Deep Learning for Constrained Utility Maximisation

Ashley Davey¹ · Harry Zheng¹

Received: 1 March 2021 / Revised: 27 August 2021 / Accepted: 25 October 2021 /
Published online: 26 November 2021
© The Author(s) 2021

Abstract

This paper proposes two algorithms for solving stochastic control problems with deep learning, with a focus on the utility maximisation problem. The first algorithm solves Markovian problems via the Hamilton Jacobi Bellman (HJB) equation. We solve this highly nonlinear partial differential equation (PDE) with a second order backward stochastic differential equation (2BSDE) formulation. The convex structure of the problem allows us to describe a dual problem that can either verify the original primal approach or bypass some of the complexity. The second algorithm utilises the full power of the duality method to solve non-Markovian problems, which are often beyond the scope of stochastic control solvers in the existing literature. We solve an adjoint BSDE that satisfies the dual optimality conditions. We apply these algorithms to problems with power, log and non-HARA utilities in the Black-Scholes, the Heston stochastic volatility, and path dependent volatility models. Numerical experiments show highly accurate results with low computational cost, supporting our proposed algorithms.

Keywords Stochastic control · Deep learning · Primal and dual BSDEs · HJB equation · Utility maximisation

Mathematics Subject Classification 93E20 · 91G80 · 90C46 · 49M29

1 Introduction

In this paper we propose algorithms that combine modern machine learning practises with theoretical stochastic control principles to solve a range of control problems with high accuracy, scalability with dimension, and low computational cost. There is a natural crossover between machine learning and stochastic control problems, as they both involve searching for features within a data set. For stochastic control the data set is a state process

Supported in part by the EPSRC (UK) Grant (EP/V008331/1)

✉ Harry Zheng
h.zheng@imperial.ac.uk

Ashley Davey
ashley.davey18@imperial.ac.uk

¹ Department of Mathematics, Imperial College, London SW7 2BZ, UK

that is influenced by a controlling process that can be chosen with the aim to optimise some objective function. Often the optimal control is a function of the underlying state process, a so-called Markovian control. This is very similar to deep learning, where the control is chosen as the output of a neural network which takes the dataset as input, and some loss function is minimised. It is in this area of crossover that this paper resides. We propose deep learning methods to solve the utility maximisation problem, as an application of a more general stochastic control solver, in a wide range of markets with arbitrary convex control constraints. Dynamic portfolio optimisation has been extensively studied within the field of mathematical finance, see Pham (2009); Karatzas and Shreve (1998) for exposition.

A typical drawback of numerical methods for stochastic control problems is that their complexity increases dramatically with increases in the state and control dimensions. Often some grid method or exhaustive search of the underlying state space is required to find optimal trajectories, which is computationally inefficient for high dimensional state spaces. It is claimed that machine learning methods provide a solution for this dimensionality problem Grohs et al. (2018). Numerical results show efficient solving of PDEs with the state space dimension in the thousands Beck et al. (2019a, 2019b); W. et al. (2017). These PDE solvers form the basis of our generic stochastic control problem solver, which is the first algorithm presented in this paper. The problems are transformed into forward-backward stochastic differential equations (FBSDEs) using the Feynman-Kac formula and the Stochastic Maximum Principle (SMP). The PDE that we solve is the Hamilton-Jacobi-Bellman (HJB) equation, a highly nonlinear PDE that is often impossible to solve analytically, or even numerically if the dimensionality is high. This paper introduces a control variable into the PDE, that is constrained to some convex set. This drastically increases the complexity and takes the problem out of the scope of W. et al. (2017); Beck et al. (2019b) unless the control can be found analytically. We show how machine learning can overcome this difficulty and produce highly accurate value approximations with comparatively low runtime.

The FBSDE based methods developed in this paper are direct extension of the first and second order deep BSDE solvers introduced in Beck et al. (2019b); W. et al. (2017). In these papers, a system of FBSDEs is solved by simulating all processes in the forward direction, driven by a neural network process which is optimised to ensure the terminal condition holds. The extension offered by this paper is an additional driving neural network process, representing a control process. The convenience of adapting this BSDE method to stochastic control problems is that the control process can be optimised using the value function approximations implied by the 2BSDE solver, via the Hamiltonian of the system. In the 2BSDE solver of Beck et al. (2019b), there are two unknown process that must be found with neural networks. This suggests that our method has three unknowns, but we can use the structure of the value problem to derive one of these processes and therefore introduce control without increasing the number of unknown processes.

There are some papers that describe generic algorithms for solving stochastic control problems in the literature. Han and W. (2016) directly maximises the gains function at the terminal time with machine learning and is applicable to many problems. However, the computation costs may be high, as one introduces a discretisation with many time steps or a high dimensional control space, and it only outputs the value at time 0. Huré et al. (2021) finds a local approximation of the value function with machine learning. However, it only holds for Markovian problems and the iterative nature of this value function approximation may lead to error propagation. Instead of directly optimising the value function, the first method presented in this paper maximises the Hamiltonian associated with the system. There are two key benefits of doing this. Firstly, it reconciles the potential errors

of Han and W. (2016); Huré et al. (2021) where the algorithm gets stuck at a local optimiser. To find the value function (given a control) we ensure that the terminal condition of the HJB equation holds, which involves minimising the least squares error of the associated processes at the terminal time. The global minimum of this error is zero by construction, ensuring that we can directly assess the quality of our approximation for any generic control problem. Secondly, the control is optimised using the Hamiltonian of the system, which is a function of only what is known at one time. This optimisation is faster than back propagating through a discretisation of the processes to compute the derivatives as in Han and W. (2016). This is even more important when the control space is high dimensional, which is normally true in portfolio optimisation problems. To implement this Hamiltonian maximisation procedure, one needs to know derivatives of the value function which can be computed accurately with deep learning based on the 2BSDE method in Beck et al. (2019b).

There are a number of papers in the literature that utilise the deep BSDE methods of W. et al. (2017); Beck and Jentzen (2019b). For example, Pereira et al. (2019); Wang et al. (2019); Exarchos and Theodorou (2018) first find an explicit formula for the optimal control in terms of the value function and its derivatives and then substitute this control into the HJB equation and solve the resulting PDE with the deep 2BSDE method. However, these algorithms cannot be applied in full generality like our method. Pereira et al. (2019); Exarchos and Theodorou (2018) are reliant on the diffusion of the state process being independent of the control, which leads to a semi-linear PDE and a representation with a first order BSDE. This framework is inappropriate for portfolio selection problems as the control naturally appears in the diffusion term. Wang et al. (2019) alleviates this restriction by deriving an explicit formula for the optimal control that is required to be unconstrained. Our approach is more general in the sense that it does not require any analytical solutions. In the literature these other algorithms are called “deep 2BSDE controllers” since the BSDE solutions are used to form the control. Our algorithm is called the deep controlled 2BSDE method as each control and state pair defines its own 2BSDE which we try to solve. A primal-dual method is introduced in Henry-Labordere (2017) for solving semi-linear PDEs. The deep 2BSDE method is used to find a lower bound for the solution. An upper bound is found by solving a maximisation control problem on the path space, that is derived using the Legendre transform. The approach in Henry-Labordere (2017) is not directly related to our notion of convex duality using martingale methods. There are also some papers on applications of deep neural networks in mathematical finance, including optimal stopping Becker et al. (2019), asset pricing with transaction cost Gonon et al. (2020), and optimal transport Eckstein and Kupper (2019).

The utility maximisation problem involves maximising concave functions, so we can naturally apply convex duality methods. Duality is a powerful tool that can be used in complete or incomplete markets with a range of convex constraints on the control process. The unique aspect of our work is the combination of deep learning methodology with convex duality. The benefits of introducing duality are threefold. Firstly, the dual problem is a minimisation problem, and solving it independently of the primal problem naturally leads to an upper bound of the value function. Sometimes the dual problem is easier to solve, since primal control constraints are characterised by penalties in the dual state process, leading the dual problem to be potentially unconstrained. Secondly, the BSDEs associated to the dual problem can be related to that of the primal problem, giving us the ability to compare our numerical results and check their accuracy. Again, as per the Hamiltonian maximisation, this benefits us since we know the error between the dual and primal processes should be zero, which allows us to easily measure the

quality of our results. We may even skip one of the problems if they are difficult to solve, as solving one problem automatically gives us the solution to the other. Thirdly, inherent in the dual formulation are certain optimality conditions that are derived from the SMP. These conditions can be used separately from the HJB equation to find the optimal control. We use these conditions as the basis for our second algorithm, called the deep SMP method, which is valid for non-Markovian control problems with convex structure. This greatly increases the range of utility maximisation problems that can be solved using the methods of this paper. In a Markovian setting we can compare the two algorithms presented in the paper. It should be noted here that results for dynamical programming principles, under which our 2BSDE scheme is formulated, exist in the literature for non-Markovian problems. Non-Markovian optimal stopping problems and their 2BSDE counterparts are studied in Karoui and Tan (2013); Possamai et al. (2018). The problem is reformulated into an optimisation in the space of probability measures, and the canonical probability space is expanded to allow a dynamical programming principle and a characterising path dependent PDE to be established, see Zhang (2017) for an exposition. However, it is not clear how a machine learning based solver could be used in this abstract setting.

The convex duality method works for complete and incomplete markets. A common incomplete market model is the Heston stochastic volatility model in which the underlying variance of the stock follows some mean-reversion SDE with its own source of randomness. This model may be seen as a 2-dimensional stock market, where only one asset is traded. This framework allows us to use our deep learning method to solve it. We compare our results to Ma et al. (2020) in which tight bounds of the value function are derived for an unconstrained optimisation problem. Ma et al. (2020) is reliant on a convenient guess of the dual control, so is not applicable to other methods where such a guess is not available, unlike our method. Another multidimensional model we consider is a path-dependent volatility model. There are several recent papers on the topic in mathematical finance with machine learning, for example Buehler et al. (2019); Liu et al. (2019); De Spiegeleer et al. (2018). These papers include model calibration, hedging and pricing, and are not directly related to our work.

The main contributions of this paper are the following. In the deep controlled 2BSDE algorithm, we present a Markovian stochastic control problem solver that combines the HJB equation with the deep 2BSDE method, but without the need for an analytical formula for the optimal control. We apply this algorithm to both the primal and dual problems which give us tight bounds on the value function. In the deep SMP method, we exploit the dual formulation of the utility maximisation problem to create a novel algorithm for non-Markovian optimisation problems using deep learning. Instead of dealing with the gains function, we solve the adjoint BSDE and use machine learning to satisfy the dual optimality conditions. To the best of our knowledge, we are the first to use this methodology in the literature.

The remainder of this paper is outlined as follows. Section 2 formulates the utility maximisation problem and defines the dual problem. Section 3 describes the general algorithm for solving Markovian stochastic control problems using deep learning and the HJB equation, and applies the algorithm to portfolio optimisation with various utilities. Section 4 introduces an algorithm for solving non-Markovian utility maximisation problems using the SMP and applies this algorithm to problems where the volatility process is not deterministic, and even path-dependent. Section 5 discusses convergence of the deep controlled 2BSDE algorithm based on empirical evidence, and compares to other algorithms. Section 6 concludes the paper.

2 The Utility Maximisation Problem

In this section we describe the standard portfolio optimisation problem and its dual problem. Let $W = (W(t))_{0 \leq t \leq T}$ be a standard m -dimensional Brownian motion on the natural filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$ augmented with all $\mathbb{P}$ -null sets, and T a fixed finite horizon. Consider an $m + 1$ dimensional stock market, consisting of a risk-free bond with interest rate $r(t)$, and a collection S of risky stocks defined by

$$dS(t) = \text{diag}(S(t))(\mu(t)dt + \sigma(t)dW(t)), \quad (1)$$

where $\text{diag}(S(t))$ is an $m \times m$ matrix with diagonal elements from $S(t)$ and all other elements zero, and r , μ and σ are uniformly bounded, progressively measurable processes valued in $\mathbb{R}$, $\mathbb{R}^m$ and $\mathbb{R}^{m \times m}$, respectively, with $\mu(t)$ and $\sigma(t)$ representing the drifts and volatilities of the m stock prices at time t , respectively. The matrix $\sigma(t)$ is invertible and σ^{-1} is uniformly bounded. An agent invests a proportion $\pi(t)$ of the wealth $X(t)$ in the risky stocks and the rest in the risk-free bond at time $t \in [0, T]$. The 1-dimensional wealth process X evolves as

$$dX(t) = X(t)(r(t) + \pi(t)^\top (\mu(t) - r(t)\mathbb{1}_m))dt + X(t)\pi(t)^\top \sigma(t)dW(t) \quad (2)$$

with $X(0) = x_0$, where $\pi(t)^\top$ is the transpose of $\pi(t)$, and $\mathbb{1}_m \in \mathbb{R}^m$ is a vector of ones. We call π an admissible control if π is a progressively measurable process valued in $K \subset \mathbb{R}^m$, a non-empty closed convex set, almost surely for almost every $t \in [0, T]$, and there exists a unique strong solution X to the SDE (2). The set of all admissible controls π is denoted by $\mathcal{A}$. The investor wishes to maximise the expected utility of the terminal wealth, that is, to find

$$V := \sup_{\pi \in \mathcal{A}} \mathbb{E}[U(X(T))],$$

where $U : [0, \infty) \rightarrow \mathbb{R}$ is a utility function that is twice continuously differentiable, strictly increasing, strictly concave and satisfies the Inada conditions $U'(0) = \infty$ and $U'(\infty) = 0$. We say that state and control processes $\hat{X}$, $\hat{\pi}$ are optimal if $V = \mathbb{E}[U(\hat{X}(T))]$.

We next construct the dual problem, see Li and Zheng (2018) [Section 2] for more details. The basis for duality is the Legendre-Fenchel transformation $\tilde{U} : (0, \infty) \rightarrow \mathbb{R}$ of the utility function U , defined by

$$\tilde{U}(y) = \sup_{x > 0} \{U(x) - xy\}, \quad y \geq 0.$$

The concavity and differentiability of U ensure that $\tilde{U}$ is well defined for all $y \geq 0$ and is twice continuously differentiable, strictly decreasing and strictly convex on $(0, \infty)$. The $\mathbb{R}$ -valued dual process Y is defined by

$$dY(t) = -Y(t)(r(t) + \delta_K(v(t)))dt - Y(t)(\theta(t) + \sigma(t)^{-1}v(t))^\top dW(t) \quad (3)$$

with $Y(0) = y$, where y is non-negative, $\theta(t) = \sigma(t)^{-1}(\mu(t) - r(t)\mathbb{1}_m)$ is the market price of risk, δ_K is the support function of the set $-K$, defined by

$$\delta_K(z) := \sup_{\pi \in K} \{-\pi^\top z\}, \quad z \in \mathbb{R}^m,$$

and v is a $\mathbb{R}^m$ -valued dual control process which is progressively measurable and ensures there exists a unique strong solution Y to the SDE (3).

Since XY is a super-martingale for any choice of π , v and y , we have the following weak duality relation [p. 5] Li and Zheng (2018)

$$V \leq \tilde{V}(x_0) := \inf_{y>0} \left\{ \inf_{v \in \mathbb{R}^m} \left\{ \mathbb{E}[\tilde{U}(Y(T)) \mid Y(0) = y] \right\} + x_0 y \right\}. \quad (4)$$

Solving the dual problem would give us an upper bound of the value function, that can be combined with a lower bound by the primal problem to produce a confidence interval for the value function. If we have an equality in (4), then we can find the value function by alternatively solving the dual problem.

The method of solving both the primal and dual problems depend on the nature of the coefficients in (2). If r , μ and σ are deterministic, then the problem is Markovian. This case is treated in Sect. 3, where we solve the HJB equation numerically using deep learning, and the control is optimised using the Hamiltonian. If the coefficients are not deterministic then the problem may not be Markovian. We treat this more general case in Sect. 4, where the adjoint BSDE is solved using deep learning, and the optimality conditions derived in Li and Zheng (2018) are used to optimise the control.

3 The Deep Controlled 2BSDE Method for the Markovian Case

The deep learning approaches for both the dual and primal Markovian problems are similar, so we present the algorithm in a more general form that covers both. We first describe the general control problem, then use the machine learning method to find the value function and optimal control. For a thorough exposition of the theory covered in this section, we direct the reader to Pham (2009) [Chapter 3]. We present the algorithm for a state dimension $d \geq 1$, as this is required in proceeding sections. For example, when we consider the Heston model with a stochastic volatility process denoted by $\sigma \in \mathbb{R}^m$, we can make the problem Markovian by introducing new $(1 + m)$ dimensional state processes (X, σ) and (Y, σ) .

Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0 \leq t \leq T}, \mathbb{P})$ be a filtered probability space with a fixed finite horizon T . Consider the $\mathcal{F}_t$ -adapted controlled diffusion process $(\mathcal{X}(t))_{0 \leq t \leq T}$ on $\mathbb{R}^d$, $d \in \mathbb{N}$, defined by

$$d\mathcal{X}(t) = b(t, \mathcal{X}(t), \alpha(t))dt + \Sigma(t, \mathcal{X}(t), \alpha(t))dW(t) \quad (5)$$

with $X(0) = x_0$, where W is an $\mathbb{R}^n$ -valued standard Brownian motion for some $n \in \mathbb{N}$ and $\alpha = (\alpha(t))_{0 \leq t \leq T}$ is a progressively measurable process, valued in $A \subset \mathbb{R}^m$ for $m \in \mathbb{N}$. It is assumed that the drift and diffusion functions $b : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}$ and $\Sigma : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}^{d \times n}$ are deterministic and measurable, such that (5) admits a strong solution. Here $\mathcal{X}$ represents the state process and α the control process. We consider only admissible controls that satisfy the integrability condition

$$\mathbb{E} \left[\int_0^T |b(t, x, \alpha(t))|^2 + |\Sigma(t, x, \alpha(t))|^2 dt \right] < \infty$$

where x is some element of the support of $\mathcal{X}$. This control is chosen to optimise the state process via maximising a reward that is defined using a terminal gain function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, which is a measurable function, either lower bounded or satisfying some quadratic growth condition. We wish to find the value function

$$u(t, x) := \sup_{a \in A} \mathbb{E} [g(\mathcal{X}(T)) \mid \mathcal{X}(t) = x]. \quad (6)$$

Remark 3.1 The utility function U is not defined on the whole space $\mathbb{R}$, but only $(0, \infty]$. We take $g(x)$ to be $U(x)$ when $x > 0$ and 0 otherwise, ensuring that quadratic growth or bounded conditions are satisfied (if they are satisfied for U). It is irrelevant what value is chosen when $x < 0$, as one can show that the state process X defined in (2) satisfies $X(T) > 0$ if and only if $X(0) > 0$ almost surely. We can similarly consider $\tilde{U}$ and Y .

Suppose that $u \in \mathcal{C}^{1,3}$. The HJB equation tells us that if u has at most quadratic growth in x then it solves the following PDE

$$\frac{\partial}{\partial t} u(t, x) = - \sup_{a \in A} F(t, x, a, D_x u(t, x), D_{xx} u(t, x)), \quad (t, x) \in [0, T) \times \mathbb{R}^d \quad (7)$$

with the terminal condition $u(T, x) = g(x)$ for $x \in \mathbb{R}^d$, where the Hamiltonian F is defined by

$$F(t, x, a, z, \gamma) = b(t, x, a)^\top z + \frac{1}{2} \text{tr}(\Sigma(t, x, a) \Sigma(t, x, a)^\top \gamma). \quad (8)$$

We can now describe the value function and its derivatives using a system of FBSDEs. This involves the generalised Hamiltonian $H : [0, T) \times \mathbb{R}^d \times A \times \mathbb{R}^d \times \mathbb{R}^{d \times n} \rightarrow \mathbb{R}$, defined by

$$H(t, x, a, z, q) = b(t, x, a)^\top z + \text{tr}(\Sigma(t, x, a)^\top q). \quad (9)$$

Theorem 3.2 Suppose that $u \in \mathcal{C}^{1,3}([0, T) \times \mathbb{R}^d) \cap \mathcal{C}^0([0, T] \times \mathbb{R}^d)$, and there exists an optimal control $\alpha \in A$ with associated controlled diffusion X , defined by (5). Then there exist continuous processes (V, Z, Γ) , valued in $\mathbb{R}, \mathbb{R}^d$, and $\mathbb{R}^{d \times d}$ respectively, solving the following 2BSDE:

$$\begin{aligned} dV(t) &= Z(t)^\top \Sigma(t, \mathcal{X}(t), \alpha(t)) dW(t) \\ dZ(t) &= -D_x H(t, \mathcal{X}(t), \alpha(t), Z(t), \Gamma(t) \Sigma(t, \mathcal{X}(t), \alpha(t))) dt + \Gamma(t) \Sigma(t, \mathcal{X}(t), \alpha(t)) dW(t) \end{aligned} \quad (10)$$

with the terminal conditions $V(T) = g(\mathcal{X}(T))$ and $Z(T) = D_x g(\mathcal{X}(T))$, where H is defined by (9) and the control α satisfies

$$F(t, \mathcal{X}(t), \alpha(t), Z(t), \Gamma(t)) = \sup_{a \in A} F(t, \mathcal{X}(t), a, Z(t), \Gamma(t)), \quad (11)$$

and F is defined by (8).

A proof is provided in the appendix. The derivative term in the dynamics of Z can be easily computed by the modern automatic differentiation method if it is unavailable analytically. This theorem is presented for a maximisation problem but easily translates to the dual problem, which is a minimisation problem, by replacing supremums with infimums.

3.1 The Deep Controlled 2BSDE Method

Now we have set up the new problem (10)–(11), we look to solve it with machine learning techniques in the spirit of Beck et al. (2019b) (the algorithm presented therein is named the “Deep 2BSDE” method, thus inspiring our adapted Deep Controlled 2BSDE method). The trick here is to simulate all processes in the forward direction, introducing new variables $v_0^0 \in \mathbb{R}$ and $z_0^0 \in \mathbb{R}^d$, then moving forward through a discretisation $(t_i)_{i=0}^N$ of $[0, T]$. We furthermore set Γ to be a neural network that at each time t_i takes the i th state position, denoted by X_i , as input since we are approximating $D_{xx}u(t, X(t))$ with this process. The control process α is also a neural network, taking in the same state X_i as input at time step t_i , which means that we are searching for a Markovian control.

Before describing the algorithms used to solve this problem in detail, we first look at what we mean by a neural network. We give a brief outline here, see for example Anthony and Bartlett (2009) for an exposition. Suppose we wish to approximate a function $\phi : \mathbb{R}^p \rightarrow \mathbb{R}^q$ for some $p, q \in \mathbb{N}$. We use the following algorithm to construct the approximating function. This neural network has a predetermined amount of ‘layers’ $L \in \mathbb{N}$ and ‘hidden nodes’ $\ell \in \mathbb{N}$. We call these the hyper-parameters of the network in the sequel.

Let $X \in \mathbb{R}^p$ be our input. For $i = 1, \dots, L$ define the linear function f_i by

$$f_i(x) = A_i x + b_i$$

We have $A_1 \in \mathbb{R}^{\ell \times p}$, $A_j \in \mathbb{R}^{\ell \times \ell}$ for $j = 2, \dots, L-1$ and $A_L \in \mathbb{R}^{q \times \ell}$ for the matrices. For the vectors we have $b_j \in \mathbb{R}^\ell$ for $j = 1, \dots, L-1$ and $b_L \in \mathbb{R}^q$. These form the unknowns of our function, and are what we need to choose to correctly approximate our function. We combine these into a set $\theta = (A_i, b_i)_{i=1}^L$ of parameters. Informally speaking, the number of hidden nodes ℓ corresponds to us ‘expanding’ the input into a higher dimensional space. One could think of each node as a piece of information we obtain from the input. The more nodes we have the more information we get, hence the better we can approximate the function. We wish to compose these linear functions in series, but we need a nonlinear function in between them. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be such a function that we apply element wise to the output of each linear function. The approximating function is therefore

$$N_\theta(X) := f_L \circ h \circ f_{L-1} \circ \dots \circ f_2 \circ h \circ f_1(X)$$

In our setting we choose $h(x) = \max(x, 0)$, called a rectified linear unit (ReLU) activation function. The choice of this function for this scheme is rather arbitrary, as there seems to be no clear preferred function here. A comparison of popular activation functions is provided in Sect. 5. This non-polynomial mapping is essential in ensuring the density of these networks within the space of continuous functions, as is shown in [Leshno et al. (1993), Theorem 1] which formalises the informal argument above about information into a density argument about the space of neural networks. Further more, if ϕ is Lipschitz then we can bound the approximation error in terms of the hyper-parameters of the network [Bach (2017), Proposition 6]. Note that no nonlinear function is applied at the final step, else we could only approximate non-negative functions! The parameter vector $\theta \in \mathbb{R}^\rho$, where $\rho = \rho(p, q)$ is defined below, is a parametrisation of the A_i and b_i matrices and is optimised at the training stage of our algorithm. We made no assumptions on the regularity of ϕ , but our approximation is almost everywhere differentiable. In practise this is sufficient for gradient descent algorithms, in which derivatives of this network with respect to its parameters must be taken. We use the Python package Tensorflow to implement this neural network procedure. This methodology reduces an optimisation over an infinite dimensional

space to a finite dimensional one. However, the dimension can still be high. For this neural network, the parameter space dimension is

$$\rho(p, q) = (L - 1)\ell^2 + \ell(L + p + q) + q, \quad (12)$$

a quadratic in the number of hidden nodes, and linear in the number of layers and the input dimension. Associated to each neural network N_θ is a loss function $\mathcal{L}(\theta)$ which is a function of the parameters forming the network. Finding the parameters that minimise this function equates to optimising the network itself. In our implementation, we use the ADAM algorithm Kingma and Ba (2014), which is an adaptation of stochastic gradient descent, implemented in `Tensorflow`. We furthermore use mini batches, which means the number of simulations we use to calculate the loss function is small, and batch normalisation, meaning that we scale our input to the neural network Ioffe and Szegedy (2015).

Remark 3.3 In this algorithm we propose a control α of the form $\alpha(t) = N_\theta(\mathcal{X}(t))$. One may then ask if this control is indeed admissible. It can be shown that a neural network is a Lipschitz function (though a global Lipschitz constant for every neural network in the training process may not be established), and therefore the required integrability properties of the state process and gains function of the control problem are satisfied.

Now we describe the algorithm to solve the control problem. We introduce parameters $\theta_i^0 \in \mathbb{R}^{\rho(d,m)}$ and $\lambda_i^0 \in \mathbb{R}^{\rho(d,d^2)}$ then use the following Euler-Maruyama Scheme. Set $\mathcal{X}_0 = x_0$, $V_0 = v_0^0$ and $Z_0 = z_0^0$ then for $i = 0, \dots, N - 1$ let $\alpha_i = N_{\theta_i^0}(\mathcal{X}_i)$, $\Gamma_i = N_{\lambda_i^0}(\mathcal{X}_i)$ and

$$\begin{aligned} \mathcal{X}_{i+1} &= \mathcal{X}_i + (t_{i+1} - t_i)b(t_i, \mathcal{X}_i, \alpha_i) + \Sigma(t_i, \mathcal{X}_i, \alpha_i)dW_i \\ V_{i+1} &= V_i + Z_i^\top \Sigma(t_i, \mathcal{X}_i, \alpha_i)dW_i \\ Z_{i+1} &= Z_i - (t_{i+1} - t_i)D_x H(t_i, \mathcal{X}_i, \alpha_i, Z_i, \Gamma_i \Sigma(t_i, \mathcal{X}_i, \alpha_i)) + \Gamma_i \Sigma(t_i, \mathcal{X}_i, \alpha_i)dW_i, \end{aligned} \quad (13)$$

where dW_i is a multivariate normal $\mathcal{N}_n(0, (t_{i+1} - t_i)\mathbb{I}_n)$ random variable. We have used minimal notation for these processes for clarity, but one should remember that these explicitly depend on the choice of parameters λ_i^0 , θ_i^0 and the initial points v_0^0 , z_0^0 . Our algorithm involves using this scheme multiple times, but with different parameters.

Remark 3.4 We use a different neural network at every time step, as opposed to a single network that takes in state and time as inputs. While this increases the number of parameters (assuming all networks use the same number of layers and hidden nodes), it also decreases the run time. This is because in the back propagation step, extra time is required to take derivatives of the network parameters, since they are involved in every time step. Even when using a high number of time steps, the increase in time outweighs the penalty of having more networks which leads to more memory usage. See Sect. 5.1 for a numerical example.

Once we have repeated this iteration $N - 1$ times, we arrive at time $t_N = T$ but realise that we have not satisfied the terminal conditions. We want to choose the neural networks and initial start points to reduce the expected error, which amounts to minimising the the loss function

$$\mathcal{L}_1(\Theta_{\text{BSDE}}^0) := \mathbb{E} \left[|V_N - g(\mathcal{X}_N)|^2 + \gamma |Z_N - D_x g(\mathcal{X}_N)|^2 \right], \quad (14)$$

where $\Theta_{\text{BSDE}}^0 \in \mathbb{R}^{1+d+\rho(d,d^2)}$ is a vector representation of $(v_0^0, z_0^0) \cup (\lambda_i^0)_{i=0}^{N-1}$. This is referred to as the BSDE loss function in the sequel. In practise, we evaluate the sample average of this loss function. The coefficient $\gamma > 0$ is some tuning parameter that can be chosen. In our case, we choose $\gamma = 0.5$.

In addition, for each time step the arbitrary control choice α is not optimal, so we must also maximise the loss functions

$$\mathcal{L}_2(\theta_i^0, i) := \mathbb{E}[F(t_i, \mathcal{X}_i, \alpha_i, Z_i, \Gamma_i)] \quad (15)$$

for $i = 0, \dots, N-1$. This is referred to as the control loss function in the sequel. Again, we evaluate a sample average of this loss function. We proceed to optimise these loss functions in turn, starting with our arbitrary 0-superscript parameters. Now suppose we have just finished step $j \in \mathbb{N}$.

The first sub-step is to improve the terminal condition loss. We generate V_N, Z_N and $\mathcal{X}_N$ using Θ_{BSDE}^j and θ_i^j for $i = 1, \dots, N-1$, which are carried forward from the previous step. We then make an improvement using one step of the ADAM algorithm, against the loss function $\mathcal{L}_1$, resulting in the updated parameters $\Theta_{\text{BSDE}}^{j+1}$. We also keep track of the moment estimates that are present in the algorithm as in Kingma and Ba (2014).

The second half of the iteration step is to improve the control process. We can first simulate the processes α and $\mathcal{X}$ using our previous parameter set $\theta^j = (\theta_i^j)_{i=0}^{N-1}$, then simulate Z and Γ using the updated parameters $\Theta_{\text{BSDE}}^{j+1}$, and finally make the following improvement using one step of the ADAM algorithm. For each $i = 0, \dots, N-1$ we update θ_i^j using the ADAM algorithm against $-\mathcal{L}_2(\cdot, i)$, resulting in θ_i^{j+1} . We use a negative here since ADAM is a minimisation algorithm, but we wish to maximise the Hamiltonian F at each time step. We track the moment estimates and move to step $j+2$ and repeat this process. We do this for each path of a batch of simulated paths of size $B \in \mathbb{N}$.

These two steps can be repeated until the new parameters are sufficiently close to the old parameters after an update, at which point we deem the algorithm to have converged. We find that making small increments at each step, ensuring that the learning rate required for the ADAM algorithm is smaller for the control step than the BSDE step, we have good convergence results.

Remark 3.5 At each time step we build 2 neural networks, which have 4 layers, with the two hidden layers consisting of $d+10$ nodes. The parameters are optimised using the ADAM algorithm, using a batch size of 64 Brownian paths. We choose to initialise each parameter near 0, though a prior guess may be used, and the rate of convergence is naturally much higher when the initial guess is closer to the true value. Substituting $p = d$, $L = 4$ and $l = d+10$ into (12) leads to a parameter set of size $\rho(d, q) = 4d^2 + 74d + 340 + q(d+11)$, where q is the output dimension of the network. The number of neural network parameters is therefore at most cubic in the state dimension and linear in the control dimension.

Remark 3.6 We use a different loss function for the BSDE and control networks in the algorithm. Another option would be combining these loss functions into one single loss function that both networks try to minimise. However, this would be slower as the back propagation step would take longer, particularly for the control network. It may also lead to false solutions, for example minimising $\mathcal{L}_1$ by making both V_N and $g(\mathcal{X}_N)$ close to zero. See Sect. 5.1 for a numerical example.

Algorithm 1 in the appendix outlines this scheme explicitly. When we use the notation $\text{ADAM}()$, we refer to one step of the ADAM algorithm outlined in [Kingma and Ba (2014), Algorithm 1]. The output is given as the value at time 0, but a byproduct of this scheme is pathwise evaluations of the value function and its derivatives, as well as the optimal control.

3.2 Extension of the Algorithm to the Dual Problem

The dual problem (4) has a similar structure, so we can adapt this method to solve it. We stay in a multidimensional setting, but this one dimensional parameter is just one component of the initial state. We define a variable y_0^0 with which the state process (which for the dual problem we denote by $\mathcal{Y}_i$ for $i = 0, \dots, N$) begins, as opposed to starting at x_0 . We only optimise one element of the initial dual state, but keep a multi dimensional state as $y_0^0 = (y, \tilde{y}) \in \mathbb{R}^d$, where $\tilde{y} \in \mathbb{R}^{d-1}$ is known. For example, when we consider the stochastic volatility model $\tilde{y}$ is the volatility process at time 0, which is known. Suppose $\mathcal{Y}$ is driven by a control β , a progressively measurable process, valued in $B \subset \mathbb{R}^m$, and has (deterministic and measurable) drift and diffusion functions $\tilde{b} : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}$ and $\tilde{\Sigma} : [0, T] \times \mathbb{R}^d \times A \rightarrow \mathbb{R}^{d \times n}$ respectively. Then, in analogue to Theorem 3.1, we have solutions (V, Z, Γ) to the 2BSDE

$$\begin{aligned} d\mathcal{Y}(t) &= \tilde{b}(t, \mathcal{Y}(t), \beta(t))dt + \tilde{\Sigma}(t, \mathcal{Y}(t), \beta(t))dW(t) \\ dV(t) &= Z(t)^\top \tilde{\Sigma}(t, \mathcal{Y}(t), \beta(t))dW(t) \\ dZ(t) &= -D_x \tilde{H}(t, \mathcal{Y}(t), \beta(t), Z(t), \Gamma(t)\Sigma(t, \mathcal{Y}(t), \beta(t)))dt + \Gamma(t)\tilde{\Sigma}(t, \mathcal{Y}(t), \beta(t))dW(t) \end{aligned} \quad (16)$$

with the terminal conditions $V(T) = \tilde{g}(\mathcal{Y}(T))$ and $Z(T) = D_y \tilde{g}(\mathcal{Y}(T))$, where $\tilde{g}$ is the dual terminal reward function. The function $\tilde{H}$ is defined as in (9) but with the dual drift and diffusion functions. We use the same notation for the solutions here for brevity, but when considering both primal and dual 2BSDEs simultaneously a distinction will be made. We then discretise and optimise in the same way, but with the addition of a third parameter set $\Theta_3 = \{y\}$ with corresponding loss function

$$\mathcal{L}_3(y) = \mathbb{E}[\tilde{g}(\mathcal{Y}_N) | \mathcal{Y}_0 = y] + yx_0. \quad (17)$$

3.3 Numerical Examples

In this section we apply the deep controlled 2BSDE algorithm to solve utility maximisation, with varying utilities and control constraints. This is a control problem (6) with $d = 1$ and $n = m$. For the primal problem, we have the process $\mathcal{X} = X$ defined by (2), controlled by $\alpha = \pi$, and the primal 2BSDE system of Theorem 3.1 denoted by (V_1, Z_1, Γ_1) , where we take $g(x) = U(x)\mathbb{1}_{x>0}$ (c.f. Remark 3.1) and define the drift and diffusion by

$$b(t, x, \pi) = x(r(t) + \pi^\top (\mu(t) - r(t)\mathbb{1}_m)), \quad \Sigma(t, x, \pi) = x\pi^\top \sigma(t).$$

For the dual problem the state process $\mathcal{Y} = Y$ is defined in (3), controlled by $\beta = v$ and starting at $y_0 > 0$, and the 2BSDE system of Theorem 3.1 is denoted by (V_2, Z_2, Γ_2) , where we take $\tilde{g}(y) = \tilde{U}(y)\mathbb{1}_{y>0}$ and define the drift and diffusion by

$$\tilde{b}(t, y, v) = -y(r(t) + \delta_K(v)), \quad \tilde{\Sigma}(t, y, v) = -y(\theta(t) + \sigma(t)^{-1}v)^\top.$$

We use the following primal-dual relations, derived in Li and Zheng (2018).

Theorem 3.7 (Li and Zheng (2018), Theorem 12). *Suppose that $\hat{y}$ and $\hat{v}$ are optimal dual controls, with the corresponding dual state process $\hat{Y}$ and 2BSDE solutions $(\hat{V}_2, \hat{Z}_2, \hat{\Gamma}_2)$. Then the primal state process $\hat{X}$ and 2BSDE solutions $(\hat{V}_1, \hat{Z}_1, \hat{\Gamma}_1)$ corresponding to the optimal primal control $\hat{\pi}$ satisfy*

$$\hat{X}(t) = -\hat{Z}_2(t), \quad \hat{V}_1(t) = \hat{V}_2(t) + \hat{Z}_2(t)^\top \hat{Y}(t), \quad \hat{Z}_1(t) = \hat{Y}(t). \quad (18)$$

The converse relation also applies, in that primal solutions can be used to derive the corresponding dual solutions. The direction described here is often the most useful, as the dual problem can be simpler to solve than the primal problem. In addition to these processes, formulas for the optimal controls are also given, but we do not use them here. To show the effectiveness of the algorithm we apply it to a series of utility problems, with certain constraints on the control space.

Example 3.8 (Unconstrained Non-HARA Utility Problem). We start with the unconstrained case $K = \mathbb{R}^m$, corresponding to a complete market. We consider the utility problem where U has the following non-HARA form

$$U(x) = \frac{1}{3}H(x)^{-3} + H(x)^{-1} + xH(x), \quad x > 0, \quad (19)$$

where $H(x) = \sqrt{2}(-1 + \sqrt{1 + 4x})^{-\frac{1}{2}}$. The dual utility is given by $\tilde{U}(y) = \frac{1}{3}y^{-3} + y^{-1}$. We have $\tilde{K} = \{0\}$, so the optimal dual control is $v(t) = 0$. Suppose that $r(s) = r$, $\sigma(s) = \sigma$, $\mu(s) = \mu$ and $\theta(s) = \theta$ are constant. It is shown in Li and Zheng (2018) that the optimiser $\hat{y}$ and processes $\hat{Y}$ and $\hat{Z}_2$ are given by

$$\begin{aligned} \hat{y} &= \frac{1}{\sqrt{2x_0}} \left[e^{(r+|\theta|^2)T} + \sqrt{e^{2(r+|\theta|^2)T} + 4x_0 e^{3(r+2|\theta|^2)T}} \right]^{\frac{1}{2}} \\ \hat{Y}(t) &= \hat{y} \exp \left(- \left(r + \frac{1}{2} |\theta|^2 \right) t - \theta^\top dW(t) \right) \\ \hat{Z}_2(t) &= a_1 S_1(t) + a_2 S_2(t), \end{aligned}$$

where $a_1 = \hat{y}^{-4} e^{3(r+2|\theta|^2)T}$, $a_2 = \hat{y}^{-2} e^{(r+|\theta|^2)T}$, $S_1(t) = e^{(r-4|\theta|^2)t + 2\theta^\top W(t)}$ and $S_2(t) = e^{rt + 2\theta^\top W(t)}$. Furthermore, the dual value function is given by

$$\tilde{u}(t, y) = \frac{1}{3} y^{-3} e^{(3r+6|\theta|^2)(T-t)} + y^{-1} e^{(r+|\theta|^2)(T-t)}.$$

From these analytical formulae, the primal processes can be derived using Theorem 3.7. For our application, we take $m = 5$, $T = 0.5$, $x_0 = 1$, $r = 0.05$, $\mu = 0.06 \mathbb{I}_5$ and set σ as a random matrix of positive numbers between 0 and 0.2 to introduce correlation of the underlying Brownian motions, with 0.2 added to the diagonal to ensure invertibility.

Figure 1 shows the development of the loss functions for the primal and dual algorithms with 50 time steps. For the control, we want to maximise the Hamiltonian, which means the derivative of the loss function should go to 0. Hence, we plot the function $j \mapsto \sum_{i=0}^{49} \sum_{\theta \in \theta_i^j} |D_\theta \mathcal{L}(\theta_i^j, i)|$ using the notation of (15), where θ_i^j denotes the parameter set

Fig. 1 Loss functions against iteration step for the primal and dual deep controlled 2BSDE methods applied to the unconstrained non-HARA utility problem

for the neural network forming π_t , at iteration step j of the algorithm. For the dual graph we plot the derivative of $\mathcal{L}_3$ with respect to y , since the above function is always zero as the dual control is forced to be zero, hence optimal, by the support function δ_K . Over the 10000 iteration steps both losses become small and settled, and therefore are deemed to have converged. In this method we choose an initial learning rate of 10^{-2} for the 2BSDE algorithm and the start control y in the dual algorithm, and 10^{-3} for the control part. Over the iterations, we divide these rates by 10 three times at regular intervals. We repeat this choice of learning rates for all applications in the sequel, though the number of iteration steps may change.

Figure 2 shows two simulations of the processes V_1 and Z_1 (labelled as primal), along with their dual counterparts given by the relations (18) using 20 time steps. We plot these, along with the explicit solution for reference. Both primal and dual approximations are very close to the solutions. The dual state process, appearing in the Z_1 graph (but very close to the solution process), is closer than the corresponding primal process since there is no error for the dynamic dual control.

Fig. 2 Two simulations of the value process and value-derivative process for the primal and dual deep controlled 2BSDE methods applied to the unconstrained non-HARA utility problem. The displayed dual processes are those implied by the duality relations (18)

Figure 3 shows the evolution of the approximation of the initial value $u(0, x_0)$ against number of time steps for the primal and dual algorithms, where $x_0 = 1$. The dual value is higher than the primal value, as expected, and the gap between these shrinks as the number of time steps increases. The final relative approximation error is $7 \times 10^{-4}\%$ for the primal problem, and $3 \times 10^{-4}\%$ for the dual problem, showing high accuracy of the proposed scheme at time 0.

Example 3.9 (Cone-Constrained Merton Problem). Now we consider the problem with power utility $U(x) = \frac{1}{p}x^p$ for some $0 < p < 1$ when the control space is the positive cone $K = \mathbb{R}_+^m$, corresponding to no short selling constraints. Let $\hat{X}$ be the primal state process and $(\hat{V}_1, \hat{Z}_1, \hat{\Gamma}_1)$ the solutions of (10) for the primal problem with optimal control $\hat{\pi}$, and let $\hat{Y}$ and $(\hat{V}_2, \hat{Z}_2, \hat{\Gamma}_2)$ be the state process and solutions respectively for the dual problem with optimal controls $\hat{\nu}$ and $\hat{y}$. The dual state process is a geometric Brownian motion given by

$$\hat{Y}(t) = \hat{Y}(0) \exp \left(- \int_0^t \left(r(s) + \frac{1}{2} |\hat{\theta}(s)|^2 \right) ds - \int_0^t \hat{\theta}(s)^\top dW(s) \right)$$

in terms of the unknown start point $\hat{y} = \hat{Y}(0)$. The corresponding dual utility function is $\tilde{U}(y) = \frac{1-p}{p} y^{\frac{p}{p-1}}$. In [Li and Zheng (2018), Subsection 4.1] the following solutions are found:

$$\begin{aligned} \hat{y} &= x_0^{p-1} \exp \left((1-p) \int_0^T \left[\frac{p}{2(p-1)^2} |\hat{\theta}(s)|^2 - \frac{p}{p-1} r(s) \right] ds \right) \\ \hat{Z}_2(t) &= -x_0 \exp \left(\int_0^t \left[r(s) + \frac{1-2p}{2(1-p)^2} |\hat{\theta}(s)|^2 \right] ds + \frac{1}{1-p} \int_0^t \hat{\theta}(s)^\top dW(s) \right), \end{aligned} \quad (20)$$

where $\hat{\theta}(t) = \theta(t) + \sigma(t)^{-1} \hat{\nu}(t)$. This deterministic optimal dual control can be found numerically using Python packages such as `Scipy`. The dual value function is given by

$$\tilde{u}(t, y) = \tilde{U}(y) \exp \left(\int_t^T \left[\frac{p}{2(p-1)^2} |\hat{\theta}(s)|^2 - \frac{p}{p-1} r(s) \right] ds \right). \quad (21)$$

(a) Value approximation.

(b) Relative error of value approximation.

Fig. 3 Approximation and relative error of the value function against number of time steps N for the primal and dual deep controlled 2BSDE methods applied to the unconstrained non-HARA utility problem. Note in the left graph that the Y values on the axis are shifted down by 2.253

For our implementation, we take $T = 0.5$, $x_0 = 1.0$, $m = 50$, $r(t) = 0.05$ and $p = \frac{1}{2}$. We define the drift as $\mu(t)_i = 0.04 + \sin(\pi t + A_i)/50$ where A_i is randomly generated for each $i = 1, \dots, 50$. We choose this function so that in the unconstrained case it would be optimal to short sell at some points. We set $\sigma(t)$ to be a constant matrix with 0.4 on the diagonal, and 0.2 everywhere else.

Remark 3.10 To ensure our controls π and v are in the correct set $K = \tilde{K} = \mathbb{R}_+^m$, we (element-wise) apply the function $x \mapsto \max(0, x)$ to the output of their respective neural networks. This function is an almost everywhere differentiable surjection from $\mathbb{R}^m$ to $\mathbb{R}_+^m$, which is sufficient for using a gradient descent method to find the optimal control. The mapping is not injective, but this is not an issue as the exact output of the neural network is not relevant. We care only about the control outputted by projecting onto the control space.

Figure 4 shows the results of applying the primal and dual 2BSDE methods to this problem. We use $N = 10$ time steps and run the algorithm for 100000 steps, notably more than for the lower dimensional unconstrained problem. This increases the run time to 5457 seconds. We start the BSDE and control learning rates at 0.01 and 0.001 respectively, then reduce these by a factor of 10 half way through training. We run both algorithms at the same time, which saves time as we can use the same generated Brownian motions for each algorithm. We see less convergent structure at the beginning of training, perhaps due to bad initial guesses for the neural networks. Indeed, we see that the value for the primal functions stays nearly constant until the loss function drops to a certain level, when the processes are accurate enough that the value approximation can move towards the analytical solution. The dual algorithm behaves similarly, but with more noise. At the end of training we have relative errors of 0.1% and 0.05% for the primal and dual algorithms respectively. This error is a few orders of magnitude higher than the unconstrained case due to the increase in dimension.

Example 3.11 (Ball Constrained Log Utility Problem). In this example we consider the log utility problem when the control process is constrained to the ball $B(0, R)$ of radius $R > 0$ in $\mathbb{R}^m$. This corresponds to the investor not wanting to invest more than a certain proportion of their wealth in risky stocks. In this case, the support function $\delta_K(z) = R|z|$ does

Fig. 4 Approximation of the value function and bsde loss against iteration step for the primal and dual deep controlled 2BSDE methods applied to the cone constrained power utility problem

not vanish for all admissible dual controls. Therefore, the SDE (3) is harder to solve. In general this problem is also difficult to solve analytically because of the dependence of the dual controls $\hat{y}$ and $\hat{v}$ on each other. However, we can exploit the multiplicity property of the log function to decouple these controls to give analytical solutions, which we describe here. To this end let $U(x) = \log(x)$ be our utility function, for which the dual utility is $\tilde{U}(y) = -(1 + \log(y))$. Since the dual process Y is a geometric Brownian motion we have

$$\mathbb{E}[\tilde{U}(Y(T))] = -1 - \log(Y(0)) + \int_0^T r(t) + R|v(t)| + \frac{1}{2}|\theta(t) + \sigma(t)^{-1}v(t)|^2 dt.$$

The measurable process

$$\hat{v}(t) := \arg \min_{v \in \mathbb{R}^m} \left\{ R|v| + \frac{1}{2}|\theta(t) + \sigma(t)^{-1}v|^2 \right\}$$

is an optimal control, and can be found using existing deterministic convex optimisation algorithms. Independently of this, the dual control $\hat{y}$ is found via

$$\hat{y} \in \arg \min_{y > 0} \left\{ x_0 y - \log(y) + \int_0^T r(t) + \delta_K(\hat{v}(t)) + \frac{1}{2}|\theta(t) + \sigma(t)^{-1}\hat{v}(t)|^2 dt \right\}.$$

yielding $\hat{y} = \frac{1}{x_0}$. The value function is then given by

$$u(0, x_0) = \log(x_0) + \int_0^T r(t) + R|\hat{v}(t)| + \frac{1}{2}|\theta(t) + \sigma(t)^{-1}\hat{v}(t)|^2 dt. \quad (22)$$

For our application, we take $R = 1.0$, $T = 0.5$, $x_0 = 5.0$, $m = 20$, $r(t) = 0.05$ and $\mu(t) = 0.07\mathbb{1}_S$. We define the volatility as a diagonal matrix with the i th diagonal element equal to $\sigma(t)_i = 0.4 + 0.2 \sin(2\pi t + A_i)$ for some arbitrarily (in our case randomly) chosen numbers A_i to shift the peaks of the sine curve for each dimension. We choose this function so that in the unconstrained case it would be reasonable to invest more than we are allowed at points when the volatility is low. The volatility is non-negative for all $t \in [0, T]$, so the problem is well posed. To ensure that our primal control is in the correct set, we use a penalty function method. In this case we do not constrain π , but instead we alter the loss function (15) to be

$$\mathcal{L}_2(\theta_i, i) := \mathbb{E}[F(t_i, X_i, \pi_i, Z_i, \Gamma_i) - \beta \max(0, |\pi_i| - 1)^2]$$

which penalises the control for straying outside of K . The parameter $\beta > 0$ is some large constant, set to 1000 for our example.

Table 1 shows the evolution of the approximated value function for the primal and dual algorithms, compared with the numerical solution 1.64755570 found using (22). The algorithm is run using $N = 10$ time steps, and the value approximation, relative error and (combined) runtime are recorded. We again perform both algorithm simultaneously using the same Brownian motions simulations. The value approximation initially moves away from the solution, but then converges. For the primal algorithm, the value moves away over the first 1000 steps, then goes to the solution. For the dual algorithm, the value does not move as far from the solution, but takes longer to do so, only starting moving towards the solution at 10000 steps. The dual value initially dips below the solution, then increases, which is why the relative error is lower at 1000 steps than at the start.

Table 1 Value approximation and runtime at certain iteration steps for the primal and dual deep controlled 2BSDE methods applied to the ball-constrained log utility problem

Step	Primal Value	Rel-Err (%)	Dual Value	Rel-Err (%)	Time (s)
1	9.92774e-01	3.97426e+01	2.50412e+00	5.19898e+01	128
1000	4.63730e-02	9.71853e+01	1.57275e+00	4.54015e+00	183
5000	2.27800e-01	8.61735e+01	1.80325e+00	9.44972e+00	405
10000	1.39359e+00	1.54147e+01	2.01337e+00	2.22035e+01	701
12000	1.62170e+00	1.56918e+00	1.98620e+00	2.05544e+01	815
18000	1.63545e+00	7.34658e-01	1.68103e+00	2.03169e+00	1149
20000	1.63543e+00	7.35726e-01	1.65361e+00	3.67465e-01	1259

4 The Deep SMP Method for the Non-Markovian Case

In this section we consider the case in which r , μ and σ may not be deterministic. The utility problem and its dual problem then become non-Markovian. This means we cannot write the value function as a measurable function of state and time nor apply the dynamical programming principle, hence the HJB equation is invalid. However, we can still employ the SMP, and it is shown in Li and Zheng (2018) that there exist certain optimality conditions for a non-Markovian utility maximisation problem, which we use to define a new algorithm, named the deep SMP method. In some of the examples we consider in this section, the problem can be made Markovian by introducing some new artificial structure to the problem. In these cases, we compare the accuracy and runtime of both algorithms.

4.1 The Deep SMP Method

Consider an agent who invests in the market (1) in full generality. We again consider the utility problem, written in this setting as

$$u(t) = \sup_{\pi \in \mathcal{A}} \mathbb{E} [U(X(T)) \mid \mathcal{F}_t].$$

We cannot use Itô's formula for this value process as it cannot be written as a measurable function of the wealth process $X(t)$, defined as in (2) and controlled by π , but we can still use the SMP. For the primal problem we have the following adjoint BSDE with unknown processes P_1 and Q_1 valued in $\mathbb{R}$ and $\mathbb{R}^m$ respectively:

$$dP_1(t) = -[(r(t) + \pi(t)^\top (\mu(t) - r(t)\mathbb{1}_m))P_1(t) + Q_1(t)^\top \sigma(t)^\top \pi(t)]dt + Q_1(t)^\top dW(t) \quad (23)$$

with the terminal condition $P_1(T) = -U'(X(T))$. Similarly, we can define the dual state Y as per (3), controlled by (y, ν) . The dual adjoint equation is the following BSDE with unknown processes P_2 and Q_2 valued in $\mathbb{R}$, $\mathbb{R}^m$ respectively:

$$dP_2(t) = [(r(t) + \delta_K(\nu(t))P_2(t) + Q_2(t)^\top (\theta(t) + \sigma(t)^{-1}\nu(t))]dt + Q_2(t)^\top dW(t) \quad (24)$$

with the terminal condition $P_2(T) = -\tilde{U}'(Y(T))$. Furthermore, we have the duality relations $P_2(t) = X(t)$ and $P_1(t) = -Y(t)$ almost surely, for almost every $t \in [0, T]$ by (18). We now state the optimality conditions for the dual problem in terms of these processes.

Theorem 4.1 (Li and Zheng (2018), Theorem 11). *Let (y, v) be an admissible dual control. Then (y, v) is optimal if and only if the solutions (P_2, Q_2) of (24) satisfy the following*

1. $P_2(0) = x_0$,
2. $\frac{(\sigma(t)^\top)^{-1}Q_2(t)}{P_2(t)} \in K$,
3. $P_2(t)\delta_K(v(t)) + Q_2(t)^\top \sigma(t)^{-1}v(t) = 0$

for all $t \in [0, T]$ almost surely.

Remark 4.2 The corresponding optimality condition for a primal control π with corresponding state process X and adjoint BSDE solutions (P_1, Q_1) is

$$[\pi(t) - a]^\top [X(t)\sigma(t)(P_1(t)\theta(t) + Q_1(t))] \geq 0$$

for all $t \in [0, T]$ and $a \in K$ almost surely. This is a much more complex condition, and the implementation would require some grid-based method to consider a sufficiently large number of $a \in K$, which would be computationally costly, especially in higher dimensions and unbounded (but still constrained) K . For this reason, we approach the dual conditions and then use the primal dual relations to derive the necessary primal processes.

The machine learning methodology to solve this problem is as follows. First, we simulate the dual processes in the forward direction using the Euler-Maruyama scheme. We start the process P_2 at $P_2(0) = x_0$ to ensure the first condition is satisfied. This condition is why it appears that the forward simulation method is naturally suited to this problem. We use neural networks to define the processes v and Q_2 , then update the values of Y and P_2 accordingly. Specifically, we set $Y(0) = y$, $P_2(0) = x_0$ then for $i = 0, \dots, N-1$ set $v(i) = N_{\theta_i^v}(Y(i))$ and

$$\begin{aligned} Q_2(i) &= P_2(i)\sigma(t_i)^\top h_K(N_{\theta_i^Q}(Y(i))) \\ Y(i+1) &= Y(i) - (t_{i+1} - t_i)Y(i)(r(t_i) + \delta_K(v(i))) - Y(i)(\theta(t_i) + \sigma(t_i)^{-1}v(i))^\top dW_i \\ P_2(i+1) &= P_2(i) + (t_{i+1} - t_i)[r(t_i)P_2(i) + Q_2(i)^\top \theta(t_i)] + Q_2(i)^\top dW_i, \end{aligned} \quad (25)$$

where $\theta_i^v, \theta_i^Q \in \mathbb{R}^{p(d,m)}$ are parameters for neural networks, dW_i is a multivariate normal $\mathcal{N}_m(0, (t_{i+1} - t_i)\mathbb{1}_m)$ random variable, and $h_K : \mathbb{R}^m \rightarrow K$ is a differentiable surjective function that ensures condition 2 in Theorem 4.1 is satisfied. The two neural networks used here both output to m -dimensional space, unlike the general algorithm outlined in Sect. 3.1, where the second neural network is valued in $\mathbb{R}^{d \times d}$. We have the variables $y, \Theta^Q := (\theta_i^Q)_{i=0}^{N-1}$ and $\Theta^v := (\theta_i^v)_{i=0}^{N-1}$, which are optimised against the following loss functions:

$$\begin{aligned} \mathcal{L}^y(y) &= \mathbb{E}[\tilde{U}(Y(N)) | Y(0) = y] + x_0 y \\ \mathcal{L}^Q(\Theta^Q) &= \mathbb{E}[|\tilde{U}'(Y(N)) + P_2(N)|^2] \\ \mathcal{L}^v(\Theta^v) &= \mathbb{E}[|P_2(i)\delta_K(v(i)) + Q_2(i)^\top \sigma(t_i)^{-1}v(i)|^2]. \end{aligned}$$

We take sample means to approximate the expected values. The first loss function represents the optimality condition of y from the definition (4) of the dual problem, the second the terminal condition of BSDE (24), and the third condition 3 of Theorem 4.1.

The optimisation process is very similar to the deep controlled 2BSDE method as outlined in Sect. 3.1, so we only provide a brief summary here. For each iteration step we simulate the processes Y , v , P_2 and Q_2 using our existing (or randomly initialised for the first step) parameter sets y , Θ^Q and θ^v and a batch of simulated Brownian motion paths. We then perform one iteration of the ADAM algorithm to update each parameter set in turn using their respective loss functions. In between the updates, we simulate the processes again using the new parameters, but do not generate new Brownian motions. When the 3 sets are updated we generate new Brownian motion paths and repeat. This process continues until the difference made by updating is sufficiently small. For our case, we start with a learning rate of 0.01 for all parameters, which decreases as the number of iterations increases.

In this algorithm we no longer have a variable representing the value at time 0. We therefore need to perform Monte Carlo simulations, forming bounds of the value function. An upper bound can be found using the dual state process Y and a lower bound the primal state process $X = P_2$, as per (18). The estimates outputted are given by

$$u_{\text{low}} := \frac{1}{M} \sum_{i=1}^M U(P_2^i(N)) \leq u(0, x_0) \leq \frac{1}{M} \sum_{i=1}^M \tilde{U}(Y^i(N)) + x_0 y =: u_{\text{high}} \quad (26)$$

for some large batch of size $M \in \mathbb{N}$, where each path is indexed by the superscript i . The start point y is not random as it is optimised using a deterministic loss function. The inequalities are valid up to discretisation and simulation errors, and as the number of time steps increases the gap $u_{\text{high}} - u_{\text{low}}$ decreases, and the optimal value is eventually contained in this gap.

4.2 Numerical Examples

Example 4.3 (Power Utility in a Stochastic Volatility Model). In this example we consider an extension of our utility model to stochastic volatility. In this model, our agent invests in a 2-dimensional stock market, out of which one asset S is traded, and is defined by

$$dS(t) = S(t)(r + Av(t))dt + S(t)\sqrt{v(t)}dW^S(t).$$

The constants $r, A > 0$ represent the risk-free rate and market price of risk respectively. The volatility process v has the dynamics

$$dv(t) = \kappa(\theta - v(t))dt + \xi\sqrt{v(t)}dW^v(t).$$

The constant $\kappa > 0$ represents the speed at which the process reverts to the long-time average volatility, given by $\theta > 0$. The constant $\xi > 0$ is a variance parameter. The Brownian motions W^S and W^v are correlated with correlation parameter $\rho \in [-1, 1]$. When the agent invests a proportion π of their wealth X in this market, their wealth process (2) evolves as

$$dX(t) = X(t)(r + \pi(t)Av(t))dt + X(t)\pi(t)\sqrt{v(t)}dW^S(t).$$

We first consider this as a 2-dimensional Markovian problem, then a random coefficient problem. In the Markovian problem the 2-dimensional process $\mathcal{X} = (X, v)$ is the state process controlled by the $K \subset \mathbb{R}$ valued process π . Consider the maximum utility problem given by

$$u(t, x, v) = \sup_{\pi \in \mathcal{A}} \mathbb{E}[U(X(T)) \mid (X(0), v(0)) = (x, v)]$$

for some utility function U . This setup fits in the framework of (6) with $d = 2$, $n = 2$, $m = 1$ and $g(x, v) := U(x)\mathbb{1}_{x>0}$ for $(x, v) \in \mathbb{R}^2$. The drift and diffusion functions are given by

$$b(t, x, v, \pi) = (x(r + \pi A v), \kappa(\theta - v))^T$$

$$\Sigma(t, x, v, \pi) = \begin{pmatrix} x\pi\sqrt{v} & 0 \\ \rho\xi\sqrt{v} & \sqrt{1 - \rho^2\xi}\sqrt{v} \end{pmatrix}.$$

It is a usual assumption that the drift and diffusion functions are Lipschitz functions of the state and control variable, but this is not the case here due to the square root appearing in the diffusion. This may invalidate certain convergence patterns that we have seen for the utility maximisation problem thus far. The dual state process consists of the volatility process v and an $\mathbb{R}$ -valued process Y that is controlled by some $\mathbb{R}^2$ -valued process (η, γ) , given by

$$dY(t) = -Y(t) \left[\left(r + \delta_K(\eta(t)) \right) dt + \left(\frac{\eta(t)}{\sqrt{v(t)}} + \rho\gamma(t) + A\sqrt{v(t)} \right) dW^s(t) - \gamma(t) dW^v(t) \right]. \quad (27)$$

As before, the dual process starts at some undetermined $y > 0$, which forms the dual control together with η and γ , and is chosen such that the process XY is a super martingale. Now consider the dual control problem

$$\tilde{u}(t, y, v) = \inf_{\gamma, \eta} \mathbb{E}[\tilde{U}(Y(T)) \mid (Y(t), v(t)) = (y, v)]. \quad (28)$$

In our numerical example we consider the unconstrained case $K = \mathbb{R}$. As before, the presence of the δ_K term means that we must have $\eta = 0$. This puts the problem in the same form as Ma et al. (2020), where only one process γ needs to be found.

This setup fits in our framework for a dual problem with $d = 2$, $m = 1$, and $n = 2$. This is not a multidimensional duality as the process v appears in both state dynamics, but a single dimensional duality between X and Y . The drift and diffusion functions are given by

$$\tilde{b}(t, y, v, \gamma) = (-yr, \kappa(\theta - v))^T$$

$$\tilde{\sigma}(t, y, v, \gamma) = \begin{pmatrix} -yA\sqrt{v} & \sqrt{1 - \rho^2 y \gamma} \\ \rho\xi\sqrt{v} & \sqrt{1 - \rho^2 \xi}\sqrt{v} \end{pmatrix}.$$

When we use the power utility $U(x) = \frac{x^p}{p}$ for $0 < p < 1$, exact solutions can be obtained for this problem. A solution is derived in Ma et al. (2020) by assuming that the value function u has the form

$$u(t, x, v) = \frac{x^p}{p} \exp(C(t) + D(t)v) \quad (29)$$

in terms of two unknown functions C and D that satisfy a system of Riccati type of ordinary differential equations for which numerical solutions can be found. We use $u(t, x, v)$ as a benchmark. For our application we take $r = 0.05$, $\rho = -0.5$, $p = 0.5$, $\kappa = 1$, $\theta = 0.05$, $\xi = 0.5$, $A = 0.5$, $x_0 = 1$, and $v_0 = 0.5$.

Table 2 shows the relative error and run time of our algorithms when we use a range of terminal times T and numbers of time steps N . We contrast this to the results of Ma et al. (2020), in which, with $T = 1.0$ and $N = 100$, upper and lower bounds for the value were found with relative errors of 0.003% in 2240 seconds and 27600 seconds respectively, albeit for slightly different parameter choices. The upper bound was found faster using a convenient guess of the form the dual control γ . We are not using 100 time steps for our algorithms due to computer memory constraints. In fact, the 50 time step algorithm is also slower than expected due to a memory bottleneck. For $T = 0.2$ the solution is 2.03289, for $T = 0.5$ it is 2.07559 and when $T = 1.0$ the solution is 2.13420. The run time for the $T = 0.5$ and $T = 1.0$ algorithms is higher than for $T = 0.2$ since twice the iteration steps were taken, 10000 compared to 5000 for $T = 0.2$. While the errors for the primal algorithm reduce as N increases, the value does not lie within the duality gap for some terminal times. This may be down to a discretisation error, and would be fixed by taking more time steps. We see convergence of the approximations to the true value as N increases and this convergence seems to be linear, as we saw in Example 3.1. This convergence is explored further in Sect. 5.

Now, let us treat v as some random process, meaning that we can only sample v , and not the underlying Brownian motion W^v . In this setting we have that both $\theta(t) = A\sqrt{v(t)}$ and $\sigma(t) = \sqrt{v(t)}$ are random processes for the stock market, and we still wish to apply Theorem 4.1. We cannot directly apply this theorem as v is not measurable with respect to the filtration generated by W^s . To circumvent this problem, we adopt a market completion method (c.f. Karatzas and Shreve (1998) for example) by first constructing an artificial stock driven by W^v , adding it to our market, and then making it unavailable for trading. To be specific, we define a new asset price S^v by

$$dS^v(t) = S^v(t)(rdt + dW^v(t)),$$

which ensures that v is adapted to the natural filtration of the Brownian motions driving the market. For simplicity we assume that $m = 1$, but this analysis extends naturally to multiple dimensions. We invest in this market with a 2-dimensional portfolio, restricted to the set $K \times \{0\}$, ensuring that our wealth process X is still defined as in (2), controlled by a

Table 2 Accuracy and runtime for the deep controlled 2BSDE method for various terminal times T and numbers of time steps N applied to the unconstrained power utility problem in a stochastic volatility model

T	N	Primal Value	Rel-Err (%)	Time (s)	Dual Value	Rel-Err (%)	Time (s)
0.2	5	2.03310	1.01654E-02	83	2.03378	4.35850E-02	98
0.2	10	2.03301	5.74161E-03	206	2.03344	2.66047E-02	239
0.2	20	2.03296	3.27670E-03	591	2.03314	1.23181E-02	604
0.2	50	2.03289	2.81206E-04	2849	2.03295	2.89678E-03	2622
0.5	5	2.07616	2.75979E-02	152	2.07969	1.97584E-01	183
0.5	10	2.07582	1.14780E-02	362	2.07747	9.08529E-02	411
0.5	20	2.07572	6.36082E-03	927	2.07670	5.38088E-02	1030
0.5	50	2.07562	1.55757E-03	5070	2.07600	1.98421E-02	4181
1.0	5	2.13328	4.32410E-02	154	2.14847	6.68466E-01	170
1.0	10	2.13377	1.99985E-02	362	2.14264	3.95457E-01	400
1.0	20	2.13399	9.92254E-03	918	2.13715	1.38171E-01	981
1.0	50	2.13403	7.97535E-03	4230	2.13548	6.01058E-02	4395

1-dimensional K -valued process π . Now the dual state process (27) is controlled by y and a 2-dimensional dynamic control (η, γ) valued in $\tilde{K} \times \mathbb{R}$. We would like to show that we can take $\gamma = 0$, as this would return the dual problem to the 1-dimensional control case independent of W^v . We want this as we cannot ‘see’ W^v , and can only sample v directly at any time. To this end, let (P_2, Q_2) be the dual adjoint solutions. Fix a time $t \in [0, T]$ and let $Q_2(t) = (q_1, q_2)$. By Theorem 4.1 (2) we have $P_2(t)^{-1}q_2 = 0$, hence $q_2 = 0$. However, this means that condition (3) becomes

$$P_2(t)\delta_K(\eta(t)) + q_1^\top \sigma(t)^{-1}\eta(t) = 0$$

which is independent of $\gamma(t)$, leading to a free control. We therefore can take $\gamma = 0$, removing the dependence of the dual state process on W^v . Note that this choice of control may not be optimal, and the decision to remove dependence on W^v may lead to a non-zero duality gap.

Table 3 shows the accuracy and runtime of the deep SMP method with a range of terminal times and step sizes, when we run it on the same problem formulation as above. Here, $K = \mathbb{R}$, so we take the constraining function h_K to be the identity. We note that the upper bounds are an order of magnitude less accurate the lower bounds for this problem. This indicates that the implied primal process approximations are more accurate than the dual approximations in this algorithm. Comparing these results to Table 2, the SMP method gives roughly the same accuracy of bounds as the dual deep controlled 2BSDE method, but with a lower computation time. In particular, the algorithm here only needed 5000 iterations steps, compared to 10000 for both the 2BSDE algorithms.

Remark 4.4 We have applied this methodology to investing in the Vasicek model where the interest rate is a random process rather than the volatility. In the same way as for the Heston model we consider it both as a 2-dimensional state process and a 1-dimensional state process with random coefficients. We find similar results as for the Heston model, so we omit this example for brevity.

Example 4.5 (Heston with non-HARA utility). We now look at a higher dimensional example with no explicit solution. Consider the same stochastic volatility market, but with multiple stocks. For simplicity we assume all stocks are independent with the same

Table 3 Accuracy and runtime for the deep SMP method for various terminal times T and numbers of time steps N applied to the unconstrained power utility problem in a stochastic volatility model

T	N	u_{low}	Rel-Diff (%)	u_{high}	Rel-Diff (%)	Time (s)
0.2	5	2.03271	9.17595E-03	2.03348	2.90360E-02	103
0.2	10	2.03280	4.40981E-03	2.03323	1.64607E-02	273
0.2	20	2.03284	2.60193E-03	2.03305	7.49044E-03	781
0.2	50	2.03289	5.76610E-05	2.03295	2.54626E-03	4788
0.5	5	2.07512	2.25897E-02	2.07974	2.00124E-01	102
0.5	10	2.07538	9.97694E-03	2.07738	8.63624E-02	301
0.5	20	2.07563	2.04008E-03	2.07656	4.67255E-02	778
0.5	50	2.07556	1.16593E-03	2.07595	1.74975E-02	3818
1.0	5	2.13025	1.84896E-01	2.14786	6.40234E-01	104
1.0	10	2.13232	8.78697E-02	2.13995	2.69313E-01	263
1.0	20	2.13489	3.25920E-02	2.13750	1.54844E-01	761
1.0	50	2.13440	9.35931E-03	2.13551	6.11910E-02	4604

parameters, but this is not a necessary restriction. We let U have the form (19). For this model we do not have a convenient ansatz for u to obtain a solution, so we simply compare our algorithms. Say we have k stocks. We would then have control processes $\pi(t) \in \mathbb{R}^k$, $\Gamma_1(t) \in \mathbb{R}^{1+k,1+k}$ for the primal 2BSDE algorithm, $v(t) \in \mathbb{R}^{2k}$, $\Gamma_2(t) \in \mathbb{R}^{1+k,1+k}$ for the dual 2BSDE, and $v(t), Q_2(t) \in \mathbb{R}^k$ for the SMP algorithm.

For this example let us take $T = 0.2$, $N = 20$, and $k = 10$. We then have $m = 10$, $d = 11$, and $n = 20$ in (5). For the 2BSDE algorithms, the combined dimensions of the unknown processes are 131 for the primal and 141 for the dual. We consider the unconstrained problem, so we can set $\eta = 0$. For the SMP algorithm we only need to find processes in 20 dimensions. We run the algorithm for 10000 steps. Table 4 shows the results of this. We give the value approximation for the 2BSDE algorithms, and the average of the upper and lower bound for the SMP. We also provide the runtime to show the efficiency of the algorithms as a whole, but the individual differences should be taken with a pinch of salt due to implementation differences, see Remark 4.6. The runtime has increased from the previous low dimensional example, but sub-linearly. The relative difference between approximations has increased to about 1%.

Remark 4.6 There are a number of factors that affect the algorithm run times. For example, we see that the dual algorithm takes twice as long as the primal, but we did not simplify the algorithm to reflect the constraints of the problem. Making this change would have reduced the run time. Additionally, the Z processes defined in (10) involve a differentiation element. This differentiation is done ‘automatically’, as it is handled by the computer, as opposed to the explicit derivative for the BSDE (24). This does not affect the accuracy, but increases the run time.

Example 4.7 (Power Utility in a Path Dependent Volatility Model). In this example not only is the control problem non-Markovian, but it cannot be made Markovian by including additional information in our state space. This means that we cannot apply our 2BSDE method.

This example is motivated by the concept that the volatility of a stock should be dependent on its value. If the stock is doing well the volatility is assumed to be lower. This may be due to infrequent trading of the stock, or the stability of the large company to which the stock is assigned. We consider an extreme example, where there are two levels of volatility. The normal higher level is maintain until the stock reaches its historical maximum, in which case it instantly transitions to the lower volatility.

Consider the market (1), where the process σ is a function of the stock history. For each $t \in [0, T]$, let $\sigma(t)$ be a diagonal matrix, with i th diagonal entry $\sigma_i(t)$ defined as

$$\sigma_i(t) = \begin{cases} \sigma_{\text{high}}, & S_i(t) < \sup_{s \leq t} S_i(s), \\ \sigma_{\text{low}}, & S_i(t) = \sup_{s \leq t} S_i(s), \end{cases}$$

Table 4 Value approximations for the 2BSDE and SMP algorithms applied to the unconstrained non-HARA utility problem in a stochastic volatility model

Method	Value	Time (s)
Primal	2.61059	2316
Dual	2.63758	5636
SMP	2.62514	1840

Table 5 Duality gap for the deep SMP method for various terminal times T and numbers of time steps N applied to the cone-constrained power utility problem with path dependent volatility

T	N	u_{low}	u_{high}	Diff	Rel-Diff (%)
0.2	5	2.01057	2.01066	8.99222E-05	4.47246E-03
0.2	10	2.01054	2.01063	8.55058E-05	4.25287E-03
0.2	20	2.01057	2.01063	6.00701E-05	2.98772E-03
0.2	50	2.01061	2.01062	1.30984E-05	6.51462E-04
0.5	5	2.02652	2.02682	3.03732E-04	1.49879E-02
0.5	10	2.02654	2.02676	2.11323E-04	1.04278E-02
0.5	20	2.02648	2.02665	1.73093E-04	8.54166E-03
0.5	50	2.02637	2.02650	1.31536E-04	6.49120E-03
1.0	5	2.05333	2.05421	8.87694E-04	4.32320E-02
1.0	10	2.05337	2.05397	5.93055E-04	2.88821E-02
1.0	20	2.05327	2.05368	4.16821E-04	2.03005E-02
1.0	50	2.05307	2.05339	3.20542E-04	1.56128E-02

where $\sigma_{\text{low}} < \sigma_{\text{high}}$ are constants and S_t is the i th stock price. In this model, the volatility of a stock is inversely proportional to the ratio of its value and the running maximum. This makes the diffusion coefficient of the wealth process X a function of the entire path of S , instead of just at one time. This path dependence ensures that the utility maximisation problem is truly non-Markovian, in the sense that it cannot be made Markovian by the introduction of a new variable. We do not need to apply an artificial stock market argument here, as the volatility is a measurable function of the adapted stock market process up to time t , hence is progressively measurable.

For our application, we consider a 2-dimensional stock market with cone constraint $K = \mathbb{R}_+^2$. The constraining function we use for this case is $h_K(x) = x^2$. We take $\sigma_{\text{low}} = 0.2$ and $\sigma_{\text{high}} = 0.3$. We then take $r(t) = 0.05$, $\mu(t) = 0.06\mathbb{1}_2$ and $p = 0.5$. Table 5 shows the duality gap when this algorithm is run for varying terminal times and numbers of time steps. We run the algorithms for 10000 iteration steps. The ‘Diff’ and ‘Rel-Diff’ columns show the absolute and relative distances between the upper and lower approximations respectively. We see in all cases that the duality gap reduces as the number of time steps increases. We do not have a benchmark to compare these results to, as the path dependence and incompleteness of the problem make it very difficult to solve analytically.

5 Methodology Comparison

First we compare the results of small changes to the methodology of the DC2BSDE algorithm, then compare the algorithms we have introduced to others existing in the literature.

5.1 Parameter Comparison

We consider the primal deep controlled 2BSDE algorithm for the unconstrained Merton problem, with $T = 1.0$ and $m = 10$ unless stated otherwise, and the other coefficients chosen as in Example 3.8. These tests were performed on an HP ZBook Studio G4 with a 2.9 GHz Intel i7 CPU, with no GPU acceleration.

Table 6 Accuracy and runtime of the primal deep controlled 2BSDE method using different activation functions, applied to the unconstrained power utility problem

Name	Function	Rel - Err (%)	Time (s)
ReLu	$x \mapsto \max(0, x)$	3.57771E-03	352.7
softplus	$x \mapsto \log(1 + e^x)$	3.57250E-03	352.0
tanh	$x \mapsto \tanh(x)$	3.44745E-03	352.8
sigmoid	$x \mapsto \frac{e^x}{e^x + 1}$	3.53815E-03	352.0

The first comparison we consider is the activation function used within the neural networks (c.f. function h in the description of the algorithm). Table 6 shows the results of this test, where we take $N = 20$. The choice of activation function is not important for this particular problem. All the functions are valid for this algorithm, since they are continuous functions that are not polynomials, see Leshno et al. (1993).

The second comparison we make is the optimisation method we use to minimise the loss function. We compare ADAM to standard gradient descent, the momentum optimiser (with momentum parameter 0.9), and the ADAGRAD optimiser, which are described in Ruder (2016). Table 7 shows the results of this test. ADAM is the best optimiser for this case and suggests also why it is the most popular in the literature. However, there is some time loss for using this method, compared to the other optimisers. A reason for the much higher convergence rate for the ADAM algorithm may be that it is better suited to a higher learning rate. Therefore, to get the same convergence results using the other optimisation algorithms one may need to run the other algorithms for longer, with a lower rate, which would cancel out the time loss.

Now we compare the running time and relative error given changes in the dimensionality of the problem. We compare dimensions $m \in \{1, 5, 10\}$ of the stock market, with numbers of time steps $N \in \{10, 20, 30, 40, 50\}$. Figure 5 shows the results of this test. The first graph shows the error of approximation against N^{-1} , and the second shows runtime against N^2 . We see the order $\mathcal{O}(N^{-1})$ convergence for each dimension, but the error increases with dimension. Therefore, to achieve the same accuracy, a larger number of time steps would be needed for a higher dimension. The dimensionality complexity of this algorithm is hence tied to the complexity with respect to the number of time steps N , which looks from the graph to be of order $\mathcal{O}(N^2)$. The time difference between the three dimensions is minimal. However, this is for a set number of iteration steps, and for even higher dimensions more steps may be needed, which would increase the runtime compared to lower dimensional models.

Finally, we consider the effect of the terminal time T on the relative error of the output. Since nothing is changing in the algorithm, we have found that changing T while keeping the number of time steps N fixed does not change the runtime of the algorithm. However, Fig. 6 shows an increasing error with T , as it increases from 0.2 to 1.0. Both graphs show the same data, the left side using a log scale for the y axis, and the right side squaring the x

Table 7 Accuracy and runtime of the primal deep controlled 2BSDE method using different optimisation algorithms, applied to the unconstrained power utility problem

Optimiser	Rel - Err (%)	Runtime (s)
Gradient Descent	1.00474E-01	358.3
ADAM	3.54173E-03	384.6
ADAGRAD	6.84528E-02	334.8
Momentum	5.16220E-02	338.1

Fig. 5 Accuracy and runtime of the primal deep controlled 2BSDE method applied to unconstrained power utility problems with different dimensions

axis, to show the order $\mathcal{O}(T^2)$ error implied by the data. We display this data using multiple time steps, from 5 to 20. It is notable here that even increasing the number of time steps by the same factor does not compensate for the error caused by increasing the terminal time. For example, the relative error when we use $(T, N) = (0.2, 5)$ is 0.05%, whereas when we use $(T, N) = (0.4, 10)$ it is 0.12%, which is just over double the error. This $\mathcal{O}(T^2)$ vs $\mathcal{O}(N^{-1})$ error may be why an assumption of T being ‘small’ is used in the convergence analysis of the first order BSDE solver, see Han and Long (2020).

Finally we consider the impact of the changes discussed in Remarks 3.4 and 3.6 for a unconstrained power utility problem with $T = 0.5$, $m = 10$ and $N = 20$. The results are shown in Table 8. We give the run time and accuracy for our method (the first row) and the two alternatives. The second row compares the accuracy and run time of using a single neural network to define the control and unknown BSDE processes, against using a separate one for each time step. We see that the error has slightly increased, but mainly the runtime increases when we use a single net. This is due to the additional time needed to take derivatives on the neural network parameters, even though there are less of them. The last row compares the accuracy and runtime if we optimise both networks using a single loss function

Fig. 6 Accuracy of the primal deep controlled 2BSDE method using different terminal times and numbers of time steps, applied to the unconstrained power utility problem. Both graphs display the same data, but the left uses a log scale for the y axis, and the right a squared scale for the x axis

Table 8 Relative error and runtime of variations of the primal DC2BSDE algorithm for the unconstrained Power Utility problem

Method	Rel-Err (%)	Time (s)
Normal	6.31965E-03	334
“One Net”	1.93321E-02	1080
“One Loss”	3.44806E-01	1689

$$\mathcal{L}(\Theta_{\text{BSDE}}^0, (\theta_i^0)_{i=0}^{N-1}) = \mathcal{L}_1(\Theta_{\text{BSDE}}^0) + \sum_{i=0}^{N-1} \mathcal{L}_2(\theta_i^0, i),$$

which is a linear combination of the original two loss functions (14) and (15). We see a much higher runtime and a lower accuracy.

5.2 Comparison to Literature

Now we compare the performance of our algorithm to two existing in the literature. We first compare to the algorithm of Han and W. (2016) where the primal gains function is optimised directly. This is not explicitly given a name by the authors, so we refer to it as the PVO (Primal Value Optimisation) algorithm. Secondly we compare to the Hybrid-Now algorithm of Huré et al. (2021) (referred to in the sequel as Hybrid), where the control is optimised using an approximation of the value function, which is also found using neural networks. First, we note that while both algorithms search for a Markovian control, and as such solve Markovian control problems, they are potentially less restrictive than our algorithm. The maximum requirement for these algorithms is satisfaction of the dynamic programming principle, as opposed to our derivation of the algorithm relying on the HJB equation. However it may be possible that our algorithm works in such a general setting as well (and resulting approximate solutions to the Hamiltonian PDE become viscosity solutions). The second algorithm is formalised in terms of successive approximations of the value function working backwards from terminal time in a discretisation. This method requires suitable knowledge of the distribution of the optimal state process at each time step, which we do not know a priori. Instead, we implement the algorithm in a similar vein to our own, generating sample paths of the state and control process and solving each optimisation problem [Bachouch et al. (2018), (2.1)] simultaneously. We run each algorithm for a fixed number of steps of gradient descent. This may not be the best way to compare the algorithms since they may converge after a different number of steps, but in practise it is difficult to identify exactly when this happens. We only compare to the primal value approximations of our DC2BSDE and DSMP methods, though it would of course be possible to use PVO and Hybrid to solve the dual problem.

Example 5.1 (Power Utility in Black Scholes Model). Consider the utility maximisation problem as of Example 3.8, but we take the power ($p = \frac{1}{2}$) utility $U(x) = 2\sqrt{x}$. We fix the terminal time $T = 0.5$ but change the stock market dimension m and the number of time steps N . The other parameters remain the same as in Example 3.8. Table 9 shows the results of this comparison. The solutions can be found (as is found in Example 3.9 but with no constraints) to be 2.04864 and 2.09756 for dimensions 5 and 20 respectively. For our methods we ran the DC2BSDE with $m = 20$ for 10000 steps, and 5000 steps for the $m = 5$ case and the DSMP.

Table 9 Comparison of the Primal DC2BSDE and (primal value approximation of the) DSMP algorithm to the PVO and Hybrid algorithms for the unconstrained Power Utility problem

m	N	DC2BSDE Value	Rel-Err (%)	Time (s)	DSMP Value	Rel-Err (%)	Time (s)
5	5	2.0483	1.4829E-02	54	2.0465	1.0476E-01	59
5	10	2.0485	9.2087E-03	128	2.0505	9.0012E-02	148
5	20	2.0486	4.4761E-03	309	2.0476	5.3657E-02	358
20	5	2.0942	1.6192E-01	92	2.0906	3.3063E-01	66
20	10	2.0958	8.2069E-02	223	2.0947	1.3682E-01	167
20	20	2.0961	6.5495E-02	538	2.0963	5.8714E-02	412
m	N	PVO Value	Rel-Err (%)	Time (s)	Hybrid Value	Rel-Err (%)	Time (s)
5	5	2.0497	4.9400E-02	79	2.0441	2.2415E-01	328
5	10	2.0478	4.3425E-02	266	2.0457	1.4347E-01	479
5	20	2.0483	1.8711E-02	950	2.0467	9.5190E-02	1630
20	5	2.0952	1.1184E-01	233	2.0431	2.5977E+00	343
20	10	2.0963	5.9899E-02	813	2.0536	2.0940E+00	813
20	20	2.0969	3.0629E-02	1824	2.0604	1.7720E+00	1747

The comparison to the PVO algorithm indicates that our method sacrifices slight accuracy, but scales much better with number of time steps and dimension. Indeed, when $m = 20$ the lower accuracy can be compensated by doubling the number of time steps without reaching the runtime of the PVO algorithm. We ran the PVO algorithm with a learning rate of 0.01 for 10000 and 20000 steps for $m = 5, 20$ respectively. The results for the Hybrid algorithm originate from our own implementation of the method. There may be a number of differences from the original implementation which could lead to less accurate results than expected. For example, we are not sure how exactly the authors used the ‘explore first, exploit later’ technique [Huré et al. (2021), Sect. 3.3] to generate sample paths of the optimal state process without first determining a control. In Bachouch et al. (2018), a 100 dimensional PDE (that is, $d = 100$ using the notation of Sect. 3) is solved with an accuracy of order 0.1% and a runtime of 2 - 4 hours, using $N = 40$ time steps and a terminal time of $T = 1.0$. This is much better than our preliminary results for this algorithm. However, we did not see any convergence like this in our implementation, in any time frame that was comparable to our methods. The hybrid algorithm is the fastest per iteration step, however we ran the algorithm for 50000 steps with a learning rate of 0.0001 to avoid instability issues, which is the cause of the large runtime.

The example we consider is a toy example with an analytical solution. In particular, the supremum term inside the HJB Eq. (7) can be written exactly and for this problem we are left with the nonlinear PDE

$$\frac{d}{dt}u(t, x) = \frac{1}{2}|\theta(t)|^2 \frac{D_x u(t, x)^2}{D_{xx} u(t, x)} - rx D_x u(t, x).$$

This PDE is control independent. In particular it is independent of the control dimension m apart from the negligible time taken to calculate $\theta(t)$. The PDE can be solved directly using

Table 10 Value, error and runtime of the Deep 2BSDE method for the unconstrained Power Utility problem

N	Value	Rel-Err(%)	Time (s)
5	1.96570E+00	6.28613E+00	363
10	2.05429E+00	2.06281E+00	865
20	2.06381E+00	1.60905E+00	1980

the Deep 2BSDE method introduced in Beck et al. (2019b). The results of solving this for $m = 20$ are given in Table 10. Due to the second derivative in the denominator, this PDE is unstable. In the algorithm we require 50000 iteration steps to see convergence results. Therefore, while this algorithm is faster per iteration step, it is slower overall. However its runtime would be comparable to our algorithm for larger value of m . It is also an order of magnitude less accurate, and can only be used in the special case where the supremum term can be written out as an explicit function of t and x .

The primal HJB equation, even in this simple case, can be very difficult to solve numerically. This further emphasises the utility of the algorithms presented in this paper, solving the equation by either finding the control without nonlinear dependence on the value function derivatives, or by solving the dual problem.

6 Conclusion

In this paper we study a variety of constrained utility maximisation problems, including some that are incomplete or non-Markovian, and introduced two algorithms, called the deep controlled 2BSDE method and the deep SMP method, for solving such problems, which involve simulating a system of BSDEs of first or second order. We use these equations, together with the optimality conditions for the optimal control, to find the value function of the problem. We apply these algorithms to a range of utility functions and closed convex control constraints, and find highly accurate solutions, with order $\mathcal{O}(N^{-1})$ convergence in terms of the number of time steps of the discretisation. We derive tight bounds of the value function at time 0 using the convex duality method that results in a primal problem and a dual problem, both can be solved with our algorithms. There remain many interesting open questions such as convergence analysis, duality theory for multidimensional state processes, the efficiency of the algorithm, etc. We leave these questions and others for future research.

Appendix

Proof of Theorem 3.2

Proof of 3.2 Let α denote the optimal control process, and $\mathcal{X}$ the corresponding state process. To make notation simpler, without loss of generality let $d = 1$, so the state space is one dimensional, and we let Σ be a vector in $\mathbb{R}^n$, rather than a matrix in $\mathbb{R}^{1 \times n}$. Define new processes on $[0, T]$ by

$$V(t) := u(t, \mathcal{X}(t)), \quad Z(t) := \frac{\partial}{\partial x} u(t, \mathcal{X}(t)), \quad \Gamma(t) := \frac{\partial^2}{\partial x^2} u(t, \mathcal{X}(t)). \quad (30)$$

By the regularity assumptions on u , these processes are continuous. Furthermore by construction (11) is satisfied. Applying Itô's Lemma to u and using the HJB equation, we get

$$\begin{aligned} dV(t) &= Z(t)d\mathcal{X}(t) + \frac{1}{2}\Gamma(t)d\langle \mathcal{X}, \mathcal{X} \rangle(t) + \frac{\partial}{\partial t} u(t, \mathcal{X}(t))dt \\ &= \left[F(t, \mathcal{X}(t), \alpha(t), Z(t), \Gamma(t)) - \sup_{a \in A} F(t, \mathcal{X}(t), a, Z(t), \Gamma(t)) \right] dt + Z(t)\Sigma(t, \mathcal{X}(t), \alpha(t))^\top dW(t), \\ &= Z(t)\Sigma(t, \mathcal{X}(t), \alpha(t))^\top dW(t), \end{aligned}$$

with the terminal condition $V(T) = g(\mathcal{X}(T))$. We assume u is three times continuously differentiable in x , so we can similarly apply Itô's lemma to $\frac{\partial}{\partial x} u$. Then

$$dZ(t) = \Gamma(t)^\top d\mathcal{X}(t) + \frac{1}{2} \frac{\partial^3}{\partial x^3} u(t, \mathcal{X}(t)) d\langle \mathcal{X}, \mathcal{X} \rangle(t) + \frac{\partial^2}{\partial x \partial t} u(t, \mathcal{X}(t)) dt$$

with terminal condition is $Z(T) = g'(\mathcal{X}(T))$, where

$$\begin{aligned} \frac{\partial^2}{\partial x \partial t} u(t, \mathcal{X}(t)) &= - \frac{\partial}{\partial x} \sup_a F(t, \mathcal{X}(t), a, Z(t), \Gamma(t)) \\ &= - \frac{\partial}{\partial x} \left[b(t, \mathcal{X}(t), \alpha(t))^\top \frac{\partial}{\partial x} u(t, \mathcal{X}(t)) + \frac{1}{2} \text{tr} \left(\frac{\partial^2}{\partial x^2} u(t, \mathcal{X}(t)) \Sigma(t, \mathcal{X}(t), \alpha(t))^\top \Sigma(t, \mathcal{X}(t), \alpha(t)) \right) \right] \\ &= - \left(\frac{\partial}{\partial x} b(t, \mathcal{X}(t), \alpha(t)) \frac{\partial}{\partial x} u(t, \mathcal{X}(t)) \right) - b(t, \mathcal{X}(t), \alpha(t)) \frac{\partial^2}{\partial x^2} u(t, \mathcal{X}(t)) \\ &\quad - \frac{\partial^2}{\partial x^2} u(t, \mathcal{X}(t)) \Sigma(t, \mathcal{X}(t), \alpha(t))^\top \frac{\partial}{\partial x} \Sigma(t, \mathcal{X}(t), \alpha(t)) - \frac{1}{2} |\Sigma(t, \mathcal{X}(t), \alpha(t))|^2 \frac{\partial^3}{\partial x^3} u(t, \mathcal{X}(t)) \end{aligned}$$

Cancellation then yields

$$\begin{aligned} dZ(t) &= - \left(Z(t) \frac{\partial}{\partial x} b(t, \mathcal{X}(t), \alpha(t)) + \Gamma(t) \Sigma(t, \mathcal{X}(t), \alpha(t))^\top \frac{\partial}{\partial x} \Sigma(t, \mathcal{X}(t), \alpha(t)) \right) dt \\ &\quad + \Gamma(t) \Sigma(t, \mathcal{X}(t), \alpha(t))^\top dW(t) \end{aligned}$$

as required.

Algorithm

Algorithm 1: Primal Deep Controlled 2BSDE (for maximisation)

Input: $x \in \mathbb{R}^d$
Result: $u(0, x)$
 Initialise $V_0^0 \sim \text{Unif}([-0.1, 0.1])$ and $Z_0^0 \sim \text{Unif}([-0.1, 0.1]^d)$;
for $i = 0, 1, 2, \dots, N - 1$ **do**
 | Initialise $\theta_i^0 \sim \text{Unif}([-0.1, 0.1]^{\rho(d, m)})$, $\lambda_i^0 \sim \text{Unif}([-0.1, 0.1]^{\rho(d, d^2)})$;
end
 Set $\Theta_1^0 = \{V_0^0, Z_0^0\} \cup \{\lambda_i^0\}_{i=0}^{N-1}$;
 Set $\Theta_2^0 = \{\theta_i^0\}_{i=0}^{N-1}$;
 Initialise $m_0 = v_0 = 0, j = 0$ and $m_0^i = v_0^i = 0$ for $i = 0, \dots, N - 1$;
while Θ_1^j, Θ_2^j not converged **do**
 Generate W_k for $k = 1, \dots, B$;
 Generate $\mathcal{X}^k, V^k, Z^k, \alpha^k, \Gamma^k$ using (3.9) and (Θ_1^j, Θ_2^j) ;
 Set $L_1^j(\Theta) = \frac{1}{B} \sum_{k=1}^B |V_N^k - g(\mathcal{X}_N^k)|^2 + \frac{1}{2} |Z_N^k - D_x g(\mathcal{X}_N^k)|^2$;
 Set $\Theta_1^{j+1}, m_{j+1}^1, v_{j+1}^1 = \text{ADAM}(\Theta_1^j, L_1^j, m_j^1, v_j^1, j)$;
 Regenerate $\mathcal{X}^k, V^k, Z^k, \alpha^k, \Gamma^k$ using (3.9) and $(\Theta_1^{j+1}, \Theta_2^j)$;
 for $i = 0, 1, 2, \dots, N - 1$ **do**
 Set $L_2^j(\theta, i) = -\frac{1}{B} \sum_{k=1}^B F(t_i, \mathcal{X}_i^k, \alpha_i^k, Z_i^k, \Gamma_i^k)$;
 Set $\theta_i^{j+1}, m_{j+1}^i, v_{j+1}^i = \text{ADAM}(\theta_i^{j+1}, L_2^j(\cdot, i), m_j^i, v_j^i, j)$;
 end
 Set $\Theta_2^{j+1} = \{\theta_i^{j+1}\}_{i=0}^{N-1}$;
 Set $j = j + 1$;
end
 Output V_0^j ;

Acknowledgements The authors are grateful to two anonymous reviewers whose constructive comments and suggestions have helped to improve the paper of the previous version.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anthony M, Bartlett PL (2009) Neural Network Learning: Theoretical Foundations. Cambridge University Press
- Bach F (2017) Breaking the curse of dimensionality with convex neural networks. *J Mach Learn Res* 18(1):629–681
- Bachouch A, Huré C, Langrené N, Pham H (in press) Deep neural networks algorithms for stochastic control problems on finite horizon: numerical applications. *Methodol Comput Appl Probab*

- Beck C, Becker S, Cheridito P, Jentzen A, Neufeld A (2019a) Deep splitting method for parabolic PDEs. [arXiv:1907.03452](#)
- Beck CEW, Jentzen A (2019b) Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science* 29(4):1563–1619
- Becker S, Cheridito P, Jentzen A (2019) Deep optimal stopping. *J Mach Learn Res* 20(74):1–25
- Buehler H, Gonon L, Teichmann J, Wood B (2019) Deep hedging. *Quantitative Finance* 19(8):1271–1291
- De Spiegeleer J, Madan DB, Reyners S, Schoutens W (2018) Machine learning for quantitative finance: Fast derivative pricing, hedging and fitting. *Quantitative Finance* 18(10):1635–1643
- Eckstein S, Kupper M (in press) Computation of optimal transport and related hedging problems via penalization and neural networks. *Appl Math Optim*
- Exarchos I, Theodorou EA (2018) Stochastic optimal control via forward and backward stochastic differential equations and importance sampling. *Automatica* 87:159–165
- Gonon L, Muhle-Karbe J, Shi X (2020) Asset pricing with general transaction costs: Theory and numerics. SSRN 3387232
- Grohs P, Hornung F, Jentzen A, Von Wurstemberger P (2018) A proof that artificial neural networks overcome the curse of dimensionality in the numerical approximation of black-scholes partial differential equations. [arXiv:1809.02362](#)
- Han J, W. E (2016) Deep learning approximation for stochastic control problems. [arXiv:1611.07422](#)
- Han J, Long J (2020) Convergence of the deep BSDE method for coupled FBSDEs. *Probability, Uncertainty and Quantitative Risk* 5(1):1–33
- Henry-Labordere P (2017) Deep primal-dual algorithm for BSDEs: Applications of machine learning to CVA and IM. SSRN 3071506
- Hur C, Pham H, Bachouch A, Langren N (2021) Deep neural networks algorithms for stochastic control problems on finite horizon: convergence analysis. *SIAM J Numer Anal* 59(1):525–557
- Ioffe S, Szegedy C (2015) Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proceedings of the 32nd International Conference on Machine Learning (ICML)* 37:448–456
- Karatzas I, Shreve SE (1998) *Methods of Mathematical Finance*. Springer
- Karoui NE, Tan X (2013) Capacities, measurable selection and dynamic programming part I: Abstract framework. [arXiv preprint arXiv:1310.3363](#)
- Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. [arXiv:1412.6980](#)
- Leshno M, Lin VY, Pinkus A, Schocken S (1993) Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Netw* 6(6):861–867
- Li Y, Zheng H (2018) Dynamic convex duality in constrained utility maximization. *Stochastics* 90(8):1145–1169
- Liu S, Oosterlee CW, Bohte SM (2019) Pricing options and computing implied volatilities using neural networks. *Risks* 7(1):16
- Ma J, Li W, Zheng H (2020) Dual control monte-carlo method for tight bounds of value function under heston stochastic volatility model. *Eur J Oper Res* 280(2):428–440
- Pereira M, Wang Z, Exarchos I, Theodorou EA (2019) Neural network architectures for stochastic control using the nonlinear Feynman-Kac lemma. [arXiv:1902.03986](#)
- Pham H (2009) *Continuous-Time Stochastic Control and Optimization with Financial Applications*. Springer
- Possamai D, Tan X, Zhou C (2018) Stochastic control for a class of nonlinear kernels and applications. *Ann Probab* 46(1):551–603
- Ruder S (2016) An overview of gradient descent optimization algorithms. [arXiv:1609.04747](#)
- W. E, Han J, Jentzen A (2017) Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics* 5(4):349–380
- Wang Z, Pereira MA, Theodorou EA (2019) Deep 2FBSDEs for systems with control multiplicative noise. [arXiv:1906.04762](#)
- Zhang J (2017) *Backward Stochastic Differential Equations*. Springer