Malliavin scheme simultaneously solving the BSDEs (3.6) and (3.14), is that through the numerical resolution of the latter, it includes a Γ process which corresponds to second-order Greeks of the associated option prices. In particular, by modeling all options in the portfolio by the One Step Malliavin scheme performed on the row-wise concatenated collection of reflected BSDEs (3.7), one simultaneously has approximations of all options' prices, Deltas and Gammas through $\widehat{Y}_n^{\pi}, \widehat{Z}_n^{\pi}$

and $\widehat{\Gamma}_n^\pi$, respectively, with the latter being recovered from the Markovian approximations $\Gamma_n^{j,\pi} = \gamma_n^{j,\pi}(X_n^\pi)$. Through a discrete time expression analogous to (3.16), we find

$$\widehat{\text{Gamma}}_n^{j,\pi} = (\sigma^T(t_n, X_n^\pi))^{-1}\left(\widehat{\Gamma}_n^{j,\pi} - \widehat{\text{Delta}}_n^{j,\pi}\nabla\sigma(t_n, X_n^\pi)\right), \tag{3.30}$$

where the approximations $\widehat{\text{Delta}}_n^{j,\pi}$ are recovered identically to (3.28). Equation (3.30) provides discrete time approximations for the right-hand side of the linear system established by the second-order conditions of delta-gamma hedging in (3.12a). Given a set of Gamma hedging instruments whose Gammas are also available in (semi-)closed form, this makes the discrete time approximation of all terms in the delta-gamma hedging portfolio (3.11) possible, after solving the following $|\mathscr{I}| \times K$ sized linear system

$$\sum_{k=1}^K \widehat{\beta}_n^{k,\pi}\partial_{li}^2 u^k(t, X_t) = \sum_{j=1}^J (\widehat{\text{Gamma}}_n^{j,\pi})^{li}, \qquad il \in \mathscr{I}, \tag{3.31a}$$

$$\widetilde{\alpha}_n^\pi = \sum_{j=1}^J \widehat{\text{Delta}}_n^{j,\pi} - \sum_{k=1}^K \widetilde{\beta}_t^k \nabla u^k(t, X_t), \qquad 1 \le i \le d. \tag{3.31b}$$

Given the deep BSDE approximations of the whole Gamma matrix, the linear system imposed by (3.31) has a solution pair $\widehat{\beta}_n^\pi, \widehat{\alpha}_n^\pi \in \mathbb{R}^K \times \mathbb{R}^m$, which solve the discrete time version of the first- and second-order conditions of delta-gamma hedging in (3.12). Consequently, after having solved the associated collection of discretely reflected FBSDEs by the OSM scheme, the investor's task at a rebalancing date is to compute the right-hand sides of (3.31), which can easily be done by evaluating the corresponding neural networks φ, ψ, χ at a given realization of the Brownian motion. Subsequently, the investor needs to buy and sell the tradeable risk factors, and the corresponding gamma-hedging securities according to the differences in the hedging weights between two rebalancing dates. The discretely rebalanced, deep BSDE approximated delta-gamma hedging portfolio evolves according to the following discrete time recursion

$$\begin{aligned}\widehat{P}_n^{\Gamma,\pi} = &-\sum_{j=0}^J \widehat{Y}_n^{j,\pi} + \sum_{i=1}^m \widehat{\alpha}_n^{i,\pi}\left(S_n^{i,\pi} + q_{t_n}^i S_n^{i,\pi}\Delta t_n\right) \\ &+ \sum_{k=1}^K \widehat{\beta}_n^{k,\pi} u^k(t_n, X_n^\pi) + B_n^\pi, \qquad \text{for } n = 0, \dots, N,\end{aligned} \tag{3.32}$$

$$\begin{aligned}B_n^\pi = &\, e^{r\Delta t_n} B_{n-1}^\pi - \sum_{i=1}^m S_n^{i,\pi}(\widehat{\alpha}_n^{i,\pi} - \widehat{\alpha}_{n-1}^{i,\pi}) \\ &- \sum_{k=1}^K (\widehat{\beta}_n^{k,\pi} - \widehat{\beta}_{n-1}^{k,\pi}) u^k(t_n, X_n^\pi), \qquad \text{for } n = 1, \dots, N,\end{aligned} \tag{3.33}$$

with $B_0^\pi = 0$. Unlike in the delta case, the replicating portfolio is rebalanced not just by buying and selling the tradeable risk factors S^i according to the approximated hedging weights $\widehat{\alpha}_n^{i,\pi}$ but also by trading each Gamma hedging instruments according to the difference $\widehat{\beta}_n^{k,\pi} - \widehat{\beta}_{n-1}^{k,\pi}$. The complete algorithm is collected in algorithm 3.

Algorithm 3 Delta-gamma hedging on the portfolio level

Require: Deep BSDE approximations of (3.7) and (3.14) for each $j = 1, \dots, J$ ▷ OSM in algorithm 2
Require: $\{W_{t_n}\}_{0 \le n \le N}$
$P_0^{\Gamma,\pi} \leftarrow 0$
for $n = 0, \dots, N-1$ **do**
  $(X_n^\pi, \hat{Y}_n^\pi, \hat{Z}_n^\pi)$ ▷ deep BSDE approximations
  compute $\widehat{\beta}_n^{\pi,k}, k = 1, \dots, K$ ▷ solution to linear system in (3.31a)
  compute $\widehat{\alpha}_n^\pi$ ▷ via (3.31b)
  if $n = 0$ **then**
    $B_n^\pi \leftarrow \widehat{Y}_n^{j,\pi} - \sum_{i=1}^m \widehat{\alpha}_n^{\pi,i} S_n^{\pi,i} - \sum_{k=1}^K \widehat{\beta}_n^{k,\pi} u^k(t_n, X_n^\pi)$
  else
    update bank account B_n^π ▷ rebalancing with (3.33)
  end if
  update portfolio value $P_{n+1}^{\Delta,\pi}$ ▷ according to (3.32)
end for
$\mathrm{PnL}^\Delta \leftarrow e^{-rT} P_N^{\Delta,\pi} / (\sum_{j=1}^J \widehat{Y}_0^{j,\pi})$ ▷ compute PnL in (3.13)

Alternative approaches to Gamma hedging. Given a differentiable function approximation of the associated option prices, one could in principle use *automatic differentiation* to approximate the corresponding Deltas and Gammas in a discrete time framework. For instance, using the deep backward dynamic programming approach of Huré et al. in [77], one has a differentiable approximation of Z in (3.6) in the form of the vector-valued function $\psi(\cdot|\widehat{\theta}^z)$. Computing the Jacobian matrix by means of automatic differentiation yields an approximation of the *"derivative Z process"*, which is analogous to Γ in (3.14) – see also (3.15). Thereafter, plugging $\nabla_x \psi(\cdot|\widehat{\theta}^z)$ in the place of Γ in the formulae (3.16), one could obtain a comparable numerical representation of the right-hand side of (3.12a). We remark that a similar approach is taken in [32], where motivated by the Feynman-Kac formula the authors approximate the Z by automatic differentiation on the parametrization of the approximation of the option prices, i.e. $\psi = \nabla_x \varphi$. The problem regarding the automatic differentiation approximations outlined above is twofold. First, as the Jacobian matrix of ψ does not form part of the loss function in (3.27), there is no guarantee that $\nabla\psi$ is an accurate approximation of $\nabla_x Z$ – even when ψ approximates Z arbitrarily well. In order to ensure for this to be the case, one would have to augment the loss function in order to account for this, which would result in a similar representation formula as (3.14) in the OSM scheme. Regardless of the lack of theoretical guarantees, one can carry out the corresponding automatic differentations and check if the corresponding results yield meaningful Gammas. We found that this in fact is not the case, rendering $\nabla_x \psi(\cdot|\widehat{\theta}^z)$ inapplicable in the context of delta-gamma hedging. According to our findings, this is already case for low-dimensional problems, and the accuracy of the automatic differentiated Gammas further decreases as d grows. In particular, we refer to figure 3 in [114], and figure 3.3b in our upcoming numerical experiments, which both demonstrate this phenomenon.

3.4.4. ABOUT THE LINEAR SYSTEM OF SECOND-ORDER CONSTRAINTS

The optimal weights with which one has to hold the Gamma hedging instruments $u^k, k = 1,\dots,K$ in (3.11) are determined by the solution of the linear system in (3.12a). The main driver of the computational complexity of gamma hedging stems from the numerical solution of this linear system, which depends not only on the number of gamma hedging instruments K contained in the replicating portfolio, but also on the type of contracts used as hedging instruments. In fact, the linear system in (3.12a) is a $|\mathscr{I}| \times K$ sized rectangular system. In the special case when one chooses to offset all upper triangular elements in the Gamma matrix, assigning a single gamma hedging instrument to each gamma and cross-gamma, this results in a system of size $(m(m+1)/2) \times (m(m+1)/2)$, depending on the number of tradable risk factors. The memory requirements of storing the coefficient matrix of one instance of such a linear system would thus scale $\mathscr{O}(m^4)$, and a naive direct solution of such a linear system would require $\mathscr{O}(m^6)$ floating point operations, depending on the number of spatial dimensions. These scaling factors grow substantially in case of large number of risk factors, especially compared to the simple matrix-vector multiplication determining the delta hedging weights in (3.29).

Nonetheless, one can choose gamma hedging instruments which preserve a special structure of the coefficient matrix in (3.12a). In fact, in the upcoming numerical experiments for each higher dimensional problem we choose the gamma hedging instruments to be European exchange options. Henceforth, given by the Margrabe formula [107], these options admit an analytical closed-form expression not only for the prices, but also for all Greeks up to second order in the Black-Scholes framework – see appendix 3.B. Consequently, the coefficient matrix in (3.12a) is computed analytically for each Brownian path. Moreover, the second order derivatives $\partial^2_{ij} u^k$ determine the shape of the coefficient matrix multiplying $\beta^k, k = 1,\dots,K$. The choice of exchange options implies that each row (or column) of the coefficient matrix in (3.12) only includes at most 3 non-zero elements, leading to a sparse linear system which can efficiently be solved by sparse numerical linear algebra methods. In what follows we use the sparse least squares method developed by [122].[2] As solving such a large linear system, for a cloud of Monte Carlo simulations – in order to be able to assess the distribution of the corresponding PnL distribution – is computationally expensive, this provides a substantial computational improvement for the delta-gamma hedging strategy in algorithm 3.

Replication errors and computational complexity. The main source of computational complexity of delta-gamma hedging is the deep BSDE optimization of the OSM scheme in algorithm 2. Nevertheless, we remark that once the replicating portfolio is constructed and the corresponding options are fixed, this step only needs to be done once. It can be done offline, and the resulting discrete time approximations can be used throughout the whole spacetime. With respect to the convergence of the discrete time approximations we refer to [77, 114] where it is shown that the deep BSDE approximations converge to the continuous solution triple of the BSDEs with an L^2 rate of $\mathscr{O}(|\pi|^{1/2})$. These results can be generalized to the vector-valued setting corresponding to the collection of BSDEs in (3.7) without any substantial difficulty. In case the forward SDE's so-

[2] In particular, the `scipy` implementation, see documentation.

lution is given in closed-form, the replication error of the corresponding delta hedging strategy coincides with those of the deep BSDE approximations. Whereas, when the forward diffusion is approximated by, e.g., an Euler-Maruyama discretization the product term in (3.2) will hamper the convergence rate by an appropriate application of Young's inequality. For a theoretical assessment of the convergence rates of the tracking errors induced by delta (3.2) and delta-gamma hedging (3.11), we refer to [64]. Therein, the authors show that the convergence rate of the hedging portfolio depends on the fractional regularity of the payoff g – at least in the European options' context. In particular, with the use of an equidistant time their results imply the convergence of the variance of the PnL distribution with a rate of $\mathcal{O}(h^{0.5})$ in case of delta-, and a rate of $\mathcal{O}(h^{0.75})$ for delta-gamma hedging of European put and call options. Consequently, including the second-order Greeks enabled by the One Step Malliavin scheme in the replicating portfolio (3.11) does not only improve replication accuracy by a constant, but in special cases may also result in a higher order convergence rate.

3.5. NUMERICAL EXPERIMENTS

In order to demonstrate the accuracy and robustness of the proposed FBSDE based delta-gamma hedging strategies, numerical experiments are presented. In what follows, we formulate a (discretely reflected) FBSDE system corresponding to each portfolio below. Subsequently, each of these systems is discretized using a fine time grid containing $N' = 100$ equally sized intervals, and solved by deep BSDE approximations as in algorithm 2, including approximations for the associated options' prices, deltas and gammas throughout the entire spacetime. This step is the most time consuming part of our approach, however, it needs to be emphasized that the training of these neural networks only needs to be done once, offline. Thereafter, the resulting approximations for prices, deltas and gammas can be simply evaluated for each rebalancing date in the updates of the hedging strategies (3.28) and (3.30), which is fast and efficient. We choose a very fine time partition for the numerical resolution of the FBSDE system in order to be able to use sub-points of it for each rebalancing frequency we consider. In order to ease the presentation, for the OSM scheme in alg. 2, we fix $\vartheta_y = 1/2$ and remark, that results are very similar in case of other choices of $\vartheta_y \in [0,1]$.

For each experiment presented below we use fully-connected, feedforward neural networks of $L = 4$ hidden layers with 50 neurons in each layer and hyperbolic tangent activations. For the networks at time step $n = N-1$, batch normalization is applied in between each pair of layers, whose parameters are thereafter frozen for preceding time-steps $n < N-1$. The historical mean and standard deviation are reset at each time step, in order to preserve adaptivity of the solutions of the BSDEs. Adam optimization is deployed with $I_{N-1} = 2^{16}$ SGD steps at the training of time step $n = N-1$ and thereafter – motivated by the transfer learning trick as in (3.26) – this is reduced to $I_n = 2^{12}$ for $n < N-1$. The OSM method in alg. 2 has been implemented in TensorFlow 2.15. The library used in this chapter is made publicly accessible on the first author's personal github repository[3]. All experiments below were run on a Dell Alienware Aurora R10 machine, equipped with an AMD Ryzen 9 3950X CPU (16 cores, 64Mb cache, 4.7GHz) and

[3] https://github.com/balintnegyesi/OSM-delta-gamma-hedging.git

	#1: Vega	#2: Gamma	#3: Vomma	#4: Vonna
T	0.3	0.4	0.3	0.25
K	10	10	9	11
R	1	1	1	1
N	60	80	60	60

Table 3.1: Example 1. Hedging instruments.

an Nvidia GeForce RTX 3090 GPU (24Gb).

For the computations of the hedging strategies in alg. 3, we use equally spaced rebalancing dates with $N = 1, 2, 5, 10, 20, 100$, which roughly correspond to *yearly*, *quarterly*, *monthly*, *fortnightly*, *weekly* and *daily* rebalancing of each portfolio. This choice ensures that each rebalance date is included in the time partition used to solve the discretely reflected FBSDEs. In each example below, we assume that options can only be exercised on rebalancing dates. In the replication of Bermudan options below, we denote the optimal stopping time at which each option is exercised by τ, which is approximated as follows

$$\tau^j := \underset{0\le n\le N}{\arg\min}\, g^j(X_n^\pi) > \hat{\hat{Y}}_n^{j,\pi}. \tag{3.34}$$

for a given path of the Brownian motion. In other words, τ^j denotes the first time the continuation value is reflected in the discretely reflected FBSDE system in (3.6). We approximate continuous Profit-and-Loss densities with a Gaussian kernel density estimate on the discrete Monte Carlo sample.[4]

3.5.1. EXAMPLE 1: TWO-DIMENSIONAL STOCHASTIC VOLATILITY MODEL

Our first example is a single Bermudan put option ($J = 1$), issued on a single asset ($m = 1$) whose price is driven by the well-known Heston model. Since multi-dimensional extensions to the Heston model are known to suffer from the phenomenon of vanishing correlations, we restrict our illustration to the case $m = 1, d = 2$, and remark that the OSM scheme would be similarly applicable in the context appropriate higher-dimensional stochastic volatility models, such as the Wishart model – see e.g. [56, 66]. We collect the corresponding first- and second-order hedging conditions in the presence of stochastic volatility in appendix 3.A – see (3.37) and (3.38) in particular. Assuming the market price of volatility risk to be $\lambda = 0$, the coefficients of the discretely reflected FBSDE system read as follows

$$\begin{aligned}
&\mu(t, x = (s;v)) = \big(\bar{\mu}s; \kappa(\bar{v} - v)\big), \qquad \sigma(t, x = (s;v)) = \begin{pmatrix} \rho\sqrt{v}s & \sqrt{1-\rho^2}\sqrt{v} \\ 0 & \eta\sqrt{v} \end{pmatrix}, \\
&f(t, x = (s;v), y, z) = -ry - (\bar{\mu} - (r-q))z\sigma^{-1}(t,x)s, \qquad l(x) \equiv g(x = (s;v)) = \max[K - s, 0],
\end{aligned} \tag{3.35}$$

with $x \in \mathbb{R}^2$. The two risk factors in (3.1) are the asset price $S = x^1$, and its stochastic volatility process $v = x^2$. The parameters are chosen according to set A in [146, sec.7],

[4] KDE is implemented by `seaborn`, using a smoothing parameter of 1.8 – see documentation.

i.e. $\bar{\mu} = r = 0.1$, $q = 0$, $\kappa = 5, \bar{\nu} = 0.16, \rho = 0.1$ and $\eta = 0.9$. The Feller condition $2\kappa\bar{\nu} \geq \eta^2$ is satisfied and hence the volatility process does not attain zero. As in [146], the initial condition of the forward SDE in (3.1) is $X_0 = (S_0, \nu_0) = (10; 0.0625)$, and set the option strike to $K = 10$. We consider a maturity of $T = 0.25$ and equidistant early exercise dates with $R = 10$, i.e. $\mathscr{R} = \{0, 0.025, 0.05, \ldots, 0.25\}$.

In order to demonstrate the method's robustness with respect to the approximation of the risk factors, we adopt a modified Euler approximation directly on the asset prices and not on log-prices. The coefficients in (3.35) are truncated according to [70, 99], in order to ensure that the discrete time approximations of the volatility process do not go below zero either

$$X_0 = (x_0, \nu_0), \qquad X^\pi_{n+1} = (S^\pi_{n+1}, \nu^\pi_{n+1}) = \left|(S^\pi_n; \nu^\pi_n) + \mu(t_n, X^\pi_n)\Delta t_n + \sigma(t_n, (S^\pi_n; \left|\nu^\pi_n\right|))\Delta W_n\right|,$$

whereas the Malliavin derivative is approximated according to (3.17). Given the approximations of the forward diffusion, we solve the discretely reflected BSDEs (3.6) and (3.14) corresponding to (3.35) by the OSM scheme in algorithm 2 once, with $N' = 50$ equally sized time intervals, which through the relations (3.28) and (3.30) provide approximations for all first- and second-order Greeks. In particular, due to stochastic volatility, the OSM scheme provides *Delta*, *Vega*, *Gamma*, *Vomma* and *Vanna* approximations. So that we can hedge the volatility, similar to (3.11), we augment the replicating portfolio with 4 additional instruments, that offset Vega, Gamma, Vomma and Vanna exposure, respectively. These instruments are all European put options issued on the same asset, with maturities and strikes as in table 3.1. In order to sketch the potential of the OSM scheme outside of the Black-Scholes framework, we use algorithm 2 as reference prices and Greeks for the coefficient matrices in the linear system described by (3.37) and (3.38). Algorithm 2 is run on each European option's associated BSDE separately, with an equidistant time grid using N intervals as in table 3.1. We consider $N = 1, 2, 5, 10, 25, 50$ rebalancing dates, for all cases.

The impact of stochastic volatility on the hedging with only first-order constraints in (3.29) is illustrated by figure 3.2a. Herein, we find that augmenting the replicating portfolio with an additional Vega hedging instrument, offsetting first-order volatility sensitivity according to (3.37) results in a significantly sharper PnL distribution. Nevertheless, as demonstrated by figures 3.2b and 3.2c, the delta-vega replicating portfolio entails substantial risk due to the neglection of second-order sensitivities. Augmenting the replicating portfolio with the additional Gamma, Vomma and Vanna hedging securities, one can further improve the replication of the Bermudan put option at τ, resulting in sharper profit and loss distributions centered around 0. As seen in fig. 3.2c, offsetting all second-order sentivities with respect to the underlying risk factors improves VaR_{95} by 30 percentage points. The impact on the choices of second-order Greeks to be offset, determined by the index set $\mathscr{I}$ in (3.38) is depicted in figure 3.3a. Three cases are compared depending on the number of second-order Greeks accounting for *(i)* Gamma hedging ($\mathscr{I} = \{11\}$); *(ii)* Gamma-Vomma hedging ($\mathscr{I} = \{11, 22\}$); *(iii)* Gamma-Vomma-Vanna hedging ($\mathscr{I} = \{11, 12, 22, 21\}$). As demonstrated by fig.3.3a, these gradually result in better replication, as more second-order Greeks are taken into account. Table 3.2 collects the mean and variance for the aforementioned replication strategies in case of $N = 25$ rebalancing dates. In line with the discussion above, we see that offsetting

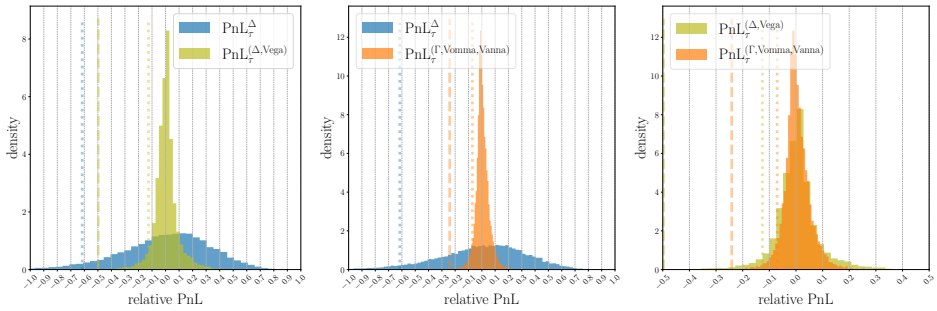

(a) Delta vs. Vega (b) Delta vs. Gamma-Vomma-Vanna (c) Delta-Vega vs. Gamma-Vomma-Vanna

Figure 3.2: Example 1. Comparison of first- and second-order hedging strategies in (3.3), (3.37) and (3.38), respectively. $N = 50$ rebalancing dates. Dashed and dotted vertical lines corresponding to VaR_{95} and ES_{99}, respectively. Hedging instruments as in table 3.1. PnL approximated through an independent Monte Carlo sample of size 2^{14}.

risk measure	PnL_τ^{Δ}	$\text{PnL}_\tau^{\Delta,\text{Vega}}$	PnL_τ^{Γ}	$\text{PnL}_\tau^{\Gamma,\text{Vomma}}$	$\text{PnL}_\tau^{\Gamma,\text{Vomma,Vanna}}$
mean	1.2e−2	1.8e−2	2.0e−2	1.9e−2	1.9e−2
variance	1.5e−1	9.0e−2	1.0e−2	8.5e−3	7.6e−3

Table 3.2: Example 1. Risk measures for different types of replicating portfolios and $N = 25$ rebalancing dates.

second-order Greeks not only with respect to the asset price but also for the volatility considerably improves the replication accuracy across all risk metrics.

Finally, as discussed in section 3.4.1, let us compare the deep BSDE approximations of all above considered option Greeks provided by the One Step Malliavin scheme in alg. 2 to the RDBDP method of [77]. The left plot of fig. 3.3b compares the delta-vega hedging strategies where the corresponding right-hand sides of (3.37) are computed by the OSM and RDBDP schemes, respectively. As we can see, even in this low-dimensional case the OSM scheme brings a marginal improvement in the replication accuracy, indicating that the corresponding first-order Greeks (Deltas and Vegas) are more accurately approximated by the OSM scheme. More importantly, the right side of fig. 3.3b demonstrates the need for the Malliavin representation formula (3.14) and its corresponding discrete time approximation by the OSM scheme in order to accurately capture second-order Greeks. Herein, the second-order hedging corresponding to Gamma-Vomma-Vanna hedging (3.38) are compared across the OSM and RDBDP schemes, where, for the latter, the right-hand side of (3.38) is computed by automatic differentiation. As we can see, automatic differentiation of [77] does not provide accurate Gammas, Vommas and Vannas, and is rendered inapplicable in the context of second-order hedging, whereas the OSM approximations bring an improvement in the replication accuracy compared to the first-order conditions.

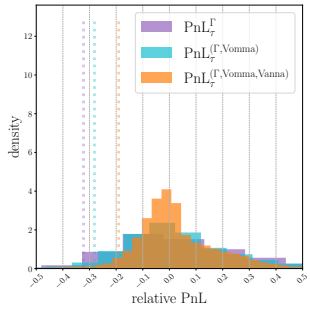

(a) Comparison on the set of second-order Greeks accounted for in (3.38) with quarterly rebalancing ($N = 2$).

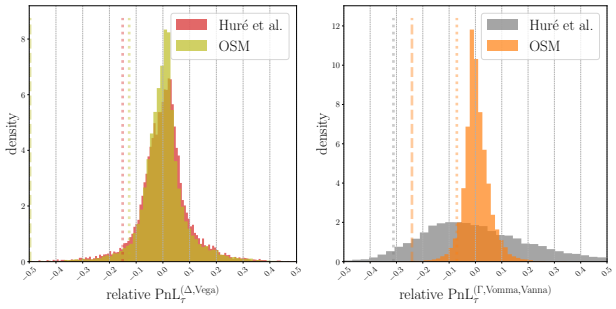

(b) Comparison between the One Step Malliavin scheme (alg.2) and the RDBDP approach of [77]. First- (3.37) and second-order (3.38) hedging on the left and right, respectively.

Figure 3.3: Example 1. Dashed and dotted vertical lines corresponding to VaR_{95} and ES_{99}, respectively. Hedging instruments as in table 3.1. PnL approximated through an independent Monte Carlo sample of size 2^{14}.

3.5.2. EXAMPLE 2: SINGLE HIGH-DIMENSIONAL OPTION

The high-dimensional examples are given on $d = m$ assets driven by Black-Scholes dynamics under the physical measure. First, we investigate a single Bermudan geometric call option ($J = 1$), with varying early exercise rights. The coefficients of the corresponding discretely reflected FBSDE system read as follows[5]

$$\mu(t,x) = \bar{\mu} \odot x, \quad \sigma(t,x) = \text{diag}(\bar{\sigma} \odot x)\Sigma, \tag{3.36}$$
$$f(t,x,y,z) = -ry - \Big(\frac{\bar{\mu} - (r-q)}{\bar{\sigma}}\Big)^T z\Sigma^{-1}, \quad l(x) \equiv g(x) = \max\left[\omega\left(\Big(\prod_{i=1}^{d} x^i\Big)^{1/d} - K\right), 0\right],$$

with $\omega \in \{-1, 1\}$ for put and call options, respectively. In the above, $\bar{\mu}, \bar{\sigma}, q \in \mathbb{R}^m$ are the drift, volatility and continuous dividend yield parameters of each asset, r is the risk free rate and $C = \Sigma^T\Sigma$ is the correlation structure between the assets with Σ being its Cholesky decomposition. This example appears often in the literature, see e.g. [32, 77], due to its special property in that the general d-dimensional problem can be reduced to a scalar one. We set the parameters in (3.36) according to [32], and consider a fixed $T = 2, X_0 = 100, r = 0.0, q = (0.02, \dots, 0.02), \omega = 1, c_{ij} = 0.75, i \neq j$.

We train a deep BSDE solver once, offline for each parameter setting with $N' = 100$ equally sized time intervals according to algorithm 2 and use the resulting approximations to recover option prices, Deltas in (3.28) and Gammas according to (3.30). We emphasize that the OSM scheme provides *all Gammas* in the high-dimensional framework, including all *cross-gammas*, i.e. the Gamma in (3.30) takes values in $\mathbb{R}^{d\times d}$. Due to the high-correlation and similarly to example 3.5.1 – see fig. 3.3a in particular –, we found that it is not sufficient to remove pure second-order sensitivities with respect to each individual asset ($\mathscr{I} = \{ii : 1 \le i \le m\}$), as there is a substantial cross-gamma exposure. Therefore, we choose to hedge the whole upper triangular part of the Hessian matrix in (3.12a) in case of delta-gamma hedging, i.e. $\mathscr{I} = \{ij : 1 \le i \le j \le d\}$. In order to be

[5] we denote element-wise multiplication by $\odot$

able to compute the coefficient matrix on the left-hand side of the second-order condition (3.12a), we augment the delta-gamma hedging portfolio (3.11) with $K = d(d+1)/2$ Gamma hedging instruments, as follows

- for the diagonal elements in the Gamma matrix ($i = j$) we choose standard European put options with the same strike and maturity $\widetilde{T} = 2T$;
- for the cross Gamma indices ($i \neq j$) we choose European exchange calls with strike $K^{ij} = 1$ and maturity $\widetilde{T} = 2T$. This in particular implies the corresponding hedging instruments' prices, Delta and Gammas are all available in closed form due to the Margrabe formula [107] – see appendix 3.B.

As discussed in section 3.4.4, these choices imply that the linear system corresponding to (3.12) has a sparse coefficient matrix, and thus the linear system in (3.31a) can be stored and solved efficiently for a large number of Monte Carlo simulations. We split the discussion of the example into four key aspects: the impact of moneyness (determined by strike K), early exercising (R), strength of volatility ($\bar{\sigma}$) and dimensionality ($d = m$).

OTM/ATM/ITM. We fix a European contract ($R = 1, \mathscr{R} = \{0, T\}$) on $d = m = 50$ assets, each with volatility $\bar{\sigma} = (0.25, \ldots, 0.25)$ and an initial condition of $X_0 = (100, \ldots, 100)$. We consider in-, at- and out-of-the-money strikes with $K = 90, 100, 110$, respectively, to investigate the impact of moneyness on the discrete replication accuracy both in case of delta and delta-gamma hedging. Results are collected in figure 3.4 and table 3.3. We see that Gamma hedging significantly outperforms delta hedging in terms of replication accuracy, irrespective of the moneyness. Across all values of moneyness, we found that a fortnightly rebalanced ($N = 10$) delta-gamma hedging strategy achieves the same replication accuracy as a daily rebalanced ($N = 100$) delta strategy. In particular, we find that once rebalanced at least weekly ($N = 20$) the delta-gamma hedging strategy achieves an expected shortfall of less than -0.4 for each moneyness, whereas delta rebalancing with the same frequency yields a much wider Profit-and-Loss distribution with more profound tail risk. All risk measures are collected in table 3.3 for fortnightly ($N = 10$) rebalanced portfolios in (3.2) and (3.11). In line with fig. 3.4, we find that the discrete replication becomes more difficult as the option is going out-of-the-money. Nonetheless, even in the OTM case, when the position is delta-gamma hedged with higher than monthly ($N = 5$) frequency, the corresponding replication achieves a VaR_{95} below -0.5, which is only attained by daily delta rebalancing. The delta-gamma hedged portfolio using alg. 2 and 3 achieves an order of magnitude higher replication accuracy across all risk measures consistently, regardless of the moneyness, in the high-dimensional setting ($d = m = 50$).

High volatility (OSM versus Huré et al. [77]). Next, we assess the impact of the strength of the volatility parameter $\bar{\sigma}$ in (3.36). In order to do so, we fix $R = 1$, consider an ATM strike of $K = 100$ and vary $\bar{\sigma}$ between $0.25, 0.5$ and 0.75 across all $d = m = 50$ assets, uniformly. The numerical results are given in figure 3.5 and table 3.4. As indicated by [114], the One Step Malliavin scheme in alg.2 outperforms the reference methods [77, 32] in the approximation accuracy of the Z process in (3.6), especially in settings of high-volatility and small time steps. This is demonstrated by figure 3.5a. We find that as the

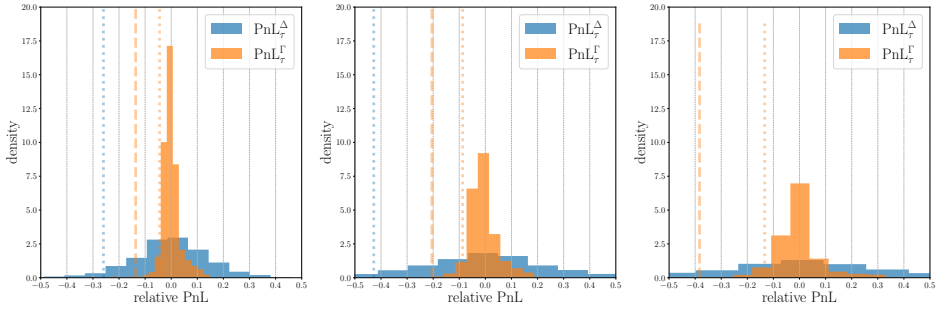

(a) $K = 90$ (ITM) (b) $K = 100$ (ATM) (c) $K = 110$ (OTM)

Figure 3.4: Example 2 in (3.36). Moneyness comparison with weekly rebalancing ($N = 20$). Example 2. Histograms for delta (3.3) and delta-gamma replication (3.12). Dotted and dashed vertical lines corresponding to VaR_{95} and ES_{99}, respectively.

	Delta			Delta-Gamma		
	$K = 90$	$K = 100$	$K = 110$	$K = 90$	$K = 100$	$K = 110$
mean	−8.8e−3	−7.4e−3	−2.7e−2	−7.1e−3	−1.4e−2	−4.1e−3
variance	4.4e−2	1.3e−1	3.3e−1	3.7e−3	7.9e−3	2.7e−2
VaR_{95}	−3.7e−1	−6.3e−1	−1.1	−7.7e−2	−1.2e−1	−2.0e−1
ES_{95}	−5.2e−1	−8.6e−1	−1.5	−1.4e−1	−1.9e−1	−3.3e−1
semivariance	2.4e−2	6.8e−2	2.0e−1	1.7e−3	2.8e−3	1.0e−2

Table 3.3: Example 2 in (3.36). Comparison of risk measures across moneyness. Fortnightly rebalancing ($N = 10$).

strength of the volatility increases the OSM scheme results in gradually more accurate Delta approximations compared to the reference method [77], which results in sharper PnL^Δ distributions. We found similar results in case of other references, such as [32]. This effect becomes more profound as $\bar{\sigma}$ increases, indicating that the One Step Malliavin scheme is not only useful in the context of second-order sensitivities, but may also result in higher first-order replication accuracy for highly volatile assets. Nonetheless, as shown by figure 3.5b, one can further improve the delta replication accuracy even by offsetting the corresponding second-order sensitivities in (3.11) – even for large $\bar{\sigma}$. We find that offsetting all option Gammas enabled by the OSM scheme, further improves the replication accuracy resulting in an approximately 30 percentage point improvement in value-at-risk and a sharper PnL distribution around 0. Unsurprisingly, as the volatility increases, the replication accuracy decreases, and one gains even more by delta-gamma hedging. All risk measures are collected in table 3.4, when the corresponding discretely rebalanced portfolios are updated daily. Comparing the delta hedging strategies between deep BSDE approximations provided by the OSM scheme and [77], we find that algorithm 2 provides higher accuracy in the Delta approximations when the volatility parameter is high, which results in higher replication accuracy. In particular, using the OSM scheme, one gains a roughly 3 times lower variance for the PnL distribution around 0. Additionally, see $\bar{\sigma} = (0.5, \ldots, 0.5)$, offsetting the associated Gammas results in an ad-

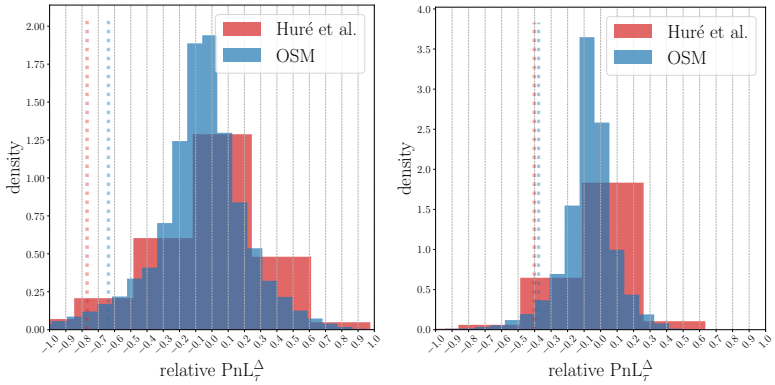

(a) Comparison between Deltas approximated by the OSM scheme (alg.2) and Huré et al. [77]. Left weekly ($N = 20$), right daily ($N = 100$) rebalancing.

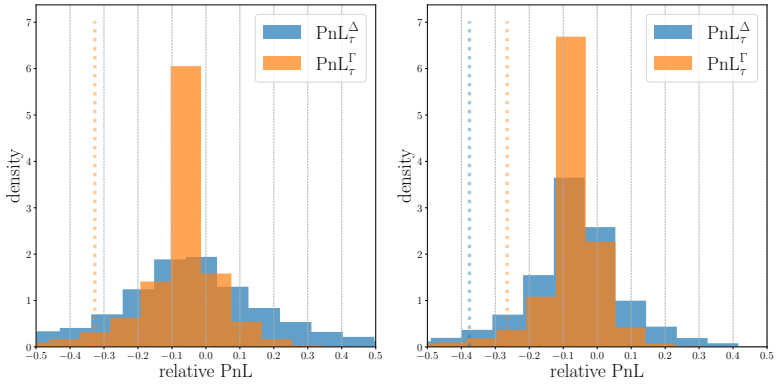

(b) Comparison of delta (3.3) and delta-gamma replication (3.12). Left weekly ($N = 20$), right daily ($N = 100$) rebalancing.

Figure 3.5: Example 2 in (3.36). Volatility impact comparison for $\bar{\sigma} = (0.75, \ldots, 0.75)$. Dotted and dashed vertical lines corresponding to VaR_{95} and ES_{99}, respectively.

ditional order of magnitude accuracy in the variance of the profit and loss distribution.

Early exercise rights. Let us investigate the impact of early exercising on the corresponding replication accuracies. We consider ATM ($K = 100$), Bermudan call options as in (3.36), issued on $d = m = 50$ underlyings, each admitting a volatility of 0.25. In order to assess the early exercise right's impact on the replication accuracy we vary the equidistant early exercise dates taking values $R = 5, 20, 100$, which can be associated with monthly, weekly and daily exercise rights, respectively. We remark that the case $R = 100$ can be thought of as a numerical approximation of the American option limit and can thus be directly compared to the results in [32]. The numerical experiments are summarized in figure 3.6 and table 3.5. As we can see, the OSM hedging strategies are robust and accurate with respect to the number of early exercise dates of the corresponding Bermudan option for both delta- and delta-gamma hedging. In fact, comparing the densities

	Huré et al. [77]		OSM (alg. 2)			
	Delta		Delta		Gamma	
	$\bar{\sigma} = 0.5$	$\bar{\sigma} = 0.75$	$\bar{\sigma} = 0.5$	$\bar{\sigma} = 0.75$	$\bar{\sigma} = 0.5$	$\bar{\sigma} = 0.75$
mean	−4.1e−3	−3.7e−2	−6.4e−3	−8.0e−2	−5.2e−3	−8.0e−2
variance	2.0e−2	1.3e−1	1.4e−2	3.8e−2	1.4e−3	2.3e−2
VaR_{95}	−2.4e−1	−4.0e−1	−2.1e−1	−3.8e−1	−6.8e−2	−2.6e−1
ES_{95}	−3.4e−1	−7.2e−1	−3.0e−1	−6.0e−1	−9.2e−2	−5.1e−1
semivariance	1.1e−2	8.5e−2	7.9e−3	3.8e−2	7.7e−4	4.4e−2

Table 3.4: Example 2 in (3.36). Comparison of risk measures across different levels of volatility. Delta hedging with OSM in alg. 2 versus Huré et al. in [77]. Daily rebalancing ($N = 100$).

	Delta			Delta-Gamma		
	$R = 5$	$R = 20$	$R = 100$	$R = 5$	$R = 20$	$R = 100$
mean	2.4e−2	7.5e−2	7.9e−2	2.9e−2	9.1e−2	9.7e−2
variance	1.3e−1	1.3e−1	1.2e−1	1.5e−2	3.5e−2	3.2e−2
VaR_{95}	−6.2e−1	−5.3e−1	−5.7e−1	−1.2e−1	−9.3e−2	−7.2e−2
ES_{95}	−9.0e−1	−8.6e−1	−8.3e−1	−2.1e−1	−1.9e−1	−1.8e−1
semivariance	7.7e−2	8.2e−2	7.4e−2	4.7e−3	5.7e−3	5.8e−3

Table 3.5: Example 2 in (3.36). Comparison of risk measures across increasing early exercise rights, from monthly to daily exercising. Fortnightly rebalancing ($N = 10$).

in fig. 3.6, we find that, even though the approximation of the corresponding discretely reflected BSDE becomes more challenging, the resulting delta-gamma hedged PnLs are still sharply distributed around 0. The scales of the vertical axes show that the replication accuracy goes down as R increases (ceteris paribus), due to the larger number of early exercised paths. Comparing the left and right columns we see that the additional second-order constraints yield a substantial improvement to delta hedging. The delta-gamma hedged portfolio achieves a similar profit and loss distribution with only fortnightly rebalancing ($N = 10$) as that of the delta hedging with daily readjusted hedging weights ($N = 100$). Risk measures for all hedging strategies are collected in table 3.5 for portfolios rebalanced every fortnight ($N = 10$). The delta-gamma hedging strategies enabled by the OSM approximations in algorithm 2 induce close to an order of magnitude improvement in the variance of the associated profit-and-loss distributions compared to mere delta-hedging, irrespective of early-exercise features, in the high-dimensional option setting with $d = m = 50$ underlyings.

Dimensionality. The main motivation behind the deep BSDE approximations for the One Step Malliavin scheme in algorithm 2 is to enable the treatment of high-dimensional basket options, issued on many (correlated) underlyings in (3.36). We fix $R = 100$ corresponding to the American option limit, and vary the number of risk factors $d = m$ in (3.36) between $1, 5, 20, 100$. Recall that these results can directly be compared to all aforementioned results given on 50 assets. The numerical results are collected in figure 3.7 and table 3.6. In figure 3.7, we see that the OSM enabled hedging strategies are robust with respect to the number of underlying assets, and both in case of delta and delta-

	Delta				Delta-Gamma			
d	1	5	20	100	1	5	20	100
mean	4.1e−3	1.7e−3	1.5e−2	3.0e−2	1.4e−3	3.9e−3	1.7e−2	3.2e−2
variance	1.0e−2	1.2e−2	1.4e−2	1.8e−2	2.6e−4	3.7e−4	1.6e−3	5.3e−3
VaR_{95}	−1.6e−1	−1.7e−1	−1.7e−1	−1.7e−1	−1.5e−2	−1.9e−2	−1.9e−2	−1.8e−2
ES_{95}	−2.3e−1	−2.3e−1	−2.5e−1	−2.5e−1	−2.5e−2	−3.3e−2	−2.8e−2	−3.3e−2
semivariance	4.7e−3	5.0e−3	6.4e−3	6.2e−3	9.9e−5	1.4e−4	1.7e−4	2.3e−4

Table 3.6: Example 2 in (3.36). Comparison of risk measures across increasing early exercise rights, from monthly to daily exercising. Fortnightly rebalancing ($N = 10$).

gamma hedging the shape of the corresponding profit-and-loss distributions are similar across $d = 5, 20, 100$. Observing the vertical scales from top to bottom, we see that both in the case of first- and second-order hedging, the peak of the distribution decreases, as the associated replication problem becomes more difficult due to dimensionality. Comparing the left and right columns indicates that including the additional second-order constraints in (3.12) brings a substantial improvement to the discrete replication accuracy. The delta-gamma hedged strategy reaches the same replication error with an order of magnitude less frequent rebalancing, as the mere delta hedging portfolio. The results above illustrate that delta-gamma replication enabled by OSM achieves a given risk tolerance with significantly less number of rebalancing dates compared to delta hedging, irrespective of the number of underlying assets. This is further demonstrated by figure 3.8, where the convergence of VaR_{95} is plotted against the number of rebalancing dates. As can be seen, for all considered number of underlying assets, the delta-gamma hedging strategy achieves a risk tolerance level of $\text{VaR}_{95} \leq 10\%$ for all rebalancing frequencies higher than $N = 10$. On the contrary, the delta hedged portfolio does not reach this accuracy, not even in case of daily rebalancing ($N = 100$). This implies that the approximation errors in alg. 2 are negligible compared to the time discretization. Moreover, in practical applications, where rebalancing is undesirable due to potential transaction costs, the delta-gamma hedging enabled by the Gamma approximations of OSM may achieve a given risk tolerance more efficiently. We collect all risk measures table 3.6 for fortnightly ($N = 10$) rebalanced replicating portfolios. For each strategy the risk measures grow as d increases. Nonetheless, both in case of delta and delta-gamma replication, the corresponding risk measures are robust with respect to the number of underlying assets. Across all values of d, the Gamma hedging strategies bring an order of magnitude improvement both in the variance and the tail risk measures of the PnL. These numbers demonstrate that the deep BSDE approximation for the OSM scheme efficiently deal with Bermudan basket options in (3.36) issued on a large number of underlyings. We emphasize that the second-order sensitivities in the OSM scheme are given by the matrix-valued Γ process in (3.18) which takes values in $\mathbb{R}^{d\times d}$, meaning that in case of $d = m = 100$ the OSM scheme accurately approximates 10^4 gammas simultaneously.

3.5.3. EXAMPLE 3: PORTFOLIO OF SEVERAL DERIVATIVES WITH DIFFERENT EARLY EXERCISE RIGHTS

In order to demonstrate the accuracy and robustness of the proposed hedging strategies in the context of a portfolio of multiple options, in our last example we investigate a high-

	contract	payoff
$j=1$	geometric put	$\max[K_1-(\prod_{i=1}^m x_i)^{1/m},0]$
$j=2$	arithmetic put on $(S_1,\dots,S_{m/2})$	$\max[K_2-2/m\sum_{i=1}^{m/2} x_i,0]$
$j=3$	call on max $(S_{m/2+1},\dots,S_m)$	$[\max_{i=m/2+1,\dots,m} x_i-K_3]^+$
$j=4$	cash or nothing	$\prod_{i=1}^m \mathbb{1}_{50\le x_i\le 150}$
$j=5$	put on min	$[K_5-\min_{i=m/2+1,\dots,m} x_i]^+$
$j=6,\dots,25$	call	$[x_i-K_j]^+$

Table 3.7: Example 3. Derivatives in the portfolio.

dimensional portfolio of derivatives. We take $d=m=20$ Black-Scholes type underlyings under the physical measure as in (3.36) , with parameters $\bar{\mu}=(0.2,0.19,\dots,0.01)$, $\bar{\sigma}=(0.4,0.25,0.2,0.15,0.1,\dots,0.4,0.25,0.2,0.15,0.1)$, $r=0.04$, $q=(0,\dots,0)$, $X_0=(100,\dots,100)$ and pairwise correlation $c_{ij}=0.25, i\neq j$. We consider a time horizon of $T=1$ year, over which $J=25$ Bermudan derivatives are held. The types of contracts are collected in table 3.7. In terms of moneyness and early exercise rights, we distinguish three different versions of the portfolio corresponding to

1. Case 1: all ATM ($K_j=100, j=1,\dots,J$), European ($R_j=1, j=1,\dots,J$) contracts;
2. Case 2:
   - varying moneyness: $K_1=100, K_2=120, K_3=80, K_5=50, K_j=150, j=6,\dots,25$,
   - Bermudan options with uniform, monthly early exercise rights: $R_j=5, j=1,\dots,J$;
3. Case 3:
   - varying moneyness: $K_1=100, K_2=120, K_3=80, K_5=50, K_j=150, j=6,\dots,25$,
   - varying early exercise rights: $R_1=20, R_2=5, R_3=2, R_4=1, R_5=10, R_6=R_7=5, R_8=R_9=10, R_{10}=\dots=R_{15}=100, R_{16}=2, R_{17}=\dots=R_{19}=20, R_{20}=\dots=R_{25}=100$.

For each choice above the solutions of the collection of discretely reflected BSDEs in (3.7) take values in $\mathbb{R}^J$, $\mathbb{R}^{J\times d}$ and $\mathbb{R}^{J\times d\times d}$ for the prices, Deltas and Gammas, respectively. Indeed, the Gamma process approximated by the One Step Malliavin scheme in alg. 3 has 10^4 elements. As this portfolio consists of multiple contracts, which may be exercised at different points in time, we provide the profit and loss distributions at maturity T, instead of τ in (3.34). This means that in case one of the options in (3.2) or (3.11) is exercised before T, the corresponding payoff is collected and the weights are computed with the remaining derivatives only.

The numerical results are collected in figures 3.9, 3.10 and table 3.8. Looking at the approximated PnL densities depicted in fig. 3.9, we find that the OSM scheme accurately approximates all Deltas and Gammas for all derivatives in the portfolio simultaneously. The replication is most accurate in the case of European contracts without early

exercise features, nevertheless, the approximations remain accurate with varying early exercise rights for each contract separately, and also for different levels of moneyness. Comparing the left and right columns, we find that for all three cases outlined above, the gamma hedging strategy yields a substantial improvement in terms of replication accuracy compared to the standard delta hedging. For both first- and second-order hedging, the PnL distributions are distributed around the origin with a decreasing variance as the number of rebalance dates increases. In particular, consistently throughout all examples the gamma hedged strategies reach the same shape of the PnL distribution with fortnightly rebalancing as the delta hedged positions with daily rebalancing. In terms of tail risk, VaR_{95} is approximately 50 percentage points better across the delta-gamma hedged portfolios.

Similar conclusions can be drawn from the histograms collected in the top row of figure 3.10. In line with the earlier results, the delta-gamma hedged OSM portfolios significantly outperform standard delta-hedging even for this portfolio of multiple derivatives. This is accentuated with less frequent rebalancing, where the delta-gamma hedging strategies achieve a much sharper PnL in the discrete time framework. These observations hold across all three cases outlined above, irrespective of the derivatives moneyness and early exercise features. This phenomenon is further demonstrated by the bottom of figure 3.10. Herein, the convergence of the 95 percentile value-at-risk against the number of rebalancing dates is collected. We find that VaR_{95} corresponding to the delta-gamma hedged PnL is consistently an order of magnitude smaller than in case of delta hedging, across all different number of rebalancing dates. In other words, any prespecified risk tolerance level – measured e.g. by value-at-risk – is achieved by the delta-gamma hedging strategy in an order of magnitude less frequent rebalancing. In particular, we find that the multi-dimensional portfolio extension of the OSM scheme yields practically identical conclusions as in case of single options, implying that the corresponding deep BSDE approximations truly excel in case of these high-dimensional equations.

Finally, risk measures corresponding to each hedging strategy are collected in table 3.8 for $N = 2$, i.e. portfolios rebalanced once every quarter. All risk measures confirm our findings above. In particular, delta-gamma hedging the portfolio of Bermudan derivatives brings an order of magnitude improvement compared to delta hedging in terms of the variance of the profit-and-loss distribution at maturity for all three cases of early exercise features and moneyness. Additionally, we find that offsetting second-order sensitivities in (3.12) enabled by the Gamma approximations of the One Step Malliavin scheme, brings a substantial improvement also in terms of tail risk measured both by value-at-risk and expected shortfall. In particular, both for the 95 and 99 percentile tails, the delta-gamma hedged distributions provide a factor of 3 improvement compared to the standard delta strategies. In light of the dimensionality considered in this last example, with a portfolio consisting of $J = 25$ derivative contracts issued on $d = m = 20$ underlyings with varying drift and volatility parameters, we can conclude that the deep BSDE approximations of the OSM scheme provide a robust and accurate approximations for all option Greeks in the portfolio up to second order.

	Delta			Delta-Gamma		
	ATM	Bermudan	Mixture	ATM	Bermudan	Mixture
mean	−2.6e−2	−2.2e−3	3.0e−2	−2.9e−4	2.9e−2	6.0e−2
variance	2.7e−2	1.2e−1	1.1e−1	1.3e−3	1.8e−2	1.7e−2
VaR_{95}	−3.0e−1	−6.5e−1	−5.9e−1	−5.9e−2	−1.8e−1	−1.4e−1
ES_{95}	−4.5e−1	−1.0	−9.3e−1	−7.2e−2	−3.1e−1	−2.6e−1
semivariance	1.7e−2	1.1e−1	1.0e−1	4.5e−4	1.2e−2	1.0e−2

Table 3.8: Example 3. Convergence of VaR_{95} against the discrete number of rebalancing dates in (3.2) and (3.11).

3.6. CONCLUSION

In this chapter we proposed a novel deep BSDE based approach for the simultaneous pricing, delta and delta-gamma hedging of large, high-dimensional portfolios of Bermudan options. First, we gave a vector-valued extension to the One Step Malliavin scheme in [114]. This way, we casted the pricing, delta- and delta-gamma hedging of a portfolio of Bermudan options into the framework of a *system of discretely reflected BSDEs*. Subsequently, we proposed a deep BSDE approach for the accurate numerical solution of this collection of equations, which is robust and efficient when the number of underlying risk factors and/or options is large. In fact, our approach includes not only prices and Deltas but also second-order Greeks, Gammas of all options in the portfolio, simultaneously approximated, throghout the entire spacetime. We demonstrated the performance of our algorithm on several examples, highlighting key features of our technique. Our findings suggest that, the hereby proposed OSM approximations outperform reference methods [77, 32] even in the context of delta hedging, when the risk factors are highly volatile. Most importantly, by performing delta-gamma hedging enabled by the Gamma approximations of the OSM scheme, we managed to improve on the discrete replication accuracy compared to standard delta hedging. Our results demonstrate that the algorithm is robust and accurate, for different levels of moneyness, high volatility, and early exercise rights up to the American option limit even in case high dimensional basket options issued on $d = 100$ assets.

3.A. BEYOND DELTA-GAMMA HEDGING

In principle, our framework is more general than mere delta-gamma hedging, as it allows for the treatment of hedging stochastic volatility. In that case, one needs to complete the market with an extra set of securities that have non-vanishing *Vegas*, such that volatility risk can be hedged, in a similar fashion to (3.11) with first-order constraints only. Hence, an appropriate delta-vega hedging strategy on the portfolio (3.11) would therefore yield

the following first-order conditions

$$\sum_{k=1}^{K} \beta_t^k \partial_l u^k(t, X_t) = \sum_{j=1}^{J} \partial_l v^j(t, X_t), \qquad l = d-m+1, \dots, d, \tag{3.37a}$$

$$\alpha_t^i = \sum_{j=1}^{J} \partial_i v^j(t, X_t) - \sum_{k=1}^{K} \beta_t^k \partial_i u^k(t, X_t), \qquad i = 1, \dots, m. \tag{3.37b}$$

Similar to the previous discussion, (3.37) only offsets first-order sensitivities of the portfolio with respect to asset prices and volatility. In order to mitigate replication errors stemming from the discrete rebalancing, one has the conditions determining the optimal hedging weights by appropriate second-order constraints including second-order sensitivities corresponding to not just asset prices (Gammas) but also to volatility such as *Vommas* and *Vannas*. Therefore, in case such second-order sensitivities are also taken into account in the presence of non-hedgeable risk factors such as stochastic volatility, one can impose the following extra second-order conditions

$$\sum_{k=1}^{K} \beta_t^k \partial_{il}^2 u^k(t, X_t) = \sum_{j=1}^{J} \partial_{il}^2 v^j(t, X_t), \quad il \in \mathscr{I}, \tag{3.38}$$

with some arbitrarily chosen index set $\mathscr{I}$. Note that in this case (3.38) and (3.37a) together form a linear system of size $(d-m+|\mathscr{I}|) \times K$. Natural choices for the index set $\mathscr{I}$ include

- $\mathscr{I} = \{il : 1 \le i = l \le m\}$ – diagonal elements of the Γ matrix are hedged, corresponding to second order sensitivities of the tradeable assets;
- $\mathscr{I} = \{il : 1 \le i, l \le m\}$ – the whole Γ matrix is hedged, with cross-gammas included;
- $\mathscr{I} = \{il : m+1 \le i, l \le d\}$ – second-order sensitivities with respect to the volatility, i.e. *Vomma*s, including cross-vommas,
- $\mathscr{I} = \{il : 1 \le i, l \le d\}$ – all second-order Greeks with respect to each underlying risk factor, e.g. including mixed partial derivatives in asset prices and volatilities *Vanna*s.

We emphasize that the focus above on volatility risk is mere illustration. In fact, the abstract FBSDE framework allows for the treatment of second-order Greeks of all types. For instance, one could consider stochastic interest rates and compute the Vera/rhova as the mixed partial derivative with respect to volatility and interest rate; as long as the corresponding risk factor forms part of the Itô diffusion in (3.1). As is made clear by the discrete backward recursion in (3.9), the only type of second-order Greeks where one needs to take additional measures is derivatives including time, e.g. *Charm*. Nonetheless, provided that the time partition used for the discrete time resolution of the discretely reflected FBSDE system associated with the portfolio is fine enough, one can approximate such Greeks by finite difference type approximations.[6] Therefore, the framework built

[6] e.g. $\text{Charm}_n^\pi = -\frac{\alpha_n^\pi - \alpha_{n-1}^\pi}{t_n - t_{n-1}}$

on the One Step Malliavin scheme allows for the hedging of all Greeks up to second order, as long as the corresponding risk factors are incorporated in associated discretely reflected FBSDE.

3.B. MARGRABE FORMULA WITH DIVIDENDS

We have the standard Black-Scholes model (3.36) in which all assets follow a geometric Brownian motion

$$\mathrm{d}S_t^i/S_t^i = (r-q^i)\mathrm{d}t + \sigma^i \mathrm{d}W_t^i, \quad 1 \le i \le m \equiv d.$$

The Brownian motions are pairwise correlated with a correlation parameter ϱ^{ij}. The Margrabe formula [107] then gives an explicit, closed-form analytical expression for the price of a European type contract whose payoff is as follows

$$g(t,S) = [S_T^k - K^{kj} S_T^j]^+,$$

at some terminal time T, where the parameter K^{kj} is an exchange strike. The price of such a contract then satisfies

$$C^{kj}(t,S) = e^{-q^k(T-t)} S_t^k \Phi(d_1^{kj}(t,S)) - e^{-q^j(T-t)} K^{kj} S_t^j \Phi(d_2^{kj}(t,S)),$$

where Φ denotes the standard normal cumulative distribution function and

$$\begin{aligned} d_1^{kj}(t,S) &= \frac{1}{\sigma^{kj}\sqrt{T-t}} \left(\ln\left(\frac{S_t^k}{K^{kj} S_t^j}\right) + (q_j - q_k + (\sigma^{kj})^2/2)(T-t) \right), \\ d_2^{kj}(t,S) &= d_1^{kj}(t,S) - \sigma^{kj}\sqrt{T-t}, \\ \sigma^{kj} &= \sqrt{(\sigma^k)^2 + (\sigma^j)^2 - 2\varrho^{kj}\sigma^k\sigma^j}. \end{aligned}$$

Straightforward computations lead to the first-order derivatives satisfying

$$\partial_k C^{kj}(t,S) = e^{-q^k(T-t)} \Phi(d_1^{kj}(t,S)), \quad \partial_j C^{kj}(t,S) = -e^{-q^j(T-t)} K^{kj} \Phi(d_2^{kj}(t,S)).$$

Further differentiation yields the second-order derivatives

$$\partial^2_{kk} C^{kj}(t,S) = \frac{e^{-q^k(T-t)}}{\sigma^{kj}\sqrt{T-t}} \frac{\phi(d_1^{kj}(t,S))}{S_t^k}, \qquad \partial^2_{jk} C^{kj}(t,S) = -\frac{e^{-q_k(T-t)}}{\sigma^{kj}\sqrt{T-t}} \frac{\phi(d_1^{kj}(t,S))}{S_t^j},$$

$$\partial^2_{kj} C^{kj}(t,S) = -\frac{e^{-q_j(T-)}}{\sigma^{kj}\sqrt{T-t}} \frac{K^{kj}\phi(d_2^{kj}(t,S))}{S_k}, \qquad \partial^2_{jj} C^{kj}(t,S) = \frac{e^{-q_j(T-t)}}{\sigma^{kj}\sqrt{T-t}} \frac{K^{kj}\phi(d_2^{kj}(t,S))}{S_t^j},$$

where ϕ is the standard normal density. Note that $\partial^2_{jk} C^{kj}(t,S) \equiv \partial^2_{kj} C^{kj}(t,S)$.

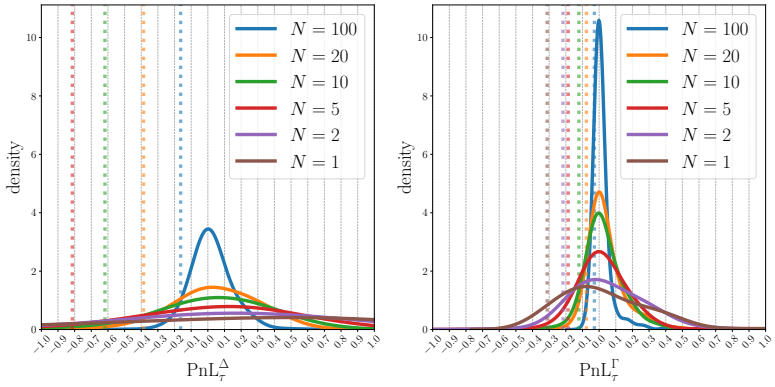

(a) monthly early exercise ($R = 5$)

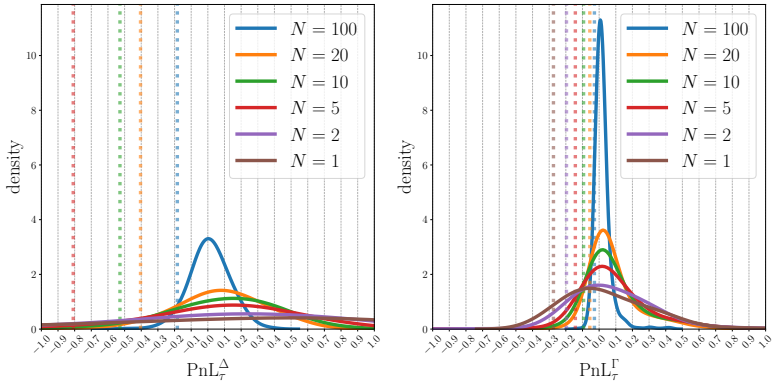

(b) weekly early exercise ($R = 20$)

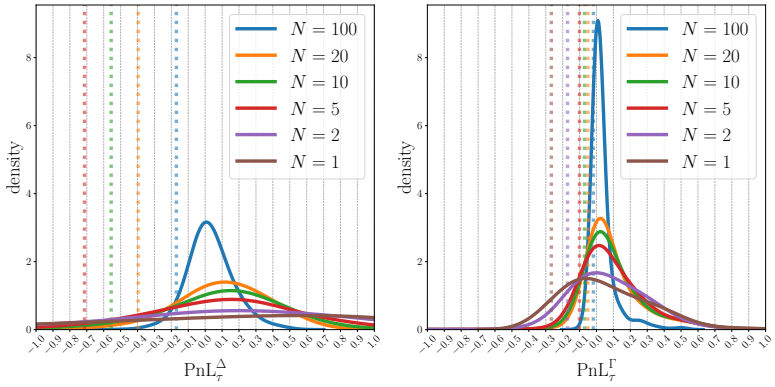

(c) daily early exercise ($R = 100$)

Figure 3.6: Example 2 in (3.36). Increasing early exercises rights, from Bermudan to American. PnL densities for several rebalancing frequencies. Left: delta hedging 3.3, right: delta-gamma hedging (3.12). Dotted vertical lines corresponding to VaR_{95}.

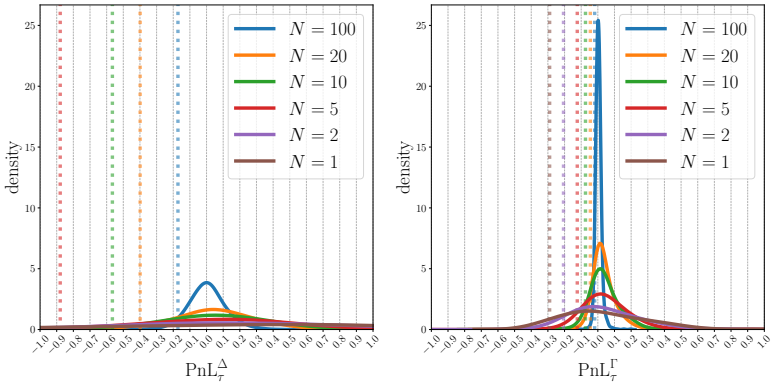

(a) $d = 5$

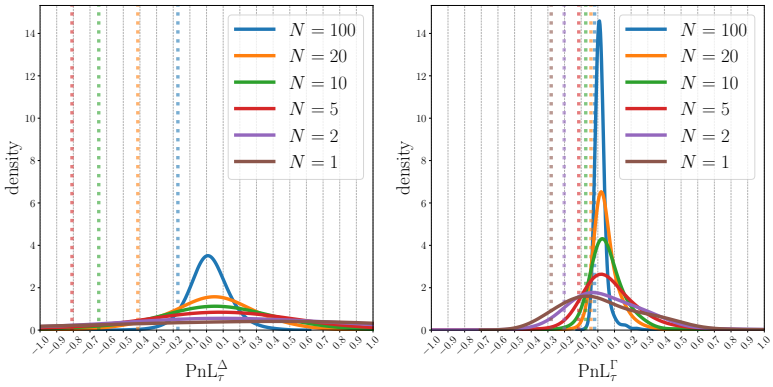

(b) $d = 20$

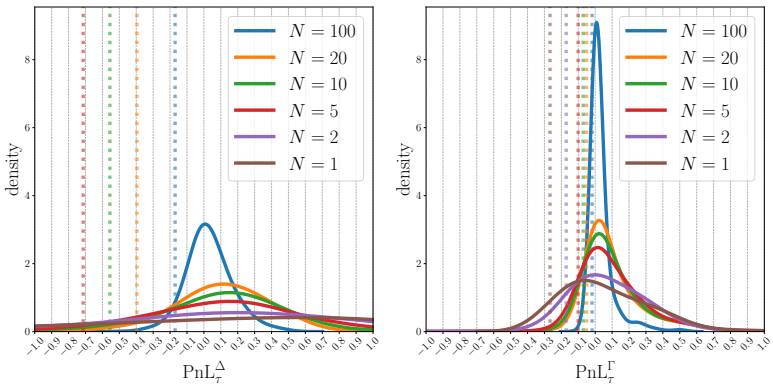

(c) $d = 100$

Figure 3.7: Example 2 in (3.36). Comparison on dimensionality, increasing number of underlying assets. PnL densities for several rebalancing frequencies. Left: delta hedging 3.3, right: delta-gamma hedging (3.12). Dotted vertical lines corresponding to VaR_{95}.

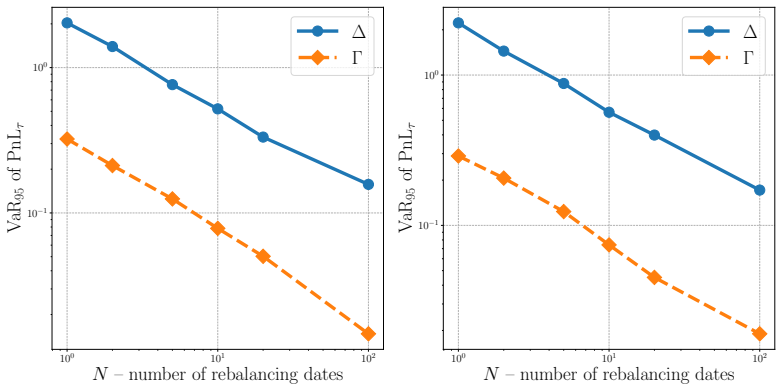

(a) $d = 1$ (b) $d = 5$

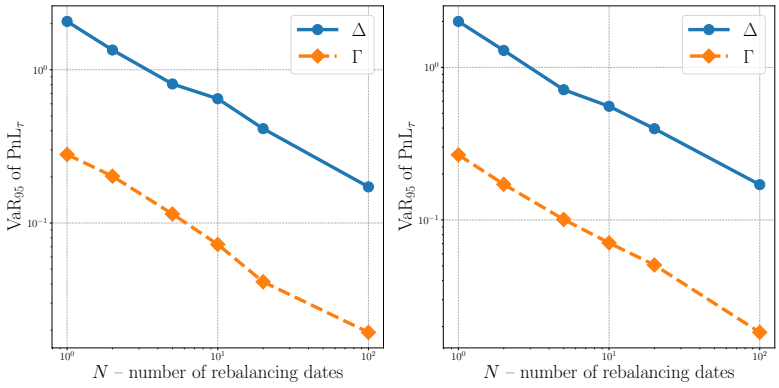

(c) $d = 20$ (d) $d = 100$

Figure 3.8: Example 2 in (3.36). Comparison on dimensionality, increasing number of underlying assets. Convergence of VaR_{95} against the discrete number of rebalancing dates in (3.2) and (3.11).

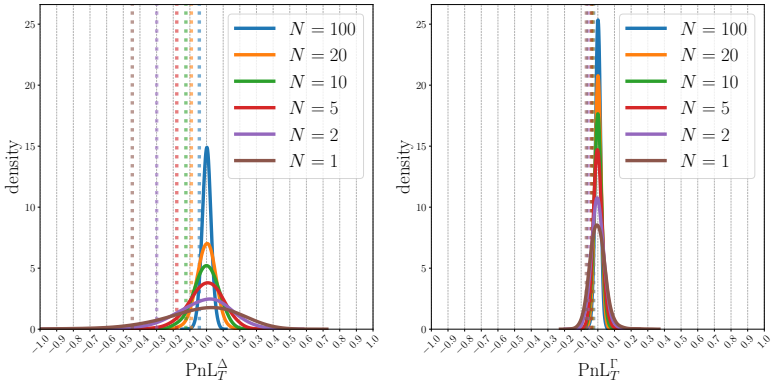

(a) Case 1: uniformly European, ATM

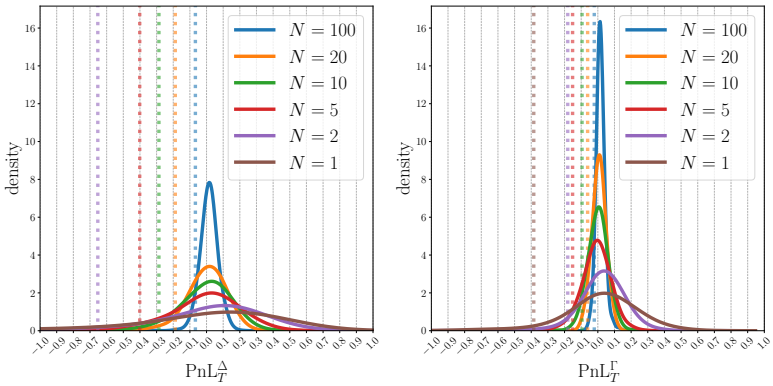

(b) Case 2: uniform monthly ($R = 5$) early exercise rights, mixed moneyness

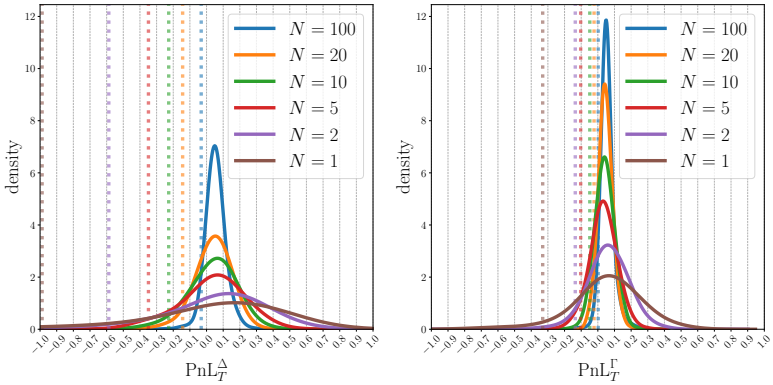

(c) Case 3: varying early exercise rights, mixed moneyness

Figure 3.9: Example 3. PnL densities for several rebalancing frequencies. Left: delta hedging (3.3), right: delta-gamma hedging (3.12). Dotted vertical lines corresponding to VaR_{95}.

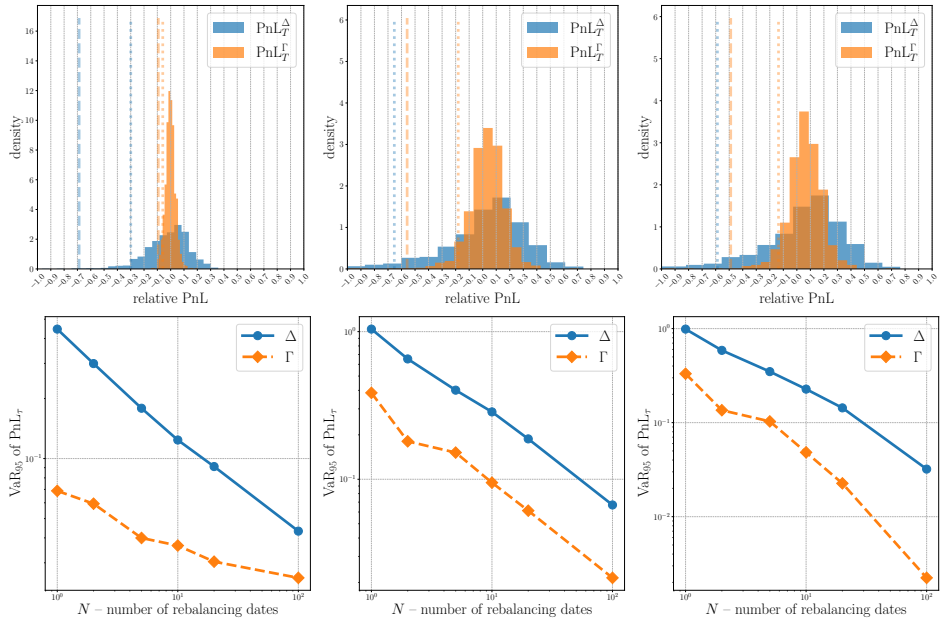

(a) Case 1 (b) Case 2 (c) Case 3

Figure 3.10: Example 3. Top: histograms for delta (3.3) and delta-gamma replication (3.12) with quarterly ($N = 2$) rebalancing. Bottom: convergence of VaR_{95} with respect to the number of rebalancing dates. Dotted and dashed vertical lines corresponding to VaR_{95} and ES_{99}, respectively.

4

GENERALIZED CONVERGENCE OF THE DEEP BSDE METHOD: A STEP TOWARDS FULLY-COUPLED FBSDES AND APPLICATIONS IN STOCHASTIC CONTROL

Minden olyan, mintha az volna ami, de mégsem.
Hogy tudsz így aludni, mikor tudod, hogy nem alszom én sem?[1]

Quimby (Álmatlan dal)

This article is based on the preprint [115]: B. Negyesi, Z. Huang, and C. W. Oosterlee. *Generalized convergence of the deep BSDE method: a step towards fully-coupled FBSDEs and applications in stochastic control.* 2025. DOI: 10.48550/arXiv.2403.18552

[1] Everything is as if it were what it is, yet it is not.
How can you sleep like this when you know that I cannot?

Quimby (Sleepless song)

4.1. INTRODUCTION

In this chapter, we are concerned with the numerical approximation of a system of coupled forward-backward stochastic differential equations (FBSDE) over a finite time interval $[0,T]$

$$\begin{cases} X_t = x_0 + \int_0^t b(s,X_s,Y_s,Z_s)\mathrm{d}s + \int_0^t \sigma(s,X_s,Y_s)\mathrm{d}W_s, \\ Y_t = g(X_T) + \int_t^T f(s,X_s,Y_s,Z_s)\mathrm{d}s - \int_t^T Z_s\mathrm{d}W_s, \end{cases} \tag{4.1}$$

where $b : [0,T]\times\mathbb{R}^d\times\mathbb{R}^q\times\mathbb{R}^{q\times m}\to\mathbb{R}^d$, $\sigma : [0,T]\times\mathbb{R}^d\times\mathbb{R}^q\to\mathbb{R}^{d\times m}$, $f : [0,T]\times\mathbb{R}^d\times\mathbb{R}^q\times\mathbb{R}^{q\times m}\to\mathbb{R}^q$ and $g : \mathbb{R}^d\to\mathbb{R}^q$ are all deterministic mappings. The equation is given on a complete probability space $(\Omega,\mathscr{F},\mathbb{P})$, over which $\{W_t\}_{0\le T}$ is a standard m-dimensional Brownian motion, $\mathscr{F} := \{\mathscr{F}_t\}_{0\le t\le T}$ its natural filtration and augmented by the usual $\mathbb{P}$-null sets. A triple of $(\mathbb{R}^d\times\mathbb{R}^q\times\mathbb{R}^{q\times m})$ valued, $\mathscr{F}_t$ adapted stochastic processes $\{(X_t,Y_t,Z_t)\}_{0\le t\le T}$ is a solution if (4.1) holds $\mathbb{P}$ almost surely and the processes satisfy natural integrability conditions, see [161, 104].

The study of FBSDEs was initiated by the seminal paper of Pardoux and Peng [126], and then extended to coupled equations by Antonelli [6]. Equations like (4.1) subsequently attracted widespread attention due to their inherent connections with systems of second-order quasi-linear partial differential equations (PDE) established by nonlinear extension to the Feynman-Kac lemma, see e.g. [104, 161] and theorem 4.3.1 below. This probabilistic representation casts FBSDEs to be the natural framework to model a wide range of problems arising in mathematical finance, physics, biology and stochastic control. The well-posedness of (4.1) has been rigorously studied and established under by now standard assumptions, see e.g. [104, 129, 39, 38] and the references therein. Most of such results rely either on classical solutions of the corresponding quasi-linear PDE with high regularity or some abstract conditions heuristically associated with monotonicity, small time duration or weak coupling. In the rest of the chapter, we consider the setting where (4.1) is well-posed and admits a unique strong solution triple.

Solving FBSDEs analytically is seldom possible and one usually has to resort to numerical approximations. In the decoupled framework where b,σ in (4.1) do not depend on Y,Z, one can detach the solution of X, and use the resulting discrete time approximation in order to approximate the solution pair of the BSDE by sequences of backward, recursive conditional expectations, we refer to [63, 24, 13, 16, 144]. In the coupled framework, things become more subtle due to the interdependence from the backward equation into the forward diffusion. Inspired by [102], classical approaches usually consist of *decoupling* the forward diffusion by means of deterministic mappings, and iteratively converging to the unique decoupling field related to the associated quasi-linear PDE, see e.g. [102, 41, 36, 40, 17, 76]. A common challenge across the aforementioned classical references is the setting of high-dimensionality. In fact, whenever either d,q or m are large, these methods suffer from the curse dimensionality and become intractable in high dimensions. Such settings arise naturally, for instance, in portfolio allocation or climate risk management. A recently emerging branch of numerical algorithms called *deep BSDE* methods and pioneered by [68, 43] addressed this gap, and has shown remarkable empirical performance in terms of tackling high-dimensional FBSDEs and associated quasi-linear PDEs. These approaches were first developed for decoupled equations, see

e.g. [68, 43, 77, 54, 114], and then extended to the coupled framework [69, 4]. For an overview, we refer to the recent survey [34].

Motivated by their outstanding empirical performance, serious efforts have been made in order to establish convergence guarantees for such deep BSDE algorithms. Originally the pioneering paper of Han and Long [69] managed to show a *posteriori bound* which depends on the objective loss functional of the machine learning algorithm. Their result was later extended in [87] to the case of non-Lipschitz continuous drift coefficients, and in [74] to the vector-valued framework in the context of stochastic optimal control. These works all have in common that they relied on the assumption of a narrower class of FBSDEs, in fact they considered a special case of (4.1) where only the Y process enters the dynamics of X, and not Z. Consequently, these convergence results were inapplicable in the context a wide range of stochastic control problems, for instance formulated through the dynamic programming principle [132], where coupling occurs in Z. In particular, these works could not provide a theoretical explanation for the phenomena observed in [4], where they found empirical evidence for the non-convergence of the deep BSDE method for FBSDEs stemming from stochastic control. More recently, the authors of [143] proved an *a-posteriori* error estimate in the framework of McKean-Vlasov FBSDEs. They give a new result for general, fully-coupled equations even a mean-field term, extending earlier analyses for decoupled equations [15]. Nonetheless, their result [143, corollary 4.5] is based on a different set of assumptions as the ones considered herein, which is concerned with the framework of [69, theorem 1].

The main motivation of the present chapter is to extend the convergence result of Han and Long [69] to drift coefficients in (4.1) that depend on the Z process as well. The main challenge is to handle the error estimate of the X process with extra Z coupling, which we control by the new estimates established in lemma 4.3.2. This enables us to derive our main result, stated in theorem 4.3.3, which is a posteriori error estimate similar to [69]. In particular, our work enjoys several relevant features:

- we recover the results of [69] in the limit case of no Z coupling;
- our result is applicable for a general class of FBSDEs, including, but not limited to, the ones obtained for stochastic control problems stemming either from dynamic programming [4] or the stochastic maximum principle [85, 74], due to the coupling of not just Y but also Z in the forward process;
- given a particular FBSDE, we can check whether or not it satisfies the convergence conditions through a directly verifiable approach.

The chapter is organized as follows. In section 4.2, we give the discrete time approximation scheme of the deep BSDE algorithm with Z coupling in the drift b. Section 4.3 contains our main result stated in theorem 4.3.3. Thereafter the abstract sufficient conditions of convergence are analyzed in section 4.4. In particular, we show that the assumptions of theorem 4.3.3 hold under heuristic interpretations such as *weak coupling* or *small time duration*. Additionally, we get earlier convergence results [69] as a limit case of our more general theory. Finally, we demonstrate our theoretical contributions by several numerical experiments on high-dimensional FBSDEs in section 4.5. These simulations confirm and showcase our theoretical findings.

4.2. THE DEEP BSDE ALGORITHM

In this section we formulate the deep BSDE algorithm for FBSDEs as in (4.1), naturally extending [68, 43, 69] to the framework of Z coupling in b. For the rest of the chapter, we denote the Frobenius norm by $\|x\|$ for any $x \in \mathbb{R}^{i\times j}$, not to be confused with the matrix 2-norm $\|x\|_2$. Without loss of generality, we work with an equidistant time partition $\pi := \{t_i, i = 0,\dots,N | 0 = t_0 < t_1 < \cdots < t_N = T\}$ with $h = T/N$ and study the following discretization

$$\inf_{\varphi_0 \in \mathcal{N}_0^Y(\theta_0^Y), \zeta_i \in \mathcal{N}_i^Z(\theta_i^Z)} \mathbb{E}\left[\left\| g\left(X_{t_N}^\pi\right) - Y_{t_N}^\pi \right\|^2\right], \tag{4.2a}$$

$$\text{s.t.} \begin{cases} X_0^\pi = x_0, \\ Y_0^\pi = \varphi_0(x_0;\theta_0^Y), \\ X_{t_{i+1}}^\pi = X_{t_i}^\pi + b\left(t_i, X_{t_i}^\pi, Y_{t_i}^\pi, Z_{t_i}^\pi\right)h + \sigma\left(t_i, X_{t_i}^\pi, Y_{t_i}^\pi\right)\Delta W_i, \\ Z_{t_i}^\pi = \zeta_i(X_{t_i}^\pi;\theta_i^Z), \\ Y_{t_{i+1}}^\pi = Y_{t_i}^\pi - f\left(t_i, X_{t_i}^\pi, Y_{t_i}^\pi, Z_{t_i}^\pi\right)h + Z_{t_i}^\pi \Delta W_i, \end{cases} \tag{4.2b}$$

for $i = 0,\dots,N-1$, where we put $\Delta W_i := W_{t_{i+1}} - W_{t_i}$. In doing so, the numerical solution of a coupled FBSDE (4.1) is reformulated into a stochastic optimization problem consisting of the minimization of an objective functional (4.2a) subject to the Euler-Maruyama discretization (4.2b). As in the continuous limit the loss functional (4.2a) attains 0 at the unique solution triple $\{(X_t, Y_t, Z_t)\}_{0\le t\le T}$ of (4.1) while also satisfying (4.2b), it is expected that for sufficiently large N and sufficiently wide function spaces $\mathcal{N}_0^Y, \mathcal{N}_i^Z$, the solution of (4.2a)-(4.2b) is an accurate discrete time approximation of (4.1). Motivated by universal approximation arguments, see e.g. [71], we set $\mathcal{N}_0^Y(\theta_0^Y), \mathcal{N}_i^Z(\theta_i^Z)$ to be spaces of deep neural networks parametrized by θ_0^Y, θ_i^Z for $i = 0,\dots,N-1$. Subsequently, the goal of the deep BSDE method is to solve this non-linear, constrained optimization problem through the training of deep neural networks. Hence, we seek to find $\varphi_0(x_0;\theta_0^Y) \in \mathcal{N}_0^Y$ and $\zeta_i(X_{t_i}^\pi;\theta_i^Z) \in \mathcal{N}_i^Z$ that approximate Y_0 and Z_{t_i} sufficiently well. The resulting pseudo-code for the complete deep BSDE method is collected in algorithm 4, its implementation is discussed in section 4.5.

4.3. CONVERGENCE ANALYSIS

This section is dedicated to our generalized convergence analysis for the deep BSDE method reviewed in section 4.2 and the discrete scheme (4.2b) for (4.1). In particular, we will show that the approximation errors of the numerical solution to the FBSDE are bounded by the simulation error of the objective function corresponding to (4.2a) which could be arbitrarily small due to the universal approximation theorem. The convergence analysis follows a similar strategy as that of [69]. We first introduce the standing assumptions and review some useful result. Throughout the chapter we use the notation $\mathbb{E}_i[\cdot] := \mathbb{E}[\cdot|\mathscr{F}_i]$.

Algorithm 4 Deep BSDE algorithm

Input: Initial parameters $(\theta_0^{Y,(0)}, \theta_i^{Z,(0)})$, learning rate η; batch size M; number of iterations K.
Data: Simulated Brownian increments $\{\Delta W_{t_i}^{(k)}\}_{0\le i\le N-1, 1\le k\le K}$
Output: Discrete time approximations $\{(\hat{X}_{t_i}^{\pi}, \hat{Y}_{t_i}^{\pi}, \hat{Z}_{t_i}^{\pi})\}_{i=0,\dots,N}$
for $k = 1$ to K **do** ▷ Euler-Maruyama (4.2b)
  $X_{t_0}^{\pi,(k)} = x_0$, $Y_{t_0}^{\pi,(k)} = \varphi_0(x_0;\theta_0^{Y,(k-1)})$
  for $i = 0$ to $N-1$ **do**
    $Z_{t_i}^{\pi,(k)} = \zeta_i(X_{t_i}^{\pi,(k)};\theta_i^{Z,(k-1)})$
    $X_{t_{i+1}}^{\pi,(k)} = X_{t_i}^{\pi,(k)} + b(t_i, X_{t_i}^{\pi,(k)}, Y_{t_i}^{\pi,(k)}, Z_{t_i}^{\pi,(k)})h + \sigma(t_i, X_{t_i}^{\pi,(k)}, Y_{t_i}^{\pi,(k)})\Delta W_{t_i}^{(k)}$
    $Y_{t_{i+1}}^{\pi,(k)} = Y_{t_i}^{\pi,(k)} - f(t_i, X_{t_i}^{\pi,(k)}, Y_{t_i}^{\pi,(k)}, Z_{t_i}^{\pi,(k)})h + Z_{t_i}^{\pi,(k)}\Delta W_{t_i}^{(k)}$
  end for
  Loss $= \frac{1}{M}\sum_{j=1}^{M} \|g(X_{t_N}^{\pi,(k)}) - Y_{t_N}^{\pi,(k)}\|^2$ ▷ empirical (4.2a)
  $(\theta_0^{Y,(k)}, \theta_0^{Z,(k)}, \dots, \theta_{N-1}^{Z,(k)}) = (\theta_0^{Y,(k-1)}, \theta_0^{Z,(k-1)}, \dots, \theta_{N-1}^{Z,(k-1)}) - \eta\nabla\text{Loss}$ ▷ SGD
end for
$(\hat{X}_i^{\pi}, \hat{Y}_{t_i}^{\pi}, \hat{Z}_{t_i}^{\pi}) = (X_{t_i}^{\pi,(K+1)}, Y_{t_i}^{\pi,(K+1)}, Z_{t_i}^{\pi,(K+1)}), \quad i = 0, \dots, N-1$

Assumption 4.3.1
There exist constants k^b and k^f, that are possibly negative, such that

$$\big(b(t,x_1,y,z) - b(t,x_2,y,z)\big)^\top \Delta x \le k^b \|\Delta x\|^2,$$
$$\big(f(t,x,y_1,z) - f(t,x,y_2,z)\big)^\top \Delta y \le k^f \|\Delta y\|^2.$$

Assumption 4.3.2
b, σ, f, g are uniformly Lipschitz continuous with respect to (x,y,z). In particular, there are non-negative constants such that

$$\big\|b(t,x_1,y_1,z_1) - b(t,x_2,y_2,z_2)\big\|^2 \le L_x^b\|\Delta x\|^2 + L_y^b\|\Delta y\|^2 + L_z^b\|\Delta z\|^2,$$
$$\big\|\sigma(t,x_1,y_1) - \sigma(t,x_2,y_2)\big\|^2 \le L_x^\sigma\|\Delta x\|^2 + L_y^\sigma\|\Delta y\|^2,$$
$$\big\|f\big(t,x_1,y_1,z_1\big) - f\big(t,x_2,y_2,z_2\big)\big\|^2 \le L_x^f\|\Delta x\|^2 + L_y^f\|\Delta y\|^2 + L_z^f\|\Delta z\|^2,$$
$$\big\|g(x_1) - g(x_2)\big\|^2 \le L_x^g\|\Delta x\|^2.$$

Assumption 4.3.3
$g(0), b(t,0,0,0), f(t,0,0,0)$ and $\sigma(t,0,0)$ are bounded for $t \in [0,T]$.

Notice that assumption 4.3.2 implies assumption 4.3.1 with $k^b, k^f \ge 0$. The reason for allowing for negativity shall be made clear by the forthcoming convergence result, see points ((3)d), ((3)e) in section 4.4 below. For convenience, we use $\mathscr{L}$ to denote the set of all constants mentioned above and assume L is the upper bound of $\mathscr{L}$.

Next, we introduce the following system of quasi-linear parabolic PDEs associated

with FBSDE (4.1),

$$\begin{cases} \partial_t v^i + \frac{1}{2}\partial_{xx} v^i : \sigma\sigma^\top(t,x,v) + \partial_x v^i\, b(t,x,v,\partial_x v\sigma(t,x,v)) + f^i(t,x,v,\partial_x v\sigma(t,x,v)) = 0, \\ v(T,x) = g(x), \quad \forall i = 1,\cdots,q, \end{cases} \tag{4.3}$$

The following assumption is needed in order to guarantee convergence of the implicit Euler-Maruyama scheme in theorem 4.3.2.

Assumption 4.3.4
The PDE (4.3) has a classical solution v with bounded derivatives $\partial_x v$ and $\partial^2_{xx} v$, and σ is bounded.

The non-linear Feynman-Kac lemma, stated below, establishes the connection between (4.3) and (4.1).

Theorem 4.3.1 (Feynman-Kac)
Under assumptions 4.3.2, 4.3.3 and 4.3.4, the FBSDE (4.1) has a unique solution (X,Y,Z), and it holds that for $t \in [0,T]$,

$$Y_t = v(t,X_t), \quad Z_t = \partial_x v(t,X_t)\,\sigma(t,X_t,v(t,X_t)).$$

The proof of this theorem can be found in [161, pp. 185-186]. Similar to other numerical methods for coupled FBSDEs, we use this theorem to decouple the original FBSDE (4.1) in order to be able to exploit standard results from the decoupled FBSDE literature. In addition to the assumptions above, we need Hölder-continuity in time for the convergence of the implicit scheme for (4.1), as stated below.

Assumption 4.3.5
b, σ, f in (4.1) are uniformly Hölder-$\frac{1}{2}$-continuous with respect to t.

Our main result in theorem 4.3.3 is concerned with an a posteriori error estimate for *discrete time approximations* of the continuous FBSDE in (4.1). A necessary ingredient in establishing this is to show that an appropriate discretization such as (4.4) below, converges in the number of time steps. This is given in the following theorem.

Theorem 4.3.2 (Convergence of the implicit scheme)
Suppose assumptions 4.3.2, 4.3.3, 4.3.4 and 4.3.5, and let a suitable set of monotonicity conditions such as in [17] or [143] hold. Then, for a sufficiently small h, the following discrete-time equation $(0 \le i \le N-1)$

$$\begin{cases} \bar{X}^\pi_0 = x_0, \\ \bar{X}^\pi_{t_{i+1}} = \bar{X}^\pi_{t_i} + b\left(t_i, \bar{X}^\pi_{t_i}, \bar{Y}^\pi_{t_i}, \bar{Z}^\pi_{t_i}\right) h + \sigma\left(t_i, \bar{X}^\pi_{t_i}, \bar{Y}^\pi_{t_i}\right) \Delta W_i, \\ \bar{Y}^\pi_T = g\left(\bar{X}^\pi_T\right), \\ \bar{Z}^\pi_{t_i} = \frac{1}{h}\mathbb{E}_i\left[\bar{Y}^\pi_{t_{i+1}} \Delta W_i^\top\right], \\ \bar{Y}^\pi_{t_i} = \mathbb{E}_i\left[\bar{Y}^\pi_{t_{i+1}} + f\left(t_i, \bar{X}^\pi_{t_i}, \bar{Y}^\pi_{t_i}, \bar{Z}^\pi_{t_i}\right) h\right], \end{cases} \tag{4.4}$$

has a solution $\{(\bar{X}^\pi_{t_i}, \bar{Y}^\pi_{t_i}, \bar{Z}^\pi_{t_i})\}_{i=0,\dots,N}$, *such that* $\bar{X}^\pi_{t_i} \in L^2\left(\Omega, \mathscr{F}_{t_i}, \mathbb{P}\right)$ *and*

$$\sup_{t\in[0,T]}\left(\mathbb{E}[\left\|X_t - \bar{X}^\pi_t\right\|^2] + \mathbb{E}[\left\|Y_t - \bar{Y}^\pi_t\right\|^2]\right) + \int_0^T \mathbb{E}[\left\|Z_t - \bar{Z}^\pi_t\right\|^2]\mathrm{d}t \le C\left(1+\mathbb{E}\|x_0\|^2\right)h, \quad (4.5)$$

with $\bar{X}^\pi_t := \bar{X}^\pi_{t_i}, \bar{Y}^\pi_t := \bar{Y}^\pi_{t_i}, \bar{Z}^\pi_t := \bar{Z}^\pi_{t_i}$ *for* $t \in [t_i, t_{i+1})$, *where* C *is a constant depending on* $\mathscr{L}$ *and* T.

Remark 4.3.1
We emphasize that the convergence of the discrete time approximation scheme in (4.4) *is only a necessary ingredient in the last step of our main result in theorem 4.3.3. The objective of the present chapter is to prove an a posteriori estimate – see* (4.24) *below – given a deep BSDE approximation, and not to analyze the convergence of an abstract time discretization. In fact, regardless of the assumptions under which theorem 4.3.2 is stated, as long as the implicit scheme admits an estimate such as* (4.5), *the conclusions of our main result remain the same. Such an estimate can be established by several different approaches in the literature. For instance, it can be shown that the weak and monotonicity conditions in [17] can be extended to our setting where the drift function b has an extra argument Z, see [75]. Alternatively, one could employ the convergence result of Reisinger et al. in [143, theorem 3.2] for the implicit Euler scheme in the framework of McKean-Vlasov FBSDEs, which would lead to a different set of monotonicity assumptions. For an overview on the literature of time discretization results, we refer to [161].*

Recall the classical Euler scheme in (4.2b). Taking conditional expectations of the discrete equation of $Y^\pi_{t_{i+1}}$, and of the same equation multiplied by $(\Delta W_i)^\top$, we obtain a formulation that does not include the objective functional (4.2a), i.e.,

$$\begin{cases} X^\pi_0 = x_0, \\ X^\pi_{t_{i+1}} = X^\pi_{t_i} + b\left(t_i, X^\pi_{t_i}, Y^\pi_{t_i}, Z^\pi_{t_i}\right)h + \sigma\left(t_i, X^\pi_{t_i}, Y^\pi_{t_i}\right)\Delta W_i, \\ Z^\pi_{t_i} = \dfrac{1}{h}\mathbb{E}_i\left[Y^\pi_{t_{i+1}}\Delta W_i^\top\right], \\ Y^\pi_{t_i} = \mathbb{E}_i\left[Y^\pi_{t_{i+1}} + f\left(t_i, X^\pi_{t_i}, Y^\pi_{t_i}, Z^\pi_{t_i}\right)h\right]. \end{cases} \quad (4.6)$$

With formulation (4.6) in hand, we can derive the following apriori estimate bounding the difference between two solutions of it.

Lemma 4.3.1
For $j = 1,2$, *suppose* $\left(\left\{X^{\pi,j}_{t_i}\right\}_{0\le i\le N}, \left\{Y^{\pi,j}_{t_i}\right\}_{0\le i\le N}, \left\{Z^{\pi,j}_{t_i}\right\}_{0\le i\le N-1}\right)$ *are two solutions of* (4.6), *with* $X^{\pi,j}_{t_i}, Y^{\pi,j}_{t_i} \in L^2\left(\Omega, \mathscr{F}_{t_i}, \mathbb{P}\right), 0 \le i \le N$. *For any* $\lambda_1 > 0, \lambda_2 > L^f_z$, *and sufficiently small* h, *denote*

$$\begin{aligned} &K_1 := 2k^b + \lambda_1 + L^\sigma_x + L^b_x h, \quad K_2 := \left(\lambda_1^{-1} + h\right)L^b_y + L^\sigma_y, \\ &K_3 := -\frac{\ln\left(1-\left(2k^f+\lambda_2\right)h\right)}{h}, \quad K_4 := \frac{L^f_x}{\left(1-\left(2k^f+\lambda_2\right)h\right)\lambda_2}, \quad C_1 := L^b_z(h+\lambda_1^{-1}). \end{aligned} \quad (4.7)$$

Let $\delta X_i := X_{t_i}^{\pi,1} - X_{t_i}^{\pi,2}$, $\delta Y_i := Y_{t_i}^{\pi,1} - Y_{t_i}^{\pi,2}$, $\delta Z_i := Z_{t_i}^{\pi,1} - Z_{t_i}^{\pi,2}$, *then we have, for* $0 \le n \le N$

$$\mathbb{E}\left[\|\delta X_n\|^2\right] \le K_2 h \sum_{i=0}^{n-1} e^{K_1(n-i-1)h}\mathbb{E}\left[\|\delta Y_i\|^2\right] + C_1 h \sum_{i=0}^{n-1} e^{K_1(n-i-1)h}\mathbb{E}\left[\|\delta Z_i\|^2\right], \tag{4.8}$$

$$\mathbb{E}\left[\|\delta Y_i\|^2\right] \le e^{K_3(N-n)h}\mathbb{E}\left[\|\delta Y_N\|^2\right] + K_4 \sum_{i=n}^{N-1} e^{K_3(i-n)h}\mathbb{E}\left[\|\delta X_i\|^2\right] h. \tag{4.9}$$

Proof. Let us define

$$\begin{aligned}
\delta b_i &:= b\left(t_i, X_{t_i}^{\pi,1}, Y_{t_i}^{\pi,1}, Z_{t_i}^{\pi,1}\right) - b\left(t_i, X_{t_i}^{\pi,2}, Y_{t_i}^{\pi,2}, Z_{t_i}^{\pi,2}\right),\\
\delta \sigma_i &:= \sigma\left(t_i, X_{t_i}^{\pi,1}, Y_{t_i}^{\pi,1}\right) - \sigma\left(t_i, X_{t_i}^{\pi,2}, Y_{t_i}^{\pi,2}\right),\\
\delta f_i &:= f\left(t_i, X_{t_i}^{\pi,1}, Y_{t_i}^{\pi,1}, Z_{t_i}^{\pi,1}\right) - f\left(t_i, X_{t_i}^{\pi,2}, Y_{t_i}^{\pi,2}, Z_{t_i}^{\pi,2}\right).
\end{aligned}$$

Then we have

$$\delta X_{i+1} = \delta X_i + \delta b_i h + \delta \sigma_i \Delta W_i, \tag{4.10}$$

$$\delta Y_i = \mathbb{E}_i\left[\delta Y_{i+1} + \delta f_i h\right], \tag{4.11}$$

and (4.6) also gives

$$\delta Z_i = \frac{1}{h}\mathbb{E}_i\left[\delta Y_{i+1}\Delta W_i^\top\right]. \tag{4.12}$$

By the martingale representation theorem, there exists an $\mathscr{F}_t$-adapted square-integrable process $\{\delta Z_t\}_{t_i \le t \le t_{i+1}}$ such that

$$\delta Y_{i+1} = \mathbb{E}_i\left[\delta Y_{i+1}\right] + \int_{t_i}^{t_{i+1}} \delta Z_t \mathrm{d}W_t,$$

which, together with (4.11), implies

$$\delta Y_{i+1} = \delta Y_i - \delta f_i h + \int_{t_i}^{t_{i+1}} \delta Z_t \mathrm{d}W_t. \tag{4.13}$$

From (4.10) and (4.13), noting that $\delta X_i, \delta Y_i, \delta b_i, \delta \sigma_i$ and δf_i are all $\mathscr{F}_{t_i}$ measurable, and $\mathbb{E}_i[\Delta W_i] = 0$, $\mathbb{E}_i[\int_{t_i}^{t_{i+1}} Z_t \mathrm{d}W_t] = 0$, we have

$$\begin{aligned}
\mathbb{E}\left[\|\delta X_{i+1}\|^2\right] &= \mathbb{E}\left[\|\delta X_i + \delta b_i h\|^2\right] + h\mathbb{E}\left[\|\delta \sigma_i\|^2\right],\\
\mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] &= \mathbb{E}\left[\left\|\delta Y_i - \delta f_i h\right\|^2\right] + \int_{t_i}^{t_{i+1}} \mathbb{E}\left[\|\delta Z_t\|^2\right] \mathrm{d}t,
\end{aligned}$$

where we also used a Fubini argument. We proceed in steps, controlling each of the terms above.

Step 1. *Estimate for δX_n.* By assumptions 4.3.1, 4.3.2, and the root-mean-square and geometric mean inequality (RMS-GM inequality), we have that for any $\lambda_1 > 0$

$$\begin{aligned}
\mathbb{E}\left[\|\delta X_{i+1}\|^2\right] &= \mathbb{E}\left[\|\delta X_i\|^2\right] + \mathbb{E}\left[\|\delta b_i\|^2\right] h^2 + h\mathbb{E}\left[\|\delta \sigma_i\|^2\right] \\
&\quad + 2h\mathbb{E}\left[\left(b\left(t_i, X_{t_i}^{\pi,1}, Y_{t_i}^{\pi,1}, Z_{t_i}^{\pi,1}\right) - b\left(t_i, X_{t_i}^{\pi,2}, Y_{t_i}^{\pi,1}, Z_{t_i}^{\pi,1}\right)\right)^\top \delta X_i\right] \\
&\quad + 2h\mathbb{E}\left[\left(b\left(t_i, X_{t_i}^{\pi,2}, Y_{t_i}^{\pi,1}, Z_{t_i}^{\pi,1}\right) - b\left(t_i, X_{t_i}^{\pi,2}, Y_{t_i}^{\pi,2}, Z_{t_i}^{\pi,2}\right)\right)^\top \delta X_i\right] \\
&\leq \left(1 + \left(2k^b + \lambda_1 + L_x^\sigma + L_x^b h\right) h\right) \mathbb{E}\left[\|\delta X_i\|^2\right] \\
&\quad + \left(\left(\lambda_1^{-1} + h\right) L_y^b + L_y^\sigma\right) \mathbb{E}\left[\|\delta Y_i\|^2\right] h + \left(L_z^b h + \lambda_1^{-1} L_z^b\right) \mathbb{E}\left[\|\delta Z_i\|^2\right] h.
\end{aligned}$$

Recalling the definition of C_1, K_1, K_2 from (4.7), we subsequently gather

$$\mathbb{E}\left[\|\delta X_{i+1}\|^2\right] \leq (1 + K_1 h)\mathbb{E}\left[\|\delta X_i\|^2\right] + K_2 h\mathbb{E}\left[\|\delta Y_i\|^2\right] + C_1 h\mathbb{E}\left[\|\delta Z_i\|^2\right].$$

Notice that $\mathbb{E}\left[\|\delta X_0\|^2\right] = 0$, and thus by induction, we have that for any $1 \leq n \leq N$,

$$\begin{aligned}
\mathbb{E}\left[\|\delta X_n\|^2\right] &\leq \prod_{i=0}^{n-1} (1 + K_1 h)\mathbb{E}\left[\|\delta X_0\|^2\right] + \sum_{i=0}^{n-1} (1 + K_1 h)^{n-1-i} K_2 \mathbb{E}\left[\|\delta Y_i\|^2\right] h \\
&\quad + \sum_{i=0}^{n-1} (1 + K_1 h)^{n-1-i} C_1 \mathbb{E}\left[\|\delta Z_i\|^2\right] h \\
&\leq K_2 h \sum_{i=0}^{n-1} e^{K_1(n-i-1)h} \mathbb{E}\left[\|\delta Y_i\|^2\right] + C_1 h \sum_{i=0}^{n-1} e^{K_1(n-i-1)h} \mathbb{E}\left[\|\delta Z_i\|^2\right],
\end{aligned}$$

where we used the inequality $(1 + x) \leq e^x$, $\forall x \in \mathbb{R}$. We remark that due to the coupling of Z in the drift, the last term of the right-hand side above is not present in [69].

Step 2. *Estimate for δY_n.* We employ a similar approach as in step 1. Using assumption 4.3.2 and the RMS-GM inequality, we obtain for any $\lambda_2 > 0$,

$$\begin{aligned}
\mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] &\geq \mathbb{E}\left[\|\delta Y_i\|^2\right] + \int_{t_i}^{t_{i+1}} \mathbb{E}\left[\|\delta Z_t\|^2\right] \mathrm{d}t \\
&\quad - 2h\mathbb{E}\left[\left(f\left(t_i, X_i^{1,\pi}, Y_i^{1,\pi}, Z_i^{1,\pi}\right) - f\left(t_i, X_i^{1,\pi}, Y_i^{2,\pi}, Z_i^{1,\pi}\right)\right)^\top \delta Y_i\right] \\
&\quad - 2h\mathbb{E}\left[\left(f\left(t_i, X_i^{1,\pi}, Y_i^{2,\pi}, Z_i^{1,\pi}\right) - f\left(t_i, X_i^{2,\pi}, Y_i^{2,\pi}, Z_i^{2,\pi}\right)\right)^\top \delta Y_i\right] \\
&\geq \mathbb{E}\left[\|\delta Y_i\|^2\right] + \int_{t_i}^{t_{i+1}} \mathbb{E}\left[\|\delta Z_t\|^2\right] \mathrm{d}t - 2k^f h\mathbb{E}\left[\|\delta Y_i\|^2\right] \\
&\quad - \left(\lambda_2 \mathbb{E}\left[\|\delta Y_i\|^2\right] + \lambda_2^{-1}\left(L_x^f \mathbb{E}\left[\|\delta X_i\|^2\right] + L_z^f \mathbb{E}\left[\|\delta Z_i\|^2\right]\right)\right) h.
\end{aligned} \tag{4.14}$$

To deal with the integral term in the last inequality, we derive the following relation via Ito's isometry, (4.13) and (4.12)

$$\delta Z_i = \frac{1}{h}\mathbb{E}_i\left[\int_{t_i}^{t_{i+1}} \delta Z_t \mathrm{d}t\right].$$

Then, by the Jensen- and Cauchy-Schwartz inequalities and the Fubini theorem, we derive a lower bound for the integral term

$$\begin{aligned}\mathbb{E}\left[\|\delta Z_i\|^2\right] h &= \sum_{j=1}^{q}\sum_{k=1}^{m}\mathbb{E}\left[(\delta Z_i)_{j,k}^2\right] h = \sum_{j=1}^{q}\sum_{k=1}^{m}\frac{1}{h}\mathbb{E}\left[\left(\mathbb{E}_i\left[\int_{t_i}^{t_{i+1}}(\delta Z_t)_{j,k}\,\mathrm{d}t\right]\right)^2\right] \\ &\leq \sum_{j=1}^{q}\sum_{k=1}^{m}\frac{1}{h}\mathbb{E}\left[\left(\int_{t_i}^{t_{i+1}}(\delta Z_t)_{j,k}\,\mathrm{d}t\right)^2\right] \\ &\leq \sum_{j=1}^{q}\sum_{k=1}^{m}\int_{t_i}^{t_{i+1}}\mathbb{E}\left[(\delta Z_t)_{j,k}^2\right]\mathrm{d}t = \int_{t_i}^{t_{i+1}}\mathbb{E}\left[\|\delta Z_t\|^2\right]\mathrm{d}t,\end{aligned} \tag{4.15}$$

where $(\cdot)_{j,k}$ denotes the (j,k)-entry of the matrix. Combining (4.14) with (4.15) gives

$$\begin{aligned}\mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] \geq \left(1-\left(2k^f+\lambda_2\right)h\right)\mathbb{E}\left[\|\delta Y_i\|^2\right] + \left(1-L_z^f\lambda_2^{-1}\right)\mathbb{E}\left[\|\delta Z_i\|^2\right] h \\ - L_x^f\lambda_2^{-1}\mathbb{E}\left[\|\delta X_i\|^2\right] h.\end{aligned} \tag{4.16}$$

Now, for any $\lambda_2 > L_z^f \geq 0$, and sufficiently small h satisfying $\left(2k^f+\lambda_2\right)h < 1$, this implies

$$\mathbb{E}\left[\|\delta Y_i\|^2\right] \leq \left(1-\left(2k^f+\lambda_2\right)h\right)^{-1}\left(\mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] + L_x^f\lambda_2^{-1}\mathbb{E}\left[\|\delta X_i\|^2\right] h\right).$$

Recalling the definitions of K_3, K_4 in (4.7), we subsequently gather by induction that for any $0 \leq n \leq N-1$

$$\mathbb{E}\left[\|\delta Y_n\|^2\right] \leq e^{K_3(N-n)h}\mathbb{E}\left[\|\delta Y_N\|^2\right] + K_4\sum_{i=n}^{N-1} e^{K_3(i-n)h}\mathbb{E}\left[\|\delta X_i\|^2\right] h.$$

We remark that this estimate coincides with the one of [69, Lemma 1].

□

Due to the coupling of Z in the drift coefficient of the forward diffusion, we need an additional estimate to handle the extra $\mathbb{E}[\|\delta Z_i\|^2]$ term in the estimate for $\mathbb{E}[\|\delta X_i\|^2]$. One of our main contributions is to establish the following lemma for this purpose.

Lemma 4.3.2
Under the setting of lemma 4.3.1, for any $\lambda_3 > 2mL_z^f$ and sufficiently small h, let us define

$$C_2 := 2((h+\lambda_3^{-1})L_y^f + \lambda_3), \quad C_3 := 2(h+\lambda_3^{-1}), \quad C_4 := (1-mC_3L_z^f)^{-1}m. \tag{4.17}$$

Then we have $C_4 > 0$, furthermore, the following estimates also hold

$$\begin{aligned}\mathbb{E}\left[\|\delta Y_i\|^2\right] \leq (1+C_2h)\mathbb{E}\left[\|\mathbb{E}_i\left[\delta Y_{i+1}\right]\|^2\right] \\ + C_3hL_x^f\mathbb{E}\left[\|\delta X_i\|^2\right] + C_3hL_z^f\mathbb{E}\left[\|\delta Z_i\|^2\right],\end{aligned} \tag{4.18}$$

$$\begin{aligned}h\sum_{i=0}^{N-1}\mathbb{E}[\|\delta Z_i\|^2] \leq C_4\left(\sum_{i=1}^{N-1} C_2h\mathbb{E}\left[\|\mathbb{E}_i\left[\delta Y_{i+1}\right]\|^2\right] + C_3hL_x^f\mathbb{E}\left[\|\delta X_i\|^2\right]\right) \\ + C_4\mathbb{E}\left[\|\delta Y_N\|^2\right].\end{aligned} \tag{4.19}$$

Proof. We take the squares of both sides of (4.11) and use the ϵ-Young inequality to get

$$\|\delta Y_i\|^2 \leq (1+\lambda_3 h)\,\|\mathbb{E}_i[\delta Y_{i+1}]\|^2 + (1+(\lambda_3 h)^{-1})\left\|h\delta f_i\right\|^2,$$

which holds for any $\lambda_3 > 0$, independent of $h > 0$. Taking expectations on both sides and using the Lipschitz continuity of f established by assumption 4.3.2 yields

$$\begin{aligned}\mathbb{E}\left[\|\delta Y_i\|^2\right] &\leq (1+\lambda_3 h)\mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right] + \left(h+\lambda_3^{-1}\right)h\mathbb{E}\left[\left\|\delta f_i\right\|^2\right] \\ &\leq (1+\lambda_3 h)\mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right] + (h+\lambda_3^{-1})h\Big(L_x^f\mathbb{E}\left[\|\delta X_i\|^2\right] + L_y^f\mathbb{E}[\|\delta Y_i\|^2] \\ &\quad + L_z^f\mathbb{E}\left[\|\delta Z_i\|^2\right]\Big).\end{aligned}$$

Therefore by a rearrangement

$$\begin{aligned}(1-(h+\lambda_3^{-1})hL_y^f)\mathbb{E}\left[\|\delta Y_i\|^2\right] \leq &(1+\lambda_3 h)\mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right] \\ &+ (h+\lambda_3^{-1})hL_x^f\mathbb{E}\left[\|\delta X_i\|^2\right] + \left(h+\lambda_3^{-1}\right)hL_z^f\mathbb{E}\left[\|\delta Z_i\|^2\right].\end{aligned}$$

Consequently, for any $\lambda_3 > 0$ and sufficiently small h, we obtain the following estimate, for $i = 0, 1, \ldots, N-1$

$$\mathbb{E}\left[\|\delta Y_i\|^2\right] \leq (1+C_2 h)\mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right] + C_3 hL_x^f\mathbb{E}\left[\|\delta X_i\|^2\right] + C_3 hL_z^f\mathbb{E}\left[\|\delta Z_i\|^2\right], \tag{4.20}$$

where we used the definitions in (4.17). This proves (4.18).

Next, we derive the estimate for Z. Recalling the definition in (4.12), we get

$$h\delta Z_i = \mathbb{E}_i\left[\delta Y_{i+1}\Delta W_i^\top\right] = \mathbb{E}_i\left[(\delta Y_{i+1} - \mathbb{E}_i[\delta Y_{i+1}])\,\Delta W_i^\top\right].$$

Taking the Frobenius norm on both sides and applying the Cauchy–Schwartz inequality then yields

$$\begin{aligned}h\,\|\delta Z_i\| &= \left\|\mathbb{E}_i\left[(\delta Y_{i+1} - \mathbb{E}_i[\delta Y_{i+1}])\,\Delta W_i^\top\right]\right\| \\ &\leq \left(\mathbb{E}_i\left[\|\delta Y_{i+1} - \mathbb{E}_i[\delta Y_{i+1}]\|^2\right]\right)^{\frac{1}{2}}\left(\mathbb{E}_i\left[\left\|\Delta W_i^\top\right\|^2\right]\right)^{\frac{1}{2}} \\ &= \left(\mathbb{E}_i\left[\|\delta Y_{i+1} - \mathbb{E}_i[\delta Y_{i+1}]\|^2\right]\right)^{\frac{1}{2}}(hm)^{\frac{1}{2}},\end{aligned}$$

which leads to

$$h\mathbb{E}\left[\|\delta Z_i\|^2\right] \leq m\mathbb{E}\left[\|\delta Y_{i+1} - \mathbb{E}_i[\delta Y_{i+1}]\|^2\right] = m\left(\mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] - \mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right]\right).$$

Summing both sides from 0 to $N-1$ and using the estimate (4.20), we gather

$$\begin{aligned}h\sum_{i=0}^{N-1}\mathbb{E}\left[\|\delta Z_i\|^2\right] &\leq m\sum_{i=0}^{N-1}\left(\mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] - \mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right]\right) \\ &= m\sum_{i=1}^{N-1}\left(\mathbb{E}\left[\|\delta Y_i\|^2\right] - \mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right]\right) + \mathbb{E}\left[\|\delta Y_N\|^2\right] - \mathbb{E}\left[\|\mathbb{E}_0[\delta Y_1]\|^2\right] \\ &\leq m\sum_{i=1}^{N-1}\left(\mathbb{E}\left[\|\delta Y_i\|^2\right] - \mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right]\right) + m\mathbb{E}\left[\|\delta Y_N\|^2\right] \\ &\leq m\sum_{i=1}^{N-1}\Big(C_2 h\mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right] + C_3 hL_x^f\mathbb{E}\left[\|\delta X_i\|^2\right] + C_3 hL_z^f\mathbb{E}\left[\|\delta Z_i\|^2\right]\Big) \\ &\quad + m\mathbb{E}\left[\|\delta Y_N\|^2\right].\end{aligned}$$

Recalling C_4 in (4.17), it is easy to check that for any $\lambda_3 > 2mL_z^f$ and sufficiently small $h > 0$, we have $C_4 > 0$ and therefore

$$h\sum_{i=0}^{N-1}\mathbb{E}\left[\|\delta Z_i\|^2\right] \le C_4\sum_{i=1}^{N-1}\left(C_2 h\mathbb{E}\left[\|\mathbb{E}_i[\delta Y_{i+1}]\|^2\right] + C_3 hL_x^f\mathbb{E}\left[\|\delta X_i\|^2\right]\right) + C_4\mathbb{E}\left[\|\delta Y_N\|^2\right].$$

□

Remark 4.3.2
As shown in the proofs of lemma 4.3.1 and 4.3.2, the constants C_j, $j = 1,2,3,4$ appear because of the Z coupling in the drift. Conversely, the constants K_j, $j = 1,2,3,4$ are present even in the less general case of only Y coupling, and they are consistent with [69]. In order to emphasize the difference, we denoted these by different letters.

4

With these auxiliary results, and particularly lemma 4.3.2, we are ready to state our main result, an a posteriori error estimate, generalizing the convergence of the deep BSDE method.

Theorem 4.3.3 (Convergence of the deep BSDE method)
Suppose assumptions 4.3.1-4.3.5 and let the conditions of theorem 4.3.2 hold. Define

$$\bar{B} := e^{\max(-\bar{K}_1 T,0)} L_x^f \bar{C}_1\bar{C}_4\bar{C}_3\frac{e^{\bar{K}_1 T}-1}{\bar{K}_1} + e^{\max(-\bar{K}_1 T,0)}\bar{C}_1\bar{C}_4 L_x^g(1+\lambda_4)e^{\bar{K}_1 T}, \tag{4.21}$$

$$\begin{aligned}\bar{A} := &\left(L_x^g(1+\lambda_4)e^{(\bar{K}_1+\bar{K}_3)T} + \frac{\bar{K}_4}{\bar{K}_1+\bar{K}_3}\left(e^{(\bar{K}_1+\bar{K}_3)T}-1\right)\right) \\ &\times\left(1-\bar{B}\right)^{-1}\left(\bar{K}_2\frac{1-e^{-(\bar{K}_1+\bar{K}_3)T}}{\bar{K}_1+\bar{K}_3} + e^{\max(-\bar{K}_1 T,0)}\bar{C}_1\bar{C}_4\bar{C}_2\frac{1-e^{-\bar{K}_3 T}}{\bar{K}_3}\right),\end{aligned} \tag{4.22}$$

where $\bar{K}_j := \lim_{h\to 0} K_j, \bar{C}_j := \lim_{h\to 0} C_j$ for $j = 1,2,3,4$. If

$$\inf_{\lambda_1>0,\lambda_2>L_z^f,\lambda_3>2mL_z^f,\lambda_4>0}\max(\bar{B},\bar{A}) < 1, \tag{4.23}$$

then there exists a constant $C > 0$, depending only on $\mathbb{E}[\|x_0\|^2], \mathscr{L}, T, \lambda_1, \lambda_2, \lambda_3$ and λ_4, such that for sufficiently small h, it holds that

$$\begin{aligned}\sup_{t\in[0,T]}\left(\mathbb{E}[\|X_t-\hat{X}_t^\pi\|^2]+\mathbb{E}[\|Y_t-\hat{Y}_t^\pi\|^2]\right) + \int_0^T\mathbb{E}[\|Z_t-\hat{Z}_t^\pi\|^2]\mathrm{d}t \\ \le C\left(h+\mathbb{E}[\|g\left(X_T^\pi\right)-Y_T^\pi\|^2]\right),\end{aligned} \tag{4.24}$$

where $\hat{X}_t^\pi := X_{t_i}^\pi, \hat{Y}_t^\pi := Y_{t_i}^\pi, \hat{Z}_t^\pi := Z_{t_i}^\pi$ for $t \in [t_i, t_{i+1})$.

Proof. Let $X_{t_i}^{\pi,1} = X_{t_i}^\pi, Y_{t_i}^{\pi,1} = Y_{t_i}^\pi, Z_{t_i}^{\pi,1} = Z_{t_i}^\pi$ given by the Euler scheme (4.2b), and $X_{t_i}^{\pi,2} = \bar{X}_{t_i}^\pi, Y_{t_i}^{\pi,2} = \bar{Y}_{t_i}^\pi, Z_{t_i}^{\pi,2} = \bar{Z}_{t_i}^\pi$ given by implicit scheme (4.4). Both of them solve (4.6), and therefore we can apply lemma 4.3.1 to bound their differences. In what follows, we use the same notations as in the proof of lemma 4.3.1.

First, using the RMS-GM inequality, for any $\lambda_4 > 0$ we get

$$\begin{aligned}\mathbb{E}\big[\|\delta Y_N\|^2\big] &\equiv \mathbb{E}\big[\|g(\bar{X}_T^\pi) - Y_T^\pi\|^2\big] \\ &\leq (1+\lambda_4^{-1})\mathbb{E}\big[\|g(X_T^\pi) - Y_T^\pi\|^2\big] + L_x^g(1+\lambda_4)\mathbb{E}\big[\|\delta X_N\|^2\big].\end{aligned} \tag{4.25}$$

Let

$$\mathscr{X} \coloneqq \max_{0\leq n\leq N} e^{-K_1 nh}\mathbb{E}\left[\|\delta X_n\|^2\right], \qquad \mathscr{Y} \coloneqq \max_{0\leq n\leq N} e^{K_3 nh}\mathbb{E}\left[\|\delta Y_n\|^2\right]. \tag{4.26}$$

4

From estimate (4.9) in lemma 4.3.1, we derive the following by multiplying with $e^{K_3 nh}$ on both sides

$$\begin{aligned}&e^{K_3 nh}\mathbb{E}\left[\|\delta Y_n\|^2\right] \\ \leq& e^{K_3 T}\mathbb{E}\left[\|\delta Y_N\|^2\right] + K_4\sum_{i=n}^{N-1} e^{K_3 ih}\mathbb{E}\left[\|\delta X_i\|^2\right]h \\ \leq& e^{K_3 T}\left(\left(1+\lambda_4^{-1}\right)\mathbb{E}\left[\|g(X_T^\pi) - Y_T^\pi\|^2\right] + L_x^g(1+\lambda_4)\mathbb{E}\left[\|\delta X_N\|^2\right]\right) + K_4\sum_{i=n}^{N-1} e^{K_3 ih}\mathbb{E}\left[\|\delta X_i\|^2\right]h \\ \leq& e^{K_3 T}(1+\lambda_4^{-1})\mathbb{E}\left[\|g(X_T^\pi) - Y_T^\pi\|^2\right] + \left(L_x^g(1+\lambda_4)e^{(K_1+K_3)T} + K_4\sum_{i=n}^{N-1} e^{(K_1+K_3)ih}h\right)\mathscr{X},\end{aligned}$$

where we used the definition of $\mathscr{X}$ in (4.26) and the estimate (4.25) in the last inequality. Maximizing over n subsequently yields

$$\begin{aligned}\mathscr{Y} \leq\ & e^{K_3 T}(1+\lambda_4^{-1})\mathbb{E}\left[\|g(X_T^\pi) - Y_T^\pi\|^2\right] \\ &+ \left(L_x^g(1+\lambda_4)e^{(K_1+K_3)T} + K_4 h\frac{e^{(K_1+K_3)T}-1}{e^{(K_1+K_3)h}-1}\right)\mathscr{X}.\end{aligned} \tag{4.27}$$

We approach δX_n in the same manner, and from (4.8) collect

$$\begin{aligned}e^{-K_1 nh}\mathbb{E}\left[\|\delta X_n\|^2\right] &\leq K_2\sum_{i=0}^{n-1} e^{-K_1(i+1)h}\mathbb{E}\left[\|\delta Y_i\|^2\right]h + C_1\sum_{i=0}^{n-1} e^{-K_1(i+1)h}\mathbb{E}\left[\|\delta Z_i\|^2\right]h \\ &\leq K_2\mathscr{Y}\sum_{i=0}^{n-1} e^{-K_1(i+1)h-K_3 ih}h + C_1\sum_{i=0}^{n-1} e^{-K_1(i+1)h}\mathbb{E}\left[\|\delta Z_i\|^2\right]h.\end{aligned}$$

Additionally, from (4.18) and (4.19) we get

$$
\begin{aligned}
&C_1 \sum_{i=0}^{n-1} e^{-K_1 h(i+1)} \mathbb{E}\left[\|\delta Z_i\|^2\right] h \\
\leq & e^{\max(-K_1 T,0)} C_1 \sum_{i=0}^{N-1} \mathbb{E}\left[\|\delta Z_i\|^2\right] h \\
\leq & e^{\max(-K_1 T,0)} C_1 C_4 \left(\sum_{i=1}^{N-1} C_2 h \mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] + \sum_{i=1}^{N-1} C_3 h L_x^f \mathbb{E}\left[\|\delta X_i\|^2\right] + \mathbb{E}\left[\|\delta Y_N\|^2\right] \right) \\
\leq & e^{\max(-K_1 T,0)} C_1 C_4 h \left(C_2 \sum_{i=1}^{N-1} e^{-K_3 (i+1) h} \mathscr{Y} + C_3 L_x^f \sum_{i=1}^{N-1} e^{K_1 i h} \mathscr{X} \right) \\
& \qquad + e^{\max(-K_1 T,0)} C_1 C_4 \mathbb{E}\left[\|\delta Y_N\|^2\right] \\
\leq & e^{\max(-K_1 T,0)} C_1 C_4 h \left(C_2 e^{-K_3 h} \frac{e^{-K_3 T}-1}{e^{-K_3 h}-1} \mathscr{Y} + C_3 L_x^f \frac{e^{K_1 T}-1}{e^{K_1 h}-1} \mathscr{X} \right) \\
& \qquad + e^{\max(-K_1 T,0)} C_1 C_4 \mathbb{E}\left[\|\delta Y_N\|^2\right],
\end{aligned}
$$

where we recall the definitions in (4.7) and (4.17). Combining these inequalities and applying estimate (4.25), we obtain the following by maximizing over n

$$
\begin{aligned}
\mathscr{X} \leq & K_2 \mathscr{Y} h e^{-K_1 h} \frac{e^{-(K_1+K_3)T}-1}{e^{-(K_1+K_3)h}-1} + e^{\max(-K_1 T,0)} C_1 C_4 h C_2 e^{-K_3 h} \frac{e^{-K_3 T}-1}{e^{-K_3 h}-1} \mathscr{Y} \\
& + e^{\max(-K_1 T,0)} C_1 C_4 h C_3 L_x^f \frac{e^{K_1 T}-1}{e^{K_1 h}-1} \mathscr{X} + e^{\max(-K_1 T,0)} C_1 C_4 L_x^g (1+\lambda_4) e^{K_1 T} \mathscr{X} \\
& + e^{\max(-K_1 T,0)} C_1 C_4 (1+\lambda_4^{-1}) \mathbb{E}\left[\left\| g\left(X_T^\pi\right) - Y_T^\pi \right\|^2 \right].
\end{aligned} \tag{4.28}
$$

In order to simplify the expressions, we define

$$
\begin{aligned}
& A_1(h) := e^{K_3 T}(1+\lambda_4^{-1}), \quad A_2(h) := L_x^g (1+\lambda_4) e^{(K_1+K_3)T} + K_4 h \frac{e^{(K_1+K_3)T}-1}{e^{(K_1+K_3)h}-1}, \\
& A_3(h) := K_2 h e^{-K_1 h} \frac{e^{-(K_1+K_3)T}-1}{e^{-(K_1+K_3)h}-1}, \quad A_4(h) := e^{\max(-K_1 T,0)} C_1 C_4 C_2 h e^{-K_3 h} \frac{e^{-K_3 T}-1}{e^{-K_3 h}-1}, \\
& A_5(h) := e^{\max(-K_1 T,0)} C_1 C_4 C_3 L_x^f h \frac{e^{K_1 T}-1}{e^{K_1 h}-1}, \quad A_6(h) := e^{\max(-K_1 T,0)} C_1 C_4 L_x^g (1+\lambda_4) e^{K_1 T}, \\
& A_7(h) := e^{\max(-K_1 T,0)} C_1 C_4 (1+\lambda_4^{-1}).
\end{aligned} \tag{4.29}
$$

Consequently, (4.27) and (4.28) read as follows

$$
\mathscr{Y} \leq A_1(h) \mathbb{E}\left[\left\| g\left(X_T^\pi\right) - Y_T^\pi \right\|^2 \right] + A_2(h) \mathscr{X}, \tag{4.30}
$$

$$
\mathscr{X} \leq A_3(h) \mathscr{Y} + A_4(h) \mathscr{Y} + A_5(h) \mathscr{X} + A_6(h) \mathscr{X} + A_7(h) \mathbb{E}\left[\left\| g\left(X_T^\pi\right) - Y_T^\pi \right\|^2 \right]. \tag{4.31}
$$

Next, we solve (4.30)-(4.31) such that $\mathscr{Y}$ and $\mathscr{X}$ are both controlled by $\mathbb{E}[\|g(X_T^\pi) - Y_T^\pi\|^2]$. Let

$$
B(h) := A_5(h) + A_6(h), \quad A(h) := A_2(h) \left(1 - A_5(h) - A_6(h)\right)^{-1} \left(A_3(h) + A_4(h)\right). \tag{4.32}
$$

Whenever $B(h) < 1$, rearranging the terms in (4.31) yields

$$\mathscr{X} \leq ((1 - A_5(h) - A_6(h))^{-1} \Big((A_3(h) + A_4(h))\mathscr{Y} + A_7(h)\mathbb{E}\Big[\big\| g\big(X_T^\pi\big) - Y_T^\pi \big\|^2 \Big] \Big). \tag{4.33}$$

Additionally, if also $A(h) < 1$, we can derive the following by substituting (4.33) into (4.30)

$$\mathscr{Y} \leq (1 - A(h))^{-1} \left(A_1(h) + A_2(h)\,(1 - B(h))^{-1}\, A_7(h) \right) \mathbb{E}\Big[\big\| g\big(X_T^\pi\big) - Y_T^\pi \big\|^2 \Big]. \tag{4.34}$$

From (4.7) and (4.17), we directly collect the limits

$$\begin{aligned} &\bar{K}_1 = 2k^b + \lambda_1 + L_x^\sigma, \quad \bar{K}_2 = \lambda_1^{-1} L_y^b + L_y^\sigma, \quad \bar{K}_3 = 2k^f + \lambda_2, \quad \bar{K}_4 = L_x^f \lambda_2^{-1}, \\ &\bar{C}_1 = \lambda_1^{-1} L_z^b, \quad \bar{C}_2 = 2(\lambda_3^{-1} L_y^f + \lambda_3), \quad \bar{C}_3 = 2\lambda_3^{-1}, \quad \bar{C}_4 = \frac{m}{1 - 2mL_z^f \lambda_3^{-1}}. \end{aligned} \tag{4.35}$$

Consequently, from (4.29) we directly have

$$\begin{aligned} &\bar{A}_1 = e^{\bar{K}_3 T}(1 + \lambda_4^{-1}), \quad \bar{A}_2 = L_x^g (1 + \lambda_4) e^{(\bar{K}_1 + \bar{K}_3)T} + \frac{\bar{K}_4}{\bar{K}_1 + \bar{K}_3} \Big(e^{(\bar{K}_1 + \bar{K}_3)T} - 1 \Big), \\ &\bar{A}_3 = \bar{K}_2 \frac{1 - e^{-(\bar{K}_1 + \bar{K}_3)T}}{\bar{K}_1 + \bar{K}_3}, \quad \bar{A}_4 = e^{\max(-\bar{K}_1 T, 0)} \bar{C}_1 \bar{C}_4 \bar{C}_2 \frac{1 - e^{-\bar{K}_3 T}}{\bar{K}_3}, \\ &\bar{A}_5 = e^{\max(-\bar{K}_1 T, 0)} \bar{C}_1 \bar{C}_4 \bar{C}_3 L_x^f \frac{e^{\bar{K}_1 T} - 1}{\bar{K}_1}, \quad \bar{A}_6 = e^{\max(-\bar{K}_1 T, 0)} \bar{C}_1 \bar{C}_4 L_x^g (1 + \lambda_4) e^{\bar{K}_1 T}, \\ &\bar{A}_7 = e^{\max(-\bar{K}_1 T, 0)} \bar{C}_1 \bar{C}_4 (1 + \lambda_4^{-1}), \end{aligned} \tag{4.36}$$

with the convention $\bar{A}_j = \lim_{h \to 0} A_j(h)$, $j = 1, \ldots, 7$. If $\bar{K}_1 < 0$ the expressions above hold only for sufficiently small h such that $K_1 < 0$. Using the definitions in (4.32), it is straightforward to check that $\lim_{h \to 0} B(h) =: \bar{B}$ and $\lim_{h \to 0} A(h) =: \bar{A}$, given by (4.21) and (4.22), respectively. From (4.26), we get

$$\max_{0 \leq n \leq N} \mathbb{E}\big[\|\delta X_n\|^2 \big] \leq e^{\max(K_1 T, 0)} \mathscr{X}, \quad \max_{0 \leq n \leq N} \mathbb{E}\big[\|\delta Y_n\|^2 \big] \leq e^{\max(-K_3 T, 0)} \mathscr{Y}, \tag{4.37}$$

with K_1, K_3 defined in (4.7), both depending on h. When $\bar{B} < 1$ and $\bar{A} < 1$, we have that for any sufficiently small h (4.34) holds true. Hence, combining (4.34) with (4.37), we derive that for any sufficiently small h

$$\max_{0 \leq n \leq N} \mathbb{E}\big[\|\delta Y_n\|^2 \big] \leq C(\lambda_1, \lambda_2, \lambda_3, \lambda_4) \mathbb{E}\Big[\big\| g\big(X_T^\pi\big) - Y_T^\pi \big\|^2 \Big], \tag{4.38}$$

with a constant independent of h, depending only on the limits defined in (4.35), (4.21), (4.22), (4.36), and thus implicitly on $\lambda_1, \lambda_2, \lambda_3, \lambda_4$. Similarly, when $\bar{B}, \bar{A} < 1$, combining (4.33) with (4.34), we also have that for any sufficiently small h

$$\max_{0 \leq n \leq N} \mathbb{E}\big[\|\delta X_n\|^2 \big] \leq C(\lambda_1, \lambda_2, \lambda_3, \lambda_4) \mathbb{E}\Big[\big\| g\big(X_T^\pi\big) - Y_T^\pi \big\|^2 \Big], \tag{4.39}$$

for a constant determined by (4.36) which is independent of h.

In order to estimate $\mathbb{E}[\|\delta Z_n\|^2]$, we consider (4.16) from the proof of lemma 4.3.1. Notice that $1 - L_z^f/\lambda_2 > 0$ since we require $\lambda_2 > L_z^f \geq 0$, then by rearranging the terms in (4.16) we obtain

$$\mathbb{E}\left[\|\delta Z_i\|^2\right] h \leq \left(1 - L_z^f \lambda_2^{-1}\right)^{-1} \left(L_x^f \lambda_2^{-1} \mathbb{E}\left[\|\delta X_i\|^2\right] h + \mathbb{E}\left[\|\delta Y_{i+1}\|^2\right] - \mathbb{E}\left[\|\delta Y_i\|^2\right] + \left(2k^f + \lambda_2\right) h \mathbb{E}\left[\|\delta Y_i\|^2\right]\right).$$

Summing from 0 to $N-1$ and taking the maximum on the right hand side, we gather

$$\begin{aligned}\sum_{i=0}^{N-1} \mathbb{E}\left[\|\delta Z_i\|^2\right] h &\leq \left(1 - L_z^f \lambda_2^{-1}\right)^{-1} \left(L_x^f \lambda_2^{-1} T \max_{0\leq n\leq N} \mathbb{E}\left[\|\delta X_n\|^2\right] \right. \\ &\qquad \left. + \left(\max\{(2k^f + \lambda_2)T, 0\} + 1\right) \max_{0\leq n\leq N} \mathbb{E}\left[\|\delta Y_n\|^2\right]\right) \\ &\leq C(\lambda_1, \lambda_2, \lambda_3, \lambda_4) \mathbb{E}\left[\left\|g\left(X_T^\pi\right) - Y_T^\pi\right\|^2\right],\end{aligned} \tag{4.40}$$

using (4.39), (4.38).

Finally, combining estimates (4.39), (4.38) and (4.40) with the convergence of the discrete time approximations such as in theorem 4.3.2, we prove our statement. □

4.4. INTERPRETATION OF THE CONDITIONS IN THEOREM 4.3.3

In this section, we apply theorem 4.3.3 to special cases of FBSDEs and discuss how the conditions change depending on the coefficients in (4.1). Furthermore, we illustrate the role of the abstract conditions imposed by (4.23), and discuss several important heuristic settings under which they are satisfied.

(1) *Decoupled FBSDE*

In this case, $L_y^b = L_z^b = L_y^\sigma \equiv 0$, which immediately implies that $\bar{B} = \bar{A} \equiv 0$, since both $\bar{C}_1 = 0$ and $\bar{K}_2 = 0$. Estimates (4.33), (4.34) then reduce to

$$\mathscr{X} = 0, \qquad \mathscr{Y} \leq e^{K_3 T}(1 + \lambda_4^{-1}) \mathbb{E}\left[\left\|g\left(X_T^\pi\right) - Y_T^\pi\right\|^2\right].$$

Consequently, the total errors in the SDE reduce to those of the Euler-Maruyama discretization from (4.4), whereas for the BSDE part an a posteriori error term remains in (4.24).

(2) *Coupled FBSDE with only Y coupling*

In this case $L_z^b = 0$ and therefore $\bar{C}_1 = 0$. We remark that we fully recover the result of [69] in this setting, in particular

$$\bar{B} \equiv 0, \quad \bar{A} = \left(L_x^g (1 + \lambda_4) e^{(\bar{K}_1 + \bar{K}_3)T} + \frac{\bar{K}_4}{\bar{K}_1 + \bar{K}_3}\left(e^{(\bar{K}_1 + \bar{K}_3)T} - 1\right)\right)\left(\bar{K}_2 \frac{1 - e^{-(\bar{K}_1 + \bar{K}_3)T}}{\bar{K}_1 + \bar{K}_3}\right),$$

where the condition $\bar{B} < 1$ becomes redundant and is automatically satisfied, whereas $\bar{A}$ has the same expression as the one derived in [69]. Moreover, we recover the weak and monotonicity conditions as in [69, remark 6], which guarantee $\bar{A} < 1$.

(3) *Coupled FBSDE in general as in* (4.1)

This is the general setting we considered throughout this chapter corresponding to (4.1). In order to guarantee that the conditions of theorem 4.3.3, and in particular (4.23) are satisfied we need certain requirements about T, the constants in $\mathscr{L}$, and choose λ_1, λ_2, λ_3 and λ_4 in an appropriate way. Recall that $\bar{B} \equiv \bar{B}(\lambda_1,\lambda_2,\lambda_3,\lambda_4)$ and $\bar{A} \equiv \bar{A}(\lambda_1,\lambda_2,\lambda_3,\lambda_4)$ are functions of all λs, defined by (4.21) and (4.22), respectively. We divide the discussion into the following five cases which all have important physical interpretations.

(a) *Small time duration.* Suppose all other constants, λ_2, λ_3 and λ_4 are fixed. If $T > 0$ is sufficiently small, we can choose, for instance, $\lambda_1 = 1/\sqrt{T}$ which implies that $\bar{B}$ is sufficiently close to zero due to the factors $\bar{C}_1$ and $e^{\bar{K}_1 T} - 1$. Similarly $\bar{A}$ is sufficiently close to zero as well, due to the scaling factors $\bar{K}_2$, $1 - e^{-(\bar{K}_1+\bar{K}_3)T}$, $\bar{C}_1$, $1 - e^{\bar{K}_3 T}$ in the last term of (4.22). Therefore (4.23) is satisfied for sufficiently small time durations T.

(b) *Weak coupling from BSDE to SDE.* Suppose all other constants, $\lambda_1,\lambda_2,\lambda_3$ and λ_4 are fixed. If $L_y^b > 0$, $L_z^b > 0$ and $L_y^\sigma > 0$ are sufficiently small, then so are the factors $\bar{C}_1$ and $\bar{K}_2$. Notice that $\bar{B}$ is scaled by $\bar{C}_1$, and for $\bar{A}$, the last term in (4.22) is scaled by both $\bar{C}_1$ and $\bar{K}_2$, and therefore both $\bar{A}$ and $\bar{B}$ are sufficiently close to zero and (4.23) holds.

(c) *Weak coupling from SDE to BSDE.* Suppose all other constants, λ_1,λ_2, λ_3 and λ_4 are fixed. If $L_x^g > 0$ and $L_x^f > 0$ are sufficiently small, and additionally $L_z^b > 0$ is sufficiently small as well, then both $\bar{B}$ and $\bar{A}$ could be sufficiently close to zero, due to the scaling factors L_x^f, $\bar{C}_1$ and L_x^g in $\bar{B}$, and L_x^g and $\bar{K}_4$ in the first term of (4.22) for $\bar{A}$. Consequently, (4.23) is satisfied.

(d) *Monotonicity in b.* Suppose all other constants, λ_3 and λ_4 are fixed, $k^b < 0$ is sufficiently negative and $L_z^b > 0$ is sufficiently small. We set $\lambda_1 = -2k^b - \epsilon > 0$ which implies $\bar{K}_1 = -\epsilon + L_x^\sigma < 0$ is fixed for any chosen $\epsilon > L_x^\sigma$. Then, $\bar{B}$ could be sufficiently close to 0 since $\bar{K}_1 < 0$ is fixed and $\bar{C}_1 = \lambda_1^{-1} L_z^b$ could be sufficiently small. For $\bar{A}$, we directly compute from (4.22)

$$\begin{aligned}\bar{A} = (1-\bar{B})^{-1}\Bigg(& e^{\max(-\bar{K}_1 T,0)}\bar{C}_1\bar{C}_4\bar{C}_2\frac{1-e^{-\bar{K}_3 T}}{\bar{K}_3}\left(L_x^g(1+\lambda_4)e^{(\bar{K}_1+\bar{K}_3)T}\right)\\ &+ e^{\max(-\bar{K}_1 T,0)}\bar{C}_1\bar{C}_4\bar{C}_2\frac{1-e^{-\bar{K}_3 T}}{\bar{K}_3}\left(\frac{\bar{K}_4}{\bar{K}_1+\bar{K}_3}\left(e^{(\bar{K}_1+\bar{K}_3)T}-1\right)\right)\\ &+ L_x^g(1+\lambda_4)\bar{K}_2\frac{e^{(\bar{K}_1+\bar{K}_3)T}-1}{\bar{K}_1+\bar{K}_3} + \bar{K}_2\bar{K}_4\frac{e^{(\bar{K}_1+\bar{K}_3)T}+e^{-(\bar{K}_1+\bar{K}_3)T}-2}{(\bar{K}_1+\bar{K}_3)^2}\Bigg).\end{aligned} \tag{4.41}$$

Let us first consider the last two terms in (4.41). Notice that we could fix some sufficiently large $\epsilon > L_x^\sigma$ and $\lambda_2 > L_z^f$ such that $\bar{K}_1 + \bar{K}_3 = -\epsilon + L_x^\sigma + 2k^f + \lambda_2 < 0$ is fixed and negative enough, and the penultimate term is sufficiently small because the fraction term is decreasing in $\bar{K}_1 + \bar{K}_3$ and λ_1 is sufficiently large. The last term could be sufficiently small as well, since $\bar{K}_1 + \bar{K}_3$ is fixed and $\bar{K}_4 = L_x^f/\lambda_2$ can be sufficiently small by choosing a large enough λ_2. For the remaining first two terms

in (4.41), notice that $\bar{K}_1 + \bar{K}_3$ is fixed, and for the previously chosen λ_2, $\bar{K}_3$ is fixed as well, then both two terms are scaled by $\bar{C}_1$ with a sufficiently small L_z^b. Combining all these arguments we have $\bar{A} < 1$ and conclude that (4.23) is verified.

(e) *Monotonicity in f.* Suppose all other constants, λ_1, λ_3 and λ_4 are fixed. If $k^f < 0$ is sufficiently negative, and $L_z^b > 0$ is sufficiently small, then it is easy to see that $\bar{B}$ could be sufficiently close to 0 due to the scaling factor $\bar{C}_1 = \lambda_1^{-1} L_z^b$ and the fact that $\bar{B}$ does not depend on k^f and λ_2. To deal with $\bar{A}$, we set $\lambda_2 = -2k^f - \epsilon$, where $\epsilon > 0$ is chosen such that $\bar{K}_1 + \bar{K}_3 = 2k^b + \lambda_1 + L_x^\sigma - \epsilon < 0$ is negative enough and the penultimate term in (4.41) is sufficiently small. Since now $\bar{K}_1 + \bar{K}_3$ and $\bar{K}_3 = -\epsilon$ are fixed, the remaining three terms in $\bar{A}$ are all scaled by $\bar{C}_1$ and therefore we have $\bar{A} < 1$ guaranteeing (4.23).

(4) *Coupled FBSDE with b only depending on Z*

In this case $L_y^b = 0$. However, all constants $\bar{K}_j$, $\bar{C}_j$ for $j = 1,2,3,4$ defined in (4.7), (4.17) are not zero in general, and therefore the conditions fall back under the general case discussed above. Even in a more special setting with $L_y^b = L_y^\sigma = L_y^f = k^f = 0$, i.e. when there is no Y coupling in neither the forward nor backward equation, the conditions do not seem to be easier to satisfy. Specifically, we have $\bar{K}_2 = 0$ in this special setting and consequently there will be one term less in $\bar{A}$ given by (4.22), but the expression (4.21) for $\bar{B}$ remains unchanged as it does not depend on $\bar{K}_2$. On one hand, this reduces some efforts due to the missing term in $\bar{A}$, for example, we do not take care of the last terms discussed in ((3)d) as they vanish. On the other hand, as $k^f = 0$ in this special case, we lost one possible way to make the conditions hold, i.e. ((3)e) does not apply anymore. In overall, we conclude that coupling in Z, even in the special cases mentioned above, induces the need to treat the conditions in theorem 4.3.3 under the general framework established by our convergence result.

Remark 4.4.1
The above five cases in (3) may be viewed as a generalization of the weak and monotonicity conditions stated in [69, 17]. One should note that, because of the extra Z coupling, we have to pay an extra price in establishing these five cases, e.g., we need to choose λs appropriately instead of fixing them as constants as in [69], and in particular, we require $L_z^b > 0$ to be sufficiently small for (3)c, (3)d and (3)e.

Remark 4.4.2
In [69, remark 2], it is claimed that it is general to consider drifts which only depend on X and Y but not on Z, since, by an appropriate change of probability measure, one can always reformulate (4.1) into an equivalent FBSDE whose drift is independent of Z yet its solution coincides with the same quasi-linear PDE. However, we find that there are several important issues with this statement both theoretically and numerically:

(1) such approach would change the probability measure under which X is simulated, consequently the training of neural networks would be carried out in a different spatial region, and therefore the algorithm may have poor accuracy around the area of the domain of interest;

(2) it is common that the reformulated FBSDE does not satisfy the theoretical assumptions needed for convergence while the original one does. For instance, a linear z term in b would result in a quadratic term in the reformulated driver, which violates the assumptions of Lipschitz continuity. We illustrate this through Example 4.5.1 in section 4.5 below. Therefore, we believe our theory is a necessary generalization to [69] and it is applicable to a wider class of FBSDEs;

(3) with the same approach, one could remove the entire drift to the driver, and simulate a reformulated FBSDE with zero drift, but this, for the two reasons above, is rarely done in practice.

Finally, let us derive a lower bound for $\bar{B}$ in (4.21) by computing the infimum of $\bar{B}$ over all possible choices of λs. Notice that $\bar{B}$ does not depend on λ_2, decreases in λ_3 and increases in λ_4, and therefore we shall mainly look at λ_1. Let

$$\bar{B}_\ell := \inf_{\lambda_1>0,\lambda_3>2mL_z^f,\lambda_4>0} \bar{B}. \tag{4.42}$$

(1) If $\lambda_1 \geq \max(-2k^b - L_x^\sigma, 0)$, then $\bar{B}$ admits a unique stationary point along the λ_1 direction, and

$$\bar{B}_\ell = mL_z^b L_x^g e^{(2k^b+T^{-1}+L_x^\sigma)T}T, \qquad \arg\inf_{\lambda_1,\lambda_3,\lambda_4} \bar{B} = (1/T, +\infty, 0).$$

(2) If $0 < \lambda_1 < \max(-2k^b - L_x^\sigma, 0)$, then $\bar{B}$ is convex in λ_1 but there is not stationary point in this range, and

$$\bar{B}_\ell = m\frac{L_z^b}{-2k^b - L_x^\sigma}L_x^g, \qquad \arg\inf_{\lambda_1,\lambda_3,\lambda_4} \bar{B} = (-2k^b - L_x^\sigma, +\infty, 0).$$

This result is particularly useful in practice and can serve as a preliminary test for the convergence of a given FBSDE. In fact, given an equation and all its relevant Lipschitz constants, if we find $\bar{B}_\ell \geq 1$, then we know that the conditions of theorem 4.3.3 cannot be satisfied and that the deep BSDE algorithm is less likely to converge. On the other hand, if $\bar{B}_\ell$ is less than 1 or in particular even close to 0, we may check if $\bar{A} < 1$, which can be solved efficiently by a numerical constrained minimization method ranging different λ_1, λ_2, λ_3 and λ_4.

4.5. NUMERICAL EXPERIMENTS

We implemented the deep BSDE method in TensorFlow 2.9. The errors correspond to the discretized version of the left hand side of (4.24) and are defined as follows

$$\text{error}(X) := \max_{0\leq n\leq N} \mathbb{E}\left[\left\|\hat{X}_n^\pi - X_{t_n}\right\|^2\right], \quad \text{error}(Y) := \max_{0\leq n\leq N} \mathbb{E}\left[\left\|\hat{Y}_n^\pi - Y_{t_n}\right\|^2\right],$$
$$\text{error}(Z) := 1/N \sum_{n=0}^{N-1} \mathbb{E}\left[\left\|\hat{Z}_n^\pi - Z_{t_n}\right\|^2\right], \quad \text{total} = \text{error}(X) + \text{error}(Y) + \text{error}(Z).$$

We also define relative L^2 approximation errors by $\text{error}(X)/\mathbb{E}[\|X_{t_n}\|^2]$ and similarly for Y, Z and total errors. In what follows, the true expectations are approximated over a Monte Carlo sample of size 2^{12}. Given a classical solution to the corresponding quasi-linear PDE used for decoupling, we gather a reference solution to the associated FBSDE (4.1) by an Euler-Maruyama simulation with $N' = 10^4$ time steps, in order to guarantee that the time discretization error of the reference solution is negligible compared to the approximation errors incurred via the deep BSDE method. In all experiments below, we use neural networks with hyperbolic tangent activation, an input layer of width d, 2 hidden layers $30 + d$ neurons wide each, and an output layer of appropriate dimensions depending on the process approximated. As an optimization strategy, we use the Adam optimizer with default initializations and a learning rate schedule of exponential decay, starting from 10^{-2} with a decay rate of 10^{-2}. For a fixed N we perform 2^{14} SGD iterations, and for each iteration we take an independent sample of 2^{10} trajectories of the underlying Brownian motion. All experiments below were run on an Dell Alienware Aurora R10 machine, equipped with an AMD Ryzen 9 3950X CPU (16 cores, 64Mb cache, 4.7GHz) and an Nvidia GeForce RTX 3090 GPU (24Gb). In order to assess the inherent stochasticity of both the regression Monte Carlo method and the SGD iterations, we run each experiment 5 times and report on the mean and standard deviations of the resulting independent approximations. All operations were carried out with single precision.

4.5.1. EXAMPLE 1

The following example is a modified version of the one in [17], where in order to demonstrate our theoretical extension we include Z coupling in the drift of the forward diffusion. The coefficients of the FBSDE system (4.1) read as follows

$$b(t,x,y,z) = \kappa_y \bar{\sigma} y \mathbf{1}_d + \kappa_z z^\top, \quad \sigma(t,x,y) = \bar{\sigma} y I_d, \quad g(x) = \sum_{i=1}^{d} \sin(x_i), \tag{4.43}$$

$$\begin{aligned} f(t,x,y,z) = &-ry + 1/2 e^{-3r(T-t)} \bar{\sigma}^2 (\sum_{i=1}^{d} \sin(x_i))^3 \\ &- \kappa_y \sum_{i=1}^{d} z_i - \kappa_z \bar{\sigma} e^{-3r(T-t)} \sum_{i=1}^{d} \sin(x_i) \sum_{i=1}^{d} \cos^2(x_i), \end{aligned}$$

with $q = 1, d = m$. The analytical solution pair to the backward equation is given by

$$y(t,x) = e^{-r(T-t)} \sum_{i=1}^{d} \sin(x_i), \quad z_i(t,x) = e^{-2r(T-t)} \bar{\sigma} \big(\sum_{j=1}^{d} \sin\big(x_j\big) \big) \cos(x_i). \tag{4.44}$$

We note that the equation above falls under the theoretical assumptions of section 4.3. In particular, we get the following set of values for the corresponding constants $L_x^g = d, L_y^b = 2(\kappa_y \bar{\sigma})^2, L_z^b = 2\kappa_z^2, L_y^\sigma = d\bar{\sigma}^2, L_x^f = 3/2 d(3\bar{\sigma}^2 d^2/2 + 2\kappa_z \bar{\sigma} d)^2, L_y^f = 18r^2, L_z^f = 3.6 d\kappa_y^2, k^f = -r, k^b = L_x^\sigma = L_z^\sigma = L_x^b = 0$. We consider the equation in $d = 10$ dimensions. The strength of coupling is determined by the values $\bar{\sigma}, r, \kappa_y, \kappa_z$. In order to satisfy the sufficient conditions of theorem 4.3.3, we put $r = 1, \bar{\sigma} = 0.1, \kappa_y = 10^{-1}, \kappa_z = 10^{-2}$ and $T = 0.25, X_0 = (\pi/4, \ldots, \pi/4)$. Convergence results are collected in figure 4.1. Figure 4.1a

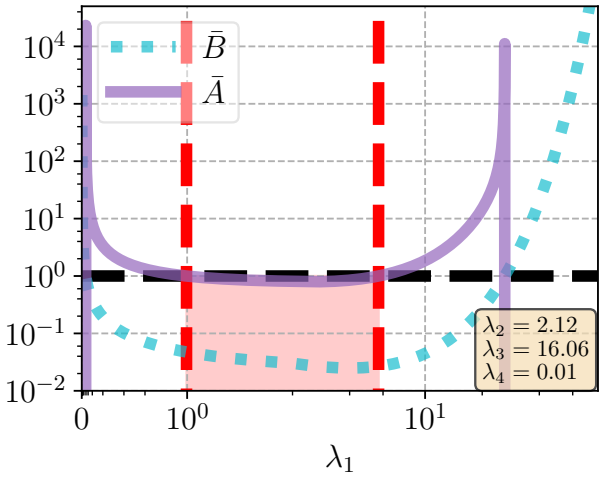

(a) $\bar{B}, \bar{A}$ as functions of λ_1 for given $(\lambda_2, \lambda_3, \lambda_4)$ in case of (4.43). Dotted vertical red lines mark the endpoints of the interval where $\bar{B}, \bar{A} < 1$, and the shaded red area the subset of the plane where the sufficient conditions of theorem 4.3.3 are satisfied.

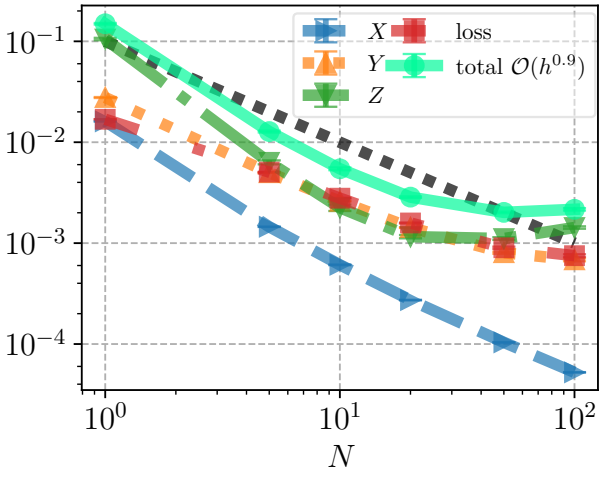

(b) Convergence in N. Empirical convergence rate in labels. Dotted black line indicates the expected $\mathcal{O}(h)$ rate predicted by theorem 4.3.3.

Figure 4.1: Example 4.5.1. $T = 0.25$, $X_0 = (\pi/4, \ldots, \pi/4)$.

shows that the conditions of theorem 4.3.3 are indeed satisfied, there exists a quadruple $(\lambda_1, \lambda_2, \lambda_3, \lambda_4)$, such that $\bar{B}, \bar{A}$ defined by (4.22), (4.21) admit $\max(\bar{B}, \bar{A}) < 1$. In particular, for the fixed $\lambda_2, \lambda_3, \lambda_4$ we mark the interval of admissible λ_1s such that the sufficient conditions are satisfied within the shaded red area. Figure 4.1b depicts the convergence of the deep BSDE method. Its most important implications are as follows. The convergence is only guaranteed in a *posteriori* sense. In fact, as can be seen the convergence only shows the expected $\mathcal{O}(h)$ behavior whenever the loss function corresponding to the last term of (4.24) is dominated by the discretization error. In particular, for $N \in \{50, 100\}$ we see that the approximation errors of Z begin to stall and the total approximation errors are dominated by the loss function. This indicates that for very fine time grids one needs to make sure that the loss is appropriately minimized when trying to recover discretization errors. Given the global minimization structure of the deep BSDE method the corresponding optimization problem becomes more difficult with an increased number of time steps. This demonstrates a clear trade-off between discretization and optimization, which is fully explained by theorem 4.3.3 and should be carefully considered in applications. Nevertheless, we get an empirical convergence rate of $\mathcal{O}(h^{0.9})$ for all time points, and accounting for the reasoning above we recover the predicted rate of our convergence analysis.

Furthermore, let us return to remark 4.4.2. In particular, Han and Long in [69, Remark 2] claim that the setting of Z independent drift is general, since, due to the connections with the associated quasi-linear PDEs, one can always move the Z dependence from the drift to the driver. In order to complement our arguments against this reasoning in remark 4.4.2, we provide a numerical demonstration of the points raised therein. One can derive a *reformulated* FBSDE system which is *decoupled in* Z and whose solution will coincide with (4.44). This equation has a modified drift and driver

$$\tilde{b}(t,x,y) = \kappa_y \bar{\sigma} y 1_d, \quad \tilde{f}(t,x,y,z) = f(t,x,y,z) + \kappa_z \|z\|^2/(\bar{\sigma} y), \tag{4.45}$$

whereas the rest of the coefficients remain the same as in (4.43). First, notice that even though (4.43) satisfies the theoretical assumptions of theorem 4.3.3, the reformulated FBSDE (4.45) does not. In particular, $\tilde{f}$ is not Lipschitz continuous in y, z which renders the results of theorem 4.3.3, or [69] as a limit case, inapplicable. This demonstrates point-(2) from remark 4.4.2. Nonetheless, as our convergence analysis only gives sufficient conditions one can still run the deep BSDE algorithm and find satisfactory results without theoretical guarantees. Table 4.1 shows that for equation (4.45) this is not case. Running the algorithm on the reformulated FBSDE (4.45) results in diverging errors. In fact, due to the singularity arising in the driver $\tilde{f}$, the backward equation blows up as N increases, which results in the forward equation also exploding due to the coupling. On the contrary, the original equation (4.43) with Z coupling converges as predicted by theorem 4.3.3 and also illustrated by figure 4.1. This observation demonstrates point-(1) from remark 4.4.2 and implies that the Lipschitz features in our analysis are crucial also in practice, in order to avoid such explosion of the coupled forward diffusion. This is in line with related results in the literature, see [17, pg.170]. Overall, we conclude that the framework of Z coupling in the drift cannot in general be circumvented neither theoretically nor numerically, and one needs to rely on our convergence result in theorem 4.3.3 instead.

Table 4.1: Comparison on the convergence of the deep BSDE algorithm between (4.43) and (4.45). Numbers correspond to the mean(std.dev.) of the total approximation errors of 5 independent runs of the algorithm. $T = 0.25, X_0 = (\pi/4, \dots, \pi/4)$.

N	1	5	10	20	50	100
total-Eq.(4.43)	1.49e−1(3e−3)	1.28e−2(3e−4)	5.5e−3(2e−4)	2.86e−3(5e−5)	2.02e−3(4e−5)	2.16e−3(6e−5)
total-Eq.(4.45)	1.52e−2(3e−3)	4e+2(2e+2)	2e+6(4e+6)	NaN	NaN	NaN

4.5.2. EXAMPLE 2

The following $d = 25$ dimensional example is related to a linear-quadratic stochastic control problem appearing in [4, example 3], which is defined by the following set of coefficients

$$
\begin{aligned}
M_x &= -\mathrm{diag}(1,2,3,1,2,3,1,2,3,1,2,3,1,2,3,1,2,3,1,2,3,1,2,3,1), \\
M_u &= (1,1,0.5,1,0,0,1,1,0.5,1,0,0,1,1,0.5,1,0,0,1,1,0.5,1,0,0,1)^\top, \\
M_c &= -M_x(-0.2,-0.1,0,0,0.1,0.2,-0.2,-0.1,0,0,0.1,0.2, \\
&\qquad -0.2,-0.1,0,0,0.1,0.2,-0.2,-0.1,0,0,0.1,0.2,-0.2)^\top, \\
\Sigma &= \mathrm{diag}(0.15,0.15,0.25,0.25,0.25,0.25,0.25,0.25,0.25,0.25,0.25, \\
&\qquad 0.25,0.15,0.15,0.25,0.25,0.25,0.25,0.25,0.25,0.25,0.25,0.25,0.25,0.25), \\
R_x &= 2\mathrm{diag}(25,1,25,1,25,1,25,1,25,1,25,1,25,1,25,1,25,1,25,1,25,1,25), R_u = 2, \\
G &= 2\mathrm{diag}(25,25,25,25,25,25,1,25,1,25,1,25,25,25,25,25,25,25,1,25,1,25,1,25,1).
\end{aligned}
\tag{4.46}
$$

One can derive an associated FBSDE system via either dynamic programming (DP) or the stochastic maximum principle (SMP), see e.g. [156]. The correspondi7eng equations' coefficients in (4.1) then take the following form in case of DP

$$
\begin{aligned}
b_{\mathrm{DP}}(t,x,y,z) &= M_x x - M_u R_u^{-1} M_u^\top (z\Sigma^{-1})^\top, \quad \sigma_{\mathrm{DP}}(t,x,y) = \Sigma, \\
f_{\mathrm{DP}}(t,x,y,z) &= 1/2(x^\top R_x x + z\Sigma^{-1}(R_u^{-1}M_u^\top)^\top M_u^\top (z\Sigma^{-1})^\top), \quad g_{\mathrm{DP}}(x) = 1/2 x^\top G x,
\end{aligned}
\tag{4.47}
$$

with $q = 1, d = m = 25$; and in case of SMP

$$
\begin{aligned}
b_{\mathrm{SMP}}(t,x,y,z) &= M_x x + M_u R_u^{-1} M_u^\top y, \quad \sigma_{\mathrm{SMP}}(t,x,y) = \Sigma, \\
f_{\mathrm{SMP}}(t,x,y,z) &= -R_x x + M_x y, \quad g_{\mathrm{SMP}}(x) = -Gx,
\end{aligned}
\tag{4.48}
$$

with $q = d = m = 25$. The main difference between the two formulations is that (4.47) leads to an FBSDE where coupling into b occurs through Z, whereas in (4.48) only through Y. Furthermore, the first equation gives a scalar-valued BSDE, whereas (4.48) is a vector-valued one. Both equations admit semi-analytical solutions given by the numerical resolution of a system of Ricatti ODEs. For details, we refer to [74] and the references therein.

Remark 4.5.1

Notice that the dynamic programming FBSDE (4.47) does not satisfy the Lipschitz con-

ditions imposed in section 4.3. In fact, g_{DP}, f_{DP} in (4.47) are quadratic in x, z. Nevertheless, one can use a localization argument, and consider the equation over a compact domain such that the corresponding coefficients become Lipschitz continuous with a constant depending on the width of the domain. We choose truncation radiuses based on upper bounds for 99% quantiles of $\|X_t\|$, $\|Z_t\|$ computed over an independent Monte Carlo simulation consisting of 2^{20} paths using the semi-analytical reference solution. This results in a negligible truncation error and truncation radiuses $r_x = 1, r_z = 10$, with which we obtain a Lipschitz continuous approximation of g_{DP}, f_{DP} for which the constants in section 4.3 can be computed even in the case of (4.47), and read as follows $L_x^g = r_x^2\|G\|_2^2/2, L_x^b = 2\|M_x\|_2^2, L_z^b = 2\|M_u R_u^{-1} M_u^\top (\Sigma^{-1})^\top\|_2^2, L_x^f = r_x^2\|R_x\|_2^2, L_z^f = r_z^2\|\Sigma^{-1}(R_u^{-1}M_u^\top)^\top M_u^\top (\Sigma^{-1})^\top\|_2^2, k^b = -1, L_y^b = L_x^\sigma = L_z^\sigma = L_y^f = k^f = 0$.

NON-CONVERGENCE OF THE DEEP BSDE METHOD

Let us first focus on the FBSDE stemming from the dynamic programming principle. In [4] it was observed that the deep BSDE method does not converge for the FBSDE defined by (4.47) with coefficients as in (4.46) and $T = 1/2$, $X_0 = (0.1, \dots, 0.1)$. Earlier convergence analyses such as [69] could not justify this phenomenon, as in (4.47) the coupling into the drift takes places via Z, which fell out of the framework of the aforementioned paper. Our generalization provided by theorem 4.3.3 enables the treatment of such equations, and in particular explains the non-convergence phenomenon. The problem lies in the *strength* of the coupling of Z into the forward diffusion. In order to demonstrate this, we consider two versions of (4.47), which differ in the coefficient M_u. One labelled as *"original"*, where the coefficients of the corresponding LQ problem are as in (4.46), and a *"rescaled"* version where M_u is replaced by $M_u/150$ and all other coefficients remain the same. Our findings are illustrated by figure 4.2. In particular, figure 4.2a depicts the contraction constants $\bar{B}, \bar{A}$ defined by (4.21), (4.22) appearing in theorem 4.3.3 for both versions. As can be seen, in case of the "original" equation one gets a lower bound $\bar{B}_\ell$ defined by (4.42) which is of $\mathcal{O}(10^{16})$. In fact, this implies that the conditions of theorem 4.3.3 can never be satisfied for the original version of (4.47). However, as is also suggested by figure 4.2a, decreasing the strength of the coupling by the given rescaling of M_u we get an equation whose $\bar{B}, \bar{A}$ satisfy the sufficient conditions (4.23). Motivated by this, we collect the convergence of the total approximation errors in figure 4.2b. In line with the discussion above, we find that the deep BSDE method does not converge in the "original" case, whereas it does converge for the "rescaled" version, once we make sure that there exist appropriate contraction constants $\bar{B}, \bar{A} < 1$ such that the sufficient conditions of theorem 4.3.3 are satisfied. The empirical rate of convergence is of $\mathcal{O}(h^{1.6})$. Note that this example illustrates the weak coupling condition described in point (3)c of section 4.4. We emphasize that $\bar{B}$ is inherent to the Z coupling and it is the main novelty of our convergence analysis.

STOCHASTIC CONTROL VIA DP OR SMP

In [74] a convergence analysis has been given in the context of solving stochastic control problems with the deep BSDE method applied on the FBSDE system derived through the stochastic maximum principle (SMP) similar to (4.48). This analysis provided a natural extension to the works of Han and Long by extending [69] to vector-valued settings.

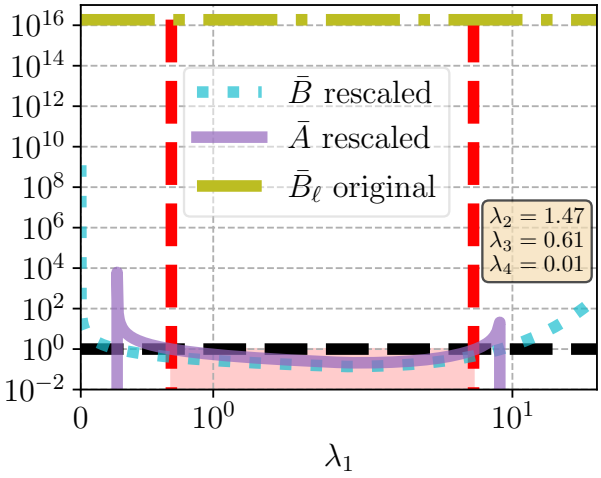

(a) $\bar{B}, \bar{A}$ as functions of λ_1 for given $(\lambda_2, \lambda_3, \lambda_4)$ in case of the "rescaled" version of (4.47). $\bar{B}_\ell$ as lower bound for $\bar{B}$ in case of the "original" version. Dotted vertical red lines mark the endpoints of the interval where $\bar{B}, \bar{A} < 1$, and the shaded red area the subset of the plane where the sufficient conditions of theorem 4.3.3 are satisfied.

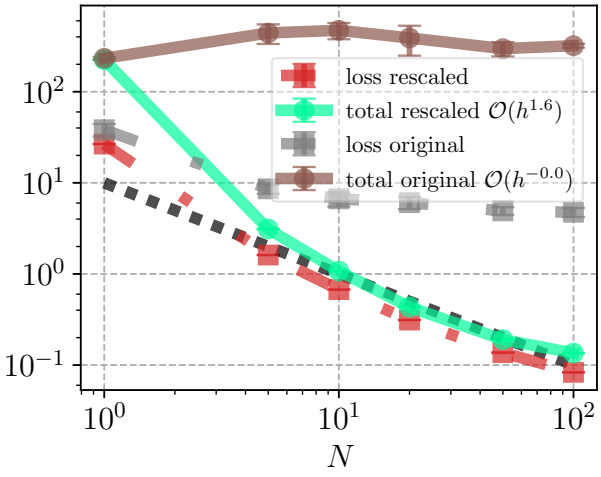

(b) Convergence in N. Empirical convergence rates in labels. Dotted black line indicates the expected $\mathcal{O}(h)$ rate predicted by theorem 4.3.3.

Figure 4.2: Example 4.5.2. Comparison of the Deep BSDE method on (4.47) between coefficients as in (4.46) (original) and M_u replaced by $M_u/150$ (rescaled). $T = 1/2$, $X_0 = (0.1, \ldots, 0.1)$.

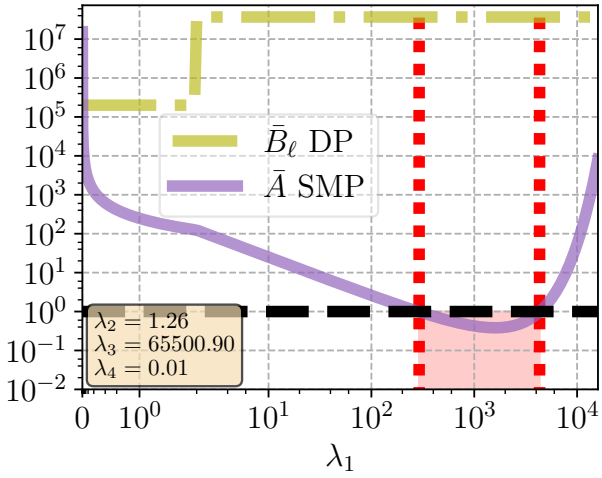

(a) $\bar{A}$ as function of λ_1 for given $(\lambda_2, \lambda_3, \lambda_4)$ in case of the SMP equation in (4.48). $\bar{B}_\ell$ as lower bound for $\bar{B}$ in case of the DP equation in (4.47). Dotted vertical red lines mark the endpoints of the interval where $\bar{A} < 1$, and the shaded red area the subset of the plane where the sufficient conditions of theorem 4.3.3 are satisfied.

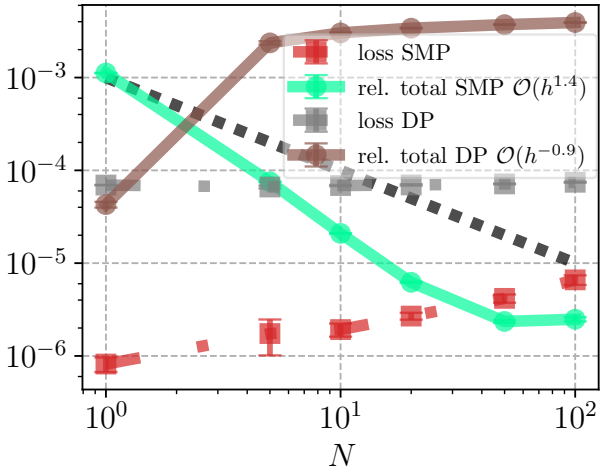

(b) Convergence in N. Empirical convergence rates in labels. Dotted black line indicates the expected $\mathcal{O}(h)$ rate predicted by theorem 4.3.3.

Figure 4.3: Example 4.5.2. Comparison between the FBSDEs derived via dynamic programming (4.47) and the stochastic maximum principle (4.48) approximated by the deep BSDE method. Coefficients as in (4.46), $T = 10^{-3}, X_0 = (0.1, \ldots, 0.1)$.

In [4] it was found that for certain FBSDEs derived in the dynamic programming framework, such as (4.47), the Deep BSDE method does not converge. On the other hand, the authors of the present chapter found in [74] that the same problems tackled by the stochastic maximum principle lead to an FBSDE (4.48) for which the deep BSDE method does converge to the unique solution triple. Our results in theorem 4.3.3 provide a natural explanation for these empirical findings. In fact, the problem lies in the Z coupling in (4.47), and in particular the value of the contraction constant $\bar{B}$ defined in (4.21). Conversely, (4.48) derived from the stochastic maximum principle has coupling only in Y, not in Z, leading to $\bar{B} \equiv 0$. In order to illustrate this, we ran the deep BSDE algorithm for both (4.47) and (4.48) with coefficients defined by (4.46), a short time horizon $T = 10^{-3}$ and $X_0 = (0.1, \dots, 0.1)$.

First, notice that only (4.48) satisfies the Lipschitz conditions imposed in section 4.3, with corresponding constants $L_x^g = \|G\|_2^2, L_x^b = 2\|M_x\|_2^2, L_y^b = 2\|M_u R_u^{-1} M_u^\top\|_2^2, L_x^f = 2\|R_x\|_2^2, L_y^f = 2\|M_x\|_2^2$, $k^f = k^b = -1, L_z^b = L_x^\sigma = L_y^\sigma = L_z^\sigma = L_z^f = 0$. Hence, only the SMP formulation guarantees direct applicability of the convergence results. Nevertheless, with the localization argument in remark 4.5.1 one can find an accurate Lipschitz continuous approximation of the DP equation (4.47) for which the sufficient conditions can also be checked. More importantly, (4.47) and (4.48) also differ in the type of coupling. Namely, in case of the former, Z couples into the forward diffusion, whereas in case of the latter only Y does. This in particular implies that besides the different Lipschitz constants, the two equations also differ in terms of the sufficient conditions of (4.24). In fact, for the SMP equation (4.48), there is no coupling in Z which implies $\bar{B} \equiv 0$, see also discussion in section 4.4. Hence (4.48) reduces to the theoretical framework of [74, 69] under which it is sufficient for $\bar{A} < 1$ to hold. On the other hand, (4.47) couples through Z, which in light of theorem 4.3.3 implies that $\bar{B}, \bar{A} < 1$ need to hold simultaneously, leading to stronger conditions to hold.

The above discussion is illustrated by figure 4.3. From figure 4.3a, we find that in case of the dynamic programming equation (4.47) $\bar{B}$ admits a lower bound $\bar{B}_\ell$ defined by (4.42) which is of $\mathcal{O}(10^5)$. In particular, this implies that the dynamic programming formulation can never satisfy the sufficient conditions imposed by theorem 4.3.3. On the contrary, under the SMP formulation we find a range of $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ such that $\bar{A} < 1$ meaning that the convergence criteria are met. Motivated by these conditions, the convergence of the deep BSDE method is collected in figure 4.3b for both equations. In line with our previous observations, we find that the method converges with an empirical rate of $\mathcal{O}(h^{1.4})$ for the SMP equation (4.48), which in this particular case is even faster than the rate predicted by [69, 74]. On the other hand, for the FBSDE derived via dynamic programming we find that the deep BSDE does not converge. As shown by figure 4.3b this is not due to the posteriori nature of the estimate in theorem 4.3.3, as the errors are growing even as the loss functional decreases. The critical phenomenon is indeed the coupling in Z, and the extra conditions it imposes as predicted by theorem 4.3.3. These findings complement our earlier convergence results in the context of stochastic control tackled by the deep BSDE method and the stochastic maximum principle [74]. In particular, our generalization in theorem 4.3.3 explains our empirical findings in [74] suggesting that for a large class of stochastic control problems, deriving an associated FBSDE through SMP leads to a system which is more tractable by deep BSDE formula-

tions. The crucial property here is the lack of Z coupling leading to milder conditions to ensure convergence according to our new convergence result in theorem 4.3.3.

4.6. CONCLUSION

In this chapter a generalized proof for the convergence of the deep BSDE method was given. Our main contributions can be summarized as follows. We extended the convergence analysis of [69] to FBSDEs with fully-coupled drift coefficients. Such an extension is essential in practice as it enables the treatment of FBSDEs stemming from stochastic optimal control problems. Our theory provides a unified framework and, in particular, includes earlier results from the literature as limit cases. Due to the extra Z coupling, the final posteriori error estimate stated in theorem 4.3.3 requires an additional condition expressed by (4.23). These sufficient conditions are directly verifiable for any equation, and as shown in section 4.4, they are satisfied under heuristic settings such as weak coupling, short time duration or monotonicity. Moreover, as demonstrated in section 4.5.2, our theory explains the non-convergence of the deep BSDE method observed in recent literature, and provides direct guidelines to avoid such issues and ensure convergence in practice. Several numerical experiments were presented for high-dimensional equations, which support and highlight key features of the theoretical findings.

5

A NUMERICAL FOURIER COSINE EXPANSION METHOD WITH HIGHER ORDER TAYLOR SCHEMES FOR FULLY COUPLED FBSDES

We are photons released from a dying star
We are fireflies a child has trapped in a jar

Nick Cave & the Bad Seeds (Fireflies)

This chapter is based on the preprint [117]: B. Negyesi and C. W. Oosterlee. *A numerical Fourier cosine expansion method with higher order Taylor schemes for fully coupled FBSDEs.* 2025. DOI: 10.48550/arXiv.2501.10988

5.1. INTRODUCTION

In this chapter we are concerned with solving the following fully-coupled system of forward-backward stochastic differential equations (FBSDE)

$$X_t = x_0 + \int_0^t \mu(s, X_s, Y_s, Z_s)\mathrm{d}s + \int_0^t \sigma(s, X_s, Y_s, Z_s)\mathrm{d}W_s, \tag{5.1a}$$

$$Y_t = g(X_T) + \int_t^T f(s, X_s, Y_s, Z_s)\mathrm{d}s - \int_t^T Z_s \mathrm{d}W_s, \tag{5.1b}$$

where all coefficients are deterministic, scalar valued functions, $T > 0$ is a finite time horizon, and $\{W_t\}_{t\in[0,T]}$ is a standard Brownian motion over an appropriate probability space. Such equations have an innate connection with second-order, parabolic, quasi-linear PDEs with terminal boundaries of the following form

$$\begin{aligned} \partial_t u(t,x) + \frac{1}{2}\widetilde{\sigma}^2(t,x,u,\partial_x u)\partial^2_{xx} u(t,x) \\ + \mu(t,x,u,\partial_x u\sigma)\partial_x u(t,x) + f(t,x,u,\partial_x u\sigma) = 0, \\ u(T,x) = g(x), \end{aligned} \tag{5.2}$$

where $\widetilde{\sigma}$ is connected to σ and the coefficients μ, f, g coincide with those in (5.1) – see e.g. [161, 52]. In particular, due to non-linear generalizations of the Feynman-Kac formula, the solution of (5.2) is related to the solution triple of (5.1) by the following relations

$$Y_t = u(t, X_t), \quad Z_t = \partial_x u(t, X_t)\sigma(t, X_t, u(t, X_t)) =: v(s, X_s), \tag{5.3}$$

at least in the case when the diffusion coefficient does not depend on Z [126, 161].

Fully-coupled FBSDEs naturally arise in applications of stochastic control, [156, 132]. Solving such equations analytically is seldom possible in an analytical fashion and one most often has to resort to numerical approximations instead. The literature in the so called *decoupled* framework, i.e. when the coefficients μ, σ do not depend on Y and Z are thoroughly studied, see e.g. [24]. However, things get more complicated in the *coupled* setting, i.e. when the coefficients μ, σ in (5.1a) take Y, Z as arguments. Due to the nature of the coupling, the forward simulation of (5.1a) is not straightforward, as one needs to have approximated the backward equation's solution pair (5.1b) in order to do so beforehand. Several approaches have been proposed in the past decades, starting with the famous *four step scheme* [102], built on the connection with the associated quasi-linear PDEs (5.3). Without the sake of completeness, we mention [17, 113, 36, 104, 161] and the references therein.

Classical techniques include Monte Carlo techniques, branching diffusions, quantization algorithms [40], Fourier cosine expansions. In particular, starting with the decoupled setting Ruijter and Oosterlee [144] proposed an algorithm for the numerical resolution of decoupled FBSDEs, built on the COS method [48], where the corresponding conditional expectations are approximated by means of Fourier cosine expansions enabled by the fact that, at least in the Markovian framework, the *characteristic function* of the transition between two time steps is known in closed form, given a suitable discretization of (5.1a). In their original work, the forward diffusion was assumed to be a standard

Arithmetic Brownian motion, giving pathwise analytical solutions for (5.1a), eliminating the need for discrete time approximations. Later on, they extended this framework in [145] to more general, state and time dependent drift and diffusion coefficients, up to discretizations by second-order Taylor schemes, including the *Euler-Maruyama* and *Milstein* schemes. In case of the former, their algorithm exhibited a strong convergence rate of order 1/2, whereas in case of the latter, this was improved to order 1. Nevertheless, these extensions were given in the decoupled setting, i.e. Y, Z did not enter the drift and diffusion coefficients in (5.1a).

In order to remedy this, following their ideas, Huijskens et al. in [76] extended their algorithm to the coupled FBSDE framework of (5.1). Therein, the forward and backward equations are first discretized by order 1/2 Euler schemes, and are subsequently decoupled by suitable *decoupling fields* which approximate the deterministic mappings in (5.3). Thereafter, they proposed three different strategies for finding the *optimal* decoupling fields: the so called *explicit method*, in which throughout the backward recursion, the decoupling relations at time t_n are replaced by the already computed approximations at time step t_{n+1}; a *local method* where the decoupling relations are obtained via Picard iterations at each time step; and a *global method* where Picard iterations take place simultaneously, at all points in time similar to [17]. They found that the explicit method, inspired by [40], is the most robust, giving strong convergence rate of order 1/2, even in the case of fully-coupled systems, where the diffusion coefficient σ depends also on Z, as in (5.1a).

The purpose of the present chapter is to combine [145] with [76] for the coupled system (5.1). Doing so, we extend the COS approximations of [76] to higher-order discretizations for fully-coupled FBSDE systems, such that strong convergence rates of order 1 and weak convergence rates of order 2 can be achieved. In particular, we generalize the discretization of the forward diffusion in (5.1a) to second-order Taylor schemes , see e.g. [90], including the *Milstein-* and *simplified order 2.0 weak Taylor* schemes. Subsequently, in order to preserve the higher-order convergence rates in the forward diffusion, we solve the corresponding backward recursions in the approximations of the backward equation by appropriate second-order schemes [164, 35], in order to obtain a suitable pair of decoupling fields (5.3). In this sense, our work is a generalization of [145], extending the approaches proposed therein from the decoupled to the fully-coupled setting. The main challenge in doing so, is that higher-order correction terms involve *derivatives* of the decoupling fields in (5.3). In order to remedy this, we capitalize on the fact that COS approximations are infinitely differentiable, and given an *explicit* scheme, these derivatives can be computed directly at each step in the backward recursion. This enables higher-order discretization schemes, up to strong convergence rates of order 1 with a *Milstein-*, and weak convergence rates of order 2 with a *simplified order 2.0 weak Taylor* scheme, when the characteristic function of the corresponding Markov transition is still available in closed form.

The rest of the chapter is organized as follows. Section 5.2 is devoted to the discrete time approximations of (5.1a). After introducing key notations and theoretical concepts, general, second-order schemes are formulated for the forward component (5.1a), followed by the standard, recursive sequence of conditional expectations related to the solution pair of the backward equation (5.1b). In order to make these discrete time ap-

proximations fully-implementable, section 5.3 explains the BCOS method in the coupled setting. Our main results are given in section 5.4, where numerical experiments are presented for a selection of coupled FBSDE systems, exhibiting the proclaimed rate of strong and weak convergence rates for a wide range of problems.

5.2. DISCRETE TIME APPROXIMATIONS OF FORWARD-BACKWARD STOCHASTIC DIFFERENTIAL EQUATIONS

Throughout the chapter we are working on a filtered probability space $(\Omega, \mathscr{F} = \cap_{t\in[0,T]}\mathscr{F}_t, \mathbb{P})$, where $\mathscr{F}_t$ is the natural filtration generated by the Brownian motion $\{W_t\}_{t\in[0,T]}$ augmented with the usual $\mathbb{P}$-null sets. The solution of (5.1) is a triplet of predictable stochastic processes $\{(X_t, Y_t, Z_t)\}_{t\in[0,T]}$ such that (5.1) is satisfied $\mathbb{P}$ almost surely, and the processes satisfy the following integrability conditions

$$\mathbb{E}\left[\sup_{t\in[0,T]}|X_t|^2\right] + \mathbb{E}\left[\sup_{t\in[0,T]}|Y_t|^2\right] + \mathbb{E}\left[\int_0^T |Z_s|^2 \mathrm{d}s\right] < \infty.$$

The well-posedness of (5.1) has been deeply studied, and established under by now classical assumptions, see e.g. [161, 126]. For the rest of chapter we consider equations such that any of such conditions is satisfied, and assume that (5.1) has a unique solution triplet $\{(X_t, Y_t, Z_t)\}_{t\in[0,T]}$.

In what follows, the following standard notations are used. Conditional expectations at $t_n \in [0,T]$ are denoted by $\mathbb{E}_n[\cdot] := \mathbb{E}[\cdot|\mathscr{F}_{t_n}]$. For a function $f : [0,T]\times\mathbb{R}\times\mathbb{R}\times\mathbb{R}\to\mathbb{R}$, we put $\partial_t f$ for the partial derivative in time, and $\partial_x f, \partial_y f, \partial_z f$ for the corresponding spatial derivatives. Second-order derivatives are denoted by $\partial^2_{yx} f = \partial_y(\partial_x f)$, and analogously for all other partial derivatives. Given a function $f : (t,x,y,z) \mapsto f$, and two deterministic mappings $\varphi : x \mapsto \varphi$, $\zeta : x \mapsto \zeta$, we put $\bar{f}^{\varphi,\zeta}(t,x) := f(t,x,\varphi(x),\zeta(x))$, consequently

$$\begin{aligned}\partial_x \bar{f}^{\varphi,\zeta}(t,x) \equiv \partial_x f(t,x,\varphi(x),\zeta(x)) + \partial_y f(t,x,\varphi(x),\zeta(x))\partial_x\varphi(x) \\ + \partial_z f(t,x,\varphi(x),\zeta(x))\partial_x\zeta(x).\end{aligned} \tag{5.4}$$

For the rest of the chapter, $\pi := \{0 = t_0 < t_1 < \cdots < t_N = T\}$, denotes a partitioning of $[0,T]$, with $\Delta t_n := t_{n+1} - t_n$, and $|\pi| := \max_{n=0,\dots,N-1}\Delta t_n$. The Brownian increments between two adjacent time points are given by $\Delta W_n := W_{t_{n+1}} - W_{t_n}$. For an equidistant time grid we use $h := T/N$, $t_n = nh$, $n = 0,\dots,N-1$. We will first discuss the discrete time approximations of the system (5.1) that are investigated in this chapter.

5.2.1. FORWARD DISCRETIZATION

The literature on numerical approximations for forward stochastic differential equations is vast, for a classical reference we refer to [90]. However, in case of coupled FBSDEs as in (5.1a) numerical approximations become more intricate due to the coefficient functions μ and σ depending on not just the state process X, but also the solution pair of the backward equation (Y,Z) in (5.1b). In order to enable the numerical approximation of forward equation, one first needs to *decouple* (5.1a) from (5.1b). Motivated by the Feynman-Kac formulae in (5.3), the standard technique is to construct a so called *decoupling field*, which is a deterministic mapping of X at each point in time. In fact,

the Markov property of the solution pair of (5.1b) together with (5.3) imply the existence of two *deterministic* mappings of time and state, such that the solution to the following forward SDE

$$X_t^{u,v} = x_0 + \int_0^t \mu(s, X_s^{u,v}, u(s, X_s^{u,v}), v(s, X_s^{u,v}))\mathrm{d}s + \int_0^t \sigma(s, X_s^{u,v}, u(s, X_s^{u,v}), v(s, X_s^{u,v}))\mathrm{d}W_s \tag{5.5}$$

coincides with that of (5.1a) pathwise, almost surely. Therefore, given the true decoupling fields u, v, one can construct a *decoupled* FBSDE system consisting of (5.5) and (5.1b), and proceed with standard approximations techniques developed for the decoupled framework in order to numerically solve the entire system in (5.1). The main difficulty in the approach described above is that the decoupling relations u, v are not known in advance, one has to provide numerical approximations of them in the discrete time framework, combined with the numerical approximation of the backward equation (5.1b).

5

To this end, let us fix two adjacent points t_n, t_{n+1} in the time partition and a decoupling pair $\varphi, \zeta : \mathbb{R} \to \mathbb{R}$. Following standard techniques of Itô-Taylor expansions, see e.g. [145] and the references therein, one can approximate the solution $\{X_t^{u,v}\}_{t\in(t_n,t_{n+1}]}$ of (5.5) by the second-order Taylor expansions of the general form

$$\begin{aligned} X_{t_n}^{u,v} &= x, \\ X_t^{t_n,x,\varphi,\zeta} &= x + \bar{m}^{\varphi,\zeta}(t,x)(t-t_n) + \bar{s}^{\varphi,\zeta}(t_n,x)(W_t - W_{t_n}) + \bar{\kappa}^{\varphi,\zeta}(t_n,x)(W_t - W_{t_n})^2, \end{aligned} \tag{5.6}$$

for $t \in (t_n, t_{n+1}]$, given $\varphi(x), \zeta(x)$ are accurate approximations of $u(t_n, x)$ and $v(t_n, x)$, respectively, for any x. We use the notation $X_{t_m}^{t_n,x,\varphi,\zeta} = X_m^{\pi,\varphi,\zeta}$ for $m = n, n+1$. The general second-order approximation in (5.6) includes famous discretization schemes such as the Euler-Maruyama scheme defined by

$$\bar{m}^{\varphi,\zeta}(t,x) = \bar{\mu}^{\varphi,\zeta}(t,x), \quad \bar{s}^{\varphi,\zeta}(t,x) = \bar{\sigma}^{\varphi,\zeta}(t,x), \quad \bar{\kappa}^{\varphi,\zeta}(t,x) = 0; \tag{5.7}$$

the Milstein scheme defined by

$$\begin{aligned} \bar{m}^{\varphi,\zeta}(t,x) &= \bar{\mu}^{\varphi,\zeta}(t,x) - \bar{\kappa}^{\varphi,\zeta}(t,x), \quad \bar{s}^{\varphi,\zeta}(t,x) = \bar{\sigma}^{\varphi,\zeta}(t,x), \\ \bar{\kappa}^{\varphi,\zeta}(t,x) &= \bar{\sigma}^{\varphi,\zeta}(t,x)\partial_x\bar{\sigma}^{\varphi,\zeta}(t,x)/2; \end{aligned} \tag{5.8}$$

or the *simplified order 2.0 weak Taylor* scheme given by

$$\begin{aligned} \bar{m}^{\varphi,\zeta}(t,x) &= \bar{\mu}^{\varphi,\zeta}(t,x) - \bar{\sigma}^{\varphi,\zeta}(t,x)\partial_x\bar{\sigma}^{\varphi,\zeta}(t,x)/2 \\ &\quad + \big(\partial_t\bar{\mu}^{\varphi,\zeta}(t,x) + \bar{\mu}^{\varphi,\zeta}(t,x)\partial_x\bar{\mu}^{\varphi,\zeta}(t,x) + \partial_{xx}^2\bar{\mu}^{\varphi,\zeta}(t,x)(\bar{\sigma}^{\varphi,\zeta}(t,x))^2/2\big)\Delta t_n/2 \\ \bar{s}^{\varphi,\zeta}(t,x) &= \bar{\sigma}^{\varphi,\zeta}(t,x) \\ &\quad + \big(\partial_x\bar{\mu}^{\varphi,\zeta}(t,x)\bar{\sigma}^{\varphi,\zeta}(t,x) + \partial_t\bar{\sigma}^{\varphi,\zeta}(t,x) + \bar{\mu}^{\varphi,\zeta}(t,x)\partial_x\bar{\sigma}^{\varphi,\zeta}(t,x) \\ &\qquad + \partial_{xx}^2\bar{\sigma}^{\varphi,\zeta}(t,x)(\bar{\sigma}^{\varphi,\zeta}(t,x))^2/2\big)\Delta t_n/2 \\ \bar{\kappa}^{\varphi,\zeta}(t,x) &= \bar{\sigma}^{\varphi,\zeta}(t,x)\partial_x\bar{\sigma}^{\varphi,\zeta}(t,x)/2. \end{aligned} \tag{5.9}$$

	Euler-Maruyama	Milstein	simplified order 2.0 weak Taylor
weak	1	1	2
strong	1/2	1	1

Table 5.1: Weak and strong convergence rates of (5.6) given different discretization schemes

It is well-known, see e.g. [90], that under standard assumptions the classical discretization schemes above converge to the continuous solution of (5.5) and thus also that of (5.1a), with strong and weak convergence rates as in table 5.1, where the convergence rates are defined as follows

$$\sup_{t\in[0,T]} (\mathbb{E}[|X_t^{u,v} - X_t^{\pi,\varphi,\zeta}|^2])^{1/2} \leq C|\pi|^{\gamma_s}, \quad |\mathbb{E}[p(X_t^{u,v})] - \mathbb{E}[p(X_t^{\pi,\varphi,\zeta})]| \leq C|\pi|^{\gamma_w}.$$

In the above $p(x)$ is any $2(\gamma_w+1)$ continuously differentiable function of polynomial growth, and C is a generic constant independent of $|\pi|$. Strong convergence implies convergence in probability to the true solution of (5.5), whereas weak convergence provides relevant information about the solution at $t=0$ whenever the initial condition of the forward SDE is deterministic, which in many financial applications is of special relevance.

Remark 5.2.1
Many higher-order Taylor expansions similar to (5.6) *could be considered such as the* order 1.5 strong Taylor *– see e.g. [90] –, including more terms of the corresponding Itô-Taylor expansion, involving powers of the Brownian increment* $W_t - W_{t_n}$ *which are higher than* 2. *The main reason why we restrict our further analysis to second-order schemes of the form* (5.6) *is due to the COS approximations that follow. In fact, when the highest power of the Brownian increment in* (5.6) *is at most* 2, *the corresponding Markov transition's characteristic function can be computed in closed form – see lemma 5.3.1 below – which is a key component of the COS method applied hereafter. However, for schemes of order higher than* 2 *this property no longer holds, and in order to be able to compute the COS formula, one would first have to numerically approximate the corresponding characteristic function.*

5.2.2. BACKWARD DISCRETIZATIONS

Let us turn to the discrete time approximations of the backward equation (5.1b). In light of (5.5), we fix a decoupling field (φ,ζ) and consider the corresponding decoupled BSDE defined by

$$Y_t^{\varphi,\zeta} = g(X_T^{\varphi,\zeta}) + \int_t^T f(s, X_s^{\varphi,\zeta}, Y_s^{\varphi,\zeta}, Z_s^{\varphi,\zeta})\mathrm{d}s - \int_t^T Z_s^{\varphi,\zeta}\mathrm{d}W_s, \tag{5.10}$$

which coincides with the solution pair of (5.1b) when $\varphi = u, \zeta = v$, established by (5.3). Given that the objective is to formulate a higher-order scheme in the fully-coupled setting, e.g. without knowing an exact decoupling pair (u,v), it is fundamental that the corresponding discrete time approximations of (5.10) admit a higher-order convergence rate than that of the classical backward Euler scheme of [24]. In what follows we consider two second-order schemes from the BSDE literature. Zhao et al. in [164] proposed

a generalized theta-scheme which approximates the Markovian solution pair of (5.10) by the following backward recursion of conditional expectations

$$\begin{aligned}
&y(t_N,x)=g(x),\quad z(t_N,x)=\partial_x g(x)\sigma(t_N,y(t_N,x),z(t_N,x)),\\
&z(t_n,x)=\frac{1}{\theta_3\Delta t_n}\big(\theta_4\Delta t_n\mathbb{E}_n^x[z(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})]+(\theta_3-\theta_4)\mathbb{E}_n^x[y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})\Delta W_n]\\
&\qquad\qquad+(1-\theta_2)\Delta t_n\mathbb{E}_n^x[f(t_{n+1},y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}),z(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}))\Delta W_n]\big),\\
&y(t_n,x)=\mathbb{E}_n^x[y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})+(1-\theta_1)\Delta t_n f(t_{n+1},y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}),z(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}))]\\
&\qquad+\theta_1\Delta t_n f(t_n,y(t_n,x),z(t_n,x)),
\end{aligned}\tag{5.11}$$

with $\theta_1,\theta_2\in[0,1]$, $\theta_3\in(0,1]$, $|\theta_4|\le\theta_3$ – see also their related works in [162, 165]. We remark that the generalized theta-scheme in (5.11) includes many classical discretization schemes such as the (implicit) backward Euler scheme of Bouchard and Touzi [24] with $\theta_1=\theta_2=\theta_3=1,\theta_4=0$; or the theta-scheme considered in [144, 76] with $\theta_2=\theta_3,\theta_4=\theta_3-1$ and also the one of [145] with the extra condition $\theta_1=\theta_2$. In [164] the authors show that the generalized theta-scheme in (5.11) has a strong convergence rate of order 2 in the decoupled framework, when $\theta_1=\theta_2=\theta_3=1/2$ and $\theta_4\le\theta_3$, given that the underlying forward diffusion is a Brownian motion, i.e. $\mu(t,x,y,z)=0,\sigma(t,x,y,z)=1$ in (5.5), and the coefficients of the BSDE g,f are sufficiently smooth with bounded derivatives. This makes the scheme (5.11) a suitable choice for the second-order approximation of (5.1), as using the estimates in (5.11) for the decoupling of (5.6) induces errors that scale with order of at most 2, preserving the convergence rates of a second-order Taylor scheme.

Alternatively, Crisan and Manolarakis in [35] proposed a second-order discretization for decoupled FBSDEs, defined by the following backward recursion

$$\begin{aligned}
&y(t_N,x)=g(x),\quad z(t_N,x)=\partial_x g(x),\\
&z(t_n,x)=\frac{4\Delta t_n+6t_n}{(\Delta t_n)^2}\mathbb{E}_n^x[(y(t_{n+1},x)+\Delta t_n f(X_{n+1}^{\pi,\varphi,\zeta},y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}),z(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}))\Delta W_n]\\
&\quad-\frac{6}{(\Delta t_n)^2}\mathbb{E}_n^x\Big[\Big(\int_{t_n}^{t_{n+1}}s\mathrm{d}W_s\Big)(y(t_{n+1},x)\\
&\qquad\qquad+\Delta t_n f(X_{n+1}^{\pi,\varphi,\zeta},y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}),z(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}))\Big],\\
&y(t_n,x)=\mathbb{E}_n^x[y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})+\Delta t_n/2 f(X_{n+1}^{\pi,\varphi,\zeta},y(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}),z(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})]\\
&\quad+\Delta t_n/2 f(x,y(t_n,x),z(t_n,x).
\end{aligned}\tag{5.12}$$

In [35], the authors prove second order convergence in $|\pi|$ for (5.12) under sufficiently smooth coefficients, generalizing second-order convergence beyond Brownian noise. Furthermore, they show that whenever those assumptions are not satisfied, their scheme still preserves the strong convergence rate of the backward Euler scheme of order 1/2.

We note that the discretization of Y is always implicit in (5.12) and also in (5.11) whenever $\theta_1>0$.

Single versus multi-step backward schemes. It is important to mention that higher-order discretizations of BSDEs have been thoroughly studied, and many important results have been established in this regard. In fact, besides the two one-step schemes

mentioned above, many multi-step approaches have been developed over the past decades, which may guarantee higher order convergence rates. Without the sake of completeness, we mention [27, 30, 166, 167] in the decoupled setting, and [163, 152] for the coupled framework. In case of the latter two, one can construct a multi-step discrete time approximation scheme for the BSDE (5.1b) which reads as follows

$$
\begin{aligned}
&y(t_N,x) = g(x), \quad z(t_N,x) = \partial_x g(x)\sigma(t_N,x),\\
&z(t_n,x) = \sum_{j=1}^{k} \alpha_j^k \mathbb{E}_n^x[(W_{t_{n+j}} - W_{t_n}) y(t_{n+j}, X_{n+j}^\pi)]\\
&y(t_n,x) = \frac{1}{\alpha_0^k}\big(-\sum_{j=1}^{k} \alpha_j^k \mathbb{E}_n^x[y(t_{n+j}, X_{n+j}^\pi)] - f(t_n,x,y(t_n,x),z(t_n,x))\big),
\end{aligned}
$$

where the $\{\alpha_j^k\}_{j=0,\dots,k}$ are known explicitly, for every $k = 1,\dots,6$. The scheme is supposed to converge with a rate of $\mathcal{O}(|\pi|^k)$, at least in the weak sense. The main difference between the thereby proposed multi-step approaches and that of the present chapter is twofold. First, the k-step multi-step schemes are not immediately implementable without appropriate approximations for the first Y, Z at the first k time steps closest to T. In other words, one has to compute $y(t_{N-j},x), z(t_{N-j},x)$ for $j = 1,\dots,k$ on either a finer time grid/using a higher-order FBSDE method, such as e.g. [164, 35], in order to be able to keep the same rate of convergence. Second, all multi-step schemes mentioned above require the numerical approximation of k conditional expectations with transitions between time step t_n and $t_{n+j}, j = 1,\dots,k$. As we shall see in the next section, in the context of the COS method this would imply the computation of the transition matrix Φ_n in (5.22), k times for each time step, which is the computationally most expensive part of the algorithm proposed therein. Therefore, a one-step scheme such as (5.11)-(5.12) is computationally preferable compared to multi-step alternatives when the spatial approximations are given by the COS method. Finally, and most importantly, the multi-step approximations for coupled FBSDEs in [163, 167, 152] are only providing higher-order convergence of the backward equation's solution pair Y, Z in (5.1b), and not that of (5.1a) which is still discretized by an Euler-Maruyama scheme. In other words, the somewhat surprising conclusion of [163, 152] is that higher-order convergence of Y, Z can be achieved even with a lower order scheme for the forward diffusion. In some applications, such as option pricing, this is sufficient as the main interest is the solution of the backward equation. However, in many other applications, e.g. stochastic optimal control, the forward diffusion is the main quantity of interest, in which case one wants to provide a higher-order scheme for the whole system in (5.1), making higher-order approximations for the forward diffusion inevitable.

For the reasons above, we discretize the backward equation by the one-step schemes (5.11)-(5.12). Both discrete time approximation schemes are only implementable given a machinery which approximates the conditional expectations on the left hand side. In our case, this will be done by the COS method, explained in the upcoming section.

5.3. COS APPROXIMATIONS

The COS method originally proposed in [48] is a Fourier based method to approximate conditional expectations, given that the underlying randomness is generated by a Markov transition whose characteristic function is known in (semi-)analytical closed form. In what follows we fix t_n, and the corresponding decoupling fields φ,ζ in (5.5). Then, we are interested in the Markov transition $X_n^{\pi,\varphi,\zeta} = x \mapsto X_{n+1}^{\pi,\varphi,\zeta}$ defined by the following second-order Taylor scheme

$$X_{n+1}^{\pi,\varphi,\zeta} = x + \bar{m}^{\varphi,\zeta}(t_n,x)\Delta t_n + \bar{s}^{\varphi,\zeta}(t_n,x)\Delta W_n + \bar{\kappa}^{\varphi,\zeta}(t_n,x)(\Delta W_n)^2. \tag{5.13}$$

The COS approximation of a conditional expectation of some function of $X_{n+1}^{\pi,\varphi,\zeta}$ then reads as follows

$$\mathbb{E}_n^x[v(t_{n+1}, X_{n+1}^{\pi,\varphi,\zeta})] \approx \sum_{k=0}^{K-1}{}' \mathcal{V}_k(t_{n+1})\,\mathrm{Re}\left\{\phi_{X_{n+1}^{\pi,\varphi,\zeta}}(k\pi/(b-a)|t_n,x)\exp(-ik\pi a/(b-a))\right\}, \tag{5.14}$$

where $\phi_{X_{n+1}^{\pi,\varphi,\zeta}}$ is the *characteristic function* of $X_{n+1}^{\pi,\varphi,\zeta}$ and the *Fourier cosine expansion coefficients* are defined by

$$\begin{aligned}\mathcal{V}_k(t_{n+1}) &:= \frac{2}{b-a}\int_a^b v(t_{n+1},x)\cos\left(\frac{k\pi}{b-a}(x-a)\right)\mathrm{d}x, \\ v(t_{n+1},x) &= \sum_{k=0}^{\infty}{}' \mathcal{V}_k(t_{n+1})\cos\left(\frac{k\pi}{b-a}(x-a)\right),\end{aligned} \tag{5.15}$$

and the notation $\sum'$ means that the 0'th term in the summation is halved.

For any given triple $\bar{m}^{\varphi,\zeta}, \bar{s}^{\varphi,\zeta}, \bar{\kappa}^{\varphi,\zeta}$, the following lemma, see [145, lemma 3.1], establishes an explicit, closed form expression for the characteristic function of the Markov transition $X_n^{\pi,\varphi,\zeta} = x \to X_{n+1}^{\pi,\varphi,\zeta}$ given a decoupling field (φ,ζ).

Lemma 5.3.1 (Characteristic function of Markov transitions, [145])
Consider (5.13)*, for any decoupling pair* φ,ζ*, the characteristic function of* $X_{n+1}^{\pi,\varphi,\zeta}$ *given* $X_n^{\pi,\varphi,\zeta} = x$ *reads as follows*

$$\begin{aligned}\phi_n^{\pi,\varphi,\zeta}(u|X_n^{\pi,\varphi,\zeta} = x) &:= \mathbb{E}_n^x\left[\exp\left(iuX_{n+1}^{\pi,\varphi,\zeta}\right)|X_n^{\pi,\varphi,\zeta} = x\right] \\ &= \frac{\exp\left(iu(x+\bar{m}^{\varphi,\zeta}(t_n,x)\Delta t_n) - \frac{u^2(\bar{s}^{\varphi,\zeta}(t_n,x))^2\Delta t_n}{2(1-2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n)}\right)}{\sqrt{(1-2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n)}}.\end{aligned} \tag{5.16}$$

Proof. The proof is analogous to that of [145, lemma 3.1]. Due to $\Delta W_n \sim \mathcal{N}(0,\Delta t_n)$, for an x such that $\bar{\kappa}^{\varphi,\zeta}(t_n,x) = 0$, the transition in (5.13) is normal with mean $x + \bar{\mu}^{\varphi,\zeta}(t_n,x)\Delta t_n$ and variance $(\bar{s}^{\varphi,\zeta}(t_n,x))^2\Delta t_n$. Equation (5.16) is then found by substituting into the characteristic function of a normal distribution.
For an x such that $\bar{\kappa}^{\varphi,\zeta}(t_n,x) \neq 0$, one can write

$$\begin{aligned}X_{n+1}^{\pi,\varphi,\zeta} &= x + \bar{\mu}^{\varphi,\zeta}(t_n,x)\Delta t_n - (\bar{s}^{\varphi,\zeta}(t_n,x))^2/(4\bar{\kappa}^{\varphi,\zeta}(t_n,x)) \\ &\quad + \bar{\kappa}^{\varphi,\zeta}(t_n,x)\left(\Delta W_n + \bar{s}^{\varphi,\zeta}(t_n,x)/(2\bar{\kappa}^{\varphi,\zeta}(t_n,x))\right)^2 \\ &\overset{\mathrm{d}}{=} x + \bar{\mu}^{\varphi,\zeta}(t_n,x)\Delta t_n - (\bar{s}^{\varphi,\zeta}(t_n,x))^2/(4\bar{\kappa}^{\varphi,\zeta}(t_n,x)) + \bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n\mathscr{X},\end{aligned} \tag{5.17}$$

where $\mathscr{X} \sim \chi'^2_\nu(\lambda(x))$ follows a noncentral χ^2 distribution with degrees of freedom $\nu = 1$ and non-centrality parameter $\lambda = (\bar{s}^{\varphi,\zeta}(t_n,x))^2/(2\bar{\kappa}^{\varphi,\zeta}(t_n,x)\sqrt{\Delta t_n})^2$, whose characteristic function is given by

$$\phi_{\chi'^2_{\nu(\lambda)}}(u|x) = \exp(i\lambda u/(1-2iu))/(1-2iu)^{-\nu/2}. \tag{5.18}$$

Combining (5.18) with (5.17), we finally obtain (5.16). □

The BCOS method, originating from [147] and later extended in [145, 76] proposes to approximate the conditional expectations in the discrete time approximations schemes of BSDEs collected in (5.11) and (5.12) by the COS formula (5.14). Indeed, by virtue of lemma 5.3.1, the conditional expectations defining the approximations of Y at time step n are fully-implementable, given the availability of Fourier cosine expansion coefficients at time step t_{n+1}. However, in order to be able to apply the COS method on the aforementioned discretizations, one also needs to solve conditional expectations of the form

$$\mathbb{E}_n^x\big[v(t_{n+1}, X_{n+1}^{\pi,\varphi,\zeta})(\Delta W_n)^k\big], \qquad k \in \mathbb{N},$$

which appear in the approximations for Z in (5.11) and (5.12). Given Fourier cosine expansion coefficients for the deterministic function $y \mapsto v(t_{n+1}, y)$ defined by (5.15), one has

$$\begin{aligned}
&\mathbb{E}_n^x[v(t_{n+1}, X_{n+1}^{\pi,\varphi,\zeta})(\Delta W_n)^k] \\
&= \sum_{l=0}^{\infty}{}' \mathcal{V}_l(t_{n+1})\mathbb{E}_n^x\left[\cos\left(l\pi\frac{X_{n+1}^{\pi,\varphi,\zeta}-a}{b-a}\right)(\Delta W_n)^k\right] \\
&= \sum_{l=0}^{\infty}{}' \mathcal{V}_l(t_{n+1})\,\mathrm{Re}\left\{\mathbb{E}_n^x\left[\exp\left(il\pi\frac{X_{n+1}^{\pi,\varphi,\zeta}}{b-a}\right)(\Delta W_n)^k\right]\exp(-il\pi a/(b-a))\right\} \\
&= \sum_{l=0}^{\infty}{}' \mathcal{V}_l(t_{n+1})\,\mathrm{Re}\big\{J_k(x|l\pi/(b-a))\exp(-il\pi a/(b-a))\big\},
\end{aligned}$$

where we put

$$J_k(x|u) := \mathbb{E}_n^x\left[\exp\left(iuX_{n+1}^{\pi,\varphi,\zeta}\right)(\Delta W_n)^k\right], \qquad k \in \mathbb{N}. \tag{5.19}$$

Given the Markov transition (5.13), the following lemma generalizes the truncated series expansion argument in [145, eq. 3.31], and is established by an integration by parts argument.

Lemma 5.3.2 (Integration by parts formulae with discretization (5.13))
Given the discretizations (5.13)*, the conditional expectations of the form* (5.19) *admit for each* $k \geq 1$

$$J_k(x|u) = \frac{iu\bar{s}^{\varphi,\zeta}(t_n,x)\Delta t_n}{1-2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n}J_{k-1}(x|u) + \mathbb{1}_{k\geq 1}(k)\frac{(k-1)\Delta t_n}{1-2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n}J_{k-2}(x|u).$$

Proof. Combining (5.19) with (5.13) we get

$$J_k(x|u) = \int_{\mathbb{R}} \exp\Big(iu(x + \bar{m}^{\varphi,\zeta}(t_n,x)\Delta t_n + \bar{s}^{\varphi,\zeta}(t_n,x)\xi + \bar{\kappa}^{\varphi,\zeta}(t_n,x)\xi^2)\Big) \times \xi^k \frac{\exp\big(-\xi^2/(2\Delta t_n)\big)}{\sqrt{2\pi\Delta t_n}} \mathrm{d}\xi,$$

as $\Delta W_n \sim \mathcal{N}(0,\Delta t_n)$. It is straightforward to check that

$$\begin{aligned}&\xi^k \exp\big(-\xi^2/(2\Delta t_n)\big)\\ &\quad = \mathbb{1}_{k\geq 1}(k)(k-1)\Delta t_n \xi^{k-2}\exp\big(-\xi^2/(2\Delta t_n)\big) - \Delta t_n \frac{\mathrm{d}(\xi^{k-1}\exp\big(-\xi^2/(2\Delta t_n)\big)}{\mathrm{d}\xi}.\end{aligned}$$

Plugging this in the above gives

$$\begin{aligned}J_k(x|u) &= \mathbb{1}_{k\geq 1}(k)(k-1)\Delta t_n J_{k-2}(x|u) \\ &\quad - \frac{\Delta t_n}{\sqrt{2\pi\Delta t_n}} \int_{\mathbb{R}} \exp\Big(iu(x + \bar{m}^{\varphi,\zeta}(t_n,x)\Delta t_n + \bar{s}^{\varphi,\zeta}(t_n,x)\xi + \bar{\kappa}^{\varphi,\zeta}(t_n,x)\xi^2)\Big) \\ &\qquad \times \frac{\mathrm{d}(\xi^{k-1}\exp\big(-\xi^2/(2\Delta t_n)\big)}{\mathrm{d}\xi}\mathrm{d}\xi \\ &=: \mathbb{1}_{k\geq 1}(k)(k-1)\Delta t_n J_{k-2}(x|u) + I_k(x|u).\end{aligned} \tag{5.20}$$

We apply integration by parts on the second term, which gives

$$\begin{aligned}I_k(x|u) &= -\Delta t_n \Bigg[\exp\Big(iu(x + \bar{m}^{\varphi,\zeta}(t_n,x)\Delta t_n + \bar{s}^{\varphi,\zeta}(t_n,x)\xi + \bar{\kappa}^{\varphi,\zeta}(t_n,x)\xi^2)\Big) \\ &\qquad \times \xi^{k-1}\frac{\exp\big(-\xi^2/(2\Delta t_n)\big)}{\sqrt{2\pi\Delta t_n}}\Bigg]_{-\infty}^{+\infty} \\ &\quad + \Delta t_n \int_{\mathbb{R}} iu\bar{s}^{\varphi,\zeta}(t_n,x)\exp\Big(iuX_{n+1}^{\pi,\varphi,\zeta}(\xi)\Big)\xi^{k-1}\frac{\exp\big(-\xi^2/(2\Delta t_n)\big)}{\sqrt{2\pi\Delta t_n}}\mathrm{d}\xi \\ &\quad + \Delta t_n \int_{\mathbb{R}} 2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\exp\Big(iuX_{n+1}^{\pi,\varphi,\zeta}(\xi)\Big)\xi^{k}\frac{\exp\big(-\xi^2/(2\Delta t_n)\big)}{\sqrt{2\pi\Delta t_n}}\mathrm{d}\xi \\ &\equiv 0 + iu\bar{s}^{\varphi,\zeta}(t_n,x)\Delta t_n J_{k-1}(x|u) + 2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n J_k(x|u),\end{aligned} \tag{5.21}$$

for any ν with sufficient radial decay in space. Combining (5.20) and (5.21) finally gives

$$J_k(x|u) = \frac{iu\bar{s}^{\varphi,\zeta}(t_n,x)\Delta t_n}{1-2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n} J_{k-1}(x|u) + \mathbb{1}_{k\geq 1(k)}\frac{(k-1)\Delta t_n}{1-2iu\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n} J_{k-2}(x|u).$$

□

Combining this with the result of lemma 5.3.2, introducing the notation

$$\Phi_n^{\pi,\varphi,\zeta}(u|x) := \phi_n^{\pi,\varphi,\zeta}(u|x)\exp(-iua), \tag{5.22}$$

we subsequently gather via recursion

$$\begin{aligned}
&\mathbb{E}_n^x[h(t_{n+1}, X_{n+1}^{\pi,\varphi,\zeta})] \\
&\quad = \sum_{l=0}^{\infty}{}' \mathscr{H}_l(t_{n+1}) \operatorname{Re}\Big\{\Phi_n^{\pi,\varphi,\zeta}(l\pi/(b-a)|x)\Big\}, \\
&\mathbb{E}_n^x[h(t_{n+1}, X_{n+1}^{\pi,\varphi,\zeta})\Delta W_n] \\
&\quad = \sum_{l=0}^{\infty}{}' \mathscr{H}_l(t_{n+1}) \operatorname{Re}\left\{\frac{il\pi\bar{s}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}{1-2il\pi\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}\Phi_n^{\pi,\varphi,\zeta}(l\pi/(b-a)|x)\right\}, \\
&\mathbb{E}_n^x[h(t_{n+1}, X_{n+1}^{\pi,\varphi,\zeta})(\Delta W_n)^2] \\
&\quad = \sum_{l=0}^{\infty}{}' \mathscr{H}_l(t_{n+1})\Bigg[\operatorname{Re}\left\{\left(\frac{il\pi\bar{s}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}{1-2il\pi\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}\right)^2\Phi_n^{\pi,\varphi,\zeta}(l\pi/(b-a)|x)\right\} \\
&\qquad\qquad \operatorname{Re}\left\{\frac{(k-1)\Delta t_n}{1-2il\pi\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}\Phi_n^{\pi,\varphi,\zeta}(l\pi/(b-a)|x)\right\}\Bigg].
\end{aligned}$$

Given the analytical expression for $\Phi_n^{\pi,\varphi,\zeta}(u|x)$ established by lemma 5.3.1, a sufficiently truncated finite cosine expansion gives the necessary COS estimates. Consequently, using lemma 5.3.2 and the expressions above, one can compute all conditional expectations arising in the scheme (5.11), provided that the Fourier cosine expansion coefficients of the deterministic functions $z(t_{n+1},x)$, $y(t_{n+1},x)$ and $\bar{f}^{z(t_{n+1},\cdot),y(t_{n+1},\cdot)}(t_{n+1},x)$ are available, or at least can be approximated.

Nonetheless, in order to make the scheme (5.12) fully-implementable in the BCOS framework, one needs to establish an additional integration by parts formula that allows the computations of conditional expectations of the form

$$\mathbb{E}_n^x\Big[\Big(\int_{t_n}^{t_{n+1}}(s-t_n)\mathrm{d}W_s\Big)v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})\Big], \tag{5.23}$$

appearing in (5.12). In the following proposition, we show that for given choices of $\theta_1,\theta_2,\theta_3,\theta_4$ the second-order scheme of [35] in (5.12) is included in the generalized theta-scheme of [164] in (5.11).

Proposition 5.3.1
For conditional expectations of the form (5.23), the following identity holds

$$\mathbb{E}_n^x\Big[\Big(\int_{t_n}^{t_{n+1}}s\mathrm{d}W_s\Big)v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})\Big] = \frac{\Delta t_n}{2}\mathbb{E}_n^x\big[\Delta W_n v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})\big]. \tag{5.24}$$

In particular, (5.12) is included in (5.11) for $\theta_1 = 1/2, \theta_3 = 1-\theta_2, \theta_4 = 0$.

Proof. Let us consider the conditional expectation given by (5.23). Similarly to lemma 5.3.2, we use an integration by parts argument to show that for $\Delta B_n := \int_{t_n}^{t_{n+1}}(s-t_n)\mathrm{d}W_s$,

$$\mathbb{E}_n^x[\Delta B_n v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})] = \frac{\mathbb{C}\mathrm{ov}[\Delta B_n,\Delta W_n]}{\mathbb{V}\mathrm{ar}[\Delta W_n]}\mathbb{E}_n^x[\Delta W_n v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta})]. \tag{5.25}$$

5

As the integrand in $\Delta B_n := \int_{t_n}^{t_{n+1}}(s-t_n)\mathrm{d}W_s$ is deterministic, we have that $\Delta B_n \sim \mathcal{N}(0,(\Delta t_n)^3/3 = \int_{t_n}^{t_{n+1}}(s-t_n)^2\mathrm{d}s)$, and due to Itô's isometry

$$\mathbb{E}[\Delta B_n]=0, \quad \mathbb{E}[(\Delta B_n)^2]=(\Delta t_n)^3/3, \quad \mathbb{C}\mathrm{ov}[\Delta B_n,\Delta W_n]=(\Delta t_n)^2/2. \tag{5.26}$$

Then the joint distribution of $(\Delta B_n,\Delta W_n)$ is a multivariate normal distribution with covariance matrix $\Sigma = \begin{pmatrix}(\Delta t_n)^3/3 & (\Delta t_n)^2/2\\ (\Delta t_n)^2/2 & \Delta t_n\end{pmatrix}$ and $\det(\Sigma) = (\Delta t_n)^4/12$. Consequently, the conditional expectation in (5.23) takes the following form

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}(\xi))\eta\frac{\exp\left(-\frac{\Delta t_n\eta^2-(\Delta t_n)^2\eta\xi+(\Delta t_n)^3\xi^2/3}{2}\right)}{2\pi(\Delta t_n)^4/12}\mathrm{d}\eta\mathrm{d}\xi. \tag{5.27}$$

By formal differentiation, it is straightforward to show that

$$\frac{\partial}{\partial\eta}\left(\frac{-\exp\left(-\frac{\Delta t_n\xi^2-(\Delta t_n)^2\eta\xi+(\Delta t_n)^3\xi^2/3}{2}\right)}{\Delta t_n}\right) = \eta\exp\left(-\frac{\Delta t_n\xi^2-(\Delta t_n)^2\eta\xi+(\Delta t_n)^3\xi^2/3}{2}\right) - \frac{\Delta t_n}{2}\xi\exp\left(-\frac{\Delta t_n\xi^2-(\Delta t_n)^2\eta\xi+(\Delta t_n)^3\xi^2/3}{2}\right),$$

and thus we can write (5.27) in the following way, by virtue of integration by parts and the Fubini theorem

$$\int_{-\infty}^{\infty}\frac{v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}(\xi))}{2\pi(\Delta t_n)^4/12}\left[\frac{-\exp\left(-\frac{\Delta t_n\xi^2-(\Delta t_n)^2\eta\xi+(\Delta t_n)^3\xi^2/3}{2}\right)}{\Delta t_n}\right]_{-\infty}^{\infty}\mathrm{d}\xi + \frac{\Delta t_n}{2}\int_{-\infty}^{\infty} v(t_{n+1},X_{n+1}^{\pi,\varphi,\zeta}(\xi))\xi\frac{\exp\left(-\frac{\Delta t_n\eta^2-(\Delta t_n)^2\eta\xi+(\Delta t_n)^3\xi^2/3}{2}\right)}{2\pi(\Delta t_n)^4/12}\mathrm{d}\eta\mathrm{d}\xi.$$

The integrand of the first term vanishes for any v with sufficient spatial radial decay, whereas the second term is identically equal to the right-hand side of (5.25) combined with (5.26). This completes the proof of identity (5.24).

Combining this with the second line of (5.12), we finally find that (5.12) is also included in the generalized θ-scheme of [164] given by (5.11), with the particular choice of $\theta_1 = 1/2, \theta_3 = 1-\theta_2, \theta_4 = 0$. □

In light of proposition 5.3.1, for the upcoming numerical experiments, we implement the theta-scheme (5.11) for general θs, which includes the second-order scheme in (5.12) for the choices above.

At this point the COS approximations of the conditional expectations in both (5.11)

and (5.12) can be obtained as follows

$$z(t_n,x) = \frac{1}{\theta_3\Delta t_n}\Bigg(\theta_4\Delta t_n \sum_{k=0}^{K-1}{}' \mathcal{Z}_k(t_{n+1})\,\mathrm{Re}\Big\{\Phi_n^{\pi,\varphi,\zeta}\big(k/(b-a)|x\big)\Big\} \tag{5.28}$$
$$+ \sum_{k=0}^{K-1}{}' \big((\theta_3-\theta_4)\mathcal{Y}_k(t_{n+1}) + (1-\theta_2)\Delta t_n \bar{\mathcal{F}}_k(t_{n+1})\big)$$
$$\times \mathrm{Re}\left\{\frac{ik\pi\bar{s}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}{1-2ik\pi\bar{\kappa}^{\varphi,\zeta}(t_n,x)\Delta t_n/(b-a)}\Phi_n^{\pi,\varphi,\zeta}(k\pi/(b-a)|x)\right\}\Bigg),$$

$$y(t_n,x) = \theta_1\Delta t_n f(t_n,x,y(t_n,x),z(t_n,x)) \tag{5.29}$$
$$+ \sum_{k=0}^{K-1}{}' (\mathcal{Y}_k(t_{n+1}) + (1-\theta_1)\Delta t_n \bar{\mathcal{F}}_k(t_{n+1}))\,\mathrm{Re}\Big\{\Phi_n^{\pi,\varphi,\zeta}\big(k/(b-a)|x\big)\Big\}.$$

However, these expressions are only fully implementable in the decoupled framework, i.e. when the drift and diffusion coefficients in (5.1a) do not depend on the solution pair of the backward equation. Indeed, in that case the characteristic function does not depend on the decoupling pair φ,ζ either – see (5.16). Nonetheless, in the coupled framework, one needs to find reasonable approximations of the true decoupling relations given in (5.5), and use those for (φ,ζ) in between two time steps t_n and t_{n+1}. In the paper of Huijskens et al. [76], using an Euler scheme for the discretization of the forward component – see (5.7), in particular $\bar{\kappa}^{\varphi,\zeta}(t_n,x)\equiv 0$ – there are three choices made in this regard. In the so called explicit method, argued by sufficient continuity in time of the solution pair of (5.1b), the decoupling relations are chosen to be the discrete time approximations at the next point in time $\varphi = y(t_{n+1},\cdot), \zeta = z(t_{n+1},\cdot)$. Note that (5.11) is a backward recursion, therefore these approximations are indeed available at the processing of time step t_n. Alternatively, [76] proposed a so-called local method, where – starting off from some initial conditions – the decoupling relations (φ,ζ) are gathered through Picard iterations at each time step t_n, assuming that the mapping $(\varphi,\zeta)\mapsto(y(t_{n+1},\cdot),z(t_{n+1},\cdot))$ defined by (5.28) is a contraction for small enough time steps. Finally, [76] also proposed a so called global approach, in which the solution pair is gathered through taking Picard iterations over the whole backward recursion in (5.11) – similar to [17]. As found in [76], the most efficient of these three options is the explicit method, as it only requires computing $\Phi_n^{\pi,\varphi,\zeta}(k\pi/(b-a)|x)$ once at every time step, and it does not require the mapping in (5.28) to be a contraction either. In our implementation, we confirmed the findings of [76] even in case of Milstein- and 2.0 weak Taylor discretizations, and found that a local method is both less stable and less accurate while creating an unnecessary computational overhead compared to the explicit decoupling. Therefore, in what follows we consider the explicit method only.

5.3.1. DERIVATIVE APPROXIMATIONS AND THE BCOS METHOD WITH HIGHER-ORDER TAYLOR SCHEMES

Nonetheless, in contrast to [76], the BCOS approximations (5.29)-(5.28) are not implementable, in case the forward diffusion is discretized with a second-order Taylor approximation, including the Milstein- and 2.0 weak Taylor approximations. Indeed, then the corresponding second-order term in (5.13) has $\bar{\kappa}^{\varphi,\zeta}\neq 0$. In particular, in case of the Mil-

stein scheme, through the total derivative of the diffusion coefficient, the derivatives of the decoupling relations also appear in the computation of $\bar{\kappa}^{\varphi,\zeta}$ which implies that one needs to choose decoupling relations carefully such that their derivatives are either available in closed form or can be accurately approximated. Our main result is built on the insight that whenever the decoupling relations in the BCOS method are chosen according to the explicit scheme – that is $\varphi = y(t_{n+1},\cdot), \zeta = z(t_{n+1},\cdot)$ – the derivatives $\partial_x\varphi$ and $\partial_x\zeta$ are analytically available provided by the smoothness of the Fourier cosine expansion (5.15). In fact, when the Fourier cosine expansion coefficients $\mathscr{Y}_k(t_{n+1}), \mathscr{Z}_k(t_{n+1})$ of $y(t_{n+1},\cdot), z(t_{n+1},\cdot)$ in (5.29) and (5.28) are known, then

$$\begin{aligned} y(t_{n+1},x) &= \sum_{k=0}^{K-1}{}' \mathscr{Y}_k(t_{n+1}) \cos\left(\frac{k\pi}{b-a}(x-a)\right), \\ z(t_{n+1},x) &= \sum_{k=0}^{K-1}{}' \mathscr{Z}_k(t_{n+1}) \cos\left(\frac{k\pi}{b-a}(x-a)\right), \end{aligned} \tag{5.30}$$

and their corresponding derivatives can analytically be computed as follows

$$\begin{aligned} \partial_x y(t_{n+1},x) &= \sum_{k=0}^{K-1}{}' -\frac{k\pi}{b-a} \mathscr{Y}_k(t_{n+1}) \sin\left(\frac{k\pi}{b-a}(x-a)\right), \\ \partial_x z(t_{n+1},x) &= \sum_{k=0}^{K-1}{}' -\frac{k\pi}{b-a} \mathscr{Z}_k(t_{n+1}) \sin\left(\frac{k\pi}{b-a}(x-a)\right). \end{aligned} \tag{5.31}$$

Similarly, the second derivatives appearing in $\bar{m}^{\varphi,\zeta}$ and $\bar{s}^{\varphi,\zeta}$ in (5.9) are given by

$$\begin{aligned} \partial^2_{xx} y(t_{n+1},x) &= \sum_{k=0}^{K-1}{}' -\left(\frac{k\pi}{b-a}\right)^2 \mathscr{Y}_k(t_{n+1}) \cos\left(\frac{k\pi}{b-a}(x-a)\right), \\ \partial^2_{xx} z(t_{n+1},x) &= \sum_{k=0}^{K-1}{}' -\left(\frac{k\pi}{b-a}\right)^2 \mathscr{Z}_k(t_{n+1}) \cos\left(\frac{k\pi}{b-a}(x-a)\right). \end{aligned} \tag{5.32}$$

For the BCOS approximation at time step t_n, we choose to decouple the Markov transition (5.13) of the forward diffusion by setting $\varphi = y(t_{n+1},\cdot)$ and $\zeta = z(t_{n+1},\cdot)$, in a similar fashion to the explicit method in [76]. In case the Markov transition is modeled by the Euler scheme (5.7), the Fourier cosine expansions in (5.30) are sufficient to compute (5.22) and thus to make the abstract scheme in (5.29)-(5.28) fully implementable. However, when the transition is approximated by a Milstein scheme then the coefficients $\bar{m}^{\varphi,\zeta}$ and $\bar{s}^{\varphi,\zeta}$ depend on the derivatives $\partial_x y(t_{n+1},\cdot)$ and $\partial_x z(t_{n+1},\cdot)$ via the total derivative of $\partial_x \bar{\sigma}^{y(t_{n+1},\cdot),z(t_{n+1},\cdot)}(t_n,x)$ according to (5.4). Nonetheless, due to (5.31), these expressions can be analytically computed for the choice $\varphi = y(t_{n+1},\cdot)$ and $\zeta = z(t_{n+1},\cdot)$, which makes the computations of the coefficients in (5.8) possible, and subsequently enables to compute $\Phi_n^{\pi,\varphi,\zeta}(u|x)$ even in the case of a Milstein transition.

Moreover, when the transition is given by a simplified order 2.0 weak Taylor approximation (5.9), then the coefficients $\bar{m}^{\varphi,\zeta}$ and $\bar{s}^{\varphi,\zeta}$ also depend on the second order derivatives of the decoupling fields φ and ζ, through $\partial^2_{xx}\bar{\sigma}^{\varphi,\zeta}(t_n,x)$ according to (5.4). Similarly to the previous reasoning, when the decoupling is chosen by $\varphi = y(t_{n+1},\cdot)$ and

$\zeta = z(t_{n+1},\cdot)$, the formulas for the second-order derivatives in (5.32) enable the computation of the coefficients in (5.9). Thereafter, (5.31) and (5.32) together enable us to compute the characteristic function in (5.22) even in case the Markov transition is modeled by a 2.0 weak Taylor approximation.

Coefficient recovery. The discussion above implicitly relied on the availability of the Fourier cosine expansion coefficients $\mathscr{Y}_k(t_{n+1})$ and $\mathscr{Z}_k(t_{n+1})$, defined by (5.15). However, during the backward recursion of the discrete approximation of the BSDE as in (5.29)-(5.28), these coefficients cannot be analytically computed. In order to make the transitions implementable, one needs to *numerically approximate* the continuous integral in (5.15), which can be achieved by a discrete Fourier cosine transform. In order to do this, we construct a discrete spatial grid partitioning the truncated integration range $[a,b]$ defined by the following points

$$x_l := a + (l+1/2)\frac{b-a}{K}, \quad l = 0,\dots,K-1, \tag{5.33}$$

given that the Fourier cosine expansion is truncated to K-many terms. We put $\Pi := \{x_l : l = 0,\dots,K-1\}$. Subsequently, the continuous integral in (5.15) can be approximated by

$$\mathscr{V}_k(t_{n+1}) \approx \frac{2}{K}\sum_{l=0}^{K-1} v(t_{n+1},x_l)\cos\left(k\pi\frac{2l+1}{2K}\right), \tag{5.34}$$

which is a Discrete Cosine Transform (DCT) approximation of type 2. Given the terminal condition of the BSDE in (5.1b), the coefficient recovery defined by (5.34) makes the backward recursion in (5.29)-(5.28) implementable for all forward transition schemes (5.7)-(5.8)-(5.9), together with (5.31)-(5.32) and the discussion above.

Picard iterations. Finally, we address the implicitness of the discrete time approximations in (5.29) in case $\theta_1 > 0$. When the time step Δt_n is sufficiently small, the implicit equation is uniquely defined for each x. In order to find the unique fixed point, we carry out Picard iterations indexed by p. We initialize the Picard iterations by the explicit approximation of (5.29) with $\theta_1 = 0$

$$h(t_n,x) = \sum_{k=0}^{K-1}{}'(\mathscr{Y}_k(t_{n+1}) + \Delta t_n \bar{\mathscr{F}}_k(t_{n+1}))\,\mathrm{Re}\left\{\Phi_n^{\pi,\varphi,\zeta}\left(k/(b-a)|x\right)\right\}, \tag{5.35}$$

and set $y^{(0)}(t_{n+1},x) = h(t_n,x)$. From thereon, the following Picard update is carried out

$$y^{(p+1)}(t_n,x) = y^{(p)}(t_n,x) + h(t_n,x) \tag{5.36}$$

until a required error tolerance level ε is reached, measured by

$$\max_{x\in\Pi}\left|y^{(p+1)}(t_n,x) - y^{(p)}(t_n,x)\right| \le \varepsilon. \tag{5.37}$$

In order to simplify the notation, we define $y_n^\pi(x) := y(t_{n+1},x)$ and $z_n^\pi(x) := z(t_{n+1},x)$ for the series expansions (5.30), with expansion coefficients recovered by DCT. With the above steps, the coupled BCOS method using higher-order Taylor discretizations as in (5.6) is now fully-implementable. The complete algorithm is given in algorithm 5.

Algorithm 5 Coupled BCOS algorithm with higher-order Taylor schemes

Require: $[a,b]$: integration range; K: number of Fourier terms; π: discrete time partition of $[0,T]$

Require: forward discretization scheme ((5.7) or (5.8) or (5.9)); $\theta_1,\theta_2,\theta_3,\theta_4$: generalized θ-scheme in (5.11)

$x_i = a + (i+1/2)(b-a)/K, i = 0,\dots,K-1$: spatial grid to compute DCT

if σ depends on Z **then**

  Solve $z_N^\pi(x_i) = \sigma(T, x_i, g(x_i), z_N^\pi(x_i))$ for all $i = 0,\dots,N-1$

end if

$\mathcal{Y}_k(T), \mathcal{Z}_k(T), \bar{\mathcal{F}}_k(T)$ ▷ Fourier coefficients at T – analytically or via DCT

for $n = N-1,\dots,0$ **do**

  $\varphi \leftarrow y_{n+1}^\pi, \zeta \leftarrow z_{n+1}^\pi$ ▷ decouple by approximations at next time step

  $\Phi_n^{\pi, y_{n+1}^\pi, z_{n+1}^\pi}$ ▷ compute characteristic function by (5.22) and (5.16) with (5.7) or (5.8) or (5.9)

  $z_n^\pi(x_i)$ ▷ compute Z approximations over spatial grid by (5.28)

  $y_n^{(0)} \leftarrow h(t_n, x)$ ▷ initialize Picard iterations by explicit approximation in (5.35)

  while ($p \le$ max. Picard iter.) **and** (tolerance in (5.37) is not reached) **do**

    $y_n^{(p+1)}(x_i)$ ▷ Picard update according to (5.36)

  end while

  $y_n^\pi(x_i) \leftarrow y_n^{p+1}(x_i)$

  $\mathcal{Y}_k(t_n), \mathcal{Z}_k(t_n), \bar{\mathcal{F}}_k(t_n)$ ▷ coefficient recovery by DCT (5.34) on $z_n(x_i), y_n^\pi(x_i), \bar{f}^{y_n^\pi, z_n^\pi}(t_n, x_i)$

end for

5.3.2. ERRORS AND COMPUTATIONAL COMPLEXITY

In what follows, we discuss the errors and computational complexity induced by the coupled BCOS method in algorithm 5.

Error analysis. The main sources of numerical errors and their contributions to the final approximation accuracy can be summarized as follows

- K – truncation of the Fourier cosine series: for smooth densities, the Fourier cosine expansion terms converge exponentially – see e.g. [48]. However, as algorithm 5 relies on Discrete Cosine Transforms to recover the coefficients $\mathcal{Y}_k(t_n), \mathcal{Z}_k(t_n), \mathcal{F}(t_n)$, the total error term only converges quadratically in the number of Fourier terms $\mathcal{O}(K^{-2})$, which is the accuracy of the numerical integration in (5.34);
- N – time discretization:
  - forward SDE (5.1a): depending on whether the forward transition is approximated by an Euler (5.7), Milstein (5.8) or 2.0 weak Taylor scheme (5.9), weak and strong errors converge according to the rates collected in table 5.1;
  - backward SDE (5.1b): considering the generalized θ-scheme of [164], the corresponding discrete time approximations in (5.11) converge with $\mathcal{O}(h^2)$ in the

strong sense when $\theta_1 = \theta_2 = \theta_3 = 1/2, \theta_4 \leq |\theta_3|$. This in particular includes the scheme of [35], for $\theta_4 = 0$ as established by (5.24);

- P – Picard iterations to approximate the implicit part of the conditional expectation in (5.29) when $\theta_1 > 0$: when the driver f is Lipschitz in its spatial arguments, the implicit mapping is contractive for small enough time steps Δt_n; then the Picard iterations converge exponentially, with the constant depending on the Lipschitz constants of the driver f in (5.1b);

- a, b – truncated integration range: for each problem the range should be chosen carefully wide enough such that its overall contribution is negligible. The precise impact of this error is difficult to quantify; a recent result on the optimal choice of a, b is given in [88], however, it is not straightforward to extend this proof to coupled FBSDEs as one can not have a-priori guarantees about the distribution of X at a given time due to the coupling. In the numerical experiments below, we choose a wide enough integration range such that the corresponding error term is negligible.

We emphasize that the derivative approximations (5.31) and (5.32), which enable second-order Taylor discretizations (5.6) of the forward SDE do not add numerical errors, as they can be computed analytically given the cosine expansion coefficients $\mathcal{Y}_k(t_n), \mathcal{Z}_k(t_n)$ at each time step. This is the key observation that extends the coupled BCOS method to Milstein and 2.0 weak Taylor approximations for the Markov transition, together with the closed-form characteristic function provided by lemma 5.3.1.

Computational complexity. The overall computational complexity of the second-order scheme proposed in algorithm 5 consists of the following components while processing time step t_n:

- $\mathcal{Y}_k(T), \mathcal{Z}_k(T), \bar{\mathcal{F}}_k(t_n)$ – terminal expansion coefficients: analytically $\mathcal{O}(K)$ or by DCT $\mathcal{O}(K \log(K))$;

- $\varphi \leftarrow y^{\pi}_{n+1}, \zeta \leftarrow z^{\pi}_{n+1}$ – computing decoupling fields: one needs to compute (5.30), (5.31) and (5.32) for each pair (x_i, k) leading to

Euler	Milstein	2.0-weak-Taylor
$\mathcal{O}(K)$	$\mathcal{O}(K^2)$	$\mathcal{O}(K^2)$

- $\Phi^{\pi,\varphi,\zeta}_n(k\pi/(b-a)|x_i)$ – computation of the characteristic function over the spatial grid x_i for each $k = 0, \ldots, K-1$: $\mathcal{O}(K^2)$;

- $y^{(p+1)}_n$ – Picard iterations for the implicit part in (5.29): $\mathcal{O}(PK)$

- $y^{\pi}_n(x_i), z^{\pi}_n(x_i)$ – BCOS approximations (5.29) and (5.28): $\mathcal{O}(K \log(K))$;

- $\mathcal{Y}_k(t_n), \mathcal{Z}_k(t_n), \bar{\mathcal{F}}_k(t_n)$ – coefficient recovery by DCT: $\mathcal{O}(K \log(K))$.

Apart from the first item, every point is repeated for all time steps $n = N-1, \ldots, 0$, and the coupled BCOS method scales linearly with respect to the number of discretization

points in time. The overall computational complexity is thus given by $\mathcal{O}(N(N+N^2+PN+N\log(N)))$. Comparing this to the computational complexity of the *explicit method* in [76], which uses the Euler discretization (5.7) in the approximation of the Markov transition (5.6), we find that our generalized higher-order BCOS method admits to the same computational complexity with a higher constant. In fact, our extensions to the method using higher-order Taylor schemes such as the Milstein (5.8) or 2.0 weak Taylor (5.9) schemes only create a computational overhead in the computations of the derivatives in (5.31) and (5.32), which scale as $\mathcal{O}(K^2)$ in both cases. As the most expensive part of the BCOS algorithm is to compute the characteristic function (5.16) at each time step, which requires $\mathcal{O}(K^2)$ operations, the additional steps for the Milstein and 2.0 weak Taylor discretizations only double and triple this, respectively. As we shall see in the numerical experiments presented below, this marginal computational overhead is justified by the significantly improved convergence rates – see table 5.1 – in the number of discretization points in time. In other words, even though the Milstein and 2.0 weak Taylor approximations require the extra computation of the derivatives of the decoupling fields, they reach a desired error tolerance level faster in N due to their higher strong and weak convergence rates, respectively.

5

Remark 5.3.1
This chapter is concerned with scalar valued FBSDE systems, i.e. the solution to (5.1) *is a triplet of scalar valued processes. Let us briefly highlight the main challenges one faces when generalizing algorithm 5 to higher dimensional equations. First, in case* (5.1a) *admits a vector-valued solution, one needs additional assumptions to preserve the convergence rates in table 5.1 for higher-order Taylor schemes. For instance, the Milstein scheme in* (5.8) *only has a strong convergence rate of* $\mathcal{O}(h)$ *for commutative noise – see [90]. Moreover, for a vector-valued* X *the corresponding decoupling fields in* (5.3) *are multivariate functions, admitting multi-dimensional Fourier cosine expansions in place of* (5.15). *In particular, this implies a Fourier expansion along* each dimension, *on top of the discretization* (5.33) *and numerical integration* (5.34) *of the multi-dimensional domain. Subsequently* $\Phi_n^{\pi,\varphi,\zeta}$ *in* (5.22) *becomes a* $K^d \times K^d$ *matrix. As discussed above, the computation of these transitional weights determined by the characteristic function is the most expensive part of algorithm 5, which would scale exponentially in the number of spatial dimensions. Regardless of this curse of dimensionality, the method could be extended up to dimension* 3 *with increasing memory constraints on* K. *We refer to [146, 139] where two-dimensional COS methods are presented, with the latter in the FBSDE framework. We refer to [146] where a two-dimensional COS methods is presented, outside of the FBSDE framework.*

5.4. NUMERICAL EXPERIMENTS

The BCOS method has been implemented in a Python library, which is openly accessible through the following github repository[1]. All computations were carried out on a Dell Alienware Aurora R10 machine equipped with an AMD Ryzen 9 3950X CPU (16 cores, 64Mb cache, 4.7 Ghz) using double precision. The solution triplet of each equation is computed by plugging the analytical solution of the corresponding quasi-linear PDE in

[1] https://github.com/balintnegyesi/coupled-BCOS

(5.3) into a simplified order 2.0 weak Taylor discretization of the forward diffusion in (5.1a) over a fine, equidistant time grid, consisting of $N' = 10^6$ equally sized intervals. This way the time discretization error of the reference solution is negligible.

Each FBSDE below is discretized by an equidistant time partition, consisting of $N+1$ points, leading to a uniform time step size $h = T/N$. The forward SDEs are discretized by (5.6), including Euler, Milstein and 2.0 weak Taylor approximations. Each BSDE below is discretized according to the generalized θ-scheme in (5.11) with $\theta_1 = \theta_2 = \theta_3 = 1/2$. In order to distinguish between the two second-order schemes of the backward equation given by (5.11) and (5.12), we consider the values $\theta_4 = -1/2$ and $\theta_4 = 0$ corresponding to (5.12) as shown by (5.24). For the COS method we specify an integration range $[a, b]$ wide enough for each problem specifically, such that the truncation error is negligible. For the implicit part of (5.29), we set the maximum number of Picard iterations in (5.36) to 100, and the tolerance level in (5.37) to $\varepsilon = 10^{-15}$.[2]

Strong L^2 errors are computed over an independently simulated Monte Carlo sample consisting of $M = 2^{10}$ paths of the Brownian motion, using the following discrete approximations

$$\begin{aligned}
\text{strong error } X &= \max_{n=0,\dots,N} \left(\frac{1}{M} \sum_{m=1}^{M} |X_n^\pi(m) - X_{t_n}(m)|^2 \right)^{1/2}, \\
\text{strong error } Y &= \max_{n=0,\dots,N} \left(\frac{1}{M} \sum_{m=1}^{M} |Y_n^\pi(m) - Y_{t_n}(m)|^2 \right)^{1/2}, \\
\text{strong error } Z &= \left(\frac{T}{NM} \sum_{n=0}^{N-1} \sum_{m=1}^{M} |Z_n^\pi(m) - Z_{t_n}(m)|^2 \right)^{1/2}.
\end{aligned} \tag{5.38}$$

The total strong error is given by the sum of the above three terms. Given the deterministic, fixed initial condition in (5.1a) for each of the equations below, the approximation errors at $t = 0$ coincide with the weak errors in mean

$$\text{error } Y_0 = |y_0^\pi(x_0) - u(0, x_0)|, \quad \text{error } Z_0 = |z_0^\pi(x_0) - v(0, x_0)|.$$

The total weak approximation error at t_0 is given by the sum of the above two terms.

5.4.1. EXAMPLE 1: DECOUPLED FBSDE

In order to demonstrate the slight differences between the discretizations (5.11) and the one applied in [145], on top of the generalization provided by lemma 5.3.2, our first example corresponds to the special case of decoupled FBSDEs. This example appears in

[2] thanks to the exponential convergence of the Picard iterations, the tolerance level is usually reached in 5 Picard iterations

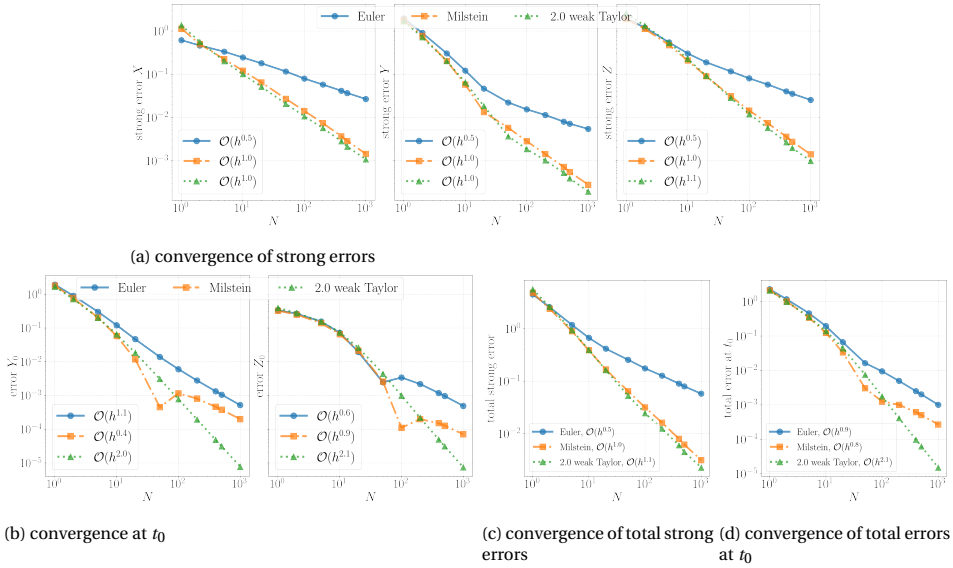

(a) convergence of strong errors

(b) convergence at t_0

(c) convergence of total strong errors

(d) convergence of total errors at t_0

Figure 5.1: Example 1 in (5.39). ($\theta_4 = 0, \theta_1 = \theta_2 = \theta_3 = 1/2$, $K = 2^9$.)

[145, example 1] and is originally from [113]. The coefficients in (5.1) read as follows

$$\begin{aligned}
&\mu(t,x,y,z) = x(1+x^2)/(2+x^2), \quad \sigma(t,x,y,z) = (1+x^2)/(2+x^2),\\
&\qquad g(x) = \exp\big(-x^2/(T+1)\big),\\
&f(t,x,y,z) = \frac{1}{t+1}\exp\big(-x^2/(t+1)\big)\left[4x^2\frac{1+x^2}{(2+x^2)^3} + \left(\frac{1+x^2}{2+x^2}\right)^2\left(1-\frac{2x^2}{t+1}-\frac{x^2}{t+1}\right)\right] \qquad (5.39)\\
&\qquad\qquad + \frac{zx}{(2+x^2)^2}\sqrt{\frac{1+y^2+\exp\big(-2x^2/(t+1)\big)}{1+2y^2}}.
\end{aligned}$$

The analytical solution pair to (5.1b) is given by the following deterministic mappings in (5.3)

$$u(t,x) = \exp\big(-x^2/(t+1)\big), \quad v(t,x) = -\frac{2x(1+x^2)}{(t+1)(2+x^2)}\exp\big(-x^2/(t+1)\big).$$

In line with [145], we choose $x_0 = 1$, $T = 10$ and for the COS method, we fix $K = 512$ Fourier coefficients, and set the domain $[a,b]$ as in[3] [145].

The numerical results are collected in figure 5.1 for all forward discretizations in (5.6), using the generalized θ-scheme (5.11) with $\theta_1 = \theta_2 = \theta_3 = 1/2, \theta_4 = 0$. In figure 5.1a, the strong convergence rates of all three processes are presented for each method, computed according to (5.38). As we can see, when the forward transition in (5.6) is approximated by the Euler discretization (5.7), each process converges with a rate of $\mathcal{O}(h^{1/2})$ in

[3] resulting in $a \approx -19.341110327048455$ and $b \approx 22.822591808529936$

K	1	2	2^2	2^3	2^4	2^5	2^6	2^7	2^8	2^9	2^{10}
Euler	0.17	0.18	0.19	0.20	0.22	0.29	0.56	1.63	5.80	23.75	107.56
Milstein	0.19	0.20	0.21	0.22	0.24	0.33	0.65	1.93	6.89	28.00	126.36
2.0 weak Taylor	0.26	0.27	0.28	0.29	0.31	0.40	0.72	2.02	7.02	28.40	121.63

Table 5.2: Example 2 in (5.40), with $\kappa_z = 10^{-2}$. CPU runtime in seconds. ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $\theta_4 = -1/2$, $N = 10^3$.)

the step size. However, this rate can be improved by the Milstein (5.8) and 2.0 weak Taylor schemes (5.9), which both show an asymptotic convergence rate of $\mathcal{O}(h)$. The convergence of the total strong errors is given in figure 5.1c, which demonstrates the improved first-order strong convergence with the Milstein and 2.0 weak Taylor discretizations. Figure 5.1b can be directly compared to fig. 5.1 in [145]. Regardless of the slight difference in the backward theta-scheme considered therein, one can draw similar conclusions about the convergence of the approximation errors at $t_0 = 0$. Given the deterministic condition in (5.1a), errors at t_0 coincide with weak errors, and the corresponding convergence rates admit weak convergence rates. This results in the Euler and Milstein discretizations exhibiting a weak convergence at t_0 with rate $\mathcal{O}(h)$; whereas the 2.0 weak Taylor scheme improves this to $\mathcal{O}(h^2)$. We remark that the latter convergence, in light of (5.9), also implies the convergence of the second-order derivatives (gammas) in (5.32). Experiments with other choices of $\theta_1, \theta_2, \theta_3$ confirmed the findings of [145], in terms of the weak convergence at t_0. In particular, by $N = 10^3$, one gains approximately two orders of magnitude accuracy at t_0 by employing the 2.0 weak Taylor scheme. In many applications – e.g. hedging, or portfolio allocation – this is of high importance. Additionally, and in line with [145], we found that in order to have a strong convergence of rate $\mathcal{O}(h)$ for the Milstein and 2.0 weak Taylor schemes, it is necessary for the backward component to be discretized by a second-order scheme.

5.4.2. EXAMPLE 2: PARTIAL COUPLING

The following coupled FBSDE system from [115] is an adaptation of [17], including Z coupling in the forward diffusion's drift coefficient. The coefficient functions in (5.1) read as follows

$$\begin{aligned} &\mu(t,x,y,z) = \kappa_y \bar{\sigma} y + \kappa_z z, \quad \sigma(t,x,y) = \bar{\sigma} y, \quad g(x) = \sin(x), \\ &f(t,x,y,z) = -ry + 1/2 e^{-3r(T-t)} \bar{\sigma}^2 \sin^3(x) - \kappa_y z - \kappa_z \bar{\sigma} e^{-3r(T-t)} \sin(x)\cos^2(x). \end{aligned} \tag{5.40}$$

The Markovian solution pair of the BSDE is given by the following deterministic mappings in (5.3) solving the corresponding quasi-linear PDE

$$u(t,x) = e^{-r(T-t)} \sin(x), \quad v(t,x) = e^{-2r(T-t)} \bar{\sigma} \sin(x)\cos(x).$$

We take $T = 1, X_0 = \pi/4$, fix $r = 0, \bar{\sigma} = 0.4$ and $\kappa_y = 10^{-1}$, which results in a numerically challenging equation with no monotonicity and strong coupling – see [17]. We fix a wide integration range by choosing $a = -3, b = 5$.

For the coupling of Z in the drift, we consider two different values for κ_z:

1. $\kappa_z = 0$: the backward equation only couples into the forward equation through the Y process;

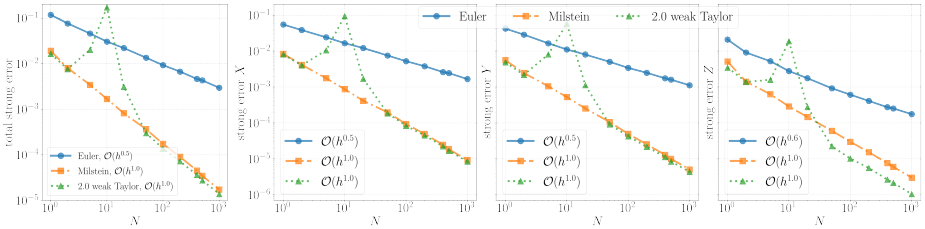

(a) $\theta_4 = -1/2$

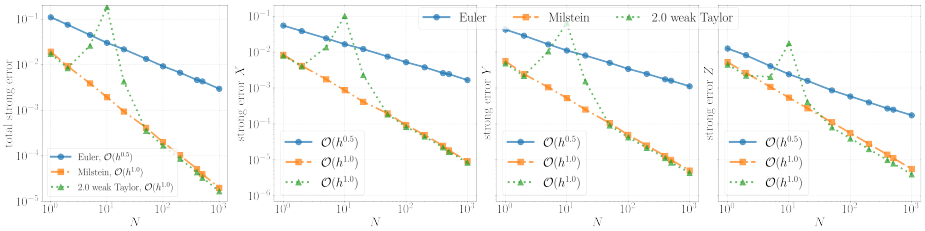

(b) $\theta_4 = 0$

Figure 5.2: Example 2 in (5.40) with $\kappa_z = 0$. Strong convergence. ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $K = 2^{10}$.)

2. $\kappa_z = 10^{-2}$: Z is also coupled into the forward SDE but only through the drift and not the diffusion coefficient.

COUPLING ONLY IN Y

We start by presenting results on the first case, when $\kappa_z = 0$, i.e. the Z process does not enter the forward diffusion (5.1a). In order to be able to neglect the Fourier truncation error term in the second-order weak convergence of the 2.0 weak Taylor discretization for very fine time grids, we use a larger number of Fourier terms $K = 2^{10}$ and remark that the rest of the results are close to identical for significantly smaller number of expansion coefficients. The strong convergence rates for each process of the solution triplet are plotted against the number of discretization points in time in figure 5.2 for $\theta_4 = 0$ (fig.5.2a) and $\theta_4 = -1/2$ (fig.5.2b). The strong errors across the different backward discretizations are comparable, with an order of magnitude gain when using $\theta_4 = -1/2$ in the approximation of Z while using Milstein or 2.0 weak Taylor forward schemes and fine time partitions. Each discretization exhibits the theoretically expected strong convergence rate predicted in table 5.1 for both θ_4 values. In fact, we recover a strong convergence rate of $\mathcal{O}(h^{1/2})$ in case of the Euler transitions, as in [76]. However, thanks to the generalization by algorithm 5 to higher order schemes, we manage to improve the strong convergence rate to $\mathcal{O}(h)$ by using the Milstein and 2.0 weak Taylor transitions in (5.6), even in the case of coupling in Y. In particular, due to the higher order discretizations enabled by algorithm 5, the Milstein and 2.0 weak Taylor approximations achieve an almost three orders of magnitude higher accuracy in the total strong approximation error than the Euler method in [76], when $N = 10^3$.

The main difference between the choices in θ_4 is illustrated by figure 5.3, where the weak convergence rates of the approximation errors at t_0 are depicted. As can be seen,

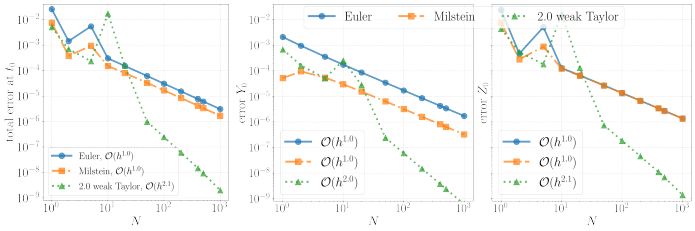

(a) $\theta_4 = -1/2$

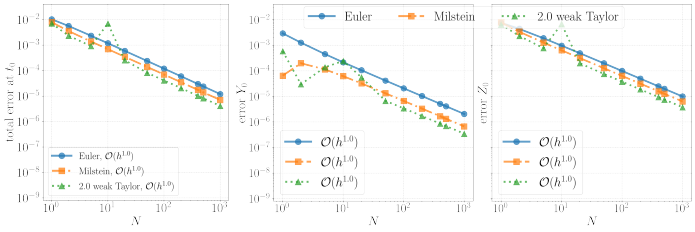

(b) $\theta_4 = 0$

Figure 5.3: Example 2 in (5.40) with $\kappa_z = 0$. Weak convergence at t_0. ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $K = 2^{10}$.)

when the forward diffusion is approximated by either an Euler or Milstein scheme, the weak errors at t_0 converge with the expected $\mathcal{O}(h)$ regardless of the value of θ_4. However, the same cannot be said about the 2.0 weak Taylor discretization. In fact, we find that the errors at t_0 only show second-order convergence when $\theta_4 = -1/2$. With $\theta_4 = 0$, the accumulating approximation errors from the backward equation result in a slower convergence of the same $\mathcal{O}(h)$ or as with the Euler and Milstein schemes. These findings suggest that in applications where the approximation accuracy at t_0 is of special relevance, $\theta_4 = -1/2$ may be a preferred choice for the discretization of the backward equation, when the forward transition is modeled with a 2.0 weak Taylor scheme.[4]

Z COUPLING IN THE DRIFT

Let us consider the second case of (5.40), corresponding to $\kappa_z = 10^{-2}$, i.e. when the Z process enters the dynamics of the forward diffusion (5.1a) but only through the drift. Similar to the previous case, we fix $K = 2^{10}$, and as the conclusion on θ_4 is verbatim, we only present results in the case $\theta_4 = -1/2$.

Numerical results are collected in figure 5.4. The strong approximation errors are depicted in fig. 5.4a. As we can see, the Milstein and 2.0 weak Taylor approximation preserve their theoretically expected $\mathcal{O}(h)$ convergence rate – improving on that of the Euler scheme – even when Z enters the forward dynamics. For fine time grids such as $N = 10^3$ this results in a more than 2 orders of magnitude improvement compared to the method employed in [76]. Moreover, similarly to the previous case, the theoretically

[4] The choice $\theta_4 = 0$ may be necessary in some applications where the solutions are highly oscillating. In such case with $\theta_4 = -1/2$ one could face stability issues, whereas the second-order scheme in (5.12) would preserve the convergence rate of the standard backward Euler scheme – see [35]. Our results suggest that even in such case, the strong convergence rates can be improved by the second-order Taylor schemes in (5.6).

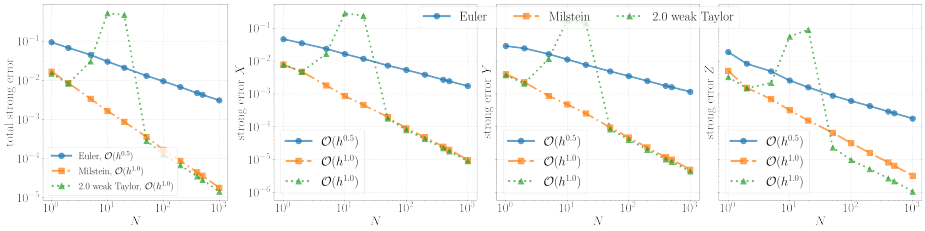

(a) strong convergence

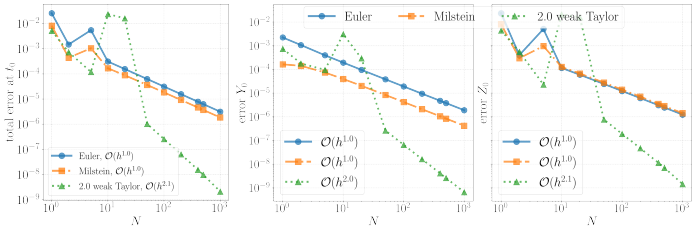

(b) weak convergence at t_0

Figure 5.4: Example 2 in (5.40), with $\kappa_z = 10^{-2}$. ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $\theta_4 = -1/2$, $K = 2^{10}$.)

expected weak convergence rates are recovered for each forward discretization for the approximation errors computed at t_0, as illustrated by fig. 5.4b. In fact, we find that whereas the Milstein scheme brings a marginal improvement in the accuracy at t_0 compared to the Euler scheme, they both exhibit first order weak convergence. On the other hand, by employing a 2.0 weak Taylor discretization in the forward diffusion's Markov transition as in (5.9), one achieves second-order convergence. This results in the approximation accuracy at t_0 reaching practically machine accuracy by $N = 10^3$, yielding 3 orders of magnitude gain compared to Euler and Milstein. As discussed in section 5.3.2, this comes with additional computational complexity driven by 2×2 matrix vector multiplications, in order to compute the corresponding first- and second-order derivatives of the decoupling relations in (5.31) and (5.32). Table 5.2 translates this additional computational complexity into CPU time by collecting the total runtime of the BCOS method for each forward discretization and different choices of the truncation of the Fourier series, with a fixed $N = 10^3$. Unsurprisingly, the Euler approximations are the fastest as they do not require the computations of the derivatives in (5.31) and (5.32). Nonetheless, the Milstein approximations are executed in merely 20% additional CPU time, while gaining an extra order in strong convergence. Moreover, the 2.0 weak Taylor approximations are the most computationally expensive, as on top of the first order derivatives of the decoupling fields, they also require the computation of the second-order derivatives in (5.32). Interestingly, the difference between the CPU times of the Milstein and 2.0 weak Taylor approximations vanishes as K increases, which is explained by the reduced number of Picard iterations in the computation of the implicit conditional expectation of y in order to reach our desired tolerance level – see discussion in section 5.3.2. We remark that the difference in computation times also shows that employing a higher-order Taylor scheme in the approximation of the forward SDE of a coupled FBSDE sys-

tem is advantegeous compared to Richardson extrapolating at t_0 with the Euler scheme as in [76, sec. 4.3]. Namely, since Richardson extrapolation would require re-running the method with a finer time grid (e.g. $h/2$) once, doing so the computation time doubles, while for smoothly converging errors at t_0 the weak convergence rate is improved to $\mathcal{O}(h^2)$. Nonetheless, Richardson extrapolation does not improve the $\mathcal{O}(h^{1/2})$ strong convergence rate of the Euler scheme , while being both more restrictive and expensive than a 2.0 weak Taylor approximation, as indicated by table 5.2.

5.4.3. EXAMPLE 3: FULLY COUPLED FBSDE, STOCHASTIC OPTIMAL CONTROL

Our final numerical example is a fully-coupled FBSDE system which is related to a linear-quadratic stochastic optimal control problem and is derived from the stochastic maximum principle – see e.g. [156, 132]. For the derivation of the corresponding FBSDE, we refer to [74] and the references therein. The coefficients in (5.1) read as follows

$$\begin{aligned}
\mu(t,x,y,z) &= \left(A - B\frac{R_{xu}}{R_u}\right)x + \frac{B^2}{R_u}y + \frac{BD}{R_u}z + \beta, \\
\sigma(t,x,y,z) &= \left(C - D\frac{R_{xu}}{R_u}\right)x + \frac{DB}{R_u}y + \frac{D^2}{R_u}z + \Sigma, \quad g(x) = -Gx, \\
f(t,x,y,z) &= \left(A - \frac{BR_{xu}}{R_u}\right)y + \left(C - \frac{DR_{xu}}{R_u}\right)z - \left(R_x - \frac{R_{xu}}{R_u}\right)x.
\end{aligned} \tag{5.41}$$

Notice that, unlike in the previous example, the diffusion coefficient of the forward SDE also takes Z as an argument. The semi-analytical solution can be obtained by numerical integration of a system of Ricatti ODEs, with practically arbitrary accuracy – see [74, eq. (45)]. We take $N' = 10^6$ steps in order to compute the reference solution over a refined time partition. As a one-dimensional version of example 1 in [74], we consider the parameter values $A = -1, B = 0.1, \beta = 0, C = 1, D = 0.01, \Sigma = 0.05, R_x = 2, R_{xu} = 0, R_u = 2, G = 2$. The truncation range in the BCOS approximations is set to $a = -5, b = 5$.

The convergence results are depicted in figure 5.5. In particular, as can be seen in fig. 5.5a, each hereby considered forward discretization exhibits its theoretically expected strong convergence rate. In case of the Euler scheme, all processes converge with a rate of $\mathcal{O}(h^{1/2})$, whereas for the Milstein and 2.0 weak Taylor approximations, this rate is improved to $\mathcal{O}(h)$. Even though the strong convergence rate with the latter two schemes is first-order in both cases, the 2.0 weak Taylor approximations admit an advantageous constant, resulting in an order of magnitude higher overall strong approximation accuracy. The weak convergence errors are collected in fig. 5.5b, from which we can draw similar conclusions as for our earlier examples. All forward discretizations preserve their theoretical weak convergence rates to the fully-coupled FBSDE setting of (5.41). In fact, the Euler and Milstein approximations agree to errors converging with a rate of $\mathcal{O}(h)$ at t_0, to which the 2.0 weak Taylor approximation brings a significant improvement by speeding up convergence at t_0 to second-order in exchange for a marginally higher total CPU time – runtimes are comparable to table 5.2 and are thus omitted.

As equation (5.41) is derived from a stochastic optimal control problem, see [74] and the references therein, the approximation accuracy for the solution of the forward SDE

is of special importance. As we can see, the methods proposed in the present chapter using Milstein and 2.0 weak Taylor approximations for the numerical resolution of the forward diffusion do not only result in a higher order strong convergence rate, but also significantly improve the approximation accuracy for a given N in the controlled forward diffusion. According to fig. 5.5a, this results in 2 orders of magnitude gain in strong approximation accuracy; and almost 4 orders of magnitude improvement in the approximation quality at t_0, when $N = 10^3$ compared to the Euler method deployed in [76].

Finally, in order to assess the influence of the truncation in the Fourier cosine expansions, we collect strong and weak approximations errors in tables 5.3 and 5.4, respectively, for different values of K and N. Comparing tables 5.3a with 5.3b, we find that both the Euler and Milstein discretizations are robust with respect to the number of Fourier coefficients K. In particular, from $K = 2^7$ the main source of strong approximation errors is the time discretization term. In line with fig. 5.5a, the Milstein approximation proposed in this chapter yields 2 orders of magnitude improvement to the total strong approximation error compared to the Euler scheme of [76], independently of the choice of K.

Similar conclusions can be drawn from the comparison of tables 5.4a and 5.4b, which collect the weak approximation errors of the Milstein and 2.0 weak Taylor schemes for the fully-coupled equation (5.41). In case of the Milstein scheme, the BCOS method is not sensitive to the choice of K and the errors at t_0 are dominated by the time discretization. On the other hand, as can be seen from tab. 5.4b, in case of the 2.0 weak Taylor discretization, due to the higher order convergence, the BCOS method more quickly reaches an error level, where the Fourier truncation becomes prominent. In particular, for the errors in Z_0, second-order convergence is not fully reached with $K = 2^7$ Fourier terms only, and one can gain an additional order of accuracy by choosing a $K = 2^{10}$ large expansion instead.

Summarizing the implications of tables 5.3 and 5.4, the BCOS method in algorithm 5 is robust with respect to the number of Fourier terms K in the cosine expansions, even in the case of fully-coupled FBSDEs. Strong approximations quickly converge to the time discretization errors in K. In terms of weak approximation errors at t_0 using the 2.0 weak Taylor scheme in the forward SDE, one may need to enlarge the truncated Fourier series in order to preserve second order convergence, as the corresponding time discretization error decays with a faster rate of $\mathcal{O}(h^2)$.

5.5. CONCLUSION

In this chapter, we extended the BCOS method of [144, 145, 76] to second-order Taylor schemes approximating the Markov transition between two time steps in the fully-coupled FBSDE setting. We presented an algorithmic framework that unifies second-order Taylor schemes for fully-coupled equations, including the Euler-, Milstein- and simplified order 2.0 weak Taylor approximations for the forward SDE; and the generalized theta-scheme of [164] for the BSDE. Building on the closed-form expression for the characteristic function of the corresponding Markov transitions in lemma 5.3.1, we extended the coupled BCOS method and gave an implementable, higher-order numerical method for the fully-coupled FBSDEs in algorithm 5. We demonstrated the robustness and accuracy of our algorithm on a wide range of equations, spanning from the decou-

	strong error X			strong error Y			strong error Z		
$N \setminus K$	128	512	1024	128	512	1024	128	512	1024
10	7.7e−3	7.7e−3	7.7e−3	1.5e−2	1.5e−2	1.5e−2	3.3e−3	3.3e−3	3.3e−3
100	2.4e−3	2.4e−3	2.4e−3	4.7e−3	4.7e−3	4.7e−3	8.3e−4	8.3e−4	8.3e−4
400	1.2e−3	1.2e−3	1.2e−3	2.4e−3	2.4e−3	2.4e−3	4.1e−4	4.1e−4	4.1e−4
1000	7.7e−4	7.7e−4	7.7e−4	1.5e−3	1.5e−3	1.5e−3	2.7e−4	2.7e−4	2.7e−4

(a) Euler with (5.7)

	strong error X			strong error Y			strong error Z		
$N \setminus K$	128	512	1024	128	512	1024	128	512	1024
10	1.7e−3	1.7e−3	1.7e−3	3.5e−3	3.5e−3	3.5e−3	1.9e−3	1.9e−3	1.9e−3
100	1.7e−4	1.7e−4	1.7e−4	3.3e−4	3.3e−4	3.3e−4	1.9e−4	1.9e−4	1.9e−4
400	4.3e−5	4.3e−5	4.3e−5	8.7e−5	8.7e−5	8.7e−5	4.9e−5	4.9e−5	4.9e−5
1000	1.7e−5	1.7e−5	1.7e−5	3.4e−5	3.4e−5	3.4e−5	2.0e−5	2.0e−5	2.0e−5

(b) Milstein with (5.8)

Table 5.3: Example 3 in (5.41). Strong approximation errors with various forward discretizations in (5.6), and for different values of K and N. ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $\theta_4 = -1/2$, $K = 2^{10}$.)

pled to the fully-coupled case, and found that the hereby proposed second-order Taylor methods bring an improved first-order strong convergence to the total approximation error of the FBSDE compared to the Euler scheme in [76], if the theta parameters of the backward discretization are chosen accordingly. Additionally, we found that the 2.0 weak Taylor discretization further improves the convergence of the approximation errors at t_0 to second-order, which is crucial in applications such as stochastic optimal control. On top of the improved accuracy, the methods also proved to be robust with respect to the number of terms in the finitely truncated Fourier expansion.

	error Y_0			error Z_0		
$N \setminus K$	128	512	1024	128	512	1024
10	8.8e−4	8.8e−4	8.8e−4	3.4e−3	3.4e−3	3.4e−3
100	8.6e−5	8.6e−5	8.6e−5	3.2e−4	3.2e−4	3.2e−4
400	2.2e−5	2.2e−5	2.2e−5	7.4e−5	7.4e−5	7.4e−5
1000	8.6e−6	8.6e−6	8.6e−6	2.4e−5	2.4e−5	2.4e−5

(a) Milstein with (5.8)

	error Y_0			error Z_0		
$N \setminus K$	128	512	1024	128	512	1024
10	1.5e−5	1.5e−5	2.1e−5	7.1e−5	6.9e−5	2.1e−4
100	8.4e−8	8.4e−8	8.4e−8	1.0e−6	7.2e−7	7.1e−7
400	5.2e−8	5.2e−8	5.2e−8	3.2e−7	6.0e−8	5.1e−8
1000	2.2e−8	2.2e−8	2.2e−8	2.8e−7	2.0e−8	1.1e−8

(b) 2.0 weak Taylor with (5.9)

Table 5.4: Example 3 in (5.41). Weak approximation errors at t_0, with various forward discretizations in (5.6), and for different values of K and N. ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $\theta_4 = -1/2$, $K = 2^{10}$.)

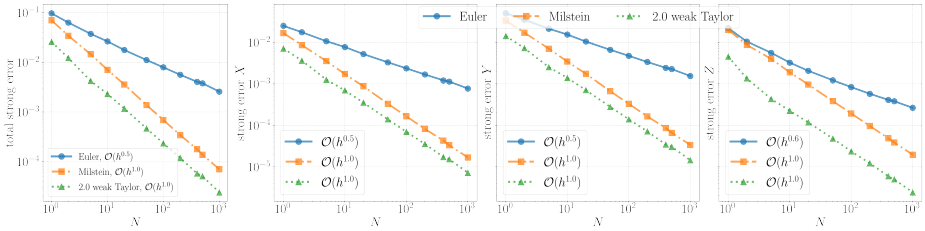

(a) strong convergence

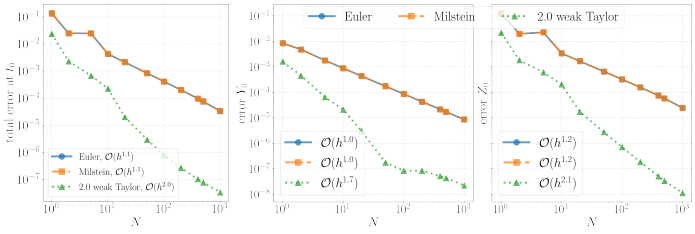

(b) weak convergence at t_0

Figure 5.5: Example 3 in (5.41). ($\theta_1 = \theta_2 = \theta_3 = 1/2$, $\theta_4 = -1/2$, $K = 2^{10}$.)

6

CONCLUSION AND OUTLOOK

I'm learning to fly but I ain't got wings
Coming down is the hardest thing

Tom Petty and the Heartbreakers (Learning to Fly)

Throughout this thesis, we investigated modern, machine learning methods for the numerical approximation of high-dimensional forward-backward stochastic differential equations. We examined the convergence of discrete time approximations in several different settings. We analyzed two different methods for the approximation of the associated conditional expectations, the deep BSDE method in high-dimensional settings, and the BCOS method in the classical, one-dimensional framework. In this concluding chapter, we summarize the main results of the thesis and point out several research directions with related interesting, open follow up questions.

6.1. MAIN RESULTS

In the first two chapters, we derived a new discretization for Malliavin differentiable BSDEs, analyzed and proved its convergence, and studied two fully-implementable approaches including a thorough regression error analysis. The main application of the One Step Malliavin scheme is the numerical resolution of delta-gamma hedging in (1.13) of high-dimensional financial options, which we carried out in chapter 3. The last two chapters were concerned with the fully-coupled FBSDE setting in (1.3)-(1.2). First, in chapter 4, we established a new, a posteriori convergence result for the approximation errors of the deep BSDE method, involving coupling in the control process. Finally, in chapter 5, we presented a higher-order numerical method for the scalar-valued case, capitalizing on differentiable approximations in a Fourier cosine expansion scheme. In what follows, we summarize the main results established along the way.

6.1.1. ONE STEP MALLIAVIN SCHEME AND ITS APPLICATIONS

Chapter 2 laid down the foundations of the One Step Malliavin scheme in the Markovian framework. Building on the representation formulae established by the theorems in section 2.3 we derived the discrete time approximations of the OSM scheme. The main ingredients were the Malliavin differentiability of the solution triplet of the FBSDE system, and the linear SDE and BSDE their Malliavin derivatives satisfy. Given a merged formulation of the Malliavin chain rule in lemma and the Feynman-Kac formulae, we derived the discrete OSM approximations, after a suitable discretization of the continuous time integrals. The crucial difference between the discretization therein, compared to, for instance, the standard Euler discretization of BSDEs, is the presence of a Γ process. This, on top of discretizing the control process through the linear BSDE governing the paths of the Malliavin derivatives, guarantees a stochastic representation of the second derivatives of the solution of the associated parabolic problem, which is in fact the key input for the delta-gamma hedging applications carried out in chapter 3.

Subsequently, we carried out a complete, discrete time approximation error analysis in section 2.4. Under standard Lipschitz assumptions, we showed that the time discretization errors induced by the OSM scheme admit an asymptotic convergence rate of order $1/2$ in natural L^2 norms. The main challenge in the proof was controlling the L^2 regularity of the Malliavin derivative of the control process and establishing sufficient estimates for the discretization error of the Γ process, that could later be combined with a discrete Grönwall lemma. The most important differences between the error estimates of the OSM scheme, and that of the convergence of the standard backward Euler

scheme of BSDEs are as follows. First, due to discretizing the control process through the Malliavin representation formula, the errors in the Z process are controlled by the maximum error over discrete time points. Furthermore, by solving an additional, linear, vector-valued BSDE of the Malliavin derivatives, the OSM scheme solves a larger problem, including second derivative estimates by means of the Γ process. Nonetheless, the One Step Malliavin scheme exhibits the same, optimal convergence rate as that of the standard Euler scheme.

The discrete time convergence analysis was followed by a thorough regression study in section 2.5, where we investigated two fully-implementable approaches approximating the backward recursion of conditional expectations in the abstract time discretization. First, we derived the corresponding *BCOS* approximations, computing the associated conditional expectations by means of Fourier cosine expansions, using the known conditional characteristic functions of the Markov transitions. This approach was particularly useful for the empirical study of regression errors stemming from the deep BSDE approximations of higher-dimensional equations – we refer to figure 2.2 in particular. Most importantly, we derived backward deep BSDE approximations of the OSM scheme in section 2.5. Argued by the martingale representation theorem, we derived two separate loss functions, depending on whether the Γ process is parameterized or computed directly by means of automatic differentiation. Sufficient minimizers of these loss functions directly approximate the conditional expectations corresponding to Z and Γ estimates. Subsequently, we derived a similar loss function for the Y part of the equation. We established the consistency of the deep BSDE approximations up to a universal approximation type argument, for both the parametrized and automatic differentiated cases of Γ. In particular, assuming perfect minimizers of the associated loss functions, we managed to control the total approximation errors by an upper bound depending on the discretization errors and the sum of regression biases. Therefore, given an appropriate UAT result, the bounds can be made arbitrarily small, assuming a converging stochastic gradient descrent optimization. The main challenge in the proof was a sufficient sandwiching of the regression errors in the presence of the Γ process.

In the final section of chapter 2, we presented several numerical experiments, spanning a wide range of high-dimensional semi-linear equations, demonstrating the robustness and accuracy of the deep BSDE approximations of the One Step Malliavin scheme up to $d = 50$ spatial dimensions. In particular, we found that the OSM approximations bring a significant improvement in the errors of the Z process for heavily control dependent equations, which are particularly relevant in stochastic optimal control.

DELTA-GAMMA HEDGING

The main motivation behind the development of the One Step Malliavin scheme was the numerical resolution of delta-gamma hedging of high-dimensional options. Capitalizing on the Γs involved in the OSM scheme, in chapter 3 we carried out this application in depth. In order to be able to treat financial options with early-exercise features, we first extended the One Step Malliavin scheme to discretely reflected BSDEs. The main challenge is the sufficient numerical treatment of the (disretized) reflection process, ensuring that the Y process stays above Markovian lower boundary over the discrete set of reflection points. Thereafter, in order to be able to hedge on the portfolio level, we generalized the OSM approximations to the vector-valued setting, allowing for the simulta-

neous treatment of systems of discretely reflected BSDEs corresponding to multiple options. For an efficient numerical resolution of this system, we proposed a vector-valued extension of the deep BSDE method for the One Step Malliavin scheme in algorithm 2. This approach involves not only option prices and Deltas, but also second-order Greeks throughout the whole time horizon. Additionally, we extended the method to stochastic volatility models.

We performed extensive numerical experiments highlighting several key features of the proposed methodology. Our findings suggest that the OSM scheme provides highly accurate Greeks up to second order, resulting in high quality discrete time replication, both in the case of delta- and delta-gamma hedging. In particular, we found that the deep BSDE approximations of the One Step Malliavin scheme surpass reference methods [77, 32], even in the context of delta hedging, when risk factors exhibit high volatility. Most importantly, we showed that the Gamma approximations induced by the OSM scheme are of high accuracy, and are applicable in the context of discretely rebalanced delta-gamma hedging for high-dimensional Bermudan options. The delta-gamma replication accuracy induced by OSM approximations significantly improved on that of sole delta hedging, even in the presence of $d = 100$ risk factors. The thorough numerical experiments indicate that our methodology is robust and accurate for different levels of moneyness, high volatility, varying early exercise rights up to the continuously reflected American option limit. Hedging on the portfolio level proved to be successful up to 25 high-dimensional basket contracts issued on $d = 100$ risk factors.

6.1.2. A POSTERIORI CONVERGENCE OF THE DEEP BSDE METHOD AND ITS APPLICATIONS

For the last two chapters of this thesis, we investigated numerical methods for the solution of fully-coupled FBSDE systems, targeting applications in stochastic control. Our first main contribution was given in chapter 4, where we generalized the *a posteriori* convergence result of the forward deep BSDE method of [69] to general drift terms in the forward diffusion, admitting Z as an argument.

After a suitable reformulation of a discretized coupled FBSDE system into a stochastic target problem, we investigated the forward deep BSDE method. Under suitable assumptions, we proved the a posteriori bound following a sequence of auxiliary lemmas, bounding the difference between two solutions to the time discretization, without imposing the terminal condition on the FBSDE. First, we gave an *a priori* estimate, controlling the difference in X and Y. Compared to the related convergence result of Han and Long [69], our conditions involved an extra constant due to the extra Z coupling. We handled this extra term in a subsequent lemma, by proving further a priori estimates. Due to the Z coupling in the drift, this induces three additional constants. The key step herein was to restrict the free constant λ_3 stemming from the application of the ϵ-Young inequality, to values $\lambda_3 > 2mL_z^f$, related to the strength of Z coupling. Building on this sequence of lemmas, our main theorem was stated and proved in theorem 4.3.3. The main condition in this theorem is related to the two coefficients $\bar{B}, \bar{A}$, and establishes that whenever these coefficients are *contractive* at the same time, then the forward deep

BSDE method, admits an a posteriori bound, which can informally be written as follows

$$\text{total approximation error} \leq \text{time discretization} \quad + \quad \text{loss function}.$$

Compared to earlier results of Han and Long [69], our extra condition related to $\bar{B}$, which stems from the coupling of the Z process.

We gave detailed interpretations for our weak coupling conditions in section 4.4. We showed that the conditions hold under many relevant applications, such as

- decoupled FBSDEs;
- short time horizons T;
- weak coupling;
- strong monotonicity.

In particular, we proved that our theory includes the setting of [69], and that in the absence of Z coupling, we recover their convergence result as a limit case of our theorem.

In the final section of chapter 4, we provided several numerical experiments supporting our theory. Among others, our theory explains the divergence of the deep BSDE method previously observed in the stochastic control literature in [4]. We showed that the reason behind this is related to the strength of Z coupling, resulting in the violation of our conditions. We demonstrated that whenever the equation is rescaled in a way that this condition is satisfied, the deep BSDE method converges for all other coefficients being equal. Ultimately, as an application on high-dimensional stochastic control problems, we compared solutions of the same optimal control problem, where the associated coupled FBSDE system was first formulated by the dynamic programming and then by the stochastic maximum principle. We found that for drift control problems, deriving an associated FBSDE system by means of the stochastic maximum principle leads to an equation more tractable in deep BSDE frameworks, due to the lack of Z coupling, which renders the extra condition of our theory automatically satisfied. Therefore, the convergence of the deep BSDE method can be ensured by milder conditions imposed only on the coupling in Y. These observations are in line with the findings of [74].

6.1.3. HIGHER-ORDER BCOS METHOD

In the final part of this thesis, chapter 5, we took a step away from high-dimensional equations, and developed a new numerical method for one-dimensional, fully-coupled FBSDEs. Distinct from most classical references, we did not only discretize the coupled forward diffusion by means of Euler-Maruyama approximations, but gave a unifying framework for second-order Taylor schemes, also including Milstein- and simplified order 2.0 weak Taylor discretizations. The latter two schemes imply an improved strong convergence rate of order 1, and in case of the last scheme, improved weak convergence of order 2. In order to preserve the higher order convergence rates from the approximation of the backward equation, we discretized the BSDE by the second order schemes of Zhao et al. [164] and Crisan and Manolarakis [35], given in (5.11) and (5.12), respectively. In proposition 5.3.1, we proved that the latter scheme can be expressed in terms of the generalized theta-scheme of Zhao et al. [164].

For the approximations in space, we employed the COS method, built on finite Fourier cosine expansions, and known conditional characteristic functions. To this end, for any fixed decoupling field (φ,ζ), we proved that the characteristic function of the Markov transition associated with second-order Taylor schemes admits a closed form expression. Subsequently, in order to approximate all conditional expectations in the discretization of the BSDE, we proved a recursive formula, establishing analytical expressions for conditional expectations where the argument is multiplied by powers of the Brownian increment. Finally, to address the derivatives of the decoupling fields (φ,ζ) appearing in the second-order Taylor approximations, we capitalized on the differentiability of COS approximations, in an explicit backward recursion. In particular, similar to the explicit scheme in Huijskens et al. [76], setting the decoupling relations $\varphi = y(t_{n+1},\cdot)$ and $\zeta = z(t_{n+1},\cdot)$ at time step t_n, one recovers analytical expressions for the corresponding spatial derivatives. This makes the computation of all coefficients in (5.6) possible. The resulting algorithm was given in algorithm 5.

Numerical experiments supported the claims on higher order convergence for a range of FBSDEs, from the decoupled to the fully-coupled settings. These results suggest that both the Milstein and simplified order 2.0 weak Taylor discretizations, approximated by the COS method, achieve a strong convergence with rate of order 1, for all levels of coupling. Additionally, in case of the latter scheme, an asymptotic weak convergence rate of order 2 can be achieved when using the generalized theta-scheme with certain parameters. This implies improved approximations at $t = 0$, in case the initial condition of the forward diffusion is deterministic, which has accentuated relevance in stochastic optimal control applications.

6.2. OUTLOOK

There are several promising directions in which the results in this thesis can be extended. As a closure to our work, we list a few of them, pointing out the main foreseeable challenges, where possible.

Convergence analysis of OSM discretization for general diffusions. In chapter 2, we proved that under certain assumptions, the discrete time approximations of the OSM scheme converge with an L^2 order of $\mathcal{O}(|\pi|^{1/2})$, as the mesh size of the corresponding time partition vanishes. A key assumption this result was based on, is the condition of additive noise in the forward diffusion. In particular, this assumption enabled us to neglect product terms from the approximation error of the merged formulation of the Malliavin chain rule and the Feynman-Kac formulae. A promising area of further research could be the relaxation of this condition, and proving convergence of the OSM discretization for more general diffusion coefficients. In doing so, the main challenge is to sufficiently control the approximation errors from the merged formulation of the Malliavin chain rule and the Feynman-Kac formula, inducing terms as in (2.41). We refer back to the discussion on this subject in section 2.4.2. Note that a necessary sub-result was shown in section 2.A, where the convergence of the Euler-Maruyama approximations for the Malliavin derivative of the forward diffusion is derived under mild conditions, capitalizing on approximations only being required between adjacent time points.

Finally, we remark that empirical results indicate that the additive noise condition can indeed be substantially relaxed – we refer to example 3 in section 2.6.

OSM for fractional diffusions and BSDEs with jumps. In chapter 2, we derived the One Step Malliavin scheme for Markovian FBSDE systems, i.e. where solutions to both the forward diffusion and the BSDE are Markov processes. A generalization of these results to non Markovian diffusion frameworks would allow the treatment of many financially relevant problems. For instance, one could consider rough volatility models, where the volatility process is driven by a *fractional Brownian motion*, see e.g. [18, 55]. Such models exhibiting memory can approximated by a sequence of Markovian approximations, and in this regard, the results in this thesis could form the basis for fractional generalizations. Due to their efficiency in handling high-dimensional problems, deep BSDE methods seem a suitable candidate for such numerical approximations – we refer to [81] for a recent study in this direction.

Similarly, another significant extension of the One Step Malliavin scheme would be the treatment of jump diffusion frameworks. Indeed, the linear BSDE governing the paths of the Malliavin derivatives of the solution pair of the backward equation, remains to be true in the general, jump diffusion settings, we refer to [57] for a recent result in this direction. Therein, the authors show that a similar Malliavin representation formula holds, as long as the equation is driven by a Lévy process. A suitable discretization of this BSDE would allow for delta-gamma hedging in the presence of jumps.

Convergence of the deep BSDE method for fully-coupled FBSDEs. In chapter 4, we proved the a posteriori convergence of the forward deep BSDE method for general drift coefficients in the forward diffusion, also admitting Z as an argument. A natural open question is what conditions would Z coupling in the diffusion have to satisfy, so that this result can be generalized to the fully-coupled case? The two main obstacles are as follows. First, one would need to ensure that the corresponding discrete time approximations admit a similar convergence result as the one employed in chapter 4. Thereafter, one would have to generalize the estimates in the relevant lemmas to allow for fully-coupled diffusions, and check how the resulting estimates would change the contractivity condition. In particular, it remains to be an open question, whether a similar result could be established without putting further restrictions on the length of the time horizon T.

Convergence analysis of BCOS. In chapter 5, we developed a BCOS method for fully-coupled FBSDEs, and showed empirically that this method admits higher order convergence rates, both in the strong- and weak senses. A rigorous, theoretical convergence proof of these observations remains an open problem, even in the case of Euler-Maruyama discretizations. The main intricacy in proving the convergence of the BCOS method is the interplay between the COS approximations and the time discretizations, in the presence of coupling. We believe that an error analysis similar to [40] could be established with some new ideas handling the COS errors terms. In particular, the truncation of both the Fourier cosine series to a finite number of terms K, and the integration range to a finite interval $[a, b]$ induce non-trivial error terms that are intertwined with the time

discretization. In case of the former, the coefficient recovery, done by discrete cosine transforms, requires tight bounds on the size of the coefficients, such that discrete time error remain meaningful in a backward recursion. For the latter, one would need to ensure that the relevant regions of Markov transitions are covered with high-probability, in between each pair of time steps. A clever combination of density estimates for FBSDEs, see e.g. [95], with recent results on the optimal integration range for the COS method [88] may provide a satisfactory answer to the these open questions.

Higher-order backward deep BSDE method using a Milstein discretization. In chapter 5, we heavily relied on the assumption that the forward diffusion in the coupled FBSDE system is one-dimensional. As is well-known in the numerical SDE literature [90], one needs to impose additional assumptions on the coefficients, in order to preserve a strong convergence rate of order 1 for the Milstein scheme of high-dimensional diffusion models. A promising, multi-dimensional extension in the spirit of chapter 5 would be to extend the backward deep BSDE method of Huré et al. [77] to forward discretizations by the Milstein scheme. In case of commutative noise, the discretized forward approximations are expected to exhibit a strong convergence rate of order 1, and it remains to be seen, how one could formulate a suitable loss function for a second-order discretization of the BSDE, in a similar way to the generalized theta-scheme of [164]. Addressing these gaps would result in a higher-order Monte Carlo method, for a special class of decoupled FBSDEs.

Comparison of deep BSDE methods with Physics Informed Neural Networks. Throughout chapters 2, 3 and 4, we considered several variants of deep BSDE methods, with the main motivation of tackling high-dimensional problems in option pricing, hedging, and stochastic optimal control. One attribute, all methods have in common is that they consider a stochastic reformulation of an associated nonlinear PDE such as (1.8), by means of a BSDE, and solve this backward SDE in a suitable machine learning formulation. There is a disjoint class of competitive numerical methods in the recent literature of scientific computing, which directly tackle the solution to a pricing PDE or HJB equation. This class of methods is often referred to as Physics Informed Neural Networks (PiNN), see e.g. [149, 142, 140]. To the best of our understanding, there is no satisfactory, in depth comparison between deep BSDE methods and PiNNs in the context of financial mathematics. Therefore, it remains to be an open question, which of these methods are more accurate in the context and applications of this thesis.

BIBLIOGRAPHY

[1] O. Aboura and S. Bourguin. "Density Estimates for Solutions to One Dimensional Backward SDE's". In: *Potential Analysis* 38.2 (2013), pp. 573–587. DOI: 10.1007/s11118-012-9287-8.

[2] S. Alanko and M. Avellaneda. "Reducing variance in the numerical solution of BSDEs". In: *Comptes Rendus. Mathématique* 351.3 (2013), pp. 135–138. DOI: 10.1016/j.crma.2013.02.010.

[3] C. Alasseur, Z. Bensaid, R. Dumitrescu, and X. Warin. *Deep learning algorithms for FBSDEs with jumps: Applications to option pricing and a MFG model for smart grids*. 2024. DOI: 10.48550/arXiv.2401.03245.

[4] K. Andersson, A. Andersson, and C. W. Oosterlee. "Convergence of a Robust Deep FBSDE Method for Stochastic Control". In: *SIAM Journal on Scientific Computing* 45.1 (2023), A226–A255. DOI: 10.1137/22M1478057.

[5] K. Andersson, A. Gnoatto, M. Patacca, and A. Picarelli. *A deep solver for BSDEs with jumps*. 2024. DOI: 10.48550/arXiv.2211.04349.

[6] F. Antonelli. "Backward-Forward Stochastic Differential Equations". In: *The Annals of Applied Probability* 3.3 (1993), pp. 777–793. DOI: 10.1214/aoap/1177005363.

[7] F. Antonelli and S. Hamadène. "Existence of the solutions of backward–forward SDE's with continuous monotone coefficients". In: *Statistics & Probability Letters* 76.14 (2006), pp. 1559–1569. DOI: 10.1016/j.spl.2006.03.018.

[8] J. L. Ba, J. R. Kiros, and G. E. Hinton. *Layer Normalization*. 2016. DOI: 10.48550/arXiv.1607.06450.

[9] V. Bally and G. Pagès. "A quantization algorithm for solving multidimensional discrete-time optimal stopping problems". In: *Bernoulli* 9.6 (2003), pp. 1003–1049. DOI: 10.3150/bj/1072215199.

[10] C. Beck, S. Becker, P. Cheridito, A. Jentzen, and A. Neufeld. "Deep Splitting Method for Parabolic PDEs". In: *SIAM Journal on Scientific Computing* 43.5 (2021), A3135–A3154. DOI: 10.1137/19M1297919.

[11] C. Beck, W. E, and A. Jentzen. "Machine Learning Approximation Algorithms for High-Dimensional Fully Nonlinear Partial Differential Equations and Second-order Backward Stochastic Differential Equations". In: *Journal of Nonlinear Science* 29.4 (2019), pp. 1563–1619. DOI: 10.1007/s00332-018-9525-3.

[12] S. Becker, P. Cheridito, and A. Jentzen. "Pricing and Hedging American-Style Options with Deep Learning". In: *Journal of Risk and Financial Management* 13.7 (2020), p. 158. DOI: 10.3390/jrfm13070158.

[13] C. Bender and R. Denk. “A forward scheme for backward SDEs”. In: *Stochastic Processes and their Applications* 117.12 (2007), pp. 1793–1812. DOI: 10.1016/j.spa.2007.03.005.

[14] C. Bender and N. Schweizer. ““Regression anytime” with brute-force SVD truncation”. In: *The Annals of Applied Probability* 31.3 (2021), pp. 1140–1179. DOI: 10.1214/20-AAP1615.

[15] C. Bender and J. Steiner. “A Posteriori Estimates for Backward SDEs”. In: *SIAM/ASA Journal on Uncertainty Quantification* 1.1 (2013), pp. 139–163. DOI: 10.1137/120878689.

[16] C. Bender and J. Steiner. “Least-Squares Monte Carlo for Backward SDEs”. In: *Numerical Methods in Finance*. Ed. by R. A. Carmona, P. Del Moral, P. Hu, and N. Oudjane. Berlin, Heidelberg: Springer, 2012, pp. 257–289. DOI: 10.1007/978-3-642-25746-9_8.

[17] C. Bender and J. Zhang. “Time discretization and Markovian iteration for coupled FBSDEs”. In: *The Annals of Applied Probability* 18.1 (2008), pp. 143–177. DOI: 10.1214/07-AAP448.

[18] L. Bergomi. *Stochastic Volatility Modeling*. New York: Chapman and Hall/CRC, 2015. 522 pp. DOI: 10.1201/b19649.

[19] J.-M. Bismut. “An Introductory Approach to Duality in Optimal Stochastic Control”. In: *SIAM Review* 20.1 (1978), pp. 62–78. DOI: 10.1137/1020004.

[20] J.-M. Bismut. “Intégrales convexes et probabilités”. In: *Journal of Mathematical Analysis and Applications* 42.3 (1973), pp. 639–673. DOI: 10.1016/0022-247X(73)90170-4.

[21] J. F. Bonnans and H. Zidani. “Consistency of Generalized Finite Difference Schemes for the Stochastic HJB Equation”. In: *SIAM Journal on Numerical Analysis* 41.3 (2003), pp. 1008–1021. DOI: 10.1137/S0036142901387336.

[22] B. C. Boonstra and C. W. Oosterlee. “Valuation of electricity storage contracts using the COS method”. In: *Applied Mathematics and Computation* 410 (2021), p. 126416. DOI: 10.1016/j.amc.2021.126416.

[23] B. Bouchard and J.-F. Chassagneux. “Discrete-time approximation for continuously and discretely reflected BSDEs”. In: *Stochastic Processes and their Applications* 118.12 (2008), pp. 2269–2293. DOI: 10.1016/j.spa.2007.12.007.

[24] B. Bouchard and N. Touzi. “Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations”. In: *Stochastic Processes and their Applications* 111.2 (2004), pp. 175–206. DOI: 10.1016/j.spa.2004.01.001.

[25] P. Briand and C. Labart. “Simulation of BSDEs by Wiener chaos expansion”. In: *The Annals of Applied Probability* 24.3 (2014), pp. 1129–1171. DOI: 10.1214/13-AAP943.

[26] G. Broux-Quemerais, S. Kaakaï, A. Matoussi, and W. Sabbagh. “Deep learning scheme for forward utilities using ergodic BSDEs”. In: *Probability, Uncertainty and Quantitative Risk* 9.2 (2024), pp. 149–180. DOI: 10.3934/puqr.2024009.

[27] J.-F. Chassagneux. "Linear Multistep Schemes for BSDEs". In: *SIAM Journal on Numerical Analysis* 52.6 (2014), pp. 2815–2836. DOI: 10.1137/120902951.

[28] J.-F. Chassagneux, J. Chen, and N. Frikha. *Deep Runge-Kutta schemes for BSDEs*. 2022. DOI: 10.48550/arXiv.2212.14372.

[29] J.-F. Chassagneux, J. Chen, N. Frikha, and C. Zhou. "A learning scheme by sparse grids and Picard approximations for semilinear parabolic PDEs". In: *IMA Journal of Numerical Analysis* 43.5 (2023), pp. 3109–3168. DOI: 10.1093/imanum/drac066.

[30] J.-F. Chassagneux and D. Crisan. "Runge–Kutta schemes for backward stochastic differential equations". In: *The Annals of Applied Probability* 24.2 (2014), pp. 679–720. DOI: 10.1214/13-AAP933.

[31] J.-F. Chassagneux and A. Richou. "Numerical simulation of quadratic BSDEs". In: *The Annals of Applied Probability* 26.1 (2016), pp. 262–304. DOI: 10.1214/14-AAP1090.

[32] Y. Chen and J. W. L. Wan. "Deep neural network framework based on backward stochastic differential equations for pricing and hedging American options in high dimensions". In: *Quantitative Finance* 21.1 (2021), pp. 45–67. DOI: 10.1080/14697688.2020.1788219.

[33] P. Cheridito, H. M. Soner, N. Touzi, and N. Victoir. "Second-order backward stochastic differential equations and fully nonlinear parabolic PDEs". In: *Communications on Pure and Applied Mathematics* 60.7 (2007), pp. 1081–1110. DOI: 10.1002/cpa.20168.

[34] J. Chessari, R. Kawai, Y. Shinozaki, and T. Yamada. "Numerical methods for backward stochastic differential equations: A survey". In: *Probability Surveys* 20 (none 2023), pp. 486–567. DOI: 10.1214/23-PS18.

[35] D. Crisan and K. Manolarakis. "Second order discretization of backward SDEs and simulation with the cubature method". In: *The Annals of Applied Probability* 24.2 (2014), pp. 652–678. DOI: 10.1214/13-AAP932.

[36] J. Cvitanic and J. Zhang. "The Steepest Descent Method for Forward-Backward SDEs". In: *Electronic Journal of Probability* 10 (none 2005), pp. 1468–1495. DOI: 10.1214/EJP.v10-295.

[37] G. Cybenko. "Approximation by superpositions of a sigmoidal function". In: *Mathematics of Control, Signals and Systems* 2.4 (1989), pp. 303–314. DOI: 10.1007/BF02551274.

[38] F. Delarue and G. Guatteri. "Weak existence and uniqueness for forward–backward SDEs". In: *Stochastic Processes and their Applications* 116.12 (2006), pp. 1712–1742. DOI: 10.1016/j.spa.2006.05.002.

[39] F. Delarue. "On the existence and uniqueness of solutions to FBSDEs in a non-degenerate case". In: *Stochastic Processes and their Applications* 99.2 (2002), pp. 209–286. DOI: 10.1016/S0304-4149(02)00085-6.

[40] F. Delarue and S. Menozzi. “A forward–backward stochastic algorithm for quasi-linear PDEs”. In: *The Annals of Applied Probability* 16.1 (2006), pp. 140–184. DOI: 10.1214/105051605000000674.

[41] J. Douglas, J. Ma, and P. Protter. “Numerical methods for forward-backward stochastic differential equations”. In: *The Annals of Applied Probability* 6.3 (1996), pp. 940–968. DOI: 10.1214/aoap/1034968235.

[42] R. M. Dudley. “Wiener Functionals as Ito Integrals”. In: *The Annals of Probability* 5.1 (1977), pp. 140–141. DOI: 10.1214/aop/1176995898.

[43] W. E, J. Han, and A. Jentzen. “Deep Learning-Based Numerical Methods for High-Dimensional Parabolic Partial Differential Equations and Backward Stochastic Differential Equations”. In: *Communications in Mathematics and Statistics* 5.4 (2017), pp. 349–380. DOI: 10.1007/s40304-017-0117-6.

[44] N. El Karoui, C. Kapoudjian, E. Pardoux, S. Peng, and M. C. Quenez. “Reflected solutions of backward SDE’s, and related obstacle problems for PDE’s”. In: *The Annals of Probability* 25.2 (1997), pp. 702–737. DOI: 10.1214/aop/1024404416.

[45] N. El Karoui, E. Pardoux, and M. Quenez. “Reflected Backward SDEs and American Options”. In: *Numerical Methods in Finance*. Ed. by D. Talay and L. C. G. Rogers. Publications of the Newton Institute. Cambridge: Cambridge University Press, 1997, pp. 215–231. DOI: 10.1017/CBO9781139173056.012.

[46] N. El Karoui, S. Peng, and M. C. Quenez. “Backward Stochastic Differential Equations in Finance”. In: *Mathematical Finance* 7.1 (1997), pp. 1–71. DOI: 10.1111/1467-9965.00022.

[47] A. Fahim, N. Touzi, and X. Warin. “A probabilistic numerical method for fully nonlinear parabolic PDEs”. In: *The Annals of Applied Probability* 21.4 (2011), pp. 1322–1364. DOI: 10.1214/10-AAP723.

[48] F. Fang and C. W. Oosterlee. “A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions”. In: *SIAM Journal on Scientific Computing* 31.2 (2009), pp. 826–848. DOI: 10.1137/080718061.

[49] F. Fang and C. W. Oosterlee. “Pricing early-exercise and discrete barrier options by fourier-cosine series expansions”. In: *Numerische Mathematik* 114.1 (2009), pp. 27–62. DOI: 10.1007/s00211-009-0252-4.

[50] F. Fang and C. W. Oosterlee. “A Fourier-Based Valuation Method for Bermudan and Barrier Options under Heston’s Model”. In: *SIAM Journal on Financial Mathematics* 2.1 (2011), pp. 439–463. DOI: 10.1137/100794158.

[51] Q. Feng, M. Luo, and Z. Zhang. “Deep signature FBSDE algorithm”. In: *Numerical Algebra, Control and Optimization* 13.3 (2023), pp. 500–522. DOI: 10.3934/naco.2022028.

[52] A. Fromm. “Theory and applications of decoupling fields for forward-backward stochastic differential equations”. PhD thesis. 2015. DOI: 10.13140/RG.2.1.5056.9129.

[53] M. Fujii, A. Takahashi, and M. Takahashi. "Asymptotic Expansion as Prior Knowledge in Deep Learning Method for High dimensional BSDEs". In: *Asia-Pacific Financial Markets* 26.3 (2019), pp. 391–408. DOI: 10.1007/s10690-019-09271-7.

[54] C. Gao, S. Gao, R. Hu, and Z. Zhu. "Convergence of the Backward Deep BSDE Method with Applications to Optimal Stopping Problems". In: *SIAM Journal on Financial Mathematics* 14.4 (2023), pp. 1290–1303. DOI: 10.1137/22M1539952.

[55] J. Gatheral, T. Jaisson, and M. Rosenbaum. *Volatility Is Rough*. Rochester, NY, 2017. DOI: 10.2139/ssrn.2509457.

[56] P. Gauthier and D. Possamaï. *Efficient Simulation of the Wishart Model*. Rochester, NY, 2009. DOI: 10.2139/ssrn.1474728.

[57] C. Geiss and A. Steinicke. "Existence, uniqueness and Malliavin differentiability of Lévy-driven BSDEs with locally Lipschitz driver". In: *Stochastics* 92.3 (2020), pp. 418–453. DOI: 10.1080/17442508.2019.1626859.

[58] M. Germain, H. Pham, and X. Warin. "Approximation Error Analysis of Some Deep Backward Schemes for Nonlinear PDEs". In: *SIAM Journal on Scientific Computing* 44.1 (2022), A28–A56. DOI: 10.1137/20M1355355.

[59] X. Glorot and Y. Bengio. "Understanding the difficulty of training deep feedforward neural networks". In: *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.

[60] A. Gnoatto, S. Lavagnini, and A. Picarelli. "Deep Quadratic Hedging". In: *Mathematics of Operations Research* (2024). DOI: 10.1287/moor.2023.0213.

[61] A. Gnoatto, A. Picarelli, and C. Reisinger. "Deep xVA Solver: A Neural Network–Based Counterparty Credit Risk Management Framework". In: *SIAM Journal on Financial Mathematics* 14.1 (2023), pp. 314–352. DOI: 10.1137/21M1457606.

[62] E. Gobet and P. Turkedjiev. "Adaptive importance sampling in least-squares Monte Carlo algorithms for backward stochastic differential equations". In: *Stochastic Processes and their Applications* 127.4 (2017), pp. 1171–1203. DOI: 10.1016/j.spa.2016.07.011.

[63] E. Gobet, J.-P. Lemor, and X. Warin. "A regression-based Monte Carlo method to solve backward stochastic differential equations". In: *The Annals of Applied Probability* 15.3 (2005), pp. 2172–2202. DOI: 10.1214/105051605000000412.

[64] E. Gobet and A. Makhlouf. "The Tracking Error Rate of the Delta-Gamma Hedging Strategy". In: *Mathematical Finance* 22.2 (2012), pp. 277–309. DOI: 10.1111/j.1467-9965.2010.00466.x.

[65] I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016. 801 pp.

[66] C. Gourieroux. "Continuous Time Wishart Process for Stochastic Risk". In: *Econometric Reviews* 25.2 (2006), pp. 177–217. DOI: 10.1080/07474930600713234.

[67] J. Han, R. Hu, and J. Long. “Learning High-Dimensional McKean–Vlasov Forward-Backward Stochastic Differential Equations with General Distribution Dependence”. In: *SIAM Journal on Numerical Analysis* 62.1 (2024), pp. 1–24. DOI: 10.1137/22M151861X.

[68] J. Han, A. Jentzen, and W. E. “Solving high-dimensional partial differential equations using deep learning”. In: *Proceedings of the National Academy of Sciences* 115.34 (2018), pp. 8505–8510. DOI: 10.1073/pnas.1718942115.

[69] J. Han and J. Long. “Convergence of the deep BSDE method for coupled FBSDEs”. In: *Probability, Uncertainty and Quantitative Risk* 5.1 (2020), p. 5. DOI: 10.1186/s41546-020-00047-w.

[70] D. J. Higham and X. Mao. “Convergence of Monte Carlo simulations involving the mean-reverting square root process”. In: *Journal of Computational Finance* (2005). DOI: 10.21314/JCF.2005.136.

[71] K. Hornik. “Approximation capabilities of multilayer feedforward networks”. In: *Neural Networks* 4.2 (1991), pp. 251–257. DOI: 10.1016/0893-6080(91)90009-T.

[72] K. Hornik, M. Stinchcombe, and H. White. “Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks”. In: *Neural Networks* 3.5 (1990), pp. 551–560. DOI: 10.1016/0893-6080(90)90005-6.

[73] Y. Hu, D. Nualart, and X. Song. “Malliavin calculus for backward stochastic differential equations and application to numerical solutions”. In: *The Annals of Applied Probability* 21.6 (2011), pp. 2379–2423. DOI: 10.1214/11-AAP762.

[74] Z. Huang, B. Negyesi, and C. W. Oosterlee. “Convergence of the deep BSDE method for stochastic control problems formulated through the stochastic maximum principle”. In: *Mathematics and Computers in Simulation* 227 (2025), pp. 553–568. DOI: 10.1016/j.matcom.2024.08.002.

[75] Z. Huang and C. W. Oosterlee. *Convergence of the Markovian iteration for coupled FBSDEs via a differentiation approach*. 2025. DOI: 10.48550/arXiv.2504.02814.

[76] T. P. Huijskens, M. J. Ruijter, and C. W. Oosterlee. “Efficient numerical Fourier methods for coupled forward–backward SDEs”. In: *Journal of Computational and Applied Mathematics* 296 (2016), pp. 593–612. DOI: 10.1016/j.cam.2015.10.019.

[77] C. Huré, H. Pham, and X. Warin. “Deep backward schemes for high-dimensional nonlinear PDEs”. In: *Mathematics of Computation* 89.324 (2020), pp. 1547–1579. DOI: 10.1090/mcom/3514.

[78] P. Imkeller and G. Dos Reis. “Path regularity and explicit convergence rate for BSDE with truncated quadratic growth”. In: *Stochastic Processes and their Applications* 120.3 (2010), pp. 348–379. DOI: 10.1016/j.spa.2009.11.004.

[79] P. Imkeller, R. L. Pellat, and O. Menoukeu-Pamen. “Differentiability of quadratic forward-backward SDEs with rough drift”. In: *The Annals of Applied Probability* 34.5 (2024), pp. 4758–4798. DOI: 10.1214/24-AAP2079.

[80] J. S. Jackson. “Some results on the well-posedness of quadratic BSDE systems”. PhD thesis. 2023. DOI: https://doi.org/10.26153/tsw/49900.

[81] A. Jacquier and M. Oumgari. “Deep Curve-Dependent PDEs for Affine Rough Volatility”. In: *SIAM Journal on Financial Mathematics* 14.2 (2023), pp. 353–382. DOI: 10.1137/19M1267805.

[82] M. Jensen and I. Smears. “On the Convergence of Finite Element Methods for Hamilton–Jacobi–Bellman Equations”. In: *SIAM Journal on Numerical Analysis* 51.1 (2013), pp. 137–162. DOI: 10.1137/110856198.

[83] A. Jentzen, B. Kuckuck, A. Neufeld, and P. von Wurstemberger. “Strong error analysis for stochastic gradient descent optimization algorithms”. In: *IMA Journal of Numerical Analysis* 41.1 (2021), pp. 455–492. DOI: 10.1093/imanum/drz055.

[84] S. Ji, S. Peng, Y. Peng, and X. Zhang. *A deep learning method for solving stochastic optimal control problems driven by fully-coupled FBSDEs*. 2024. DOI: 10.48550/arXiv.2204.05796.

[85] S. Ji, S. Peng, Y. Peng, and X. Zhang. “Solving Stochastic Optimal Control Problem via Stochastic Maximum Principle with Deep Learning Method”. In: *Journal of Scientific Computing* 93.1 (2022), p. 30. DOI: 10.1007/s10915-022-01979-5.

[86] S. Ji, S. Peng, Y. Peng, and X. Zhang. “Three Algorithms for Solving High-Dimensional Fully Coupled FBSDEs Through Deep Learning”. In: *IEEE Intelligent Systems* 35.3 (2020), pp. 71–84. DOI: 10.1109/MIS.2020.2971597.

[87] Y. Jiang and J. Li. “Convergence of the Deep BSDE method for FBSDEs with non-Lipschitz coefficients”. In: *Probability, Uncertainty and Quantitative Risk* 6.4 (2021), pp. 391–408. DOI: 10.3934/puqr.2021019.

[88] G. Junike and K. Pankrashkin. “Precise option pricing by the COS method—How to choose the truncation range”. In: *Applied Mathematics and Computation* 421 (2022), p. 126935. DOI: 10.1016/j.amc.2022.126935.

[89] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. Vol. 113. Graduate Texts in Mathematics. New York, NY: Springer, 1998. DOI: 10.1007/978-1-4612-0949-2.

[90] P. E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Berlin, Heidelberg: Springer, 1992. DOI: 10.1007/978-3-662-12616-5.

[91] S. Kremsner, A. Steinicke, and M. Szölgyenyi. “A Deep Neural Network Algorithm for Semilinear Elliptic PDEs with Applications in Insurance Mathematics”. In: *Risks* 8.4 (2020), p. 136. DOI: 10.3390/risks8040136.

[92] W. La Cruz, J. Martínez, and M. Raydan. “Spectral residual method without gradient information for solving large-scale nonlinear systems of equations”. In: *Mathematics of Computation* 75.255 (2006), pp. 1429–1448. DOI: 10.1090/S0025-5718-06-01840-0.

[93] J. Liang, Z. Xu, and P. Li. "Deep learning-based least squares forward-backward stochastic differential equation solver for high-dimensional derivative pricing". In: *Quantitative Finance* 21.8 (2021), pp. 1309–1323. DOI: 10.1080/14697688.2021.1881149.

[94] S. Liang, S. W. Jiang, J. Harlim, and H. Yang. "Solving PDEs on unknown manifolds with machine learning". In: *Applied and Computational Harmonic Analysis* 71 (2024), p. 101652. DOI: 10.1016/j.acha.2024.101652.

[95] R. Likibi Pellat and O. Menoukeu Pamen. "Density analysis for coupled forward–backward SDEs with non-Lipschitz drifts and applications". In: *Stochastic Processes and their Applications* 173 (2024), p. 104359. DOI: 10.1016/j.spa.2024.104359.

[96] R. Likibi Pellat, O. Menoukeu Pamen, and Y. Ouknine. "A class of quadratic forward-backward stochastic differential equations". In: *Journal of Mathematical Analysis and Applications* 514.2 (2022), p. 126100. DOI: 10.1016/j.jmaa.2022.126100.

[97] C. Liu, A. Papapantoleon, and A. Saplaouras. *Convergence rates for Backward SDEs driven by Lévy processes*. 2024. DOI: 10.48550/arXiv.2402.01337.

[98] F. A. Longstaff and E. S. Schwartz. "Valuing American Options by Simulation: A Simple Least-Squares Approach". In: *The Review of Financial Studies* 14.1 (2001), pp. 113–147. DOI: 10.1093/rfs/14.1.113.

[99] R. Lord, R. Koekkoek, and D. V. Dijk. "A comparison of biased simulation schemes for stochastic volatility models". In: *Quantitative Finance* 10.2 (2010), pp. 177–194. DOI: 10.1080/14697680802392496.

[100] Q. Lü and X. Zhang. "Well-posedness of backward stochastic differential equations with general filtration". In: *Journal of Differential Equations* 254.8 (2013), pp. 3200–3227. DOI: 10.1016/j.jde.2013.01.010.

[101] G. Ma and S.-P. Zhu. "Revisiting the Merton Problem: from HARA to CARA Utility". In: *Computational Economics* 59.2 (2022), pp. 651–686. DOI: 10.1007/s10614-021-10102-z.

[102] J. Ma, P. Protter, and J. Yong. "Solving forward-backward stochastic differential equations explicitly — a four step scheme". In: *Probability Theory and Related Fields* 98.3 (1994), pp. 339–359. DOI: 10.1007/BF01192258.

[103] J. Ma, Z. Wu, D. Zhang, and J. Zhang. "On well-posedness of forward–backward SDEs—A unified approach". In: *The Annals of Applied Probability* 25.4 (2015), pp. 2168–2214. DOI: 10.1214/14-AAP1046.

[104] J. Ma and J. Yong. *Forward-Backward Stochastic Differential Equations and their Applications*. Red. by J.-M. Morel, F. Takens, and B. Teissier. Vol. 1702. Lecture Notes in Mathematics. Berlin, Heidelberg: Springer, 2007. DOI: 10.1007/978-3-540-48831-6.

[105] J. Ma and J. Zhang. "Representation theorems for backward stochastic differential equations". In: *The Annals of Applied Probability* 12.4 (2002), pp. 1390–1418. DOI: 10.1214/aoap/1037125868.

[106] J. Ma and J. Zhang. “Representations and regularities for solutions to BSDEs with reflections”. In: *Stochastic Processes and their Applications* 115.4 (2005), pp. 539–569. DOI: 10.1016/j.spa.2004.05.010.

[107] W. Margrabe. “The Value of an Option to Exchange One Asset for Another”. In: *The Journal of Finance* 33.1 (1978), pp. 177–186. DOI: 10.1111/j.1540-6261.1978.tb03397.x.

[108] T. Mastrolia. “Density analysis of non-Markovian BSDEs and applications to biology and finance”. In: *Stochastic Processes and their Applications* 128.3 (2018), pp. 897–938. DOI: 10.1016/j.spa.2017.06.009.

[109] T. Mastrolia, D. Possamaï, and A. Réveillac. “Density analysis of BSDEs”. In: *The Annals of Probability* 44.4 (2016), pp. 2817–2857. DOI: 10.1214/15-AOP1035.

[110] T. Mastrolia, D. Possamaï, and A. Réveillac. “On the Malliavin differentiability of BSDEs”. In: *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques* 53.1 (2017), pp. 464–492. DOI: 10.1214/15-AIHP723.

[111] R. C. Merton. “Optimum consumption and portfolio rules in a continuous-time model”. In: *Journal of Economic Theory* 3.4 (1971), pp. 373–413. DOI: 10.1016/0022-0531(71)90038-X.

[112] R. C. Merton. “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case”. In: *The Review of Economics and Statistics* 51.3 (1969), pp. 247–257. DOI: 10.2307/1926560.

[113] G. N. Milstein and M. V. Tretyakov. “Numerical Algorithms for Forward-Backward Stochastic Differential Equations”. In: *SIAM Journal on Scientific Computing* 28.2 (2006), pp. 561–582. DOI: 10.1137/040614426.

[114] B. Negyesi, K. Andersson, and C. W. Oosterlee. “The One Step Malliavin scheme: new discretization of BSDEs implemented with deep learning regressions”. In: *IMA Journal of Numerical Analysis* 44.6 (2024), pp. 3595–3647. DOI: 10.1093/imanum/drad092.

[115] B. Negyesi, Z. Huang, and C. W. Oosterlee. *Generalized convergence of the deep BSDE method: a step towards fully-coupled FBSDEs and applications in stochastic control*. 2025. DOI: 10.48550/arXiv.2403.18552.

[116] B. Negyesi and C. W. Oosterlee. *A deep BSDE approach for the simultaneous pricing and delta-gamma hedging of large portfolios consisting of high-dimensional multi-asset Bermudan options*. 2025. DOI: 10.48550/arXiv.2502.11706.

[117] B. Negyesi and C. W. Oosterlee. *A numerical Fourier cosine expansion method with higher order Taylor schemes for fully coupled FBSDEs*. 2025. DOI: 10.48550/arXiv.2501.10988.

[118] B. Négyesi. “A Novel Method for Solving High-Dimensional Backward Stochastic Differential Equations Using Malliavin Calculus and Deep Learning”. MSc thesis. 2020.

[119] D. Nualart. *The Malliavin Calculus and Related Topics*. Probability, its Applications. Berlin/Heidelberg: Springer-Verlag, 2006. DOI: 10.1007/3-540-28329-3.

[120] G. D. Nunno and P. D. Lozano. *Deep Operator BSDE: a Numerical Scheme to Approximate the Solution Operators*. 2024. DOI: 10.48550/arXiv.2412.03405.

[121] B. Øksendal. *Stochastic Differential Equations*. Universitext. Berlin, Heidelberg: Springer, 2003. DOI: 10.1007/978-3-642-14394-6.

[122] C. C. Paige and M. A. Saunders. "LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares". In: *ACM Trans. Math. Softw.* 8.1 (1982), pp. 43–71. DOI: 10.1145/355984.355989.

[123] A. Papapantoleon, D. Possamaï, and A. Saplaouras. "Existence and uniqueness results for BSDE with jumps: the whole nine yards". In: *Electronic Journal of Probability* 23 (none 2018), pp. 1–68. DOI: 10.1214/18-EJP240.

[124] A. Papapantoleon, D. Possamaï, and A. Saplaouras. "Stability of backward stochastic differential equations: the general Lipschitz case". In: *Electronic Journal of Probability* 28 (none 2023), pp. 1–56. DOI: 10.1214/23-EJP939.

[125] A. Papapantoleon, D. Possamaï, and A. Saplaouras. "Stability results for martingale representations: The general case". In: *Transactions of the American Mathematical Society* 372.8 (2019), pp. 5891–5946. DOI: 10.1090/tran/7880.

[126] E. Pardoux and S. Peng. "Backward stochastic differential equations and quasilinear parabolic partial differential equations". In: *Stochastic Partial Differential Equations and Their Applications*. Ed. by B. L. Rozovskii and R. B. Sowers. Berlin, Heidelberg: Springer, 1992, pp. 200–217. DOI: 10.1007/BFb0007334.

[127] E. Pardoux and S. G. Peng. "Adapted solution of a backward stochastic differential equation". In: *Systems & Control Letters* 14.1 (1990), pp. 55–61. DOI: 10.1016/0167-6911(90)90082-6.

[128] E. Pardoux and A. Rascanu. *Stochastic Differential Equations, Backward SDEs, Partial Differential Equations*. Vol. 69. Stochastic Modelling and Applied Probability. Cham: Springer International Publishing, 2014. DOI: 10.1007/978-3-319-05714-9.

[129] E. Pardoux and S. Tang. "Forward-backward stochastic differential equations and quasilinear parabolic PDEs". In: *Probability Theory and Related Fields* 114.2 (1999), pp. 123–150. DOI: 10.1007/s004409970001.

[130] S. Peng. "Open Problems on Backward Stochastic Differential Equations". In: *Control of Distributed Parameter and Stochastic Systems: Proceedings of the IFIP WG 7.2 International Conference, June 19–22, 1998 Hangzhou, China*. Ed. by S. Chen, X. Li, J. Yong, and X. Y. Zhou. Boston, MA: Springer US, 1999, pp. 265–273. DOI: 10.1007/978-0-387-35359-3_32.

[131] S. Peng and Z. Wu. "Fully Coupled Forward-Backward Stochastic Differential Equations and Applications to Optimal Control". In: *SIAM Journal on Control and Optimization* 37.3 (1999), pp. 825–843. DOI: 10.1137/S0363012996313549.

[132] H. Pham. *Continuous-time Stochastic Control and Optimization with Financial Applications*. Vol. 61. Stochastic Modelling and Applied Probability. Berlin, Heidelberg: Springer, 2009. DOI: 10.1007/978-3-540-89500-8.

[133] H. Pham, X. Warin, and M. Germain. “Neural networks-based backward scheme for fully nonlinear PDEs”. In: *SN Partial Differential Equations and Applications* 2.1 (2021), p. 16. DOI: 10.1007/s42985-020-00062-8.

[134] A. Pinkus. “Approximation theory of the MLP model in neural networks”. In: *Acta Numerica* 8 (1999), pp. 143–195. DOI: 10.1017/S0962492900002919.

[135] A. Popier and C. Zhou. “Second-order BSDE under monotonicity condition and liquidation problem under uncertainty”. In: *The Annals of Applied Probability* 29.3 (2019), pp. 1685–1739. DOI: 10.1214/18-AAP1435.

[136] D. Possamaï. “Second order backward stochastic differential equations under a monotonicity condition”. In: *Stochastic Processes and their Applications* 123.5 (2013), pp. 1521–1545. DOI: 10.1016/j.spa.2013.01.002.

[137] D. Possamaï and M. Rodrigues. “Reflections on BSDEs”. In: *Electronic Journal of Probability* 29 (2024), pp. 1–82. DOI: 10.1214/24-EJP1123.

[138] D. Possamaï and C. Zhou. “Second order backward stochastic differential equations with quadratic growth”. In: *Stochastic Processes and their Applications* 123.10 (2013), pp. 3770–3799. DOI: 10.1016/j.spa.2013.05.007.

[139] M. Pou, M. R. Ruijter, and C. W. Oosterlee. “Extension of a Fourier-Cosine Method to Solve BSDEs with Higher Dimensions”. In: *Progress in Industrial Mathematics at ECMI 2014*. Ed. by G. Russo, V. Capasso, G. Nicosia, and V. Romano. Cham: Springer International Publishing, 2016, pp. 75–101. DOI: 10.1007/978-3-319-23413-7_12.

[140] M. Raissi, P. Perdikaris, and G. E. Karniadakis. “Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations”. In: *Journal of Computational Physics* 378 (2019), pp. 686–707. DOI: 10.1016/j.jcp.2018.10.045.

[141] M. Raissi. “Forward-Backward Stochastic Neural Networks: Deep Learning of High-Dimensional Partial Differential Equations”. In: *Peter Carr Gedenkschrift*. WORLD SCIENTIFIC, 2023, pp. 637–655. DOI: 10.1142/9789811280306_0018.

[142] M. Raissi, P. Perdikaris, and G. E. Karniadakis. *Physics Informed Deep Learning (Part I): Data-driven Solutions of Nonlinear Partial Differential Equations*. 2017. DOI: 10.48550/arXiv.1711.10561.

[143] C. Reisinger, W. Stockinger, and Y. Zhang. “A posteriori error estimates for fully coupled McKean–Vlasov forward-backward SDEs”. In: *IMA Journal of Numerical Analysis* 44.4 (2024), pp. 2323–2369. DOI: 10.1093/imanum/drad060.

[144] M. J. Ruijter and C. W. Oosterlee. “A Fourier Cosine Method for an Efficient Computation of Solutions to BSDEs”. In: *SIAM Journal on Scientific Computing* 37.2 (2015), A859–A889. DOI: 10.1137/130913183.

[145] M. J. Ruijter and C. W. Oosterlee. “Numerical Fourier method and second-order Taylor scheme for backward SDEs in finance”. In: *Applied Numerical Mathematics* 103 (2016), pp. 1–26. DOI: 10.1016/j.apnum.2015.12.003.

[146] M. J. Ruijter and C. W. Oosterlee. "Two-Dimensional Fourier Cosine Series Expansion Method for Pricing Financial Options". In: *SIAM Journal on Scientific Computing* 34.5 (2012), B642–B671. DOI: 10.1137/120862053.

[147] M. Ruijter, C. Oosterlee, and R. Aalbers. "On the Fourier cosine series expansion method for stochastic control problems". In: *Numerical Linear Algebra with Applications* 20.4 (2013), pp. 598–625. DOI: 10.1002/nla.1866.

[148] R. L. Schilling, L. Partzsch, B. Böttcher, R. L. Schilling, and B. Böttcher. *Brownian Motion: An Introduction to Stochastic Processes*. Berlin/Boston, GERMANY: De Gruyter, Inc., 2012. DOI: 10.1515/9783110278989.

[149] J. Sirignano and K. Spiliopoulos. "DGM: A deep learning algorithm for solving partial differential equations". In: *Journal of Computational Physics* 375 (2018), pp. 1339–1364. DOI: 10.1016/j.jcp.2018.08.029.

[150] H. M. Soner, N. Touzi, and J. Zhang. "Wellposedness of second order backward SDEs". In: *Probability Theory and Related Fields* 153.1 (2012), pp. 149–190. DOI: 10.1007/s00440-011-0342-y.

[151] A. Takahashi, Y. Tsuchida, and T. Yamada. "A new efficient approximation scheme for solving high-dimensional semilinear PDEs: Control variate method for Deep BSDE solver". In: *Journal of Computational Physics* 454 (2022), p. 110956. DOI: 10.1016/j.jcp.2022.110956.

[152] L. Teng and W. Zhao. "High-order Combined Multi-step Scheme for Solving Forward Backward Stochastic Differential Equations". In: *Journal of Scientific Computing* 87.3 (2021), p. 81. DOI: 10.1007/s10915-021-01505-z.

[153] N. Touzi. *Optimal Stochastic Control, Stochastic Target Problems, and Backward SDE*. Vol. 29. Fields Institute Monographs. New York, NY: Springer, 2013. DOI: 10.1007/978-1-4614-4286-8.

[154] P. Turkedjiev. "Two algorithms for the discrete time approximation of Markovian backward stochastic differential equations under local conditions". In: *Electronic Journal of Probability* 20 (2015), pp. 1–49. DOI: 10.1214/EJP.v20-3022.

[155] H. Wang, H. Chen, A. Sudjianto, R. Liu, and Q. Shen. *Deep Learning-Based BSDE Solver for Libor Market Model with Application to Bermudan Swaption Pricing and Hedging*. 2018. DOI: 10.48550/arXiv.1807.06622.

[156] J. Yong and X. Y. Zhou. *Stochastic Controls*. New York, NY: Springer, 1999. DOI: 10.1007/978-1-4612-1466-3.

[157] B. Yu, X. Xing, and A. Sudjianto. *Deep-Learning Based Numerical BSDE Method for Barrier Options*. Rochester, NY, 2019. DOI: 10.2139/ssrn.3366314.

[158] Y. Yu, N. Ganesan, and B. Hientzsch. "Backward Deep BSDE Methods and Applications to Nonlinear Problems". In: *Risks* 11.3 (2023), p. 61. DOI: 10.3390/risks11030061.

[159] B. Zhang and C. W. Oosterlee. "Efficient Pricing of European-Style Asian Options under Exponential Lévy Processes Based on Fourier Cosine Expansions". In: *SIAM Journal on Financial Mathematics* 4.1 (2013), pp. 399–426. DOI: 10.1137/110853339.

[160] J. Zhang. “A numerical scheme for BSDEs”. In: *The Annals of Applied Probability* 14.1 (2004), pp. 459–488. DOI: 10.1214/aoap/1075828058.

[161] J. Zhang. *Backward Stochastic Differential Equations*. Vol. 86. Probability Theory and Stochastic Modelling. New York, NY: Springer, 2017. DOI: 10.1007/978-1-4939-7256-2.

[162] W. Zhao, L. Chen, and S. Peng. “A New Kind of Accurate Numerical Method for Backward Stochastic Differential Equations”. In: *SIAM Journal on Scientific Computing* 28.4 (2006), pp. 1563–1581. DOI: 10.1137/05063341X.

[163] W. Zhao, Y. Fu, and T. Zhou. “New Kinds of High-Order Multistep Schemes for Coupled Forward Backward Stochastic Differential Equations”. In: *SIAM Journal on Scientific Computing* 36.4 (2014), A1731–A1751. DOI: 10.1137/130941274.

[164] W. Zhao, Y. Li, and G. Zhang. “A generalized theta-scheme for solving backward stochastic differential equations”. In: *Discrete and Continuous Dynamical Systems - B* 17.5 (2012), pp. 1585–1603. DOI: 10.3934/dcdsb.2012.17.1585.

[165] W. Zhao, J. Wang, and S. Peng. “Error estimates of the theta-scheme for backward stochastic differentialequations”. In: *Discrete and Continuous Dynamical Systems - B* 12.4 (2009), pp. 905–924. DOI: 10.3934/dcdsb.2009.12.905.

[166] W. Zhao, G. Zhang, and L. Ju. “A Stable Multistep Scheme for Solving Backward Stochastic Differential Equations”. In: *SIAM Journal on Numerical Analysis* 48.4 (2010), pp. 1369–1394. DOI: 10.1137/09076979X.

[167] W. Zhao, W. Zhang, and L. Ju. “A Numerical Method and its Error Estimates for the Decoupled Forward-Backward Stochastic Differential Equations”. In: *Communications in Computational Physics* 15.3 (2014), pp. 618–646. DOI: 10.4208/cicp.280113.190813a.

CURRICULUM VITÆ

Bálint NÉGYESI

23-10-1993 Born in Budapest, Hungary.

EDUCATION AND RESEARCH EXPERIENCE

2012–2016 BSc in Economic and Financial Mathematical Analysis
Corvinus University of Budapest
Thesis: Equilibrium, stability and control of dynamical systems
Advisor: Prof. dr. Z. Kánnai

2013–2017 BSc in Physics
Eötvös Loránd University
Thesis: Analysis of one-dimensional Einstein's box
Advisor: Prof. dr. Zs. Papp & Prof. dr. L. Orosz

2018–2020 MSc in Applied Mathematics
Delft University of Technology
Thesis: A novel method for solving high-dimensional backward stochastic differential equations using Malliavin calculus and deep learning
Advisor: Prof. dr. ir. C. W. Oosterlee

2020-2025 PhD in Applied Mathematics
Delft University of Technology
Thesis: Modern machine learning methods for high-dimensional forward-backward SDEs with applications in mathematical finance
Promotors: Prof. dr. ir. C. W. Oosterlee & Prof. dr. ir. M. C. Veraar

AWARDS

2024 Travel award, 12th Bachelier World Congress
2023 Third prize of Best Poster Award, Woudschoten conference
2020 Peter Paul Peterich Scholarship

LIST OF PUBLICATIONS

PUBLISHED

[74] Z. Huang, B. Negyesi, and C. W. Oosterlee. "Convergence of the deep BSDE method for stochastic control problems formulated through the stochastic maximum principle". In: *Mathematics and Computers in Simulation* 227 (2025), pp. 553–568. DOI: 10.1016/j.matcom.2024.08.002

[114] B. Negyesi, K. Andersson, and C. W. Oosterlee. "The One Step Malliavin scheme: new discretization of BSDEs implemented with deep learning regressions". In: *IMA Journal of Numerical Analysis* 44.6 (2024), pp. 3595–3647. DOI: 10.1093/imanum/drad092

UNDER REVIEW

[116] B. Negyesi and C. W. Oosterlee. *A deep BSDE approach for the simultaneous pricing and delta-gamma hedging of large portfolios consisting of high-dimensional multi-asset Bermudan options*. 2025. DOI: 10.48550/arXiv.2502.11706

[117] B. Negyesi and C. W. Oosterlee. *A numerical Fourier cosine expansion method with higher order Taylor schemes for fully coupled FBSDEs*. 2025. DOI: 10.48550/arXiv.2501.10988

[115] B. Negyesi, Z. Huang, and C. W. Oosterlee. *Generalized convergence of the deep BSDE method: a step towards fully-coupled FBSDEs and applications in stochastic control*. 2025. DOI: 10.48550/arXiv.2403.18552

List of Academic Activities

Talks given at selected conferences

January 2025	22nd Winter school on Mathematical Finance, Soesterberg, the Netherlands
July 2024	12th Bachelier World Congress, Rio de Janeiro, Brazil
April 2024	International Conference on Computational Finance, Amsterdam, the Netherlands
June 2023	SIAM Conference on Financial Mathematics and Engineering, Philadelphia, USA
June 2022	9th Colloquium on BSDEs and Mean Field Systems, Annecy, France
June 2022	International Conference on Computational Finance, Wuppertal, Germany

Poster presentation

September 2023	B. Negyesi. *One Step Malliavin schemes: a BSDE approach for delta-gamma hedging.* Woudschoten conference (Dutch-Flemish Scientific Computing Society), Woudschoten, the Netherlands.[1]

Peer reviewing for Journals

Applied Mathematics & Optimization
International Journal of Computer Mathematics
Journal of Scientific Computing (Springer)

Supervision

2024	Lisanne van Wijk (MSc, UU): On Stochastic Control Theory for Dynamic Carbon Emission Reduction
2023	Jord van Eldik (BSc, UU): Lookback Option Pricing with the COS Method
2023	Stijn Vermeulen (MSc, UU): Solving Stochastic Optimal Control Problems Using Fully Coupled FBSDEs: and its applications to Pension Funds
2022	Martijn Brouwer (BSc, UU): Applying Physics-informed Neural Networks to Chaotic Nonlinear Systems of Ordinary Differential Equations

[1] Third prize, Best Poster Award.