



null

Calypso Herrera, Florian Krach, Pierre Ruysen, Josef

Teichmann

2023/02/02



null

abstract

This paper presents new machine learning approaches to approximate the solutions of optimal stopping problems. The key idea of these methods is to use neural networks, where the parameters of the hidden layers are generated randomly and only the last layer is trained, in order to approximate the continuation value. Our approaches are applicable to high dimensional problems where the existing approaches become increasingly impractical. In addition, since our approaches can be optimized using simple linear regression, they are easy to implement and theoretical guarantees are provided. Our randomized reinforcement learning approach and randomized recurrent neural network approach outperform the state-of-the-art and other relevant machine learning approaches in Markovian and non-Markovian examples, respectively. In particular, we test our approaches on Black-Scholes, Heston, rough Heston and fractional Brownian motion. Moreover, we show that they can also be used to efficiently compute Greeks of American options.

Keywords: optimal stopping, least squares Monte Carlo, reinforcement learning, randomized neural networks, reservoir computing, Greeks of American options

Introduction

The optimal stopping problem consists in finding the optimal time to stop in order to maximize an expected reward. This problem is found in areas of statistics, economics, and in financial mathematics. Despite significant advances, it remains one of the most challenging problems in optimization, in particular when more than one factor affects the expected reward. A common provable and widely used approach is based on Monte Carlo simulations, where the stopping decision is estimated via backward induction (Tsitsiklis and Van Roy, 2001; Longstaff and Schwartz, 2001), which is an (approximate) dynamic programming approach. Another provable approach is based on reinforcement learning (RL) (Tsitsiklis and Van Roy, 1997, 2001; Yu and Bertsekas, 2007; Li et al., 2009; Chen et al., 2020). Both approaches are based on the ordinary least squares approximation which involves choosing basis functions. There are many different sets of basis functions available that are commonly used, however, it can be difficult to choose a good set for the considered problem. Moreover, the number of basis functions often increases polynomially or even exponentially (Longstaff and Schwartz, 2001, Section 2.2) in the dimension of the underlying process, making those algorithms impractical for high dimensions.

A relatively new approach consists in replacing the basis functions by a neural network and performing gradient descent instead of ordinary least squares (Kohler et al., 2010; Lapeyre and Lelong, 2021;

Becker et al., 2019, 2020). The big advantage is that the basis functions do not need to be chosen but are learned instead. Compared to using a polynomial basis, neural networks have the advantage to be dense in any space $L^p(\mu)$, for $1 \leq p < \infty$ and finite measure μ (Hornik, 1991), while for polynomials this is only true under certain additional conditions on the measure (Bakan, 2008). Moreover, in many cases the neural network overcomes the curse of dimensionality, which means that it can easily scale to high dimensions. However, as the neural network is a non convex function with respect to its parameters, the gradient descent does not necessarily converge to the global minimum, while this is the case for the ordinary least squares minimization. Hence, the main disadvantage of those methods is that there are no convergence guarantees without strong and unrealistic assumptions.

In this paper, we propose two neural network based algorithms to solve the optimal stopping problem for Markovian settings: a backward induction and a reinforcement learning approach. The idea is inspired by randomized neural networks (Cao et al., 2018; Huang et al., 2006). Instead of learning the parameters of all layers of the neural network, those of the hidden layers are randomly chosen and fixed and only the parameters of the last layer are learned. Hence, the non convex optimization problem is reduced to a convex problem that can be solved with linear regression. The hidden layers form random feature maps, which can be interpreted as random basis functions. In particular, in this paper we show that there is actually no need for complicated or a large number of basis functions. Our algorithms are based on the methods proposed by Longstaff and Schwartz (2001) (backward induction approach) and Tsitsiklis and Van Roy (2001) (reinforcement learning approach). The difference is that we use a randomized neural network instead of a linear combination of basis functions. However, a randomized neural network can also be interpreted as a linear combination of random basis functions. On the other hand, our algorithms can also be interpreted as the neural network extensions of these methods, where not the entire neural network but only the last layer is trained.

In addition, we provide a randomized recurrent neural network approach for nonMarkovian settings. We compare our algorithms to the most relevant baselines in terms of accuracy and computational speed in different option pricing problems. With only a fraction of trainable parameters compared to existing methods, we achieve high quality results considerably faster.

Optimal Stopping via Randomized Neural Networks

One of the most popular and most studied applications of optimal stopping is the pricing of American options. Hence, we explain our approach in this context.

American and Bermudan Options

An American option gives the holder the right but not the obligation to exercise the option associated with a non-negative payoff function g at any time up to the maturity. An American option can be approximated by a Bermudan option, which can be exercised only at some specific dates $t_0 < t_1 < t_2 < \dots < t_N$, transforming the continuous-time problem to a discrete one. If the time grid is chosen small enough, the American option is well approximated by the Bermudan option. In the case of a Rough Heston model, the convergence rate of the Bermudan option price to the American option price was shown in (Chevalier et al., 2021, Theorem 4.2). For equidistant dates we simply write $0, 1, 2, \dots, N$ instead of $t_0 < t_1 < t_2 < \dots < t_N$.

Option Price and Optimal Stopping.

For $d \in \mathbb{N}$, we assume to have a d -dimensional Markovian stochastic process $(X_t)_{t \geq 0}$ describing the stock prices. With respect to a fixed (pricing) probability measure \mathbb{Q} , the (superhedging seller's) price of the discretized American option can be expressed through the Snell envelope described by

$$\begin{aligned} U_N &:= g(X_N), \\ U_n &:= \max(g(X_n), \mathbb{E}[\alpha U_{n+1} | X_n]), \quad 0 \leq n < N, \end{aligned}$$

where α is the step-wise discounting factor and $g(X_n)$ is assumed to be square integrable for all n . Then the (superhedging seller's) price of the option at time n is given by U_n and can equivalently be expressed as the optimal stopping problem

$$U_n = \sup_{\tau \in \mathcal{T}_n} \mathbb{E}[\alpha^{\tau-n} g(X_\tau) | X_n],$$

where \mathcal{T}_n is the set of all stopping times $\tau \geq n$. The smallest optimal stopping time is given by

$$\begin{aligned} \tau_N &:= N, \\ \tau_n &:= \begin{cases} n, & \text{if } g(X_n) \geq \mathbb{E}[\alpha U_{n+1} | X_n], \\ \tau_{n+1}, & \text{otherwise.} \end{cases} \end{aligned}$$

In particular, at maturity N , the holder receives the final payoff, and the value of the option U_N is equal to the payoff $g(X_N)$. At each time prior to the maturity, the holder decides whether to exercise or not, depending on whether the current payoff $g(X_n)$ is greater than the continuation value $c_n(X_n) := \mathbb{E}[\alpha U_{n+1} | X_n]$. Combining expression (1), (2) and (3), we can write the price at initial time as

$$U_0 = \max (g (X_0), \mathbb{E} [\alpha^{\tau_1} g (X_{\tau_1})]).$$

In the following we approximate the price U_0 and continuation values $c_n (X_n)$ which are defined theoretically but cannot be computed directly.

Monte Carlo Simulation and Backward Recursion

We assume to have access to a procedure to sample discrete paths of X under \mathbb{Q} . A standard example is that X follows a certain stochastic differential equation (SDE) with known parameters. Therefore, we can sample m realizations of the stock price paths, where the i -th realization is denoted by $x_0, x_1^i, x_2^i, \dots, x_N^i$, with the fixed initial value x_0 . For each realization, the cash flow realized by the holder when following the stopping strategy (3) is given by the backward recursion

$$p_N^i := g (x_N^i),$$

$$p_n^i := \begin{cases} g (x_n^i), & \text{if } g (x_n^i) \geq c_n (x_n^i) \\ \alpha p_{n+1}^i, & \text{otherwise} \end{cases}$$

As p_1^i are samples of $\alpha^{\tau_1-1} g (X_{\tau_1})$, we have by the strong law of large numbers that almost surely

$$U_0 = \max \left(g (X_0), \lim_{m \rightarrow \infty} \frac{1}{m} \sum_{i=1}^m \alpha p_1^i \right)$$

Randomized Neural Network Approximation of the Continuation Value

For each path i in $\{1, 2, \dots, m\}$ and each date n in $\{1, 2, \dots, N-1\}$, the continuation value is $c_n (x_n^i) = \mathbb{E} [\alpha U_{n+1} | X_n = x_n^i]$, where $c_n : \mathbb{R}^d \rightarrow \mathbb{R}$ describes the expected value of the discounted price αU_{n+1} if we keep the option until next exercising date $n+1$, knowing the current values of the stocks X_n . We approximate this continuation value function by a neural network, where only the parameters of the last layer are learned. We refer to such a network as a randomized neural network. Even though the architecture of the neural network can be general, we present our algorithm with a simple dense shallow neural network, where the extension to deep networks is immediate. We call $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ the activation function. A common choice is $\sigma(x) = \tanh(x)$, however, there are many other suitable alternatives. For $K \in \mathbb{N}$, we define $\sigma : \mathbb{R}^{K-1} \rightarrow \mathbb{R}^{K-1}$, $\sigma(x) = (\sigma(x_1), \dots, \sigma(x_{K-1}))^\top$ for $x \in \mathbb{R}^{K-1}$. Let $\vartheta := (A, b) \in \mathbb{R}^{(K-1) \times d} \times \mathbb{R}^{K-1}$ be the parameters of the hidden layer which are randomly and identically sampled and not optimized. In general, A and b can be sampled from different distributions that are continuous and cover all of \mathbb{R} . The distributions and their parameters are hyperparameters

of the randomized neural network that can be tuned. For simplicity we use a standard Gaussian distribution. Let us define

$$\phi : \mathbb{R}^d \rightarrow \mathbb{R}^K, x \mapsto \phi(x) = \left(\sigma(Ax + b)^\top, 1 \right)^\top.$$

and let $\theta_n := \left((A_n)^\top, b_n \right) \in \mathbb{R}^{K-1} \times \mathbb{R}$ be the parameters that are optimized. Then for each n the continuation value is approximated by

$$c_{\theta_n}(x) := \theta_n^\top \phi(x) = A_n^\top \sigma(Ax + b) + b_n.$$

Least Squares Optimization of Last Layer's Parameters θ_n

While the parameters ϑ of the hidden layer are set randomly, the parameters θ_n of the last layer are found by minimizing the squared error of the difference between conditional expectation of the discounted future price and the approximation function. This is equivalent to finding θ_n which minimizes $\mathbb{E}_{\mathbb{Q}} \left[(c_{\theta_n}(x_n^i, n) - \alpha U_{n+1})^2 \mid X_n = x_n^i \right]$ for each time n in $\{1, 2, \dots, N-1\}$. The backward recursive Monte Carlo approximation of this expectation at time n yields the loss function

$$\psi_n(\theta_n) = \sum_{i=1}^m \left(c_{\theta_n}(x_n^i) - \alpha p_{n+1}^i \right)^2.$$

As the approximation function c_{θ_n} is linear in the parameters θ_n , the minimizer can be found by ordinary least squares. It is given by the following closed form expression, which is well defined under the standard assumptions (see Theorems 5 and 6)

$$\theta_n = \alpha \left(\sum_{i=1}^m \phi(x_n^i) \phi^\top(x_n^i) \right)^{-1} \cdot \left(\sum_{i=1}^m \phi(x_n^i) p_{n+1}^i \right).$$

Splitting the Data Set into Training and Evaluation Set

The parameters θ_n are determined using 50% of the sampled paths (training data). Given θ_n , the remaining 50% of the sampled paths (evaluation data) are used to compute the option price. By definition, the continuation value c_n is a conditional expectation, which is not allowed to depend on the future values X_{n+1} . If the data set was not split, this might not be satisfied, since the loss function (5) uses the future values X_{n+1} . In particular, the neural network can suffer from overfitting to the training data, by memorizing the paths, instead of learning the continuation value. This is related to the maximization bias discussed in (Sutton and Barto, 2018, Section 6.7).

Algorithm

We first sample $2m$ paths and then proceed backwards as follows. At maturity, the pathwise option price approximation is equal to the payoff, which means that $p_N^i = g(x_N^i)$. For each time n in $\{N-1, N-2, \dots, 0\}$, we first determine θ_n as described before using the paths $\{1, 2, \dots, m\}$. For all paths $i \in \{1, 2, \dots, 2m\}$ we then compare the exercise value $g(x_n^i)$ to the continuation value $c_{\theta_n}(x_n^i)$ and determine the path-wise option price approximation at time n as

$$p_n^i = \underbrace{g(x_n^i)}_{\text{payoff}} \underbrace{\mathbf{1}_{\{g(x_n^i) \geq c_{\theta_n}(x_n^i)\}}}_{\text{exercise}} + \underbrace{\alpha p_{n+1}^i}_{\text{discounted future price}} \underbrace{\mathbf{1}_{\{g(x_n^i) < c_{\theta_n}(x_n^i)\}}}_{\text{continue}}.$$

Finally, the second half of the paths $\{m+1, \dots, 2m\}$ is used to compute the option price approximation $p_0 = \max\left(g(x_0), \frac{1}{m} \sum_{i=m+1}^{2m} \alpha p_1^i\right)$. We call this algorithm, which is presented in Algorithm 1, randomized least squares Monte Carlo (RLSM).

Algorithm 1 Optimal stopping via randomized least squares Monte Carlo (RLSM)

Input: discount factor α , initial value x_0

Output: price p_0

- 1: sample a random matrix $A \in \mathbb{R}^{(K-1) \times d}$ and a random vector $b \in \mathbb{R}^{K-1}$
- 2: simulate $2m$ paths of the underlying process (x_1^i, \dots, x_N^i) for $i \in \{1, \dots, 2m\}$
- 3: for each path $i \in \{1, \dots, 2m\}$, set $p_N^i = g(x_N^i)$
- 4: for each time $n \in \{N-1, \dots, 1\}$
 - a: for each path $i \in \{1, \dots, 2m\}$, set $\phi(x_n^i) = \left(\sigma(Ax_n^i + b)^\top, 1\right)^\top \in \mathbb{R}^K$
 - b: set $\theta_n = \alpha \left(\sum_{i=1}^m \phi(x_n^i) \phi(x_n^i)^\top\right)^{-1} \left(\sum_{i=1}^m \phi(x_n^i) p_{n+1}^i\right)$
 - c: for each path $i \in \{1, \dots, 2m\}$, set $p_n^i = g(x_n^i) \mathbf{1}_{g(x_n^i) \geq \theta_n^\top \phi(x_n^i)} + \alpha p_{n+1}^i \mathbf{1}_{g(x_n^i) < \theta_n^\top \phi(x_n^i)}$
- 5: set $p_0 = \max\left(g(x_0), \frac{1}{m} \sum_{i=m+1}^{2m} \alpha p_1^i\right)$

Guarantees of Convergence

We present results that guarantee convergence of the price computed with our algorithm to the correct price of the discretized American option. The formal results with precise definitions and proofs are given in Section A. In contrast to comparable results for neural networks (Lapeyre and Lelong, 2021; Becker et al., 2019, 2020), our results do not need the assumption that the optimal weights are found

by some optimization scheme like stochastic gradient descent. Instead, our algorithms imply that the optimal weights are found and used.

Theorem 1 (informal). As the number of sampled paths m and the number of random basis functions K go to ∞ , the price p_0 computed with Algorithm 1 converges to the correct price of the Bermudan option.

Possible Extensions

When the set of pricing measures \mathcal{Q} has more than one element (in case of an incomplete market), the option price is given by $\sup_{\mathbb{Q} \in \mathcal{Q}} U_0^{\mathbb{Q}}$, where $U^{\mathbb{Q}}$ is defined as in (1). Assuming that we can sample from a finite subset $\mathcal{Q}_1 \subset \mathcal{Q}$, this price can be approximated by first computing the price for each measure in \mathcal{Q}_1 and taking the maximum of them.

For simplicity we assume that the payoff function only takes the current price as input, however, all our methods and results stay valid if $g(X_n)$ is replaced by a square integrable \mathcal{F}_n -measurable random variable Z_n , where \mathcal{F}_n denotes the information available up to time n . In the case that $(Z_n)_{1 \leq n \leq N}$ is Markov, Algorithm 1 and Algorithm 2 (Section 3) can be used, otherwise Algorithm 3 (Section 4) has to be used, to deal with the path dependence. In the following sections we stick to the notation $g(X_n)$ for the payoff, keeping in mind that the extension to a general Z_n is also valid there.

Optimal Stopping via Randomized Reinforcement Learning

In order to avoid approximating the continuation value at each single date $n \in \{1, \dots, N-1\}$ with a different function, as it is done in Section 2, we can directly learn the continuation function which also takes the time as argument. Hence, instead of having a different function $c_{\theta_n}(x_n^i)$ for each date n , we learn one function which is used for all dates n . As previously, we define $\vartheta := (A, b) \in \mathbb{R}^{(K-1) \times (d+2)} \times \mathbb{R}^{K-1}$ the parameters of the hidden layer which are randomly chosen and not optimized, and $\phi : \mathbb{R}^{d+2} \rightarrow \mathbb{R}^K$, $\phi(n, x) = \left(\sigma(A\tilde{x}_n + b)^\top, 1 \right)^\top$, where $\tilde{x}_n = \left(n, N-n, x_n^\top \right)^\top$. Let $\theta \in \mathbb{R}^K$ define the parameters that are optimized, then the continuation value is approximated by

$$c_\theta(n, x) = \theta^\top \phi(n, x).$$

Instead of having a loop backward in time with N steps, we iteratively improve the approximation c_θ . More precisely, we start with some (random) initial weight θ_0 and then iteratively improve it by minimizing the difference between the continuation function c_{θ_ℓ} and the prices p computed with the previous weight $\theta_{\ell-1}$. Moreover, differently than in Section 2, we use the continuation value for the decision whether to continue and for the approximation of the discounted future price, as in (Tsitsiklis

and Van Roy, 2001). This second algorithm can be interpreted as a randomized fitted Q-iteration (RFQI) and is presented in Algorithm 2. It is a very simple type of reinforcement learning, where the agent has only two possible actions and the agent's decision does not influence the transitions of the state. In particular the agent's decision does not influence the evolution of the underlying stocks. As a reinforcement learning method, it is based on the assumption that the optimization problem can be modeled by a Markov decision process.

Algorithm 2 Optimal stopping via randomized fitted Q-Iteration (RFQI)

Input: discount factor α , initial value x_0

Output: price p_0

- 1: sample a random matrix $A \in \mathbb{R}^{(K-1) \times (d+2)}$ and a random vector $b \in \mathbb{R}^{K-1}$
 - 2: simulate $2m$ paths of the underlying process (x_1^i, \dots, x_N^i) for $i \in \{1, \dots, 2m\}$
 - 3: initialize weights $\theta_0 = 0 \in \mathbb{R}^K$ and set $\ell = 0$
 - 4: until convergence of θ_ℓ
 - a: for each path $i \in \{1, \dots, 2m\}$
 - i: set $p_N^i = g(x_N^i)$
 - ii: for each date $n \in \{1, \dots, N-1\}$,
 - set $\phi(n, x_n^i) = (\sigma(A\tilde{x}_n^i + b), 1) \in \mathbb{R}^K$
 - set $p_n^i = \max(g(x_n^i), \phi(n, x_n^i)^\top \theta_\ell)$
 - b: set $\theta_{\ell+1} = \alpha \left(\sum_{n=1}^N \sum_{i=1}^m \phi(n, x_n^i) \phi^\top(n, x_n^i) \right)^{-1} \cdot \left(\sum_{n=1}^N \sum_{i=1}^m \phi(n, x_n^i) p_{n+1}^i \right) \in \mathbb{R}^K$
 - c: set $\ell \leftarrow \ell + 1$
 - 5: set $p_0 = \max\left(g(x_0), \frac{1}{m} \sum_{i=m+1}^{2m} \alpha p_1^i\right)$
-

Guarantees of Convergence.

We present results that guarantee convergence of the price computed with our algorithm to the correct price of the discretized American option. The formal results with precise definitions and proofs are given in Section B.

Theorem 2 (informal). As the number of iterations L , the number of sampled paths m and the number of random basis functions K go to ∞ , the price p_0 computed with Algorithm 2 converges to the correct price of the Bermudan option.

Optimal Stopping via Randomized Recurrent Neural Networks for Non-Markovian Processes

For non-Markovian processes, for each date n , the continuation function is no longer a function of the last stock price, $c_n(X_n)$, but a function depending on the entire history $c_n(X_0, X_1, \dots, X_{n-1}, X_n)$.

More precisely, the continuation value is now defined by $c_n := \mathbb{E}[\alpha g(X_{n+1}) \mid \mathcal{F}_n]$ where \mathcal{F}_n denotes the information available up to time n . Therefore, we replace the randomized feed-forward neural network by a randomized recurrent neural network (randomized RNN), which can utilize the entire information of the path up to the current time $(x_0, x_1, \dots, x_{n-1}, x_n)$. In particular, we define $\vartheta := (A_x, A_h, b) \in \mathbb{R}^{(K-1) \times d} \times \mathbb{R}^{(K-1) \times (K-1)} \times \mathbb{R}^{K-1}$, the parameters of the hidden layer which are randomly sampled and not optimized. However, their distributions and parameters, which don't have to be the same for A_x and A_h , are hyperparameters that can be tuned. Those tuning parameters are more important in this case, as they determine the interplay between past and new information. Moreover, we define

$$\phi : \mathbb{R}^d \times \mathbb{R}^K \rightarrow \mathbb{R}^{K+1}, \quad (x, h) \mapsto \phi(x, h) = \left(\sigma(A_x x + A_h h + b)^\top, 1 \right)^\top$$

and $\theta_n := ((A_n)^\top, b_n) \in \mathbb{R}^{K-1} \times \mathbb{R}$, the parameters that are optimized. Then for each n , the continuation value is approximated by

$$\begin{cases} h_n = \sigma(A_x x_n + A_h h_{n-1} + b) \\ c_{\theta_n}(h_n) = A_n^\top h_n + b_n = \theta_n^\top \phi(x_n, h_{n-1}) \end{cases}$$

with $h_{-1} = 0$. We call this algorithm, which is presented in Algorithm 3, randomized recurrent least squares Monte Carlo (RRLSM).

Algorithm 3 Optimal stopping via randomized recurrent neural network (RRLSM)

Input: discount factor α , initial value x_0 , initial latent variable $h_{-1} = 0$

Output: price p_0

1: sample random matrices $A_x \in \mathbb{R}^{(K-1) \times d}$, $A_h \in \mathbb{R}^{(K-1) \times (K-1)}$ and a random vector $b \in \mathbb{R}^{K-1}$

2: simulate $2m$ paths of the underlying process (x_1^i, \dots, x_N^i) for $i \in \{1, \dots, 2m\}$

3: for each path $i \in \{1, \dots, 2m\}$, set $p_N^i = g(x_N^i)$

4: for each date $n \in \{0, \dots, N-1\}$

a: for each path $i \in \{1, \dots, 2m\}$

set $h_n^i = \sigma(A_x x_n^i + A_h h_{n-1}^i + b)$

5: for each date $n \in \{N-1, \dots, 1\}$

a: for each path $i \in \{1, \dots, 2m\}$

set $\phi_n^i = ((h_n^i)^\top, 1)^\top \in \mathbb{R}^K$

b: set $\theta_n = \alpha \left(\sum_{i=1}^m \phi_n^i (\phi_n^i)^\top \right)^{-1} \left(\sum_{i=1}^m \phi_n^i p_{n+1}^i \right)$

c: for each path $i \in \{1, \dots, 2m\}$, set $p_n^i = g(x_n^i) \mathbf{1}_{g(x_n^i) \geq \theta_n^\top \phi_n^i} + \alpha p_{n+1}^i \mathbf{1}_{g(x_n^i) < \theta_n^\top \phi_n^i}$

6: set $p_0 = \max(g(x_0), \frac{1}{m} \sum_{i=m+1}^{2m} \alpha p_1^i)$

Guarantees of Convergence

We present results that guarantee convergence of the price computed with our algorithm to the correct price of the discretized American option. The formal results with precise definitions and proofs are given in Section C.

Theorem 3 (informal). As the number of sampled paths m and the number of random basis functions K go to ∞ , the price p_0 computed with Algorithm 3 converges to the correct price of the Bermudan option.

Related Work

In this section we present the most relevant approaches for the optimal stopping problem: backward induction either with basis functions or with neural networks and reinforcement learning. Moreover, we explain the connection of our algorithms to randomized neural networks and reservoir computing techniques.

Optimal Stopping

Numerous works studied the optimal stopping problem via different approaches. A common approach consists in using a regression based method to estimate the continuation value (Tilley, 1993; Barraquand and Martineau, 1995; Carriere, 1996; Tsitsiklis and Van Roy, 1997, 2001; Longstaff and Schwartz, 2001; Schweizer, 2002; Boyle et al., 2003; Broadie and Glasserman, 2004; Kolodko and Schoenmakers, 2004; Egloff et al., 2007; Jain and Oosterlee, 2015), or the optimal stopping boundary (Pham, 1997; Andersen, 1999; Garcia, 2003). A different approach uses quantization (Bally and Pagès, 2003; Bally et al., 2005). A dual approach was developed and extended in (Rogers, 2002; Haugh and Kogan, 2004; Rogers, 2010). Bank and Besslich (2019) studied Lenglart's Theory of Meyer-sigma-fields and El Karoui's Theory of Optimal Stopping (El Karoui, 1981). An in depth review of the different methods is given in (Bouchard and Warin, 2012; Pagès, 2018).

Optimal Stopping Via Backward Induction

One of the most popular approaches are the backward induction methods introduced by Tsitsiklis and Van Roy (2001) and Longstaff and Schwartz (2001). Tsitsiklis and Van Roy (2001) use the approximated continuation value to estimate the current price, by using the following backward recursion

$$p_n^i = \max \left(g \left(x_n^i \right), c_{\theta_n} \left(x_n^i \right) \right).$$

Instead, Longstaff and Schwartz (2001) use the continuation value only for the decision to stop or to continue.

$$p_n^i = \begin{cases} g(x_n^i), & \text{if } g(x_n^i) \geq c_{\theta_n}(x_n^i) \\ \alpha p_{n+1}^i, & \text{otherwise} \end{cases}$$

The second algorithm is more robust, as the approximation is only used for the decision and not for the estimation of the price. Hence, the method proposed by Longstaff and Schwartz (2001) is the most used method in the financial industry and can be considered state-of-the-art. In both papers, the approximation $c_{\theta}(x_n^i) = \theta^\top \phi(x_n^i)$ is used, where $\phi = (\phi_1, \dots, \phi_K)$ is a set of K basis functions and $\theta \in \mathbb{R}^K$ are the trainable weights. Possible choices for the basis functions proposed in Longstaff and Schwartz (2001) are Laguerre, Hermite, Legendre, Chebyshev, Gegenbauer, and Jacobi polynomials. While they have the advantage of having convergence guarantees, both algorithms do not easily scale to high dimensional problems since the number of basis functions usually grows polynomially or even exponentially (Longstaff and Schwartz, 2001, Section 2.2) in the number of stocks. One direction of research to overcome this problem is to apply dimension reduction techniques (Bayer et al., 2021).

Optimal Stopping via Backward Induction using Neural Networks.

Another idea to overcome this issue was proposed by Kohler et al. (2010), which consists in approximating the continuation value by a neural network

$$f_{\theta}(x_n^i) \approx c_{\theta}(x_n^i).$$

That way, the features are learned contrary to the basis functions which must be chosen. While Kohler et al. (2010) use the backward recursion (7) introduced by Tsitsiklis and Van Roy (2001), both Lapeyre and Lelong (2021) and Becker et al. (2020) use the backward recursion (8) suggested by Longstaff and Schwartz (2001). Instead of approximating the continuation value, Becker et al. (2019) suggested to approximate the whole indicator function presented in (8) by a neural network $f_{\theta_n}(x_n^i) \approx \mathbf{1}_{\{g(x_n^i) \geq c(x_n^i)\}}$. Therefore the current price can be estimated by

$$p_n^i = g(x_n^i) \underbrace{f_{\theta_n}(x_n^i)}_{\text{stop}} + \alpha p_{n+1}^i \underbrace{(1 - f_{\theta_n}(x_n^i))}_{\text{continue}}.$$

Moreover, instead of minimizing the loss function (5) in order to find a good approximation of the continuation function, Becker et al. (2019) optimize the parameters by directly maximizing the option price $\psi_n(\theta_n) = \frac{1}{m} \sum_{i=1}^m \alpha p_n^i$.

All those methods use a stochastic gradient based method to optimize the parameters of the neural networks. They have to find the parameters of $N - 1$ neural networks (using a different neural network for each date). Since they use stochastic gradient methods with a non-convex loss function they cannot provide theoretical convergence guarantees, without the strong assumption that they find the optimal parameters.

Optimal Stopping Via Reinforcement Learning

By its nature, reinforcement learning is closely related to the dynamic programming principle as shown in (Sutton and Barto, 2018; Bertsekas and Tsitsiklis, 1996). Moreover the optimal stopping problem is well studied as an application of reinforcement learning (Tsitsiklis and Van Roy, 1997, 2001; Yu and Bertsekas, 2007; Li et al., 2009). In all those methods, a linear approximator is used (linear combination of basis functions), similarly to the LSM method (Longstaff and Schwartz, 2001). If a standard set of basis functions that grows polynomially in the dimension is used, then these methods suffer from the curse of dimensionality. In particular, they cannot practically be scaled to high dimensions as can be seen in our numerical results. To the best of our knowledge, our approach constitutes the first time that randomized neural networks are used to approximate the value function in reinforcement learning.

Randomized Neural Networks and Reservoir Computing

In RLSM and RFQI we use a neural network with randomly sampled and fixed hidden layers, where only the last layer is reinitialized and trained at each time $n \in \{N - 1, \dots, 1\}$. The architecture used at each time can be interpreted as a neural network with random weights (NNRW) studied and reviewed in (Cao et al., 2018), where a universality result was provided in (Huang et al., 2006). Randomized neural networks as approximation functions were also studied by Gorban et al. (2016).

Randomized recurrent neural networks are an extension of randomized neural networks. A recurrent neural network (RNN) where the parameters are randomly generated and fixed and only the readout map is trained, is known as reservoir. Reservoir computing not only reduces the computation time, but also outperforms classical fully trained RNNs in many tasks (Schrauwen et al., 2007; Verstraeten et al., 2007; Lukoševičius and Jaeger, 2009; Gallicchio et al., 2017). Similarly as in reservoir computing, in our randomized recurrent neural network algorithm RRLSM the parameters of the hidden layers are randomly sampled and fixed thereafter. However, while reservoir computing trains only one readout map which has the same parameters for all times, we train a different readout map for each single time $n \in \{N - 1, \dots, 1\}$ similarly to RLSM.

Backward Induction versus Reinforcement Learning

Backward induction is an (approximate) dynamic programming (ADP) approach. While Sutton and Barto (2018) regard ADP as a class of RL algorithms, we distinguish these two approaches in this work, because of their different algorithmic structure and their different ways of using the training data. In particular, backward recursion computes the approximation of the continuation value for each date sequentially. More precisely, it starts at the final date and goes backward in time. For the approximation at each date, only the data of this date is used. In contrast to this, RL starts with an initial approximation that is applied for all dates and iteratively improves this approximation. This way, the data of all dates is used to improve the approximation of all dates. In comparison to backward recursion, this can be interpreted as a type of transfer learning between the dates.

Experiments

There are numerous ways to empirically evaluate optimal stopping approaches. Therefore, we choose the most studied settings that were considered in the American option pricing literature. In particular, we only consider synthetic data. Applications to real data involve model calibration which is an independent problem and finally results in applying the optimal stopping algorithm to synthetically generated data again.

Besides our algorithms, we also implemented the baselines and provided all of them at <https://github.com/HeKrRuTe/OptStopRandNN>.

Experimental Setup

The evaluation of all the algorithms was done on the same computer, a dedicated machine with 2× Intel Xeon CPU E5-2697 v2 (12 Cores) 2.70GHz and 256GiB of RAM.

Baselines (LSM, NLSM, DOS And FQI)

We compare RLSM and RFQI to three backward induction algorithms and one reinforcement learning approach. First, the state-of-the-art least squares Monte Carlo (LSM) (Longstaff and Schwartz, 2001). Second, the algorithm proposed by Lapeyre and Lelong (2021), where the basis functions are replaced by a deep neural network (NLSM). Third, the deep optimal stopping (DOS) (Becker et al., 2019), where instead of the continuation value the whole indicator function of the stopping decision is approximated by a neural network. Finally, the fitted Q-iteration (FQI) presented as the second algorithm in (Tsitsiklis and Van Roy, 1997). Li et al. (2009) studied and compared two reinforcement learning based methods

(FQI and LSPI) to solve the optimal stopping problem. Since FQI always worked better in our experiments, we only show comparisons to this algorithm. Our aim is to compare the main concepts of all the different algorithms in a fair way, hence we leave away certain (more sophisticated) particularities unique to each of them.

Choice of Basis Functions for the Baselines

There are many possible choices for the set of basis functions. Longstaff and Schwartz (2001) proposed to use the first three weighted Laguerre polynomials for LSM and Li et al. (2009) added three additional basis functions of the date for FQI. While the size of this set of basis functions scales linearly with the dimension, it does not include any interaction terms. The classical polynomial basis functions up to the second order are the easiest way to include coupling terms in the basis. To deal with the time dependence of FQI, the relative date t/T and $1 - t/T$ are added as additional coordinates to the d -dimensional stock vector. The size of this basis grows quadratically in the dimension d , i.e. it has $1 + 2d + d(d - 1)/2$ elements for LSM and for FQI d is replaced by $d + 2$. The results obtained with the classical polynomials up to degree two were better than with the weighted Laguerre polynomials for LSM and FQI, therefore we only present these results in our Tables. For large d the computations of LSM and FQI did not terminate within a reasonable amount of time (several hours) and therefore were aborted.

No Regularization for LSM And FQI

While increasing the number of hidden nodes without applying penalization led to overfitting for RLSM and RFQI, this was not observed for LSM and FQI. In particular, for LSM Ridge regression (L^2 -penalisation) was tested without leading to better results than standard linear regression. Moreover, comparing the results of FQI, RFQI and DOS for growing dimensions shows that overfitting does not become a problem when more basis functions are used. Therefore, also for FQI standard linear regression was used as suggested by Tsitsiklis and Van Roy (1997).

Architecture of Neural Networks

In order to have a fair comparison in terms of accuracy and in terms of computation time, we use the same number of hidden layers and nodes per layer for all the algorithms.

- As we observed that one hidden layer was sufficient to have a good accuracy (an increase of the number of the hidden layers did not lead to better accuracy), we use one hidden layer. Therefore, NLSM, DOS, and all algorithms that we proposed have only one hidden layer.

- We use 20 nodes for the hidden layer. For RFQI the number of nodes is set to the minimum between 20 and the number of stocks for stability reasons.
- Leaky ReLU is used for RLSM and RFQI and tanh for the randomized recurrent neural network RRLSM. For NLSM and DOS, we use the suggested activation functions, Leaky ReLU for NLSM and ReLU and sigmoid for DOS.
- The parameters (A, b) of the random neural networks of RLSM and RFQI are sampled using a standard normal distribution with mean 0 and standard deviation 1. Different hyper-parameters were tested, but they didn't have a big influence on the results so we kept the standard choice. - For the randomized recurrent neural network of RRLSM, we use a standard deviation of 0.0001 for A_x and 0.3 for A_h . Also here different hyper-parameters were tested, and the best performing were chosen and used to present the results. The same holds for tested path-dependent versions of RFQI, however, none of the hyper-parameters performed very well as shown below.
- Some of the reference methods suggest to use the payoff as additional input, while others do not or leave this open. Therefore, we tested using the payoff as input and not using it for each method in each experiment. We came to the conclusion that the backward induction algorithms (LSM, DOS, NLSM, RLSM) usually work slightly better with, while the reinforcement learning algorithms (FQI, RFQI) usually work slightly better without the payoff. Hence, we show these results.
- As suggested by the authors, we used batch normalization for the implementation of DOS.

The Markovian Case - Bermudan Option Pricing

First we evaluate RLSM and RFQI in the standard Markovian setting of Bermudan option pricing with different stock price models and payoff functions.

Stock Models (Black-Scholes and Heston)

We test our algorithm on two multidimensional stochastic models, Black-Scholes and Heston with fixed parameters. For each model we sample $m = 20'000$ paths on the time interval $[0, 1]$ using the Euler-scheme with $N = 10$ equidistant dates. As explained in Section 2.6, we use half of the paths as training set and the second half to compute the approximated price using the trained continuation value respectively decision function.

The Black-Scholes model for a max call option is a widely used example in the literature (Longstaff and Schwartz, 2001; Lapeyre and Lelong, 2021; Becker et al., 2019). The Stochastic Differential Equation (SDE) describing this model is

$$dX_t = (r - \delta)X_t dt + \sigma X_t dW_t$$

with $X_0 = x_0$ where $(W_t)_{t \geq 0}$ is a d -dimensional Brownian motion. If not stated differently, we choose the rate $r = 0\%$, the dividend rate $\delta = 0\%$, the volatility $\sigma = 20\%$ and the initial stock price $x_0 \in \{80, 100, 120\}$.

To increase the complexity, we also compare the algorithms on the Heston model (Heston, 1993), which is also used in (Lapeyre and Lelong, 2021). The SDE describing this model is

$$\begin{aligned} dX_t &= (r - \delta)X_t dt + \sqrt{v_t}X_t dW_t, \\ dv_t &= -\kappa(v_t - v_\infty) dt + \sigma\sqrt{v_t}dB_t \end{aligned}$$

with $X_0 = x_0$ and $v_0 = \nu_0$, where $(W_t)_{t \geq 0}$ and $(B_t)_{t \geq 0}$ are two d -dimensional Brownian motions correlated with coefficient $\rho \in (-1, 1)$. Here, X is the stock price and v the stochastic variance process. If not stated differently, we choose the drift $r = 0\%$, the dividend rate $\delta = 0\%$, the volatility of volatility $\sigma = 20\%$, the long term variance $v_\infty = 0.01$, the mean reversion speed $\kappa = 2$, the correlation $\rho = -30\%$, the initial stock price $x_0 = 100$ and the initial variance $\nu_0 = 0.01$. Since the Heston model is Markovian only if the price and the variance (X_t, v_t) are observed simultaneously, we give both values as inputs to the algorithms here, and denote this below by ‘‘Heston (with variance)’’.

Payoffs (Max Call, Geometric Put, Basket Call and Min Put)

We test our algorithms on three different types of options: the max call, the geometric put and the basket call. First, we consider the max call option as it is a classical example used in optimal stopping (Lapeyre and Lelong, 2021; Becker et al., 2019). The payoff of a max call option is defined by $g(x) = (\max(x_1, x_2, \dots, x_d) - K)_+$ for any $x = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d$. Moreover, we also consider the geometric put option as used in (Lapeyre and Lelong, 2021). The payoff of the geometric put option is defined by $g(x) = \left(K - \left(\prod_{i=1}^d x_i\right)^{1/d}\right)_+$. We also test our approach on a basket call option (Hanbali and Linders, 2019), where the payoff is given by $g(x) = \left(\frac{1}{d} \sum_{i=1}^d x_i - K\right)_+$ and a min put option with payoff $g(x) = (K - \min(x_1, x_2, \dots, x_d))_+$. For all these payoffs, the strike K is set to 100.

Reference Prices

In some cases reference prices can be computed, to which the prices computed with the different algorithms can be compared. All call options where the underlying stocks have a rate $r \geq 0$ and dividend $\delta = 0$ are optimally executed at maturity. Therefore, the price of the American option and

of the corresponding European option are the same under these constraints (Föllmer and Schied, 2016). For all examples where this is the case, we compute the European option price (EOP) as an approximation of the correct American option price.

Moreover, as explained in (Lapeyre and Lelong, 2021), geometric put options on d dimensional stocks following Black-Scholes are equivalent to one-dimensional put options on a 1-dimensional stock following Black-Scholes with adjusted parameters. The 1-dimensional problem can be priced efficiently with the CRR binomial-tree method (B) (Cox et al., 1979). With the adjusted parameters $\hat{\sigma} = \frac{\sigma}{\sqrt{d}}$ and $\hat{\delta} = \delta + \frac{\sigma^2 - \hat{\sigma}^2}{2}$ the binomial-tree model is

defined with factors for the stock price going up and down $up = \exp(\hat{\sigma}\sqrt{T/N})$, $down = \frac{1}{up}$, probabilities to go up and down $p = \frac{\exp((r-\hat{\delta})T/N) - down}{up - down}$, $1 - p$ and step-wise discounting factor $\exp(-rT/N)$. Cox et al. (1979) have shown that the price computed with this method converges to the correct price under the Black-Scholes model as $N \rightarrow \infty$. Hence, this method yields good approximations of the correct option price for large N . Whenever applied, we use $N = 10'000$ for the binomial-tree method. While in the first case of call options, the optimal stopping problem has an easy solution, i.e. to wait until maturity, this is not the case here, where the optimal stopping problem is harder.

The remaining options, i.e. call options with $\delta > 0$, put options with $r > 0$ and geometric put option with underlying stocks following a Heston model, also constitute more complex stopping problems and no efficient methods to compute the (approximately) correct price are available. Therefore, we evaluate the performance of the algorithms by comparing the approximated prices directly. Since these prices are computed on unseen paths for all algorithms, where at each time, the algorithm can only decide whether to exercise or not, higher prices imply better performance of the algorithms.

Results and Discussion

All algorithms are run 10 times in parallel and the mean and standard deviation (in parenthesis) of the prices respectively the median of the corresponding computation times are reported. In particular, the computation times do not include the time for generating the stock paths, since the main interest is in the actual time the algorithms need to compute prices and paths can be generated offline and stored. In the following, we always compare computation times for large d , since random machine influences have less impact there.

d	x_0	price							duration						
		LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP
5	80	5.08 (0.05)	4.99 (0.06)	4.92 (0.06)	5.11 (0.10)	5.11 (0.07)	5.12 (0.11)	5.12 (0.06)	11s	9s	0s	0s	2s	0s	0s
	100	24.76 (0.09)	24.79 (0.16)	24.58 (0.19)	24.83 (0.14)	24.86 (0.25)	24.89 (0.14)	24.91 (0.11)	11s	8s	2s	0s	2s	0s	0s
	120	49.50 (0.12)	49.37 (0.23)	49.09 (0.21)	49.55 (0.27)	49.73 (0.18)	49.55 (0.11)	49.63 (0.20)	11s	7s	2s	0s	2s	0s	0s
10	80	8.97 (0.13)	8.91 (0.09)	8.67 (0.19)	8.95 (0.15)	8.99 (0.09)	8.99 (0.08)	8.98 (0.07)	28s	7s	1s	0s	6s	0s	0s
	100	33.88 (0.21)	33.85 (0.12)	33.23 (0.22)	33.97 (0.15)	33.95 (0.18)	33.85 (0.12)	33.90 (0.15)	29s	7s	2s	0s	7s	0s	0s
	120	60.61 (0.09)	60.57 (0.14)	60.16 (0.23)	60.67 (0.24)	60.81 (0.16)	60.78 (0.27)	60.72 (0.17)	29s	7s	2s	0s	6s	0s	0s
50	80	21.44 (0.12)	22.51 (0.12)	21.01 (0.43)	21.29 (0.12)	22.64 (0.11)	22.49 (0.11)	22.69 (0.08)	8m39s	8s	2s	0s	6m28s	1s	0s
	100	52.56 (0.13)	53.17 (0.12)	50.86 (0.50)	51.48 (0.23)	53.18 (0.17)	53.18 (0.16)	53.28 (0.10)	8m42s	8s	3s	0s	6m57s	1s	0s
	120	83.17 (0.10)	83.84 (0.22)	81.49 (0.77)	81.81 (0.18)	83.79 (0.22)	83.76 (0.22)	83.97 (0.14)	8m46s	9s	3s	0s	7m 4s	1s	0s
100	80	24.14 (0.17)	28.71 (0.08)	25.74 (0.77)	27.35 (0.08)	28.48 (0.12)	28.69 (0.14)	28.86 (0.07)	39m44s	13s	3s	0s	1h23m39s	1s	0s
	100	55.95 (0.24)	60.71 (0.15)	57.46 (0.88)	59.18 (0.10)	60.54 (0.16)	60.84 (0.11)	61.04 (0.09)	40m42s	13s	4s	0s	1h23m28s	1s	0s
	120	86.97 (0.19)	92.89 (0.14)	88.91 (0.87)	91.01 (0.11)	92.74 (0.18)	92.99 (0.16)	93.27 (0.13)	40m25s	13s	4s	0s	1h22m15s	1s	0s
500	80	-	41.34 (0.10)	37.09 (1.12)	41.12 (0.11)	-	42.16 (0.06)	42.38 (0.06)	-	53s	11s	1s	-	1s	0s
	100	-	76.72 (0.14)	71.90 (1.12)	76.39 (0.09)	-	77.75 (0.07)	77.95 (0.09)	-	53s	12s	1s	-	1s	0s
	120	-	111.95 (0.13)	106.29 (0.99)	111.68 (0.10)	-	113.30 (0.20)	113.47 (0.11)	-	53s	12s	1s	-	1s	0s
1000	80	-	46.38 (0.13)	43.11 (0.87)	46.78 (0.09)	-	47.74 (0.07)	47.89 (0.08)	-	1m34s	20s	2s	-	1s	0s
	100	-	82.97 (0.12)	78.98 (1.21)	83.50 (0.10)	-	84.76 (0.14)	84.89 (0.07)	-	1m35s	20s	3s	-	1s	0s
	120	-	119.53 (0.13)	114.21 (1.45)	120.17 (0.12)	-	121.62 (0.11)	121.88 (0.09)	-	1m36s	19s	2s	-	1s	0s
2000	80	-	49.92 (0.09)	49.00 (0.29)	52.37 (0.08)	-	53.15 (0.06)	53.37 (0.07)	-	2m57s	34s	5s	-	2s	0s
	100	-	87.38 (0.15)	86.40 (0.67)	90.41 (0.11)	-	91.54 (0.10)	91.77 (0.06)	-	3m 2s	39s	5s	-	2s	0s
	120	-	124.98 (0.14)	123.45 (0.71)	128.53 (0.10)	-	129.88 (0.14)	130.09 (0.09)	-	2m57s	37s	4s	-	2s	0s

Table 1: Max call option on Black-Scholes for different number of stocks d and varying initial stock price x_0 .

d	price							duration						
	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP
5	8.34 (0.08)	8.36 (0.07)	8.22 (0.09)	8.37 (0.07)	8.25 (0.03)	8.33 (0.07)	8.23 (0.04)	31s	6s	3s	0s	8s	0s	0s
10	11.81 (0.06)	11.83 (0.07)	11.51 (0.12)	11.83 (0.02)	11.79 (0.06)	11.83 (0.05)	11.79 (0.07)	1m30s	6s	3s	0s	28s	0s	0s
50	16.85 (0.07)	20.01 (0.06)	18.60 (0.32)	19.31 (0.05)	20.05 (0.06)	20.09 (0.05)	20.04 (0.04)	39m37s	8s	4s	0s	1h22m45s	1s	0s
100	-	23.49 (0.06)	21.75 (0.41)	22.90 (0.02)	-	23.69 (0.06)	23.66 (0.04)	-	14s	6s	0s	-	1s	0s
500	-	31.31 (0.06)	29.93 (0.32)	31.35 (0.06)	-	32.14 (0.06)	32.13 (0.07)	-	1m19s	24s	3s	-	2s	0s
1000	-	34.23 (0.08)	33.79 (0.29)	35.09 (0.06)	-	35.82 (0.06)	35.86 (0.04)	-	2m59s	41s	6s	-	4s	0s
2000	-	35.18 (0.14)	37.76 (0.23)	38.84 (0.05)	-	39.63 (0.08)	39.60 (0.05)	-	13m11s	1m28s	13s	-	7s	0s

Table 2: Max call option on Heston (with variance) for different number of stocks d .

d	price							duration						
	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP
5	3.60 (0.05)	3.57 (0.05)	3.49 (0.06)	3.58 (0.03)	3.61 (0.03)	3.62 (0.06)	3.59 (0.02)	13s	6s	2s	0s	2s	0s	0s
10	2.54 (0.04)	2.52 (0.03)	2.45 (0.06)	2.54 (0.04)	2.53 (0.03)	2.53 (0.03)	2.54 (0.01)	30s	6s	1s	0s	7s	0s	0s
50	0.94 (0.01)	1.12 (0.01)	0.83 (0.03)	1.06 (0.01)	1.13 (0.01)	1.15 (0.01)	1.14 (0.01)	8m51s	8s	1s	0s	7m 3s	1s	0s
100	0.51 (0.01)	0.78 (0.01)	0.55 (0.01)	0.75 (0.01)	0.80 (0.01)	0.81 (0.01)	0.81 (0.01)	38m59s	13s	2s	0s	1h21m59s	1s	0s
500	-	0.33 (0.01)	0.24 (0.00)	0.34 (0.00)	-	0.36 (0.00)	0.36 (0.00)	-	1m 7s	7s	1s	-	1s	0s
1000	-	0.22 (0.00)	0.17 (0.00)	0.24 (0.00)	-	0.25 (0.00)	0.26 (0.00)	-	2m24s	14s	2s	-	2s	0s
2000	-	0.13 (0.00)	0.12 (0.01)	0.17 (0.00)	-	0.18 (0.00)	0.18 (0.00)	-	5m35s	25s	7s	-	3s	0s

Table 3: Basket call options on Black-Scholes for different number of stocks d .

In all cases RLSM and RFQI are the fastest algorithms while achieving at least similar prices to the best performing baselines. Their biggest strength are high dimensional problems ($d \geq 500$), where this speed-up becomes substantial.

model	d	price							duration						
		LSM	DOS	NLSM	RLSM	FQI	RFQI	B	LSM	DOS	NLSM	RLSM	FQI	RFQI	B
BlackScholes	5	3.34 (0.04)	3.31 (0.03)	3.29 (0.06)	3.33 (0.04)	3.31 (0.05)	3.35 (0.04)	3.35 (nan)	11s	6s	1s	0s	2s	0s	3m12s
	10	2.37 (0.04)	2.42 (0.02)	2.33 (0.02)	2.40 (0.04)	2.39 (0.03)	2.40 (0.03)	2.40 (nan)	28s	6s	1s	0s	7s	0s	3m12s
	20	1.65 (0.02)	1.71 (0.04)	1.57 (0.04)	1.65 (0.03)	1.73 (0.04)	1.72 (0.02)	1.71 (nan)	1m31s	6s	1s	0s	32s	1s	3m12s
	50	0.91 (0.01)	1.07 (0.02)	0.80 (0.02)	1.03 (0.01)	1.09 (0.02)	1.09 (0.02)	1.09 (nan)	8m26s	10s	2s	0s	7m24s	1s	3m31s
	100	0.50 (0.01)	0.76 (0.01)	0.54 (0.01)	0.73 (0.01)	0.77 (0.01)	0.77 (0.01)	0.78 (nan)	38m37s	16s	2s	0s	1h23m50s	1s	3m31s
Heston	5	2.45 (0.03)	2.44 (0.03)	2.30 (0.06)	2.44 (0.02)	2.44 (0.04)	2.43 (0.03)	-	11s	6s	1s	0s	2s	0s	-
	10	2.00 (0.02)	2.00 (0.02)	1.75 (0.04)	2.00 (0.03)	2.00 (0.02)	2.01 (0.02)	-	29s	6s	2s	0s	7s	0s	-
	20	1.68 (0.02)	1.69 (0.02)	1.21 (0.05)	1.62 (0.05)	1.72 (0.02)	1.71 (0.01)	-	1m31s	7s	2s	0s	32s	1s	-
	50	1.33 (0.02)	1.47 (0.01)	0.83 (0.03)	1.24 (0.01)	1.49 (0.01)	1.48 (0.01)	-	8m31s	7s	3s	0s	7m13s	1s	-
	100	0.88 (0.01)	1.39 (0.01)	0.71 (0.02)	1.18 (0.01)	1.41 (0.01)	1.40 (0.01)	-	41m34s	15s	4s	0s	1h24m11s	1s	-

Table 4: Geometric put options on Black-Scholes and Heston (with variance) for different number of stocks d . Here $r = 2\%$ is used as interest rate.

d	x_0	price						duration					
		LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI
5	80	35.49 (0.07)	35.48 (0.06)	35.21 (0.12)	35.46 (0.07)	35.53 (0.08)	35.54 (0.05)	11s	10s	3s	0s	3s	0s
	100	19.98 (0.09)	19.96 (0.09)	19.68 (0.07)	19.96 (0.14)	19.97 (0.10)	19.95 (0.09)	11s	9s	3s	0s	3s	0s
	120	7.46 (0.10)	7.36 (0.08)	7.25 (0.07)	7.39 (0.10)	7.45 (0.11)	7.38 (0.10)	11s	6s	1s	0s	2s	0s
10	80	40.22 (0.05)	40.17 (0.05)	39.91 (0.10)	40.21 (0.07)	40.31 (0.07)	40.30 (0.04)	28s	6s	2s	0s	9s	0s
	100	25.74 (0.09)	25.74 (0.10)	25.36 (0.12)	25.76 (0.09)	25.79 (0.10)	25.83 (0.13)	28s	6s	3s	0s	6s	0s
	120	11.98 (0.07)	11.92 (0.09)	11.62 (0.14)	11.94 (0.13)	11.96 (0.10)	12.03 (0.07)	28s	5s	1s	0s	6s	0s
50	80	48.08 (0.05)	48.27 (0.04)	47.03 (0.19)	47.72 (0.03)	48.36 (0.05)	48.34 (0.04)	8m25s	8s	3s	0s	5m20s	1s
	100	35.57 (0.07)	35.80 (0.08)	34.27 (0.40)	35.11 (0.04)	35.91 (0.07)	35.87 (0.08)	8m35s	8s	3s	0s	6m57s	1s
	120	22.93 (0.08)	23.33 (0.07)	21.40 (0.41)	22.50 (0.05)	23.41 (0.06)	23.42 (0.10)	8m28s	8s	2s	0s	6m52s	1s
100	80	49.71 (0.06)	50.93 (0.04)	48.78 (0.26)	50.48 (0.04)	50.93 (0.04)	50.99 (0.03)	39m57s	13s	3s	0s	1h22m58s	1s
	100	37.63 (0.07)	39.11 (0.05)	36.42 (0.80)	38.55 (0.05)	39.11 (0.06)	39.22 (0.04)	40m12s	12s	3s	0s	1h23m26s	1s
	120	25.52 (0.08)	27.31 (0.05)	24.13 (0.52)	26.64 (0.03)	27.28 (0.07)	27.42 (0.05)	40m40s	12s	3s	0s	1h22m53s	1s
500	80	-	55.71 (0.03)	51.51 (0.46)	55.66 (0.03)	-	56.04 (0.03)	-	54s	13s	1s	-	1s
	100	-	45.14 (0.05)	40.06 (1.04)	45.05 (0.02)	-	45.51 (0.03)	-	53s	13s	1s	-	1s
	120	-	34.53 (0.05)	28.39 (0.75)	34.45 (0.05)	-	34.99 (0.02)	-	54s	12s	1s	-	1s
1000	80	-	57.40 (0.03)	53.50 (0.64)	57.52 (0.03)	-	57.84 (0.02)	-	1m36s	21s	3s	-	2s
	100	-	47.24 (0.05)	42.35 (0.63)	47.40 (0.05)	-	47.76 (0.03)	-	1m37s	21s	3s	-	2s
	120	-	37.04 (0.03)	31.10 (0.87)	37.25 (0.04)	-	37.68 (0.03)	-	1m34s	20s	3s	-	2s
2000	80	-	58.59 (0.04)	55.21 (0.67)	59.21 (0.02)	-	59.50 (0.02)	-	3m 1s	30s	6s	-	3s
	100	-	48.72 (0.04)	44.37 (0.61)	49.49 (0.04)	-	49.83 (0.03)	-	2m56s	31s	6s	-	3s
	120	-	38.84 (0.06)	33.31 (0.73)	39.79 (0.04)	-	40.18 (0.04)	-	3m 0s	31s	6s	-	3s

Table 5: Min put option on Black-Scholes for different number of stocks d and varying initial stock price x_0 . Here $r = 2\%$ is used as interest rate.

d	price						duration					
	LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI
5	18.83 (0.17)	18.66 (0.11)	18.62 (0.18)	18.83 (0.12)	18.43 (0.10)	18.83 (0.16)	22s	7s	3s	0s	2s	0s
10	26.67 (0.14)	26.72 (0.17)	26.35 (0.14)	26.60 (0.12)	26.56 (0.13)	26.77 (0.09)	46s	7s	3s	0s	8s	0s
50	43.86 (0.10)	44.52 (0.13)	43.27 (0.33)	43.37 (0.11)	44.66 (0.14)	44.78 (0.12)	10m 5s	10s	4s	0s	7m23s	1s
100	46.62 (0.19)	51.71 (0.09)	49.31 (0.56)	50.61 (0.10)	51.79 (0.17)	52.16 (0.09)	49m27s	15s	5s	0s	1h21m26s	1s
500	-	67.12 (0.09)	62.82 (0.72)	67.02 (0.08)	-	68.48 (0.13)	-	59s	14s	2s	-	2s
1000	-	73.37 (0.12)	69.25 (0.83)	73.85 (0.10)	-	75.31 (0.08)	-	1m52s	26s	4s	-	2s
2000	-	78.17 (0.11)	76.57 (0.83)	80.54 (0.07)	-	81.96 (0.16)	-	5m26s	47s	8s	-	3s

Table 6: Max call option on Black-Scholes for different number of stocks d . Here $r = 5\%$ is used as interest rate and $\delta = 10\%$ as dividend rate.

In these high dimensional problems, RLSM outperforms all baselines in terms of prices, even though RLSM has much less trainable parameters than DOS and NLSM. Moreover, RFQI achieves the highest prices there, and therefore works best, while having considerably less trainable parameters, since only one neural network (with a random hidden layer) of

d	price						duration					
	LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI
5	12.34 (0.05)	12.31 (0.05)	12.16 (0.11)	12.29 (0.06)	12.35 (0.09)	12.37 (0.09)	30s	6s	3s	0s	8s	0s
10	16.48 (0.07)	16.52 (0.08)	16.09 (0.13)	16.55 (0.06)	16.64 (0.07)	16.61 (0.08)	1m31s	6s	3s	0s	28s	0s
50	22.86 (0.05)	25.56 (0.04)	24.03 (0.42)	24.85 (0.08)	25.72 (0.03)	25.71 (0.07)	39m57s	9s	4s	0s	1h21m59s	1s
100	-	29.13 (0.04)	27.30 (0.46)	28.50 (0.06)	-	29.33 (0.07)	-	16s	6s	0s	-	1s
500	-	36.26 (0.05)	34.74 (0.31)	36.28 (0.04)	-	36.95 (0.05)	-	1m21s	24s	3s	-	2s
1000	-	38.62 (0.08)	38.19 (0.20)	39.32 (0.03)	-	39.93 (0.05)	-	3m18s	45s	6s	-	4s
2000	-	39.22 (0.13)	41.05 (0.21)	42.22 (0.04)	-	42.81 (0.04)	-	12m51s	1m37s	13s	-	8s

Table 7: Min put option on Heston (with variance) for different number of stocks d . Here $r = 2\%$ is used as interest rate.

d	price						duration					
	LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI
5	4.88 (0.03)	4.89 (0.03)	4.69 (0.05)	4.83 (0.04)	4.37 (0.06)	4.59 (0.08)	31s	5s	3s	0s	8s	0s
10	7.19 (0.06)	7.20 (0.04)	6.90 (0.07)	7.17 (0.04)	6.63 (0.07)	6.84 (0.06)	1m33s	5s	2s	0s	27s	0s
50	11.68 (0.05)	13.99 (0.07)	12.93 (0.28)	13.70 (0.05)	13.72 (0.09)	13.71 (0.04)	41m 6s	8s	3s	0s	1h22m14s	1s
100	-	17.04 (0.07)	15.94 (0.29)	16.80 (0.03)	-	16.97 (0.05)	-	11s	5s	0s	-	1s
500	-	24.05 (0.05)	22.95 (0.40)	24.35 (0.05)	-	24.70 (0.05)	-	1m19s	23s	3s	-	2s
1000	-	26.86 (0.05)	26.47 (0.39)	27.71 (0.04)	-	28.08 (0.05)	-	2m48s	41s	6s	-	4s
2000	-	28.01 (0.11)	30.12 (0.18)	31.13 (0.05)	-	31.55 (0.07)	-	12m56s	1m30s	14s	-	7s

Table 8: Max call option on Heston (with variance) for different number of stocks d . Here $r = 5\%$ is used as interest rate and $\delta = 10\%$ as dividend rate.

d	N	price						duration							
		LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP
10	10	33.88 (0.21)	33.85 (0.12)	33.23 (0.22)	33.97 (0.15)	33.95 (0.18)	33.85 (0.12)	33.90 (0.15)	29s	7s	2s	0s	7s	0s	0s
	50	34.13 (0.12)	34.14 (0.20)	33.96 (0.11)	33.98 (0.08)	34.25 (0.21)	34.15 (0.13)	34.23 (0.11)	2m44s	32s	20s	0s	46s	6s	0s
	100	34.15 (0.14)	34.14 (0.26)	33.98 (0.24)	34.05 (0.15)	34.29 (0.16)	33.97 (0.11)	34.28 (0.10)	5m18s	1m 6s	32s	1s	1m27s	7s	0s
50	10	52.56 (0.13)	53.17 (0.12)	50.86 (0.50)	51.48 (0.23)	53.18 (0.17)	53.18 (0.16)	53.28 (0.10)	8m42s	8s	3s	0s	6m57s	1s	0s
	50	52.82 (0.13)	53.94 (0.18)	53.24 (0.26)	50.85 (0.18)	54.31 (0.14)	53.74 (0.08)	54.46 (0.11)	48m36s	41s	18s	1s	21m15s	7s	0s
	100	52.74 (0.11)	54.09 (0.15)	53.61 (0.18)	50.42 (0.17)	54.15 (0.10)	53.77 (0.12)	54.33 (0.14)	1h37m48s	1m36s	37s	2s	41m 8s	16s	0s
100	10	55.95 (0.24)	60.71 (0.15)	57.46 (0.88)	59.18 (0.10)	60.54 (0.16)	60.84 (0.11)	61.04 (0.09)	40m42s	13s	4s	0s	1h23m28s	1s	0s
	50	-	61.88 (0.06)	60.72 (0.24)	58.68 (0.13)	-	61.66 (0.14)	62.48 (0.07)	-	1m15s	22s	1s	-	8s	0s
	100	-	62.07 (0.11)	61.19 (0.15)	58.26 (0.16)	-	61.79 (0.11)	62.46 (0.04)	-	2m23s	44s	3s	-	15s	0s
500	10	-	76.72 (0.14)	71.90 (1.12)	76.39 (0.09)	-	77.75 (0.07)	77.95 (0.09)	-	53s	12s	1s	-	1s	0s
	50	-	79.14 (0.08)	75.63 (1.07)	76.68 (0.05)	-	79.23 (0.07)	80.44 (0.09)	-	4m59s	1m 4s	8s	-	9s	0s
	100	-	79.44 (0.09)	76.46 (0.41)	76.33 (0.05)	-	79.34 (0.08)	80.47 (0.10)	-	10m12s	2m13s	18s	-	19s	0s

Table 9: Max call option on Black-Scholes for different number of stocks d and higher number of exercise dates N .

the respective size is used for all exercise dates. In particular, RFQI has only 21 trainable parameters, compared to more than $20dN$ for DOS and NLSM.

Comparing the achieved prices of LSM and FQI, we can confirm the claim of Li et al. (2009), that reinforcement learning techniques usually outperform the backward induction in the Markovian setting. RFQI, achieving similar prices as FQI, therefore naturally outperforms RLSM which achieves similar prices as LSM. A possible explanation for the outperformance of the reinforcement learning algorithm is the following. The backward induction algorithms have approximately N times the number of trainable parameters used in the reinforcement learning algorithms, since a different network is trained for each discretisation date. Moreover, for the backward induction algorithms, a different continuation value function is approximated for each date, hence, only the data of this date is used to

d	N	price						duration					
		LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI
10	10	26.67 (0.14)	26.72 (0.17)	26.35 (0.14)	26.60 (0.12)	26.56 (0.13)	26.77 (0.09)	46s	7s	3s	0s	8s	0s
	50	26.65 (0.12)	26.56 (0.20)	26.42 (0.17)	26.61 (0.13)	26.51 (0.07)	26.69 (0.18)	2m21s	44s	53s	2s	44s	4s
	100	26.68 (0.19)	26.44 (0.14)	26.42 (0.19)	26.59 (0.13)	26.55 (0.14)	26.65 (0.16)	4m46s	2m46s	1m45s	1s	1m25s	8s
50	10	43.86 (0.10)	44.52 (0.13)	43.27 (0.33)	43.37 (0.11)	44.66 (0.14)	44.78 (0.12)	10m 5s	10s	4s	0s	7m23s	1s
	50	43.42 (0.09)	44.50 (0.08)	44.13 (0.19)	42.26 (0.10)	44.68 (0.15)	44.72 (0.12)	43m 5s	1m14s	52s	3s	16m46s	9s
	100	43.21 (0.14)	44.45 (0.15)	44.30 (0.13)	41.89 (0.31)	44.64 (0.17)	44.60 (0.14)	1h25m 4s	2m 0s	1m45s	2s	36m10s	17s
100	10	46.62 (0.19)	51.71 (0.09)	49.31 (0.56)	50.61 (0.10)	51.79 (0.17)	52.16 (0.09)	49m27s	15s	5s	0s	1h21m26s	1s
	50	-	51.73 (0.10)	50.89 (0.21)	49.36 (0.09)	-	51.91 (0.12)	-	52s	33s	1s	-	8s
	100	-	51.72 (0.15)	51.27 (0.12)	48.90 (0.08)	-	51.84 (0.11)	-	1m48s	46s	3s	-	19s
500	10	-	67.12 (0.09)	62.82 (0.72)	67.02 (0.08)	-	68.48 (0.13)	-	59s	14s	2s	-	2s
	50	-	67.48 (0.11)	64.75 (0.50)	65.70 (0.07)	-	68.03 (0.12)	-	2m56s	1m 4s	7s	-	10s
	100	-	67.50 (0.15)	65.55 (0.34)	65.22 (0.07)	-	67.91 (0.10)	-	5m48s	1m52s	16s	-	20s

Table 10: Max call option on Black-Scholes for different number of stocks d and higher num-

ber of exercise dates N . Here $r = 5\%$ is used as interest rate and $\delta = 10\%$ as dividend rate.




Figure 1: At the money max call option without dividend on Black-Scholes.

learn the parameters. In contrast, the reinforcement learning methods train their parameters using the data of all dates. Hence, the reinforcement learning methods use N times the number of data to train $1/N$ times the number of parameters, which seems to lead to better approximations.

We first give a detailed discussion of results for the easy optimal stopping problems, where it is optimal to exercise the option at maturity. Although these optimal stopping problems

0. Due to memory overflow issues, FQI could only be run with 5 instead of 10 parallel runs for larger N , hence the computation times are smaller than they would otherwise be, due to more CPU power per run.




Figure 2: At the money max call option with dividend on Black-Scholes.

are less complex, they are still interesting, because a minimal requirement for the algorithms should be that they perform well in these basic examples. Moreover, a comparison to the reference price is possible. In Table 1 and Figure 1 we show results of a max call option on Black-Scholes. For high dimensions ($d \geq 500$), RLSM is about 8 times faster than the fastest baseline NLSM and about 30 times faster than DOS. Moreover, RFQI is about twice as fast as RLSM. For $d = 100$ we also see the large difference in computation time between LSM (respectively FQI), where the number of basis functions grows quadratically in d , and RLSM (respectively RFQI), where the number of basis functions does not grow in d . The computed prices of RLSM are at most 2.1% smaller than those of LSM and the prices of RFQI are at most 0.6% smaller than those of FQI. For $d \leq 100$ the maximal relative errors compared to the reference prices are 16.3% for LSM, 2.4% for DOS, 10.8% for NLSM, 6.1% for RLSM, 1.3% for FQI and 0.8% for RFQI. For $d \geq 500$ these errors are 6.4% for DOS, 12.5% for NLSM, 2.9% for RLSM and 0.5% for RFQI. The results of Table 2 (max call on Heston with variance and payoff as additional input) and Table 3 (basket call on Black-Scholes) are similar, except that relative errors become larger in Table 3 for growing d , since the prices become very small.

In the remaining examples, it is in general not optimal to exercise the options at maturity, making the

stopping decisions harder and therefore more challenging for the algorithms.

For the geometric put options (Table 4), we do not present dimensions larger than 100, because prices cannot be computed numerically any more. In the Black-Scholes case, the maximal relative errors compared to the reference price are 35.7% for LSM, 2.1% for DOS, 29.9% for NLSM, 6% for RLSM, 1.1% for FQI and 0.5% for RFQI. Again, the prices computed with RLSM (respectively RFQI) are never much smaller than those of LSM (FQI); 0.3% (0.1%) for Black-Scholes and 6% (0.6%) for Heston (with variance). On the Heston model, RFQI, FQI and DOS achieve the highest prices that never deviate more than 1.5% from each other.

For the min put option on Black-Scholes (Table 5) RLSM is about 7 times faster than NLSM and more than 30 times faster than DOS for high dimensions. Furthermore RFQI is again about twice as fast as RLSM. For $d \leq 50$ all algorithms yield very similar prices and for larger d the highest prices are always achieved by RFQI, whereby the prices computed with RFQI never deviate more than 1% from those computed with FQI. Moreover, the prices computed with RLSM are never more than 1.9% smaller than those computed with LSM. In addition, RLSM achieves the second highest prices for high dimensions. For the max call option with dividends on Black-Scholes (Table 6 and Figure 2), the situation is similar. However, the highest prices are always achieved by RFQI and the prices computed with RLSM are at most 1.1% smaller than those of LSM. For the min put option on Heston (with variance and payoff as additional input) (Table 7) we have similar results as on Black-Scholes, but the prices computed with RLSM (RFQI) are at most 0.5% (0.2%) smaller than those computed with LSM (FQI).

For the max call option with dividend on Heston (with variance) (Table 8), RLSM is about 7 times faster than NLSM and more than 26 times faster than DOS for high dimensions. RFQI is again about twice as fast as RLSM. For $d \in \{5, 10\}$ DOS yields the highest prices, RLSM deviates at most than 1.2% from them and RFQI at most 6%. FQI yields lower prices than RFQI. For $d \in \{50, 100\}$, DOS, RLSM and RFQI yield very similar prices deviating at most 2% from each other. For higher dimensions of $d \geq 500$, RFQI yields the highest and RLSM the second highest prices.

When increasing the number of exercise dates for the maxcall option on Black-Scholes from $N = 10$ to $N \in \{50, 100\}$ (Table 9) the Bermudan option price should become closer to the American option price. The highest prices are achieved either by RFQI, FQI or DOS, with a maximum deviation of less than 1.4% between their results and a maximum deviation from the reference prices of 1.6% for DOS and 1.5% for RFQI. RFQI is more than 30 times faster than DOS for high dimensions. Increasing the number of dates further, the computation time can become a limiting factor for DOS, while this is not the case for RFQI. We see similar results for the more complex maxcall option on Black-Scholes with dividends (Table 10), where RFQI always achieves the highest price.

Computation of Greeks

The Greeks are the sensitivities of the option price to a small change in a given underlying parameter. In particular, they are partial derivatives of the option prices with respect to different parameters, such as the spot price, time, rate and volatility. We provide experiments

(and the code) where we compute the most popular Greeks: delta $\left(\frac{\partial p_0}{\partial x_0}\right)$, gamma $\left(\frac{\partial^2 p_0}{\partial x_0^2}\right)$, theta $\left(\frac{\partial p_0}{\partial t}\right)$, rho $\left(\frac{\partial p_0}{\partial r}\right)$ and vega $\left(\frac{\partial p_0}{\partial \sigma}\right)$. The straight forward method to compute them is via the finite difference (FD) method. For theta, rho and vega, the standard forward finite difference method can be used with our algorithms, however, they turn out to be unstable for NLSM and DOS. Therefore, we use the central finite difference method, where the exercise boundary is frozen to be the one of the central point and report results only with this method. For

K	algo	price		delta		gamma		theta	rho	vega
		FD	regr.	FD	regr.	PDE	regr.			
36	B	0.9192 (-)		-0.1982 (-)		0.0389 (-)		-0.7152 (-)	-6.7085 (-)	10.9100 (-)
36	LSM	0.9024 (0.0086)	0.9006 (0.0094)	-0.1919 (0.0019)	-0.1888 (0.0027)	0.0368 (0.0004)	0.0381 (0.0009)	-0.6615 (0.0068)	-6.8267 (0.0497)	10.7107 (0.0918)
36	RLSM	0.9020 (0.0069)	0.9073 (0.0110)	-0.1940 (0.0019)	-0.1947 (0.0026)	0.0371 (0.0003)	0.0379 (0.0010)	-0.6665 (0.0063)	-6.8241 (0.0639)	10.7590 (0.0629)
36	FQI	0.8504 (0.0067)	0.8642 (0.0116)	-0.1777 (0.0016)	-0.1770 (0.0015)	0.0329 (0.0003)	0.0329 (0.0011)	-0.5753 (0.0047)	-7.7168 (0.0676)	10.3836 (0.0841)
36	RFQI	0.9005 (0.0087)	0.8766 (0.0138)	-0.1924 (0.0029)	-0.1847 (0.0041)	0.0368 (0.0005)	0.0369 (0.0009)	-0.6612 (0.0105)	-6.8282 (0.0891)	10.7093 (0.0880)
36	NLSM	0.8948 (0.0238)	0.8817 (0.0151)	-0.2010 (0.0174)	-0.1905 (0.0038)	0.0368 (0.0013)	0.0380 (0.0006)	-0.6427 (0.0124)	-6.9383 (0.0625)	10.6464 (0.1207)
36	DOS	0.9068 (0.0093)	0.9081 (0.0106)	-0.1957 (0.0024)	-0.1940 (0.0033)	0.0359 (0.0013)	0.0378 (0.0012)	-0.6251 (0.0393)	-7.0119 (0.3001)	10.7249 (0.1419)
40	B	2.3196 (-)		-0.4047 (-)		0.0611 (-)		-0.8446 (-)	-11.2405 (-)	14.7517 (-)
40	LSM	2.2916 (0.0107)	2.2715 (0.0116)	-0.3930 (0.0022)	-0.4049 (0.0053)	0.0579 (0.0003)	0.0637 (0.0009)	-0.7711 (0.0051)	-11.6289 (0.0573)	14.6886 (0.0548)
40	RLSM	2.2897 (0.0095)	2.2970 (0.0135)	-0.3984 (0.0052)	-0.4076 (0.0039)	0.0583 (0.0006)	0.0610 (0.0012)	-0.7721 (0.0067)	-11.6357 (0.1092)	14.7026 (0.0467)
40	FQI	2.2593 (0.0121)	2.2078 (0.0161)	-0.3661 (0.0017)	-0.3934 (0.0038)	0.0541 (0.0003)	0.0655 (0.0013)	-0.7179 (0.0057)	-12.6145 (0.0856)	14.7479 (0.0670)
40	RFQI	2.2239 (0.0407)	2.1252 (0.0374)	-0.3623 (0.0140)	-0.3654 (0.0123)	0.0529 (0.0024)	0.0570 (0.0024)	-0.6897 (0.0415)	-12.9713 (0.6541)	14.6794 (0.0831)
40	NLSM	2.2586 (0.0149)	2.2599 (0.0236)	-0.3830 (0.0136)	-0.4034 (0.0035)	0.0565 (0.0013)	0.0620 (0.0012)	-0.7529 (0.0166)	-11.7895 (0.2469)	14.6545 (0.0970)
40	DOS	2.2884 (0.0102)	2.2963 (0.0099)	-0.4031 (0.0037)	-0.4071 (0.0039)	0.0587 (0.0004)	0.0611 (0.0006)	-0.7727 (0.0055)	-11.5870 (0.1041)	14.7045 (0.0579)
44	B	4.6629 (-)		-0.6654 (-)		0.0779 (-)		-0.6169 (-)	-11.8974 (-)	12.8541 (-)
44	LSM	4.6141 (0.0175)	4.6177 (0.0120)	-0.6476 (0.0037)	-0.6693 (0.0051)	0.0743 (0.0003)	0.0676 (0.0012)	-0.5468 (0.0055)	-12.8537 (0.1337)	13.1799 (0.1129)
44	RLSM	4.6167 (0.0205)	4.6407 (0.0178)	-0.6541 (0.0067)	-0.6699 (0.0036)	0.0746 (0.0005)	0.0641 (0.0022)	-0.5400 (0.0034)	-12.7787 (0.2212)	13.0670 (0.1566)
44	FQI	4.5366 (0.0221)	4.5476 (0.0137)	-0.5962 (0.0039)	-0.6766 (0.0054)	0.0695 (0.0005)	0.0710 (0.0013)	-0.5216 (0.0066)	-15.2857 (0.1248)	14.3877 (0.0471)
44	RFQI	4.3469 (0.0708)	4.1557 (0.0415)	-0.5463 (0.0073)	-0.5791 (0.0059)	0.0607 (0.0023)	0.0687 (0.0019)	-0.3688 (0.0534)	-19.9917 (1.0327)	15.6835 (0.1199)
44	NLSM	4.5820 (0.0211)	4.5996 (0.0531)	-0.6553 (0.0150)	-0.6713 (0.0048)	0.0742 (0.0008)	0.0658 (0.0029)	-0.5268 (0.0182)	-13.1123 (0.9381)	12.9567 (0.5430)
44	DOS	4.6206 (0.0126)	4.6536 (0.0113)	-0.6580 (0.0046)	-0.6695 (0.0042)	0.0747 (0.0003)	0.0641 (0.0016)	-0.5350 (0.0050)	-12.8331 (0.2236)	13.0304 (0.1295)

Table 11: Prices and Greeks computed for different strikes K of a 1-dimensional put option on Black-Scholes. For the binomial (B) algorithm, the spacing of the FD method is set to $\varepsilon = 10^{-9}$, which is also used for the other algorithms for delta, theta, rho and vega. For the regression method, $\epsilon = 5$ and a polynomial basis up to degree 9 are used.

computing delta we use the same method, since the others are unstable for all algorithms. Moreover, the computation of gamma, as a second derivative, turns out to be unstable when computed with the second order finite difference method, even when using the same technique as for delta. Therefore, we use two alternative ways to circumvent this instability. The first one (PDE method) is specific to the case of an underlying Black-Scholes model, where the Black-Scholes PDE

$$\frac{\partial p_0}{\partial t} + \frac{1}{2}\sigma^2 x_0^2 \frac{\partial^2 p_0}{\partial x_0^2} + r x_0 \frac{\partial p_0}{\partial x_0} - r p_0 = 0$$

can be used to express gamma in terms of the price, delta and theta. The second one (regression method) is the “naive method” suggested in (Létourneau and Stentoft, 2019, Section 3.1). It fits a

polynomial regression to option prices achieved when distorting the initial price x_0 by a noise term $\xi \sim N(0, \epsilon^2)$. Then the price, delta and gamma can easily be computed by evaluating the fitted regression and its first and second derivative (which are easily computed, since polynomial regression is used) at the initial price x_0 . The parameter ϵ controls the variance-bias trade-off and has to be chosen by hand. However, the authors also suggested a 2-step method that reduces variance and bias, where this parameter is chosen automatically.

For comparability, we compute the Greeks for the same example as in (Létourneau and Stentoft, 2019). In particular, we consider a put option on $d = 1$ stock following a Black-Scholes model with initial price $x_0 = 40$, strike $K \in \{36, 40, 44\}$, rate $r = 6\%$, volatility $\sigma = 20\%$, $N = 10$ equidistant dates and $m = 100'000$ paths. The models are run 10 times and mean and standard deviations are reported in Table 11. The price, delta and gamma are computed with both, the finite difference (respectively PDE) and the regression method. As reference we use the binomial model with $N = 50'000$ equidistant dates, for which only the finite difference (respectively PDE) method is used. The hidden size was set to 10 to account for the smaller input dimension and the payoff was not used as input except for DOS, where it improved the results considerably. For RLSM the activation function was changed to Softplus, since this worked best, although all other tested activation functions did also yield good results. Overall, RLSM and DOS with the regression method achieve the best results. Furthermore, we highlight, that the time advantage of RLSM and RFQI also comes into play for the computation of Greeks, when increasing the dimension d .

Empirical Convergence Study

We confirm the theoretical results of Theorem 3 (Figure 3 left) and Theorem 2 (Figure 3 right) by an empirical convergence study for a growing number of paths m . For RLSM we also increase the number of hidden nodes K , while they are fixed for RFQI since $d = 5$ is used. For each combination of the number of paths m and the hidden size K , the algorithms are run 20 times and their mean prices with standard deviations are shown. For small m , we see that smaller hidden sizes achieve better prices. This is due to overfitting to the training paths when using larger networks. Regularization techniques like L^1 - or L^2 -penalization could be used to reduce overfitting for larger networks. However, our results suggest that restricting the hidden size is actually the simplest and best regularization technique, since it additionally leads to lower training times.




Figure 3: Mean \pm standard deviation (bars) of the price for a max call on 5 stocks following the Black-Scholes model for RLSM (left) and RFQI (right) for varying the number of paths m and varying for RLSM the number of neurons in the hidden layer K .

The Non-Markovian Case - Optimally Stopping Fractional Brownian Motions

In order to compare our algorithms on a problem where the underlying process is nonMarkovian, we take the example of the fractional Brownian motion $(W_t^H)_{t \geq 0}$ as in (Becker et al., 2019). Unlike classical Brownian motion, the increments of fractional Brownian motion need not be independent. Fractional Brownian motion is a continuous centered Gaussian process with covariation function $E(W_t^H W_s^H) = \frac{1}{2} (|t|^{2H} + |s|^{2H} - |t - s|^{2H})$ where $H \in (0, 1]$ is called the Hurst parameter. When the Hurst parameter $H = 0.5$, then W^H is a standard Brownian motion; when $H \neq 0.5$, the increments of $(W_t^H)_{t \geq 0}$ are correlated (positively if $H > 0.5$ and negatively if $H < 0.5$) which means that for $H \neq 0.5$, $(W_t^H)_{t \geq 0}$ is not Markovian (Bayer et al., 2016; Livieri et al., 2018; Gatheral et al., 2018; El Euch et al., 2018; Abi Jaber and El Euch, 2019).

Stock Model, Payoffs and Baselines

In this section we use a d -dimensional fractional Brownian motion, with independent coordinates all starting at $X_0 = 0$, as the underlying process $X_t = W_t^H$. In contrast to the price processes we used before, this process can become negative. In the one-dimensional case, we use the identity as “payoff” function $g = \text{id}$ as in (Becker et al., 2019), which can lead to negative “payoff” values. Moreover, we use the maximum $g(x) = \max(x_1, x_2, \dots, x_d)$ for any $x = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d$ and the mean $g(x) = 1/d \sum_{i=1}^d x_i$ as “payoffs” for higher dimensions, which can also yield negative values. In particular, this setting leads to an optimal stopping problem outside of the standard discretized American option pricing setting. We compare RLSM and RRLSM to DOS and the path-version of DOS

(denoted pathDOS for our implementation of it and pathDOS-paper for results reported from (Becker et al., 2019)), where the entire path until the current date is used as input (Becker et al., 2019). Moreover, we test RFQI and its recurrent and path-version in this setting.

For two values of the Hurst parameter the optimal value can be computed explicitly. In particular, for $H = 0.5$ we have a Brownian motion and therefore the optimal value is 0 and for $H = 1$ we have a fully correlated process (i.e. all information is known after the first step), where the optimal value is approximately 0.39495 (Becker et al., 2019).

Results and Discussion

For $d = 1$, we clearly see the outperformance of the algorithms processing information of the path compared to the ones using only the current value as input (Figure 4 top left). Moreover, this application highlights the limitation of reinforcement learning techniques when applied in non-Markovian settings as discussed in (Kaelbling et al., 1996). In particular, RFQI, the randomized RNN version of it and its path-version do not work well in this example (Figure 4 top right). This poor performance was consistent under varying hyper-parameters.

RRLSM achieves very similar results to those reported for pathDOS in (Becker et al., 2019) with an MSE of 0.0005 between their reported values and ours, while using only 20 K instead of 4M paths (Figure 4 bottom). RRLSM needs only 1s to be trained in contrast to 430s reported in (Becker et al., 2019). The longer training times can partly be explained by the larger amount of paths used. However, our implementation of pathDOS using the same number of 20 hidden nodes as RRLSM and also being trained on 20K paths (hence completely comparable to the training of RRLSM) takes approximately 175 s and achieves slightly worse results than RRLSM (Figure 4 top left) with an MSE of 0.0018. The exact prices displayed in Figure 4 are provided in Appendix C.1.

For higher dimensions, we use the small Hurst parameter $H = 0.05$ for which a big difference between the standard and the path dependent algorithms was visible in the onedimensional case. RFQI yields very similar prices as DOS and RRLSM yields very similar prices as pathDOS. However, RFQI and RRLSM are considerably faster than DOS and pathDOS (Table 12).




Figure 4: Top left: algorithms processing path information outperform. Top right: reinforcement learning algorithms do not work well in non-Markovian cases. Bottom: RRLSM achieves similar results as reported in (Becker et al., 2019), while using only 20 K paths instead of 4M for training which took only 1s instead of the reported 430s.

payoff	d	price				duration			
		DOS	pathDOS	RLSM	RRLSM	DOS	pathDOS	RLSM	RRLSM
Identity	1	0.67(0.02)	1.24(0.01)	0.65(0.01)	1.24(0.01)	1 m15 s	3 m1 s	0 s	1 s
	5	1.96(0.01)	2.15(0.01)	2.00(0.01)	2.16(0.01)	3 m8 s	21 m46 s	4 s	1 s
	10	2.34(0.01)	2.43(0.01)	2.40(0.01)	2.43(0.02)	3 m49 s	37 m46 s	4 s	2 s
Mean	5	0.29(0.01)	0.53(0.00)	0.28(0.01)	0.52(0.01)	3 m40 s	21 m8 s	3 s	1 s
	10	0.20(0.01)	0.36(0.00)	0.21(0.01)	0.33(0.01)	3 m39 s	36 m1 s	5 s	1 s

Table 12: Identity, maximum and mean on the fractional Brownian motion with $H = 0.05$ and different number of stocks d .

Non-Markovian Stock Models

In Section 6.3 we saw that the RL based algorithms do not perform well on problems which are highly path dependent. In this section, we consider “intermediate” problems of typical non-Markovian stock models, where a path dependence exists, but where this path dependence is not very strong.

Heston Without Variance as Input

First, we revisit the Heston model (9), but this time without feeding the algorithms the variance, which makes it a non-Markovian problem. For the max call (Table 16), min put (Table 17) and max call with dividend (Table 18) options on Heston without variance, all the algorithms yield very similar prices as on Heston with variance (Tables 2, 7 and 8), therefore we do only show the tables in Appendix C.2. In particular, this suggests that even though the Heston model is not Markovian without providing the current variance, this doesn’t make a difference for option pricing.

Rough Heston

Moreover, we test on the rough Heston model, where the variance itself is path-dependent. This model recently became a very popular choice for modelling financial markets (El Euch and Rosenbaum, 2018; El Euch et al., 2019; Gatheral et al., 2020). The rough Heston model (El Euch and Rosenbaum, 2018) is defined as

$$dX_t = (r - \delta)X_t dt + \sqrt{v_t}X_t dW_t$$

$$v_t = v_0 + \int_0^t \frac{(t-s)^{H-1/2}}{\Gamma(H+1/2)} \kappa (v_\infty - v_s) ds + \int_0^t \frac{(t-s)^{H-1/2}}{\Gamma(H+1/2)} \sigma \sqrt{v_s} dB_s$$

where $X_0 = x_0$, the Hurst parameter $H \in (0, 1/2)$ and $(W_t)_{t \geq 0}$ and $(B_t)_{t \geq 0}$ are two d dimensional Brownian motions correlated with coefficient $\rho \in (-1, 1)$. We choose the drift $r = 5\%$, the dividend rate $\delta = 10\%$, the volatility of volatility $\sigma = 20\%$, the long term variance $v_\infty = 0.01$, the mean reversion speed $\kappa = 2$, the correlation $\rho = -30\%$, the initial stock price $x_0 = 100$ and the initial variance $v_0 = 0.01$ and consider a max call option on the stock price X .

As for the Heston model, also for the rough Heston model there is no significant difference between the computed prices with and without providing the current variance, therefore we only show prices where the current variance was also fed to the algorithms, which is still a non-Markovian setting. For the max call option on the rough Heston model (with variance) (Table 13), we see again that the reinforcement learning based algorithms FQI and RFQI do not work well.

Overall, DOS, NLSM, RLSM and RRLSM achieve very similar prices, never deviating more than 2.2% from each other. In particular, we do not see a better performance of the path dependent algorithms pathDOS and RRLSM compared to DOS and RLSM.

d	price								duration							
	LSM	DOS	pathDOS	NLSM	RLSM	RRLSM	FQI	RFQI	LSM	DOS	pathDOS	NLSM	RLSM	RRLSM	FQI	RFQI
5	8.78 (0.08)	8.78 (0.10)	8.76 (0.08)	8.69 (0.06)	8.81 (0.10)	8.74 (0.07)	8.33 (0.11)	8.01 (0.29)	29s	7s	14s	3s	0s	0s	7s	0s
10	12.99 (0.09)	13.00 (0.06)	13.00 (0.11)	12.94 (0.07)	13.02 (0.13)	12.74 (0.09)	12.22 (0.08)	10.68 (0.74)	1m31s	7s	18s	3s	0s	0s	29s	0s
50	22.55 (0.11)	25.67 (0.07)	25.06 (0.11)	26.08 (0.10)	25.64 (0.19)	25.77 (0.19)	22.92 (0.72)	17.30 (0.29)	40m17s	10s	1m27s	4s	0s	0s	1h24m	6s
100	-	32.52 (0.15)	31.24 (0.14)	32.39 (0.23)	32.60 (0.13)	32.71 (0.16)	-	24.46 (0.19)	-	17s	2m37s	6s	0s	0s	-	1s

Table 13: Max call option on Rough-Heston for different number of stocks d . The interest rate is $r = 5\%$ and the dividend rate is $\delta = 10\%$.

Discussion on the Sensitivity of the Randomness of the Hidden Layers

We perform a test specifically designed to study the model's sensitivity to the randomness of the weights of the hidden layers. In our previous tests in this paper we performed 10 runs, where a different set of paths and different weights of the hidden layer were chosen for each run. In order to test the sensitivity to the randomness of the weights, we perform an experiment with 10 runs, where only the set of hidden weights are different for each run, while the paths are the same.

We compare RLSM and NLSM in the setting of a 1-dimensional Black-Scholes call option with spot $x_0 = 100$ and strike $K = 100$, where we use 100 K paths and 10 exercises dates with 20 hidden nodes and either 10, 30 or 50 epochs of training for NLSM.

In order to have a fair comparison, we do not fix the initial weights of NLSM, as it would be equivalent to reusing the same random weights for RLSM in each run, with the possibility of having a good or bad initialisation. Hence, similar to RLSM's sensitivity to the randomness of the weights in the hidden layer, NLSM is sensitive to the randomness in the initialization of the weights (of the hidden layer). In order to reduce this sensitivity in the algorithms, one should always take the average of several runs with different sets of weights (and paths). This can be easily done in parallel in order to reduce the computation time. The results of this sensitivity analysis are given in Table 14 . We see that the sensitivity of NLSM to the randomness of the initialization depends on the number of epochs of the training, becoming smaller with longer training.

In order to further reduce the sensitivity of RLSM to the randomness of the hidden layer weights we propose a variant of it, which we call RLSMreinit. Instead of using the same random weights for each date, we use different ones, which have an averaging effect and therefore reduce the variance in multiple runs.

algo	#epochs	price	delta	gamma	theta	rho	vega	comp. time median
NLSM	10	8.8961 (0.0356)	0.5845 (0.0018)	0.0194 (0.0001)	-4.8682 (0.0142)	46.0934 (1.2538)	39.4202 (0.0613)	7.86s
NLSM	30	8.9175 (0.0171)	0.5836 (0.0008)	0.0194 (0.0001)	-4.8723 (0.0206)	47.2922 (1.2718)	39.3781 (0.0640)	19.53s
NLSM	50	8.9178 (0.0160)	0.5835 (0.0006)	0.0194 (0.0001)	-4.8780 (0.0166)	46.6717 (1.6730)	39.3755 (0.1004)	35.03s
RLSM		8.9439 (0.0179)	0.5828 (0.0004)	0.0195 (0.0001)	-4.8952 (0.0153)	48.0459 (0.7537)	39.3425 (0.0951)	1.37s
RLSMreinit		8.9425 (0.0081)	0.5828 (0.0002)	0.0195 (0.0000)	-4.8887 (0.0078)	47.7329 (0.3765)	39.3477 (0.0573)	1.52s

Table 14: Prices and Greeks for NLSM (with different number of training epochs), RLSM and RLSMreinit with standard deviations computed over 10 runs with different initializations on the same set of paths.

Conclusion

Based on a broad study of machine learning based approaches to approximate the solution of optimal stopping problems, we introduced two simple and powerful approaches, RLSM and RFQI. As state-of-the-art algorithms, they are very simple to implement and have convergence guarantees. Moreover, similarly to the neural network methods, they are easily scalable to high dimensions and there is no need to choose basis functions by hand. Furthermore, in our empirical study we saw that RLSM and RFQI are considerably faster than existing algorithms for high dimensional problems. In particular, up to 2400 (and 4800) times faster than LSM (and FQI respectively) with basis functions of order 2, 5 to 16 times faster than NLSM and 20 to 66 times faster than DOS.

In our Markovian experiments, RFQI often achieves the best results and if not, usually is very close to the best performing baseline method under consideration, reconfirming that reinforcement learning methods surpass backward induction methods.

In our non-Markovian experiments on fractional Brownian Motion, our randomized recurrent neural network algorithm RRLSM achieves similar results as the path-version of DOS, while requiring less training data and being much faster. However, this example also brought up the limitations of reinforcement learning based approaches, in particular of RFQI, which do not work well in those non-Markovian experiments.

In our non-Markovian experiments on rough Heston, we concluded that there is no need of using a recurrent neural network, since RLSM has similar results as RRLSM. This is also the case with DOS and pathDOS.

Overall, the speed of our algorithms is very promising for applications in high dimensions and with many discretization times, where existing methods might become impractical and where our methods show very reliable performance. To summarize, we suggest to use RFQI for Markovian problems, RLSM for non-Markovian processes which do not have a strong path-dependence, as the stock price of rough Heston and finally RRLSM for non-Markovian processes which have a strong path-dependence like fractional Brownian Motion.

Acknowledgments

The authors would like to thank Sebastian Becker, Patrick Cheredito, Blanka Horvath, Arnulf Jentzen, Hartmut Maennel and Louis Paulot for helpful feedback and discussions. In addition, we would like to warmly thank the quant team of Crédit Agricole CIB, and in particular Arthur Semin, Ryan Kurniawan and Wail El Allali for the great collaboration we had which considerably improved our paper. Thanks to this collaboration, in this current version, we provide the computation of the greeks, we improved the sensitivity to the randomness of the hidden layers of RLSM and we improved the proof of convergence of RLSM. Moreover, the authors would like to acknowledge support for this project from the Swiss National Science Foundation (SNF grant 179114).

References

Eduardo Abi Jaber and Omar El Euch. Multifactor Approximation of Rough Volatility Models. *SIAM Journal on Financial Mathematics*, 10(2):309-349, 2019.

Leif Andersen. A Simple Approach to the Pricing of Bermudan Swaptions in the Multi-Factor Libor Market Model. *Mathematical Finance*, 3(2):5-32, 1999.

Andrew Bakan. Representation of Measures with Polynomial Denseness in $L_p(\mathbb{R}, d\mu)$, $0 < p < \infty$, and its Application to Determinate Moment Problems. *Proceedings of the American Mathematical Society*, 136(10):3579-3589, 2008.

Vlad Bally and Gilles Pagès. A Quantization Algorithm for Solving Multi-Dimensional Discrete-Time Optimal Stopping Problems. *Bernoulli*, 9(6):1003-1049, 2003.

Vlad Bally, Gilles Pagès, and Jacques Printems. A Quantization Tree Method for Pricing and Hedging Multidimensional American Options. *Mathematical Finance*, 15(1):119-168, 2005.

Peter Bank and David Besslich. On Lenglart's Theory of Meyer-sigma-fields and El Karoui's Theory of Optimal Stopping. *arXiv:1810.08485*, 2019.

Jérôme Barraquand and Didier Martineau. Numerical Valuation of High Dimensional Multivariate American Securities. *The Journal of Financial and Quantitative Analysis*, 30 (3):383-405, 1995.

Christian Bayer, Peter Friz, and Jim Gatheral. Pricing under Rough Volatility. *Quantitative Finance*, 16(6):887-904, 2016.

Christian Bayer, Martin Eigel, Leon Sallandt, and Philipp Trunschke. Pricing HighDimensional Bermudan Options with Hierarchical Tensor Formats. *arXiv:2103.01934*, 2021.

Sebastian Becker, Patrick Cheredito, and Arnulf Jentzen. Deep Optimal Stopping. *Journal of Machine Learning Research*, 20:74, 2019.

Sebastian Becker, Patrick Cheridito, and Arnulf Jentzen. Pricing and Hedging American-Style Options with Deep Learning. *Journal of Risk and Financial Management*, 2020.

Dimitri P. Bertsekas and John N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, 1996. Bruno Bouchard and Xavier Warin. Monte-Carlo Valuation of American Options: Facts and New Algorithms to Improve Existing Methods. *Numerical Methods in Finance*, pages 215 – 255, 2012.

Phelim P. Boyle, Adam W. Kolkiewicz, and Ken Seng Tan. An improved simulation method for pricing high-dimensional american derivatives. *Mathematics and Computers in Simulation*, 62(3) : 315 – 322, 2003.

Mark Broadie and Paul Glasserman. A stochastic mesh method for pricing high-dimensional american options. *Journal of Computational Finance*, 7(4):35-72, 2004.

Weipeng Cao, Xizhao Wang, Zhong Ming, and Jinzhu Gao. A Review on Neural Networks with Random Weights. *Neurocomputing*, 275:278-287, 2018.

Jacques F. Carriere. Valuation of the Early-Exercise Price for Options using Simulations and Nonparametric Regression. *Insurance: Mathematics and Economics*, 19(1):19-30, 1996.

Shuhang Chen, Adithya M Devraj, Ana Bušić, and Sean Meyn. Zap Q-Learning for Optimal Stopping. *IEEE American Control Conference (ACC)*, pages 3920-3925, 2020.

Etienne Chevalier, Sergio Pulido, and Elizabeth Zúñiga. American Options in the Volterra Heston Model. *arXiv:2103.11734*, 2021.

Emmanuelle Clément, Damien Lamberton, and Philip Protter. An Analysis of the LongstaffSchwartz Algorithm for American Option Pricing. Technical report, Cornell University Operations Research and Industrial Engineering, 2001.

John C Cox, Stephen A Ross, and Mark Rubinstein. Option Pricing: A Simplified Approach. *Journal of financial Economics*, 1979.

Daniel Egloff. Monte Carlo Algorithms for Optimal Stopping and Statistical Learning. *The Annals of Applied Probability*, 15(2):1396-1432, 2005.

Daniel Egloff, Michael Kohler, and Nebojsa Todorovic. A Dynamic Look-Ahead Monte Carlo Algorithm for Pricing Bermudan Options. *The Annals of Applied Probability*, 17(4): 1138 – 1171, 2007.

Omar El Euch and Mathieu Rosenbaum. Perfect Hedging in Rough Heston Models. *The Annals of Applied Probability*, 28(6):3813-3856, 2018.

Omar El Euch, Masaaki Fukasawa, and Mathieu Rosenbaum. The Microstructural Foundations of Leverage Effect and Rough Volatility. *Finance and Stochastics*, 22(2):241-280, 2018.

Omar El Euch, Jim Gatheral, and Mathieu Rosenbaum. Roughening Heston. *Risk Management & Analysis in Financial Institutions eJournal*, 2019.

Nicole El Karoui. Les Aspects Probabilistes du Contrôle Stochastique. Ecole d'Été de Probabilités de Saint-Flour IX-1979, 1981. Hans Föllmer and Alexander Schied. Stochastic Finance: An Introduction in Discrete Time. De Gruyter, 2016.

Claudio Gallicchio, Alessio Micheli, and Luca Pedrelli. Deep Reservoir Computing: A Critical Experimental Analysis. *Neurocomputing*, 268:87 - 99, 2017.

Diego Garcia. Convergence and Biases of Monte Carlo Estimates of American Option Prices using a Parametric Exercise Rule. *Journal of Economic Dynamics and Control*, 27(10): 1855 – 1879, 2003.

Jim Gatheral, Thibault Jaisson, and Mathieu Rosenbaum. Volatility is Rough. *Quantitative Finance*, 18(6):933-949, 2018.

Jim Gatheral, Paul Jusselin, and Mathieu Rosenbaum. The Quadratic Rough Heston Model and the Joint S&P 500/VIX Smile Calibration Problem. *arXiv:2001.01789*, 2020.

Emmanuel Gobet, Jean-Philippe Lemor, and Xavier Warin. A Regression-Based Monte Carlo Method to Solve Backward Stochastic Differential Equations. *The Annals of Applied Probability*, 15(3):2172-2202, 2005.

L. Gonon and J. Ortega. Reservoir Computing Universality With Stochastic Inputs. *IEEE Transactions on Neural Networks and Learning Systems*, 31:100-112, 2020.

Alexander N. Gorban, Ivan Yu. Tyukin, Danil V. Prokhorov, and Konstantin I. Sofeikov. Approximation with Random Bases: Pro et Contra. *Information Sciences*, 364-365:129-145, 2016.

Hamza Hanbali and Daniel Linders. American-Type Basket Option Pricing: a Simple Two-Dimensional Partial Differential Equation. *Quantitative Finance*, 19(10):1689-1704, 2019.

Martin B. Haugh and Leonid Kogan. Pricing American Options: A Duality Approach. *Operations Research*, 52(2):258-270, 2004.

Steven L. Heston. A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The review of financial studies*, 6(2):327-343, 1993.

Kurt Hornik. Approximation Capabilities of Multilayer Feedforward Networks. *Neural networks*, 4(2):251-257, 1991.

Guang-Bin Huang, Lei Chen, Chee Kheong Siew, et al. Universal Approximation using Incremental Constructive Feedforward Networks with Random Hidden Nodes. *IEEE Transactions on Neural Networks*, 17(4):879-892, 2006.

Shashi Jain and Cornelis W. Oosterlee. The Stochastic Grid Bundling Method: Efficient Pricing of Bermudan Options and their Greeks. *Applied Mathematics and Computation*, 269:412-431, 2015.

Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement Learning: A Survey. *Journal of artificial intelligence research*, 4:237-285, 1996. Michael Kohler, Adam Krzyżak, and N

- Zaklina Todorovic. Pricing of High-Dimensional American Options by Neural Networks. Wiley-Blackwell: Mathematical Finance, 2010.
- Anastasia Kolodko and John Schoenmakers. Iterative Construction of the Optimal Bermudan Stopping Time. *Finance and Stochastic*, 10:27-49, 2004.
- Bernard Lapeyre and Jérôme Lelong. Neural Network Regression for Bermudan Option Pricing. *Monte Carlo Methods and Applications*, 2021.
- Pascal Létourneau and Lars Stentoft. Simulated Greeks for American Options. Available at SSRN 3503889, 2019.
- Yuxi Li, Csaba Szepesvari, and Dale Schuurmans. Learning Exercise Policies for American Options. In *International Conference on Artificial Intelligence and Statistics*, 2009.
- Giulia Livieri, Saad Mouti, Andrea Pallavicini, and Mathieu Rosenbaum. Rough Volatility: Evidence from Option Prices. *IJSE transactions*, 50(9):767-776, 2018.
- Francis A. Longstaff and Eduardo S. Schwartz. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Review of Financial Studies*, 2001.
- Mantas Lukoševičius and Herbert Jaeger. Reservoir Computing Approaches to Recurrent Neural Network Training. *Computer Science Review*, 3(3):127-149, 2009.
- Gilles Pagès. *Numerical Probability: An Introduction with Applications to Finance*. Springer, 2018.
- Huyên Pham. Optimal Stopping, Free Boundary, and American Option in a Jump-Diffusion Model. *Applied Mathematics and Optimization*, 1997.
- Chris Rogers. Monte Carlo Valuation of American Options. *Mathematical Finance*, 12(3): 271 – 286, 2002.
- Chris Rogers. Dual Valuation and Hedging of Bermudan Options. *SIAM Journal on Financial Mathematics*, 1(1):604-608, 2010.
- Anton Maximilian Schäfer and Hans Georg Zimmermann. Recurrent Neural Networks are Universal Approximators. In *Artificial Neural Networks - ICANN 2006*, pages 632-640. Springer Berlin Heidelberg, 2006.
- Benjamin Schrauwen, David Verstraeten, and Jan Van Campenhout. An Overview of Reservoir Computing: Theory, Applications and Implementations. In *European Symposium on Artificial Neural Networks*, pages 471-482, 2007.
- Martin Schweizer. On bermudan options. *Advances in Finance and Stochastics: Essays in Honour of Dieter Sondermann*, 2002.
- Lars Stentoft. Convergence of the Least Squares Monte Carlo Approach to American Option Valuation. *Management Science*, 50(9):1193-1203, 2004. Richard S. Sutton and Andrew G. Barto. Reinforcement

Learning: An Introduction. MIT press, 2018.

James A. Tilley. Valuing American Options in a Path Simulation Model. In Transactions of the Society of Actuaries, 1993.

John Tsitsiklis and Benjamin Van Roy. Optimal Stopping of Markov Processes: Hilbert Space Theory, Approximation Algorithms, and an Application to Pricing High-Dimensional Financial Derivatives. IEEE Transactions on Automatic Control, 44:1840-1851, 1997.

John Tsitsiklis and Benjamin Van Roy. Regression Methods for Pricing Complex American Style Options. IEEE Transactions on Neural Networks, 12(4):694-703, 2001.

David Verstraeten, Benjamin Schrauwen, Michiel d'Haene, and Dirk Stroobandt. An Experimental Unification of Reservoir Computing Methods. Neural Networks, 20(3):391-403, 2007.

Huizhen Yu and Dimitri P. Bertsekas. Q-Learning Algorithms for Optimal Stopping Based on Least Squares. European Control Conference, 2007.

Daniel Zanger. Quantitative Error Estimates for a Least-Squares Monte Carlo Algorithm for American Option Pricing. Finance and Stochastics, 17, 072013.

Daniel Z Zanger. Convergence of a Least-Squares Monte Carlo Algorithm for Bounded Approximating Sets. Applied Mathematical Finance, 16(2):123-150, 2009.

Daniel Z Zanger. Convergence of a Least-Squares Monte Carlo Algorithm for American Option Pricing with Dependent Sample Data. Mathematical Finance, 28(1):447-479, 2018.

Daniel Z. Zanger. General Error Estimates for the Longstaff-Schwartz Least-Squares Monte Carlo Algorithm. 45(3), 2020.

Rui Zhang, Yuan Lan, Guang-Bin Huang, and Zong-Ben Xu. Universal Approximation of Extreme Learning Machine With Adaptive Growth of Hidden Nodes. IEEE Transactions on Neural Networks and Learning Systems, 23(2):365-371, 2012.

Appendix A. Convergence of the Randomized Least Square Monte Carlo (RLSM)

We first introduce some technical notation that will be helpful for the proofs. Then we describe the steps from the theoretical idea of RLSM to its implementable version that was presented in Section 2.7. These descriptions and proofs are based on (Tsitsiklis and Van Roy, 2001; Clément et al., 2001), in particular, our theoretical results are a direct consequence of these works and the universal approximation theorem of Zhang et al. (2012). Nevertheless, we give a detailed description here for completeness.

A.1 Definitions

We assume to have a sequence of infinitely many random basis functions $\phi = (\phi_k)_{k \geq 1}$, where each ϕ_k is of the form

$$\phi_k : \mathbb{R}^d \rightarrow \mathbb{R}, x \mapsto \phi_k(x) := \sigma \left(\alpha_k^\top x + \beta_k \right),$$

with σ a bounded activation function, $\alpha_k \in \mathbb{R}^d$ and $\beta_k \in \mathbb{R}$. The parameters α_k and β_k have i.i.d. entries with a standard Gaussian distribution, hence the name random basis functions. With $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{P}})$ we denote the probability space on which the random weights are defined. For each $K \in \mathbb{N}$ we define the operator Φ_K acting on $\theta = (\theta_1, \dots, \theta_K) \in \mathbb{R}^K$ by

$$(\Phi_K \theta)(x) := \theta^\top \phi(x) := \sum_{k=1}^K \theta_k \phi_k(x).$$

In particular, Φ_K is the operator producing a linear combination of the first K random basis functions. We assume to have a Markovian, discrete time stochastic process $X = (X_0, \dots, X_N)$ defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_n)_{n=0}^N, \mathbb{P})$. In particular, each X_n is a \mathcal{F}_{n-} -measurable random variable. We assume that there exists an absolutely continuous measure $\mathbb{Q} \ll \mathbb{P}$, the pricing measure, and that the distribution of X_n under \mathbb{Q} is π_n . For expectations with respect to these random variables under \mathbb{Q} , we write $\mathbb{E}[\cdot]$. For $0 \leq n \leq N$ we use the norm

$$\|f\|_{\pi_n}^2 := \mathbb{E} \left[|f(X_n)|_2^2 \right] = \int_{\mathbb{R}} |f(x)|_2^2 d\pi_n(x),$$

where $|\cdot|_2$ is the Euclidean norm and f a measurable function. We introduce the operators E_n and Π_n^K defined by

$$\begin{aligned} (E_n J)(x) &:= \mathbb{E} [J(X_{n+1}) \mid X_n = x], \\ (\Pi_n^K J) &:= \arg \min_{\Phi_K \theta} \|J - \Phi_K \theta\|_{\pi_n} \end{aligned}$$

for $J \in L^2(\pi_n)$. With \hat{E}_n we denote the one-sample approximation of E_n , i.e. $(\hat{E}_n J)(X_n) = J(X_{n+1})$, which is better understood in terms of a realization of $x = (x_0, \dots, x_N)$ of X as $(\hat{E}_n J)(x_n) = J(x_{n+1})$. Moreover, $\hat{\Pi}_n^K$ is the Monte Carlo approximation of Π_n^K , i.e. if x_n^1, \dots, x_n^m are i.i.d. samples of π_n , then $(\hat{\Pi}_n^K J) := \arg \min_{\Phi_K \theta} \frac{1}{m} \sum_{i=1}^m (J(x_n^i) - (\Phi_K \theta)(x_n^i))^2$. In the following, we write Π_n and $\hat{\Pi}_n$ whenever K is fixed.

The payoff at any exercise time n is given by $g(X_n)$ and we assume that they are square integrable, i.e. $\|g(X_n)\|_{\pi_n} < \infty$.

A.2 Theoretical Description of RLSM

We first introduce the exact algorithm to compute the continuation value and then give definitions of the 2-step approximation of this exact algorithm. The first step is to introduce projections on the subspace of functions spanned by Φ_K , while assuming that (conditional) expectations can be computed exactly. We call this the idealized algorithm. We remark that also the projection itself is based on minimizing an expectation. The second step is to introduce Monte Carlo and one-sample approximations of the projections and (conditional) expectations using m sample paths. This we call the implementable algorithm, since it can actually be implemented. Our goal is then to show that the price computed with those two approximation steps converges to the true price, when K and m increase to infinity.

A.2.1 Exact Algorithmic

The continuation value is the expected discounted payoff at the current time conditioned on a decision not to exercise the option now. The exact algorithmic definition of the continuation value is defined backwards step-wise as in (Tsitsiklis and Van Roy, 2001) as

$$\begin{cases} Q_{N-1} & := \alpha E_{N-1} g \\ Q_n & := \alpha E_n \max(g, Q_{n+1}). \end{cases}$$

A.2.2 Idealized Algorithm

Our idealized algorithm to compute the continuation value, written similar as in (Tsitsiklis and Van Roy, 2001), is defined for fixed K as

$$\begin{cases} \tilde{Q}_{N-1}^K & := \alpha E_{N-1} P_N^K \\ \tilde{Q}_n^K & := \alpha E_n P_{n+1}^K \end{cases}$$

where

$$\begin{cases} P_N^K & := g \\ P_n^K & := g \mathbf{1}_{g \geq \alpha \Pi_n^K E_n P_{n+1}^K} + \alpha E_n P_{n+1}^K \mathbf{1}_{g < \alpha \Pi_n^K E_n P_{n+1}^K}. \end{cases}$$

In particular, P_n^K can be interpreted as the choice of the algorithm at time step n , to either execute and take the payoff or to continue with the expected discounted future payoff. We drop the superscript K

whenever it is clear from the context which K is meant. We see from this equation, that the difference from the idealized algorithm in (Tsitsiklis and Van Roy, 2001, described in (1) and before Theorem 1) is, that we use the \tilde{Q}_{n+1} instead of its linear approximation with the random basis functions $\Pi_n \tilde{Q}_{n+1}$, if we decide to continue. However, the decision to continue or to stop, is still based on the approximation $\Pi_n \tilde{Q}_{n+1}$ as it is also the case in the idealized algorithm (Tsitsiklis and Van Roy, 2001). If the linear approximation is exact, both algorithms produce the same output, but if it is not exact, our algorithm uses a better approximation of the continuation value.

A.2.3 Implementable Algorithm

Finally, we define our implementable algorithm to compute the continuation value, which is an approximation of the idealized algorithm using the approximations \hat{E}_n and $\hat{\Pi}_n^K$ as

$$\begin{cases} \hat{Q}_{N-1}^K & := \alpha \hat{E}_{N-1} \hat{P}_N^K \\ \hat{Q}_n^K & := \alpha \hat{E}_n \hat{P}_{n+1}^K \end{cases}$$

where

$$\begin{cases} \hat{P}_N^K & := g \\ \hat{P}_n^K & := g \mathbf{1}_{g \geq \alpha \hat{\Pi}_n^K \hat{E}_n \hat{P}_{n+1}^K} + \alpha \hat{E}_n \hat{P}_{n+1}^K \mathbf{1}_{g < \alpha \hat{\Pi}_n^K \hat{E}_n \hat{P}_{n+1}^K} \end{cases}$$

Also here we drop the superscript K whenever it is clear from the context which K is meant.

A.3 Preliminary Results

The following result is proven in (Zhang et al., 2012, Theorem 3) and states, that the error of the approximation of any integrable function by randomized neural networks converges $\tilde{\mathbb{P}}$ -a.s. to 0 as the number of hidden nodes goes to infinity, where $\tilde{\mathbb{P}}$ is the probability measure associated with the random weights.

Theorem 4. Let $0 \leq n \leq N - 1$ and J be an square integrable function, i.e. $\|J\|_{\pi_n} < \infty$, then

$$\left\| \Pi_n^K J - J \right\|_{\pi_n} \xrightarrow[K \rightarrow \infty]{\tilde{\mathbb{P}}\text{-a.s.}} 0.$$

A.4 Convergence Results

The price of the Bermudan approximation of the American option can be expressed with the exact algorithm as

$$U_0 := \max(g(X_0), Q_0(X_0)),$$

the price computed with the idealized algorithm is

$$U_0^K := \max(g(X_0), \tilde{Q}_0^K(X_0))$$

and the price computed with the implementable algorithm is

$$U_0^{K,m} := \max\left(g(X_0), \frac{1}{m} \sum_{i=1}^m \hat{Q}_0^K(x_0, x_1^i, \dots, x_N^i)\right).$$

We provide two different convergence results with two different assumptions. The first result is based on (Clément et al., 2001) which needs a technical assumption that might not be satisfied in general. The second result is based on (Zanger, 2020), which replaces this assumption by a stronger integrability assumption on the payoff. A.4.1 Convergence Results based on Clément Et al. (2001)

Combining the following two results, convergence of U_0^{K,m_K} to U_0 as $K \rightarrow \infty$ can be established by choosing a suitable sequence $(m_K)_{K \geq 1}$, under the assumption that $g(X_n)$ is square integrable for all $0 \leq n \leq N$

Theorem 5. The idealized price U_0^K converges to the correct price U_0 $\tilde{\mathbb{P}}$ -a.s. as $K \rightarrow \infty$.

Theorem 6. We assume that $\mathbb{Q}[\alpha \Pi_n^K E_n P_{n+1}^K(X_n) = g(X_n)] = 0$ for all $0 \leq n \leq N - 1$. Then the implementable price $U_0^{K,m}$ converges almost surely to the idealized price U_0^K as $m \rightarrow \infty$.

The proofs are a direct consequence of Clément et al. (2001).

Proof [Theorems 5 and 6] The proofs are implied by the results presented in (Clément et al., 2001, Section 3). We only need to establish that their assumption A_1 is satisfied. The assumption A_2 is actually not needed, as explained below.

Assumption A_1 is that $(\phi_k(X_n))_{k \geq 1}$ is total in $L^2(\sigma(X_n))$ for every $1 \leq n \leq N - 1$, which is used to show that $\left\| \prod_n^K Q_n - Q_n \right\|_{\pi_n}$ converges to 0. We replace this assumption by our Theorem 4, which therefore yields \mathbb{P} -almost sure convergence in the result.

Assumption A_2 is that for every $1 \leq n \leq N$ and every $K > 0$, if $\sum_{k=0}^K \lambda_k \phi_k(X_n) = 0$ almost surely, then all $\lambda_k = 0$. This assumption is actually only needed for the projection weights to be uniquely defined, such that they can be expressed by the closed-form ordinary least squares formula. Otherwise, if this assumption is not satisfied, there exist several weight vectors θ , which all define the same projection $\Phi_K \theta$ minimizing the projection objective. By Gram-Schmidt, we can generate an orthonormal basis $(\tilde{\phi}_k)_{1 \leq k \leq \tilde{K}(K)}$ of the linear subspace of L^2 that is spanned by $(\phi_k)_{1 \leq k \leq K}$, with $\tilde{K}(K) \leq K$. By its definition, $(\tilde{\phi}_k)_{1 \leq k \leq \tilde{K}(K)}$ satisfies assumption A_2 and therefore, the results of (Clément et al., 2001, Section 3) can be applied. Finally, we note that the projections are the same, no matter whether $(\tilde{\phi}_k)_{1 \leq k \leq \tilde{K}}$ or $(\phi_k)_{1 \leq k \leq K}$ are used to describe the space that is spanned. We are interested in the convergence of the price. Considering the definition (12), we see that the price depends only on the projection but not on the used weights. Therefore, we can conclude that the same statements hold with our originally defined random basis functions $(\phi_k)_{1 \leq k \leq K}$.

The technical assumption that $\mathbb{Q} \left[\alpha \prod_n^K E_n P_{n+1}^K(X_n) = g(X_n) \right] = 0$ for all $0 \leq n \leq N - 1$ of the result of Clément et al. (2001) that shows up in Theorem 6 is not always satisfied. In particular, it is easy to construct examples of finite probability spaces, where this is not the case. Indeed, consider the easiest possible case of probability space which is a singleton, with a (deterministic) constant stock price without discounting, then $\mathbb{Q} \left[\prod_n^K E_n P_{n+1}^K(X_n) = g(X_n) \right] = 1$. Therefore, in the next section, we provide a different proof based on the work of Zanger (2020), which replaces this assumption by a slightly stronger integrability assumption on the payoff process.

A.4.2 Convergence Results Based On Zanger (2020)

After the work of Clément et al. (2001), improved theoretical guarantees to the original Least Squares Monte Carlo algorithm (LSM) have been proposed, such as (Stentoft, 2004; Egloff, 2005; Gobet et al., 2005). An important improvement of the convergence results is done in (Zanger, 2009, 2013, 2018, 2020). In particular, Zanger (2009) raised the issue that Clément et al. (2001) have additional restrictions on the law of the underlying Markov process such as the assumption in Theorem 6 mentioned above. Zanger (2009) proposed a generalized LSM algorithm and provides a proof of convergence in probability (Zanger, 2009, Theorem 5.1). In this theorem, he does not need the condition of Clément et al. (2001),

but needs the payoff to be almost surely bounded (Zanger, 2009, Definition 5.1 and 5.2). Zanger (2013) provides error estimates (convergence rates), even when the underlying process and payoff process are not necessarily in L^∞ . Later, Zanger (2018) provides a convergence result (Zanger, 2018, Corollary 5.5) without the assumption in Theorem 6 of Clément et al. (2001), but with a bounded payoff process. However, this time, he provides almost sure convergence instead of convergence in probability. Finally, in his last paper, Zanger (2020) replaces the assumption of having a bounded payoff process by a condition on its moments (Zanger, 2020, Corollary 1). We use this last result to prove our second convergence theorem.

Theorem 7. Assume that there exists some $2 < p \leq \infty$ such that

$$M_p := \max_{1 \leq n \leq N} \|g(X_n)\|_{L^p}^p < \infty$$

and that all payoffs are non-negative. Moreover, assume that we use the truncated versions of the payoffs $g(X_n)$ in Algorithm 1 as well as the truncated versions of the randomized neural networks, with truncation level $1 \leq \lambda < \infty$. Then

$$\mathbb{E} \left[|U_0^{K,m} - U_0| \right] \xrightarrow[K, m \rightarrow \infty]{\mathbb{P}\text{-a.s.}} 0,$$

when choosing $\lambda = m^{1/8}$.

The proof is a direct consequence of (Zanger, 2020, Corollary 1).

Proof Let us fix the number of paths m and the number of random basis functions K . Then (Zanger, 2020, Corollary 1) implies that

$$\begin{aligned} \mathbb{E} \left[|U_0^{K,m} - U_0| \right] &\leq 6^N \left(\frac{C\lambda^2 \left(\sqrt{\nu c_0} \log^{\frac{1}{2}}(m) + \log^{\frac{1}{2}}(C_0) \right)}{\sqrt{m}} \right. \\ &\quad \left. + 4\sqrt{\varepsilon} + \max_{n=1, \dots, N-1} \left(\inf_{f \in \mathcal{B}_n^{K, \lambda}} \|f - Q_n\|_{\pi_n} \right) + \left(\frac{8M_p \lambda^{(2-p)}}{p-2} \right)^{1/2} \right), \end{aligned}$$

where $C_0 = C(c_0\nu + 1)^4 (C\lambda^4)^{2\nu(1+c_0)}$, $c_0 = 2(N+1) \log_2(e(N+1))$ and C is a numerical constant with $1 \leq C < \infty$, and $\varepsilon \geq 0$ as defined in (Zanger, 2020, Equation 13). Here, ν is the Vapnik-Chervonenkis (VC) dimension of the set of randomized neural networks, which is finite according to (Zanger, 2020, Remark 8). For each exercise time $1 \leq n \leq N-1$ the set $\mathcal{B}_n^{K, \lambda}$ is defined to be the set of all λ -truncated randomized neural networks using the first K random basis functions (i.e. any truncated version of a linear combinations of the basis functions $(\phi_k)_{1 \leq k \leq K}$). In particular

$$\mathcal{B}_n^{K,\lambda} = \{\mathcal{T}_\lambda f \mid f \in \text{span}\{\phi_1, \dots, \phi_K\}\},$$

where \mathcal{T}_λ is the operator truncating a function at λ . Note that for any function f , we have that

$$\|(\mathcal{T}_\lambda f - Q_n) 1_{\{|Q_n| < \lambda\}}\|_{\pi_n} \leq \|(f - Q_n) 1_{\{|Q_n| < \lambda\}}\|_{\pi_n}$$

Therefore,

$$\begin{aligned} \inf_{f \in \mathcal{B}_n^{K,\lambda}} \|(f - Q_n) 1_{\{|Q_n| < \lambda\}}\|_{\pi_n} &\leq \inf_{f \in \text{span}\{\phi_1, \dots, \phi_K\}} \|(f - Q_n) 1_{\{|Q_n| < \lambda\}}\|_{\pi_n} \\ &\leq \inf_{f \in \text{span}\{\phi_1, \dots, \phi_K\}} \|f - Q_n\|_{\pi_n} = \|\Pi_n^K Q_n - Q_n\|_{\pi_n} \end{aligned}$$

Hence, we can now bound the approximation error with truncated randomized neural networks by

$$\begin{aligned} \inf_{f \in \mathcal{B}_n^{K,\lambda}} \|f - Q_n\|_{\pi_n} &\leq \inf_{f \in \mathcal{B}_n^{K,\lambda}} \left(\|(f - Q_n) 1_{\{|Q_n| < \lambda\}}\|_{\pi_n} + \|(f - Q_n) 1_{\{|Q_n| \geq \lambda\}}\|_{\pi_n} \right) \\ &\leq \inf_{f \in \mathcal{B}_n^{K,\lambda}} \|(f - Q_n) 1_{\{|Q_n| < \lambda\}}\|_{\pi_n} + \sup_{f \in \mathcal{B}_n^{K,\lambda}} \|(f - Q_n) 1_{\{|Q_n| \geq \lambda\}}\|_{\pi_n} \\ &\leq \|\Pi_n^K Q_n - Q_n\|_{\pi_n} + 2 \|Q_n 1_{\{|Q_n| \geq \lambda\}}\|_{\pi_n}, \end{aligned}$$

where in the last inequality we used that functions in $\mathcal{B}_n^{K,\lambda}$ are truncated at λ implying that they are bounded by $|Q_n|$ on the set $\{|Q_n| \geq \lambda\}$. Moreover, we can choose $\varepsilon = 1/m$, replace $\lambda = m^{1/8}$ and simplify all expressions by using one common constant \tilde{C} to rewrite

$$\begin{aligned} \mathbb{E} \left[|U_0^{K,m} - U_0| \right] &\leq \tilde{C} \left(\frac{\log^{\frac{1}{2}}(m)}{m^{1/4}} + \frac{1}{\sqrt{m}} \right. \\ &\quad \left. + \max_{n=1, \dots, N-1} \left(\|\Pi_n^K Q_n - Q_n\|_{\pi_n} + 2 \|Q_n 1_{\{|Q_n| \geq m^{1/8}\}}\|_{\pi_n} \right) + m^{\frac{2-p}{16}} \right). \end{aligned}$$

Now it suffices to note that the terms $\|\Pi_n^K Q_n - Q_n\|_{\pi_n}$ converge to 0 as $K \rightarrow \infty$ by Theorem 4, the terms $\|Q_n 1_{\{|Q_n| \geq m^{1/8}\}}\|_{\pi_n}$ converge to 0 as $m \rightarrow \infty$ by dominated convergence and the remaining terms trivially converge to 0 as $m \rightarrow \infty$.

Appendix B. Convergence of the Randomized Fitted Q-Iteration (RFQI)

Similar as in Section A, we first introduce some additional technical notation needed for the proofs. Then we describe the steps from the theoretical idea of RFQI to its implementable version that was presented in Section 3. In contrast to Section A, the algorithms described here are applied simultaneously for all times. Again, the proof is a direct consequence of (Tsitsiklis and Van Roy, 2001) and the universal approximation theorem of Zhang et al. (2012), but is given in detail for completeness.

B.1 Definitions

In Section 6, Tsitsiklis and Van Roy (2001) introduced a reinforcement learning version of their optimal stopping algorithm, where a stopping function is learned that generalizes over time. In particular, instead of learning a different function for each time step, a single function that gets the time as input is learned with an iterative scheme. In accordance with this, the random basis functions are redefined such that they also take time as input

$$\begin{aligned} \phi_k \mathbb{R}^d \times \{0, \dots, N-1\} &\rightarrow \mathbb{R}, \\ (x, n) &\mapsto \phi_k(x, n) := \sigma \left(\alpha_k^\top(x, n)^\top + \beta_k \right), \end{aligned}$$

with $\alpha_k \in \mathbb{R}^{d+1}$ and $\beta_k \in \mathbb{R}$. For $0 \leq n \leq N-1$ let $\Phi_{K,n}$ be defined similarly to before as

$$(\Phi_{K,n}\theta)(x) := \theta^\top \phi(x, n) := \sum_{k=1}^K \theta_k \phi_k(x, n),$$

for $\theta \in \mathbb{R}^K$ and $x \in \mathbb{R}^d$. Moreover, let $\Phi_k := (\Phi_{K,0}, \dots, \Phi_{K,N-1})$, such that

$$\Phi_k \theta := (\Phi_{K,0}\theta, \dots, \Phi_{K,N-1}\theta).$$

In the following, we consider the product space $(L^2)^N := L^2(\pi_0) \times \dots \times L^2(\pi_{N-1})$, which is the space on which the functions for all time steps can be defined concurrently. For $J = (J_0, \dots, J_{N-1}) \in (L^2)^N$ we define the norm

$$\|J\|_\pi := \frac{1}{N} \sum_{n=0}^{N-1} \|J_n\|_{\pi_n},$$

where $\|\cdot\|_{\pi_n}$ is as defined in Section A. Let us define the projection operator Π^K as

$$\left(\Pi^K J\right) := \arg \min_{\Phi_K \theta} \|\Phi_K \theta - J\|_\pi,$$

for $J = (J_0, \dots, J_{N-1}) \in (L^2)^N$. Finally, we define the operator

$$H : (L^2)^N \rightarrow (L^2)^N, \quad \begin{pmatrix} J_0 \\ \vdots \\ J_{N-2} \\ J_{N-1} \end{pmatrix} \mapsto \begin{pmatrix} \alpha E_0 \max(g, J_1) \\ \vdots \\ \alpha E_{N-2} \max(g, J_{N-1}) \\ \alpha E_{N-1} g \end{pmatrix},$$

where E_n and g are as defined in the previous sections.

B.2 Theoretical Description of the Algorithm

Based on the definitions in Section A.2, we first introduce the exact algorithm and then give the two-step approximation with the idealized and implementable algorithm.

B.2.1 Exact Algorithm

Let Q_n as defined in (10), then $Q := (Q_0, \dots, Q_{N-1})$ satisfies $Q = HQ$ by definition. In particular, Q is a fixed point of H . It was shown in (Tsitsiklis and Van Roy, 2001, Section 6) that H is a contraction with respect to the norm $\|\cdot\|_\pi$ with contraction factor α . Hence, the Banach fixed point theorem implies that there exists a unique fixed point, which therefore has to be Q , and that for any starting element $J^0 \in (L^2)^N$, J^i converges to Q in $(L^2)^N$ as $i \rightarrow \infty$, where $J^{i+1} := HJ^i$. This yields a way to find the exact algorithm Q iteratively.

B.2.2 Idealized Algorithm

The combined operator $\Pi^K H$ is a contraction on the space $\Pi^K (L^2)^N$, since the projection operator is a non-expansion as outlined in (Tsitsiklis and Van Roy, 2001, Section 6). The idealized algorithm is then defined as the unique fixed point \tilde{Q}^K of $\Pi^K H$, which can again be found by iteratively applying this operator to an arbitrary starting point. Since any element in $\Pi^K (L^2)^N$ is given as $\Phi_K \theta$ for some weight vector $\theta \in \mathbb{R}^K$, this iteration can equivalently be given as iteration on the weight vectors. To do this, let us assume without loss of generality that $(\phi_k)_{1 \leq k \leq K}$ are independent (if not, see the strategy

in Proof of Theorem 5 and 6). Then, given some starting weight vector θ_K^0 , the iterative application of $\Pi^K H$ defines the weight vectors

$$\begin{aligned} \theta_K^{i+1} := & \alpha \left(\mathbb{E} \left[\sum_{n=0}^{N-1} \phi_{1:K}^\top (X_n, n) \phi_{1:K} (X_n, n) \right] \right)^{-1} \\ & \cdot \mathbb{E} \left[\sum_{n=0}^{N-1} \phi_{1:K}^\top (X_n, n) \cdot \max \left(g(X_{n+1}), \left(\Phi_{K,n+1} \theta_K^i \right) (X_{n+1}) \right) \right], \end{aligned}$$

where $\phi_{1:K} = (\phi_1, \dots, \phi_K)$. This closed-form solution is exactly the ordinary least squares (OLS) formula and this result was shown in (Tsitsiklis and Van Roy, 2001, Section 6).

B.2.3 Implementable Algorithm

An implementable version of this iteration is defined by the Monte Carlo approximation of the weight vectors. In particular, we assume that m realizations $(x_0^j, \dots, x_N^j)_{1 \leq j \leq m}$ of X are sampled and fixed for all iterations. Then for $\hat{\theta}_{K,m}^0 = \theta_K^0$ we iteratively define

$$\begin{aligned} \hat{\theta}_{K,m}^{i+1} := & \alpha \left(\sum_{j=1}^m \sum_{n=0}^{N-1} \phi_{1:K}^\top (x_n^j, n) \phi_{1:K} (x_n^j, n) \right)^{-1} \\ & \cdot \sum_{j=1}^m \sum_{n=0}^{N-1} \phi_{1:K}^\top (x_n^j, n) \cdot \max \left(g(x_{n+1}^j), \left(\Phi_{K,n+1} \hat{\theta}_{K,m}^i \right) (x_{n+1}^j) \right), \end{aligned}$$

which in turn defines $\hat{Q}^{K,m,i} := \Phi_K \hat{\theta}_{K,m}^i$. As explained in (Tsitsiklis and Van Roy, 2001, Section 6), this implementable iteration can equivalently be described as iteratively applying the operator $\widehat{\Pi^K H}$. Here $\widehat{\Pi^K H}$ is identical to $\Pi^K H$, but with the measures π_n replaced by the empirical measures $\hat{\pi}_n$ arising from the sampled trajectories $(x_0^j, \dots, x_N^j)_{1 \leq j \leq m}$. Hence, $\widehat{\Pi^K H}$ is also a contraction and Banach's fixed point theorem implies convergence to the unique fixed point

$$\hat{Q}^{K,m,i} \xrightarrow{i \rightarrow \infty} \hat{Q}^{K,m} =: \Phi_K \hat{\theta}_{K,m}^*.$$

We note that this also implies that $\hat{\theta}_{K,m}^i \xrightarrow{i \rightarrow \infty} \hat{\theta}_{K,m}^*$.

B.3 Convergence Result

In the following, we show that prices of Bermudan options computed with the two approximation steps of the exact algorithm converge to the correct price, as $K, m \rightarrow \infty$. The prices are defined similarly as in Section A.4. Hence, it is enough to show that $\hat{Q}^{K, m_i, i}$ converges to \tilde{Q}^K as $i \rightarrow \infty$ and that \tilde{Q}^K converges to Q as $K \rightarrow \infty$.

Theorem 8. \tilde{Q}^K converges $\tilde{\mathbb{P}}$ -a.s. to Q as $K \rightarrow \infty$, i.e.

$$\left\| \tilde{Q}^K - Q \right\|_{\pi} \xrightarrow[\pi, K \rightarrow \infty]{\tilde{\mathbb{P}}\text{-a.s.}} 0.$$

Proof First, let us recall (Tsitsiklis and Van Roy, 2001, Theorem 3), which states that for $0 < \kappa < 1$ the contraction factor of $\Pi^K H$, we have

$$\left\| \tilde{Q}^K - Q \right\|_{\pi} \leq \frac{1}{\sqrt{1 - \kappa^2}} \left\| \Pi^K Q - Q \right\|_{\pi}.$$

Now remark that since Π^K is a non-expansion and H a contraction with factor α , we have $\kappa \leq \alpha < 1$. Therefore, for every K we have

$$\left\| \tilde{Q}^K - Q \right\|_{\pi} \leq \frac{1}{\sqrt{1 - \alpha^2}} \left\| \Pi^K Q - Q \right\|_{\pi}.$$

Finally, we remark that Theorem 4 holds equivalently for the norm $\| \cdot \|_{\pi}$, since the universal approximation theorem can equivalently be applied to the functions with the combined input (x, n) . Hence, the right hand side of (14) converges to 0 $\tilde{\mathbb{P}}$ -a.s. as $K \rightarrow \infty$.

We recall that the weight vectors $\hat{\theta}_{K, m}^i$ are random variables since they depend on the m sampled trajectories of X .

Lemma 9. For any fixed $i \in \mathbb{N}$ we have that $\hat{\theta}_{K, m}^i$ converges to θ_K^i \mathbb{Q} -a.s. as $m \rightarrow \infty$.

Proof The proof follows the proof of (Tsitsiklis and Van Roy, 2001, Theorem 2). We introduce the intermediate weight as

$$\begin{aligned} \tilde{\theta}_{K, m}^i &:= \alpha \left(\sum_{j=1}^m \sum_{n=0}^{N-1} \phi_{1:K}^{\top} \left(x_n^j, n \right) \phi_{1:K} \left(x_n^j, n \right) \right)^{-1} \\ &\cdot \sum_{j=1}^m \sum_{n=0}^{N-1} \phi_{1:K}^{\top} \left(x_n^j, n \right) \cdot \max \left(g \left(x_{n+1}^j \right), \left(\Phi_{K, n+1} \theta_K^{i-1} \right) \left(x_{n+1}^j \right) \right). \end{aligned}$$

Then it is clear that $\tilde{\theta}_{K, m}^i$ converges to θ_K^i \mathbb{Q} -a.s. as $m \rightarrow \infty$, by the strong law of large numbers. Hence,

$\delta_i(m) := \left| \tilde{\theta}_{K,m}^i - \theta_K^i \right|_2$ converges to 0 \mathbb{Q} -a.s. Moreover, for suitably chosen random variables $A_i(m)$ that remain bounded as $m \rightarrow \infty$, we have

$$\hat{\theta}_{K,m}^i - \tilde{\theta}_{K,m}^i = A_i(m) \left| \hat{\theta}_{K,m}^{i-1} - \theta_K^{i-1} \right|_2.$$

Therefore we have by the triangle inequality

$$\left| \hat{\theta}_{K,m}^i - \theta_K^i \right|_2 \leq \delta_i(m) + A_i(m) \left| \hat{\theta}_{K,m}^{i-1} - \theta_K^{i-1} \right|_2.$$

Since (by our choice) we start with the same weight vector $\hat{\theta}_{K,m}^0 = \theta_K^0$, we can conclude by induction that

$$\left| \hat{\theta}_{K,m}^i - \theta_K^i \right|_2 \xrightarrow[m \rightarrow \infty]{\mathbb{Q}\text{-a.s.}} 0.$$

However, we remark that this proof only works as long as i is fixed, but not in the limit $i \rightarrow \infty$, since their induction would lead to an infinite sum.

Theorem 10. Let $K \in \mathbb{N}$ be fixed. Then there exists a random sequence $(m_i)_{i \geq 0}$ such that $\hat{Q}^{K,m_i,i}$ converges \mathbb{Q} -a.s. to \tilde{Q}^K as $i \rightarrow \infty$, i.e.

$$\left\| \hat{Q}^{K,m_i,i} - \tilde{Q}^K \right\|_{\pi} \xrightarrow[i \rightarrow \infty]{\mathbb{Q}\text{-a.s.}} 0.$$

Proof Let us define $\theta_K^* \in \mathbb{R}^K$ to be the weight vector of the unique fixed point \tilde{Q}^K of $\Pi^K H$, i.e. $\tilde{Q}^K = \Phi_K \theta_K^*$. From Banach's fixed point theorem we know that $\left| \theta_K^i - \theta_K^* \right|_2 \rightarrow 0$ as $i \rightarrow \infty$.

With Lemma 9 we know that for every $i \in \mathbb{N}$ there exists $\Omega_i \subset \Omega$ with $\mathbb{Q}(\Omega_i) = 1$ such that $\hat{\theta}_{K,m}^i(\omega)$ converges to θ_K^i for all $\omega \in \Omega_i$. Let $\Omega_\infty := \bigcap_{i=1}^{\infty} \Omega_i$ be the set on which this convergence holds for all $i \in \mathbb{N}$, then $\mathbb{Q}(\Omega_\infty) = 1$. Fix $\omega \in \Omega_\infty$. Now let us choose $m_0 = 0$ and for every $i > 0$, $m_i > m_{i-1}$ such that $\left| \hat{\theta}_{K,m_i}^i(\omega) - \theta_K^i \right|_2 \leq 1/i$. Therefore, we obtain that

$$\left| \hat{\theta}_{K,m_i}^i(\omega) - \theta_K^* \right|_2 \leq \left| \hat{\theta}_{K,m_i}^i(\omega) - \theta_K^i \right|_2 + \left| \theta_K^i - \theta_K^* \right|_2 \leq \frac{1}{i} + \left| \theta_K^i - \theta_K^* \right|_2,$$

which converges to 0 when i tends to infinity.

Appendix C. Convergence of the Randomized Recurrent Least Square Monte Carlo (RRLSM)

In this section, we extend the results of Section A to the non-Markovian setting, where we assume that the path up to the current time is a Markov process. In particular, given a discrete time stochastic process $X = (X_0, \dots, X_N)$ as before, we assume that its extension $Z = (Z_0, \dots, Z_N)$ with $Z_n = (X_n, X_{n-1}, \dots, X_0, 0, \dots, 0)$ taking values in $\mathbb{R}^{N+1 \times d}$ for all $0 \leq n \leq N$ is a Markov process. Hence, all results of Section A hold similarly up to replacing X by Z and they also hold for payoff functions that depend on the entire path of X up to the current time. In particular, this immediately implies that RLSM with the path input Z approximates the correct price of the Bermudan option arbitrarily well as $K \rightarrow \infty$. Therefore, it is only left to show that an equivalent result to Theorem 4 holds for our randomized recurrent neural network (6), which takes X as input instead of Z , but makes use of a latent variable in which information about the past is stored.

Fix some $1 \leq n \leq N - 1$ and let π_n now be the distribution of Z_n under \mathbb{Q} . Moreover, let the basis functions $\phi^n = (\phi_k^n)_{k \geq 1}$ be now given by the n -th latent variable h_n of (6). In particular, we define ϕ_k^n as the function mapping $z_n = (x_n, x_{n-1}, \dots, x_0, 0, \dots, 0)$ to the k -th coordinate of

$$h_n = \sigma(A_x x_n + A_h h_{n-1} + b),$$

where $h_{-1} = 0$. By abuse of notation, for growing k we let the matrices grow by adding now rows of random elements to b , A_x and A_h and filling up the new columns of previous rows of A_h with zeros. Like this, ϕ_k^n is well defined for all $k \geq 1$. The operator Π_n^K is defined similarly as before, but with this new set of basis functions, defined on the set of π_n -integrable function J . Then we have to show that the following results are true, so that the assumptions for Theorem 5 and 6 are satisfied. The remainder of their proof works as before.

Proposition 11. If the activation function σ is invertible then for all $0 \leq n \leq N - 1$

$$\left\| \Pi_n^K Q_n - Q_n \right\|_{\pi_n} \xrightarrow[K \rightarrow \infty]{\tilde{\mathbb{P}}\text{-a.s.}} 0.$$

Before we start with the proof, we remark that standard results for the approximation of dynamical systems with RNNs (Schäfer and Zimmermann, 2006) and reservoir computing systems (Gonon and Ortega, 2020) do not apply here, since the dynamical system to approximate $Q = (Q_0, \dots, Q_{N-1})$ is

not time-invariant (in the language of Gonon and Ortega (2020)).

Proof Firstly, we note that it is enough to show that for any $\epsilon > 0$ there exists some size $K \in \mathbb{N}$ and weight matrices b, A_x, A_h such that the corresponding neural network approximation $\tilde{\Pi}_K^n Q_n$ satisfies $\|\tilde{\Pi}_K^n Q_n - Q_n\|_{\pi_n} < \epsilon$ for all $0 \leq n \leq N - 1$. Indeed, if this is true, the convergence (15) follows by the same arguments as in (Zhang et al., 2012, Theorem 3).

Secondly, we note that it is enough to show the statement above for any fixed n separately, i.e. that for each $0 \leq n \leq N - 1$ and $\epsilon > 0$ there exist $K^n \in \mathbb{N}$ and weight matrices b^n, A_x^n, A_h^n such that the corresponding neural network approximation $\tilde{\Pi}_K^n Q_n$ satisfies $\|\tilde{\Pi}_K^n Q_n - Q_n\|_{\pi_n} < \epsilon$. Indeed, if this is true, the stronger statement follows immediately by setting

$$A_x = \begin{pmatrix} A_x^0 \\ \vdots \\ A_x^{N-1} \end{pmatrix}, \quad A_h = \begin{pmatrix} A_h^0 & & \\ & \ddots & \\ & & A_h^{N-1} \end{pmatrix} \text{ and } b = \begin{pmatrix} b^0 \\ \vdots \\ b^{N-1} \end{pmatrix}.$$

Hence, let us fix some $\epsilon > 0$ and n and let us assume that $d = 1$ for simplicity of notation, while the extension to $d > 1$ is immediate. Then we know from Theorem 4 that there exists some neural network f such that $\|f - Q_n\|_{\pi_n} < \epsilon$. The difference between the approximation $\tilde{\Pi}_K^n$ and f is that $\tilde{\Pi}_K^n$ gets a recurrent input, while f gets the entire path as input. However, since n is fixed and finite, we can simply accumulate the same path information in h_n by setting $b^0 = 0, A_x^0 = (1, 0, \dots, 0)^\top \in \mathbb{R}^n$ and

$$A_h^0 = \begin{pmatrix} 0 & \dots & & 0 \\ 1 & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & & \\ 0 & \dots & 0 & 1 & 0 \end{pmatrix} \in \mathbb{R}^{n \times n}.$$

Indeed, with this choice we have $h_{n-1}^0 = (\sigma(x_{n-1}), \sigma(\sigma(x_{n-2})), \dots, \sigma^{(n)}(x_0))^\top$. Let us define the function

$$\varphi : (x_n, \dots, x_0, 0, \dots, 0) \mapsto (x_n, \sigma(x_{n-1}), \dots, \sigma^{(n)}(x_0))$$

Under the assumption that σ is invertible also φ is and there exists a function \tilde{Q}_n such that $\tilde{Q}_n \circ \varphi = Q_n$. Since Q_n is integrable with respect to π_n , the change of variables formula implies that \tilde{Q}_n is integrable with respect to $\varphi^{-1} \circ \pi_n$ and $\mathbb{E}^{(\varphi^{-1} \circ \pi_n)}[\tilde{Q}_n] = \mathbb{E}^{\pi_n}[\tilde{Q}_n \circ \varphi] = \mathbb{E}^{\pi_n}[Q_n]$. Therefore, there exists a neural network $\tilde{f} = \tilde{\beta}\sigma(\tilde{A} \cdot + \tilde{b})$ such that

$$\|\tilde{f} \circ \varphi - Q_n\|_{\pi_n} = \|(\tilde{f} - \tilde{Q}_n) \circ \varphi\|_{\pi_n} = \|\tilde{f} - \tilde{Q}_n\|_{\varphi^{-1} \circ \pi_n} < \epsilon.$$

By extending b^0, A_x^0, A_h^0 to

$$b = \begin{pmatrix} b^0 \\ \tilde{b} \end{pmatrix}, \quad A_x = \begin{pmatrix} A_x^0 \\ \tilde{A}_1 \end{pmatrix}, \quad A_h = \begin{pmatrix} A_h^0 & 0 \\ \tilde{A}_{2:n+1} & 0 \end{pmatrix},$$

where $\tilde{A} = (\tilde{A}_1, \tilde{A}_{2:n+1})$, we get

$$h_n = \begin{pmatrix} \sigma(A_x^0 x_n + A_h^0 h_{n-1}^0 + b^0) \\ \sigma(\tilde{A}\varphi(z_n) + \tilde{b}) \end{pmatrix}.$$

Therefore, we can conclude the proof, since the corresponding approximation $\tilde{\Pi}_K^n$ satisfies $\|\tilde{\Pi}_K^n Q_n - Q_n\|_{\pi_n} \leq \|\tilde{f} \circ \varphi - Q_n\|_{\pi_n} \leq \epsilon$

Remark 12. The idea of the proof is to use the recurrent structure only to recover the path-wise input z_n for which the standard feed-forward neural network approximation results can be used. This is clearly less efficient than using the path-wise input directly. However, in practice, the recurrent neural network approach is usually more efficient than the path-wise approach, finding better ways to store and process the past information than the one given in the proof. This is in line with our empirical findings.

C.1 Stopping of Fractional Brownian Motion Table

The results shown in the plots of Section 6.3 are given in Table 15.

H	price				duration										
	DOS	pathDOS	RLSM	RRLSM	FQJ	RFQJ	RRFQJ	pathRFQJ	pathRRFQJ						
0.01	0.85 (0.02)	1.48 (0.01)	0.84 (0.01)	1.45 (0.01)	0.79 (0.01)	0.78 (0.02)	0.85 (0.07)	1.09 (0.08)	1m15s	2m59s	0s	1s	9s	5s	18s
0.05	0.67 (0.02)	1.24 (0.01)	0.65 (0.01)	1.24 (0.01)	0.68 (0.01)	0.67 (0.02)	0.71 (0.04)	0.99 (0.07)	1m15s	3m 1s	0s	1s	9s	4s	19s
0.1	0.50 (0.02)	0.99 (0.01)	0.49 (0.01)	1.02 (0.01)	0.57 (0.01)	0.55 (0.01)	0.56 (0.05)	0.83 (0.03)	1m12s	2m58s	0s	1s	10s	4s	19s
0.15	0.37 (0.02)	0.77 (0.01)	0.38 (0.01)	0.82 (0.01)	0.47 (0.02)	0.45 (0.02)	0.47 (0.05)	0.65 (0.03)	1m13s	2m59s	0s	1s	9s	4s	19s
0.2	0.28 (0.01)	0.60 (0.01)	0.31 (0.01)	0.64 (0.01)	0.38 (0.01)	0.35 (0.09)	0.31 (0.02)	0.53 (0.02)	1m15s	2m58s	1s	1s	9s	4s	19s
0.25	0.23 (0.01)	0.44 (0.01)	0.25 (0.01)	0.49 (0.01)	0.29 (0.01)	0.26 (0.05)	0.26 (0.04)	0.39 (0.02)	1m14s	2m58s	1s	1s	9s	4s	18s
0.3	0.18 (0.01)	0.30 (0.01)	0.20 (0.01)	0.36 (0.01)	0.21 (0.01)	0.17 (0.01)	0.15 (0.01)	0.27 (0.01)	1m13s	2m57s	1s	1s	9s	4s	18s
0.35	0.13 (0.01)	0.19 (0.01)	0.15 (0.01)	0.25 (0.01)	0.14 (0.01)	0.13 (0.02)	0.12 (0.03)	0.17 (0.01)	1m15s	2m57s	1s	1s	9s	4s	18s
0.4	0.08 (0.01)	0.10 (0.01)	0.10 (0.01)	0.14 (0.01)	0.09 (0.01)	0.06 (0.01)	0.06 (0.01)	0.10 (0.02)	1m15s	2m59s	1s	1s	9s	4s	18s
0.45	0.04 (0.01)	0.03 (0.01)	0.05 (0.01)	0.06 (0.01)	0.04 (0.01)	0.02 (0.01)	0.03 (0.01)	0.05 (0.01)	1m14s	2m58s	0s	1s	9s	4s	18s
0.5	0.00 (0.00)	0.01 (0.01)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.01)	0.01 (0.01)	0.00 (0.01)	1m14s	2m57s	0s	1s	9s	4s	18s
0.55	0.03 (0.01)	0.02 (0.01)	0.03 (0.01)	0.05 (0.01)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	1m16s	3m 0s	1s	1s	9s	4s	17s
0.6	0.07 (0.00)	0.09 (0.01)	0.08 (0.01)	0.10 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)	1m12s	2m56s	1s	1s	9s	4s	17s
0.65	0.10 (0.01)	0.14 (0.01)	0.12 (0.01)	0.16 (0.01)	0.00 (0.00)	0.01 (0.01)	0.00 (0.00)	0.00 (0.00)	1m13s	2m59s	1s	1s	9s	4s	17s
0.7	0.14 (0.01)	0.19 (0.01)	0.16 (0.01)	0.20 (0.01)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	1m13s	2m57s	1s	1s	9s	4s	18s
0.75	0.18 (0.01)	0.23 (0.01)	0.19 (0.00)	0.23 (0.01)	0.00 (0.00)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)	1m15s	2m55s	1s	1s	9s	4s	18s
0.8	0.22 (0.01)	0.26 (0.01)	0.23 (0.01)	0.26 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)	1m15s	2m58s	1s	1s	9s	4s	17s
0.85	0.26 (0.00)	0.29 (0.01)	0.27 (0.01)	0.29 (0.01)	0.00 (0.01)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)	1m16s	2m55s	1s	1s	9s	4s	18s
0.9	0.30 (0.01)	0.33 (0.01)	0.30 (0.00)	0.32 (0.00)	0.00 (0.01)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	1m14s	2m55s	1s	1s	9s	4s	18s
0.95	0.34 (0.01)	0.35 (0.00)	0.34 (0.01)	0.35 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	1m 9s	2m55s	1s	1s	9s	4s	18s
0.999	0.38 (0.01)	0.39 (0.01)	0.38 (0.01)	0.38 (0.00)	0.00 (0.00)	0.01 (0.01)	0.00 (0.00)	0.00 (0.00)	1m18s	2m45s	1s	1s	9s	4s	18s

Table 15: Results of stopping fractional Brownian Motion for different Hurst parameters.

C.2 Non-Markovian Stock Models - Additional Tables

d	price							duration						
	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP	LSM	DOS	NLSM	RLSM	FQI	RFQI	EOP
5	8.34 (0.07)	8.29 (0.09)	8.17 (0.06)	8.31 (0.07)	8.23 (0.04)	8.34 (0.08)	8.23 (0.04)	11s	7s	3s	0s	3s	0s	0s
10	11.83 (0.07)	11.81 (0.09)	11.39 (0.16)	11.83 (0.07)	11.77 (0.04)	11.82 (0.05)	11.79 (0.05)	29s	6s	3s	0s	6s	0s	0s
50	19.60 (0.07)	20.04 (0.04)	18.14 (0.37)	19.32 (0.05)	20.05 (0.06)	20.08 (0.06)	20.06 (0.03)	8m50s	7s	3s	0s	6m36s	1s	0s
100	20.51 (0.09)	23.57 (0.07)	21.29 (0.46)	22.87 (0.04)	23.56 (0.07)	23.67 (0.05)	23.67 (0.05)	40m44s	9s	3s	0s	1h21m35s	1s	0s
500	-	31.62 (0.06)	28.38 (0.55)	31.33 (0.04)	-	32.09 (0.06)	32.14 (0.02)	-	44s	8s	1s	-	1s	0s
1000	-	34.99 (0.08)	33.03 (0.50)	35.06 (0.04)	-	35.83 (0.05)	35.84 (0.03)	-	1m16s	15s	2s	-	1s	0s
2000	-	37.77 (0.07)	36.77 (0.32)	38.83 (0.06)	-	39.64 (0.07)	39.61 (0.04)	-	2m17s	25s	4s	-	2s	0s

Table 16: Max call option on Heston for different numbers of stocks d .

d	price							duration						
	LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI		
5	12.29 (0.07)	12.26 (0.06)	12.12 (0.08)	12.25 (0.07)	12.38 (0.08)	12.34 (0.07)	12s	6s	3s	0s	2s	0s		
10	16.55 (0.06)	16.54 (0.10)	16.03 (0.19)	16.50 (0.06)	16.63 (0.09)	16.64 (0.06)	30s	6s	3s	0s	10s	0s		
50	25.24 (0.07)	25.66 (0.07)	23.67 (0.35)	24.87 (0.04)	25.71 (0.07)	25.68 (0.04)	8m42s	8s	3s	0s	7m34s	1s		
100	26.84 (0.09)	29.22 (0.07)	26.47 (0.62)	28.45 (0.03)	29.26 (0.06)	29.32 (0.07)	42m26s	12s	4s	0s	1h24m 4s	1s		
500	-	36.47 (0.05)	33.80 (0.65)	36.26 (0.05)	-	36.93 (0.04)	-	56s	13s	1s	-	1s		
1000	-	39.25 (0.04)	37.01 (0.34)	39.33 (0.02)	-	39.93 (0.04)	-	1m49s	23s	2s	-	2s		
2000	-	41.45 (0.03)	39.92 (0.26)	42.25 (0.05)	-	42.78 (0.04)	-	3m58s	43s	5s	-	2s		

Table 17: Min put option on Heston for different numbers of stocks d and varying initial stock price x_0 . Here $r = 2\%$ is used as interest rate.

d	price							duration						
	LSM	DOS	NLSM	RLSM	FQI	RFQI	LSM	DOS	NLSM	RLSM	FQI	RFQI		
5	4.82 (0.03)	4.78 (0.04)	4.68 (0.04)	4.75 (0.04)	4.29 (0.12)	4.57 (0.06)	12s	5s	3s	0s	2s	0s		
10	7.20 (0.06)	7.16 (0.04)	6.92 (0.06)	7.13 (0.05)	6.60 (0.14)	6.76 (0.16)	29s	6s	3s	0s	8s	0s		
50	13.48 (0.05)	13.98 (0.03)	12.44 (0.18)	13.69 (0.04)	13.79 (0.03)	13.72 (0.07)	8m34s	8s	3s	0s	7m 7s	1s		
100	14.63 (0.07)	17.13 (0.06)	15.19 (0.32)	16.83 (0.04)	16.97 (0.07)	16.99 (0.04)	39m49s	12s	6s	0s	1h23m 4s	1s		
500	-	24.31 (0.08)	21.83 (0.63)	24.37 (0.04)	-	24.69 (0.05)	-	54s	12s	1s	-	1s		
1000	-	27.42 (0.07)	25.64 (0.55)	27.73 (0.03)	-	28.08 (0.06)	-	1m39s	23s	2s	-	2s		
2000	-	30.10 (0.08)	29.27 (0.36)	31.09 (0.04)	-	31.50 (0.06)	-	3m47s	43s	5s	-	2s		

Table 18: Max call option on Heston for different numbers of stocks d . Here $r = 5\%$ is used as interest rate and $\delta = 10\%$ as dividend rate.

