Pricing American Options by Exercise Rate Optimization

Christian Bayer[a], Raúl Tempone[b], Sören Wolfers[b*]

[a] Weierstrass Institute for Applied Analysis and Stochastics (WIAS)

[b] King Abdullah University of Science and Technology (KAUST), CEMSE

October 17, 2018

Abstract We present a novel method for the numerical pricing of American options based on Monte Carlo simulation and optimization of exercise strategies. Previous solutions to this problem either explicitly or implicitly determine so-called optimal *exercise regions*, which consist of points in time and space at which the option is exercised. In contrast, our method determines *exercise rates* of randomized exercise strategies. We show that the supremum of the corresponding stochastic optimization problem provides the correct option price. By integrating analytically over the random exercise decision, we obtain an objective function that is differentiable with respect to perturbations of the exercise rate even for finitely many sample paths. Starting in a neutral strategy with constant exercise rate then allows us to globally optimize this function in a gradual manner. Numerical experiments on vanilla put options in the multivariate Black–Scholes model and preliminary theoretical analysis underline the efficiency of our method both with respect to the number of time-discretization steps and the required number of degrees of freedom in the parametrization of exercise rates. Finally, the flexibility of our method is demonstrated by numerical experiments on max call options in the Black–Scholes model and vanilla put options in Heston model and the non-Markovian rough Bergomi model.

Keywords Computational finance, American option pricing, stochastic optimization problem, Monte Carlo, multivariate approximation, rough volatility

2010 Mathematics Subject Classification 91G60, 91G20, 49M20, 90C90, 65K10, 65C05

1 Introduction

American options on $d \geq 1$ underlying assets $S_t = (S_{1,t}, \ldots, S_{d,t})$ may be exercised by their holder at any time t before a given expiration time $T \in \mathbb{R}_+ := [0, \infty)$, upon which the holder receives the payoff $g(t, S_t)$, for some previously agreed function $g\colon [0,T] \times \mathbb{R}^d_+ \to \mathbb{R}_+$.

If the underlying market is Markovian and has a security with interest rate $r > 0$, then the arbitrage-free value of an American option under an equivalent martingale measure $\mathbb{Q}$ is uniquely determined by the current asset values. Moreoever, if we assume for simplicity that the underlying assets do not pay dividends, then the corresponding *value function* $v\colon \mathbb{R}^d_+ \to \mathbb{R}_+$ is given by

$$v(s_0) = \sup_{\tau \in \mathcal{S}} \mathbb{E}_{\mathbb{Q}}[Y_{\tau \wedge T} | S_0 = s_0], \quad \forall s_0 \in \mathbb{R}^d_+, \tag{1}$$

where $Y_t := \exp(-rt) g(t, S_t)$, $t \geq 0$ is the *discounted payoff process* and $\mathcal{S}$ denotes the set of all stopping times with respect to the filtration generated by $(S_t)_{0 \leq t \leq T}$ [19, Theorem 5.3]. In the remainder of this work, all expectations are taken with respect to the same equivalent martingale measure and denoted by $\mathbb{E}$.

Most state-of-the-art methods for American option pricing – including all variants of the Longstaff–Schwartz [21], PDE [1], binomial tree [12] and stochastic mesh [9] methods, to give a just a few references – exploit the dynamic

*Corresponding author. Email address: soeren.wolfers@kaust.edu.sa

programming principle to determine the value function directly in a backwards iteration scheme. Further approaches are based on dual problems [23, 3], policy iteration [7], or (quasi-)analytic solutions [5, 20]. For options on multiple underlying assets, many methods break down or become prohibitively expensive with the Monte Carlo simulation based Longstaff–Schwartz algorithm and its relatives possibly being most efficient and most popular among practitioners.

In this work, we propose a method that is based on the following variation of Equation (1), which states that one may restrict the optimization to *hitting times* instead of general stopping times:

$$v(s_0) = \sup_{E \in \mathcal{B}([0,T]\times\mathbb{R}^d_+)} \mathbb{E}[Y_{\tau_E \wedge T} | S_0 = s_0], \quad \forall s_0 \in \mathbb{R}^d_+. \tag{2}$$

Here, the supremum is taken over Borel measurable subsets of $E \subset [0,T] \times \mathbb{R}^d_+$, whose hitting time is given by $\tau_E := \inf\{t \geq 0 : (t, S_t) \in E\}$. Given the value function, an optimal *early exercise region* in Equation (2) may be obtained by $E^* := \{(t,s) : v(t,s) = g(t,s)\}$. As a practical alternative to compute $v(s_0)$, we solve a relaxation of the stochastic optimization problem in Equation (2) that amounts to the determination of an optimal *randomized exercise strategy*. To be precise, both Equation (1) and Equation (2) require some technical conditions on the processes $(Y_t)_{0\leq t\leq T}$ and $(S_t)_{0\leq t\leq T}$ [24, Corollary 2, Section 3.3.1]. Throughout this work we tacitly assume such conditions hold and restrict our attention to the solution of Equation (2).

To the best of our knowledge, optimization of the exercise region in Equation (2) was first proposed in [16] and developed in [2, 13, 18, 14] but has not found its way into the canon of numerical algorithms for American option pricing. In [16], exercise strategies were determined separately for each exercise date of an American Asian option in a backwards iteration. The optimization at each step was done in a brute force fashion, which explains why only two parameters were allowed in the parametrization of the exercise regions. In [13, 14], ad-hoc parametrizations exploiting known behavior of the optimal exercise regions were used to optimize exercise regions as subsets of time-space without applying backwards iteration.

In general, optimization of the exercise region faces two challenges. First, as mentioned in [14], it is not obvious how to parametrize all possible exercise regions in the multi-dimensional setting or even in a one-dimensional settings that goes beyond vanilla options in the Black–Scholes model. Second, it is not obvious how to find the global optimum once a parametrization has been found, as mentioned in both [13] and [14]. Indeed, once the expectation is replaced by an empirical average for numerical approximations of the expected payoff, the quantity to be maximized depends highly irregularly on the exercise region E (see Figure 1b below). Furthermore, even if a large number of sample paths is used to reduce the small scale oscillatory behavior, the resulting surface may still be non-concave and exhibit isolated local optima, as reported in [13].

To address these challenges, Section 2 introduces a relaxation of the optimization problem in Equation (2) that replaces exercise regions $E \subset [0,T] \times \mathbb{R}^d_+$ by *exercise rates* $f\colon [0,T] \times \mathbb{R}^d_+ \to \mathbb{R}_+$, which define *randomized exercise strategies* where options are exercised with an infinitesimal probability depending on the current time and asset values. [1] We show that the supremum of the resulting relaxed optimization problem coincides with the correct option price. To find the supremum numerically, we parametrize the space of exercise rates using an arbitrary (truncated) basis of functions on $[0,T] \times \mathbb{R}^d$. By integrating analytically with respect to the exponential distribution that underlies the random exercise decision, we obtain a smooth objective function even for finitely many sample paths. We may then use gradient based optimization routines to determine an optimal coefficient vector. Furthermore, we may start this optimization from an exercise rate that has a constant non-zero value across time and space and then let the optimization routine gradually refine this neutral strategy towards an optimal one with marked variations in the exercise rate. Intuitively, this facilitates the search of a global optimum without requiring an informed initial guess close to the optimum. Details of the numerical implementation are discussed in Section 2.1. There, we also briefly discuss how the accuracy of our method depends on the various involved discretization parameters. In particular, we provide heuristic bounds on the number of degrees of freedom in the exercise rate required for satisfactory randomized exercise strategies. The main message here is that such bounds are determined by the smoothness of the optimal exercise boundary as a manifold, not as a function of time.

[1] This notion of randomization is not related to that of [10], where the option is replaced by a similar one with random expiry, which can be priced more easily.

Finally, Section 3 presents numerical experiments for various market models and options. In Sections 3.1 and 3.2, we consider vanilla put options in the classical Black–Scholes model.

In the case of a single underlying, the exercise boundary of an American put option, whose payoff function is given by $g(t,s) := g(s) := (K-s)^+$ for some strike $K > 0$, can be written as a function of time with asymptotic behavior $s(t) \approx K - C_1\sqrt{(T-t)\log(T-t)}$ for some $C_1 > 0$ as $t \to T$. Despite the square root singularity near expiry, the experiments presented in Section 3.1 show that low-degree polynomials suffice to capture the optimal exercise boundary well. This can be explained by the fact that the graph of the similar function $\tilde{s}(t) = K - C_1\sqrt{(T-t)}$ is smooth as a one-dimensional manifold in $\mathbb{R}^2$ and indeed coincides with the zero level set (intersected with $x < K$) of the quadratic polynomial $f(t,s) := (K-s)^2 - C_1^2(T-t)$, whose scalar multiples therefore constitute close-to-optimal exercise rates. In fact, we obtain a relative error of less than 0.1% with quadratic polynomials. Although we solve non-concave maximization problems, we are able to find global optima starting from a constant exercise rate. Furthermore, since we perform global optimization without backwards iteration in time, our method is insensitive to the number of exercise dates – in contrast to exponential blow up of the error in terms of the number of exercise dates in conventional Longstaff–Schwartz algorithms [25]. Furthermore, in Section 3.2 we show that our algorithm outperforms the Longstaff–Schwartz algorithm with respect to the required polynomial degree for the pricing of basket put options, which is crucial when the number of underlying asset is large.

In Section 3.3, we consider call options on the maximum of a number of underlying assets, $g(s) = \max_{i=1}^d (s_i - K)^+$. Numerical algorithms for the pricing of such max call options were previously discussed in [3, 22]. They pose a challenge to the direct determination of exercise regions because optimal exercise regions are disconnected [8]. Still, our results show that polynomials of low degree suffice to obtain highly accurate estimates despite the nontrivial topology of the optimal exercise region.

In Section 3.4, we consider the Heston model, in which the underlying asset and its stochastic volatility form a joint Markov process. Since our method involves the market model only for the generation of random sample paths, its application in this scenario is straightforward. Finally, we consider the non-Markovian rough Bergomi model of [6] in Section 3.5. Here, to recover Markovianity, we must extend our process by its past values. In practice, using a large but finite number of past values leads to very high-dimensional approximation problems. However, our experiments indicate that exercise strategies depending only on the spot values of the underlying asset and its volatility achieve near-optimal performance.

2 Exercise rate optimization

For the remainder of this work, we aim to compute $v(s_0)$ for fixed $s_0 \in \mathbb{R}_+^d$. We let $\mathcal{T} := [0,T]$ and assume throughout that $(S_t)_{t\in\mathcal{T}}$ is conditioned on $S_0 = s_0$.

Definition 2.1. *For any $f\colon \mathcal{T} \times \mathbb{R}_+^d \to \mathbb{R}_+$, the* randomized exercise strategy with exercise rate f *is given by early exercise at the time*

$$\tau_f := \inf\{t \geq 0 : \int_0^t \lambda_u \,\mathrm{d}u \geq X\}, \tag{3}$$

where $\lambda_t := f(t, S_t)$, $t \in \mathcal{T}$ and X is a standard exponential random variable that is independent of $(S_t)_{t\in\mathcal{T}}$.

The exercise time τ_f equals the first jump time of a Poisson process with rate $(\lambda_t)_{t\in\mathcal{T}}$. In other words, the exercise rate f determines the time- and space dependent infinitesimal probability with which the American option is exercised in a infinitesimal time interval $\mathrm{d}t$.

With Equation (2) in mind, we are interested in the expected payoff under a randomized exercise strategy with early exercise time τ_f, which we denote by

$$\psi(f) := \mathbb{E}[Y_{\tau_f \wedge T}]. \tag{4}$$

Since $\int_0^t \lambda_u \,\mathrm{d}u$ is a deterministic function of the asset path until t and X is independent of $(S_u)_{u\in\mathcal{T}}$, we have

$$\mathbb{P}(\tau_f \geq t \mid (S_u)_{u\in\mathcal{T}}) = \mathbb{P}(X > \int_0^t \lambda_u \,\mathrm{d}u \mid (S_u)_{u\in\mathcal{T}}) = \exp\left(-\int_0^t \lambda_u \,\mathrm{d}u\right) =: U_t$$

and

$$\mathbb{P}(\tau_f \in \mathrm{d}t \mid (S_u)_{u\in\mathcal{T}}) = -\,\mathrm{d}U_t = \lambda_t U_t\,\mathrm{d}t.$$

Hence, we obtain

$$\phi(f,(S_u)_{u\in\mathcal{T}}) := \mathbb{E}[Y_{\tau_f\wedge T} \mid (S_u)_{u\in\mathcal{T}}] = \int_0^T Y_t\lambda_t U_t\,\mathrm{d}t + Y_T U_T.$$

By the law of total expectation, which we may apply because all involved random variables are nonnegative, we may thus deduce the formula

$$\psi(f) = \mathbb{E}[\phi(f,(S_u)_{u\in\mathcal{T}})] = \mathbb{E}\left[\int_0^T Y_t\lambda_t U_t\,\mathrm{d}t + Y_T U_T\right]. \tag{5}$$

The following proposition shows that, in theory, exercise rate optimization yields the correct option value.

Proposition 2.2. *We have*

$$v(s_0) = \sup_{f\colon [0,T]\times\mathbb{R}_+^d\to\mathbb{R}_+} \psi(f). \tag{6}$$

Proof. For any $E \in \mathcal{B}(\mathcal{T}\times\mathbb{R}_+^d)$, we may formally insert the indicator function

$$f_E(t,s) := \begin{cases} +\infty, & (t,s)\in E \\ 0, & (t,s)\notin E \end{cases}$$

in Equation (3) to obtain $\tau_{f_E} = \tau_E$. If we replace $+\infty$ by large numbers diverging to $+\infty$, we may conclude from Equation (2) that $\sup_{f\colon [0,T]\times\mathbb{R}_+^d\to\mathbb{R}_+}\psi(f) \geq v(s_0)$ by first applying Fatou's lemma and then taking the supremum over E.

Conversely, the law of total expectation shows for any $f\colon [0,T]\times\mathbb{R}_+^d\to\mathbb{R}_+$ that

$$\psi(f) = \mathbb{E}[Y_{\tau_f\wedge T}] = \mathbb{E}\left[\mathbb{E}\left[Y_{\tau_f\wedge T} \mid X\right]\right].$$

Because τ_f conditioned on X is a stopping time and $(S_t)_{t\in\mathcal{T}}$ is independent of X, Equation (1) implies that $\mathbb{E}\left[Y_{\tau_f\wedge T} \mid X\right] \leq v(s_0)$ almost surely, hence $\psi(f)\leq v(s_0)$. □

2.1 Numerical algorithm

To determine optimal exercise rates numerically, we

(i) replace the time continuous model of the stochastic process $(S_t)_{t\in\mathcal{T}}$ by a discretization with $N<\infty$ time steps, for example using the Euler–Maruyama scheme;

(ii) replace the expectation in Equation (5) by an average over $M<\infty$ fixed sample paths $(S_t^{(m)})_{t\in\mathcal{T},1\leq m\leq M}$ [2];

(iii) introduce a B-dimensional, $B<\infty$ parametrization $\mathbb{R}^B \ni \boldsymbol{c}\mapsto f_{\boldsymbol{c}}$ of the space of exercise rates;

(iv) maximize the surrogate function

$$\begin{aligned}\overline{\psi}\colon \mathbb{R}^B &\to \mathbb{R} \\ \boldsymbol{c} &\mapsto \frac{1}{M}\sum_{m=1}^M \phi(f_{\boldsymbol{c}},(S_t^{(m)})_{t\in\mathcal{T}}).\end{aligned}$$

[2] We assume for notational simplicity that each path is interpolated beyond the N nodes of the time-discretization scheme, for example in a piecewise constant fashion, and that all subsequent expressions are evaluated using this interpolation.

Parametrization To address step (iii), we work with the logarithmic asset values $x_i := \log(s_i)$, $1 \leq i \leq d$ and let

$$F_{\mathcal{P}} := \left\{ f_p(t,x) := 1_{g(t,s)>0} \exp(p(t,x)) \;\middle|\; p \in \mathcal{P} \right\}$$

for any finite-dimensional linear space $\mathcal{P}$ of functions on $\mathcal{T} \times \mathbb{R}^d$. After choosing a basis of $\mathcal{P}$, we then obtain the desired parametrization $\boldsymbol{c} \mapsto f_{\boldsymbol{c}}$. In particular, we work with spaces $\mathcal{P}_k$ of polynomials of degree less than or equal to $k \geq 0$ in $d+1$ variables in the remainder of this manuscript.

Optimization Concerning step (iv), due to lack of concavity of $\overline{\psi}$ it is not clear that globally optimal coefficients, which may even lie at infinity, can be found numerically. However, in our numerical experiments, we found the L-BFGS-B algorithm as implemented in Python's SciPy library[3] to perform well and not get stuck in local maxima when started from a constant exercise rate. For example, a simple gradient descent algorithm could be used to maximize the function in Figure 1a, which shows the dependence of $\overline{\psi}$ on the coefficient $c_{(0,0)}$ of the constant polynomial $p_{(0,0)} \equiv 1$ for a one-dimensional put option. Note that this is not possible for the function in Figure 1b, which arises in the optimization of deterministic exercise regions as in [14] and requires the use of finite difference stochastic gradient algorithms.

Differentiability of ϕ, and thus ψ and $\overline{\psi}$, with respect to f is easy to show. Using the fact that $\lambda_t U_t \,\mathrm{d}t = -\,\mathrm{d}U_t$, we obtain the simple gradient formula

$$\langle \nabla_f \phi(f, (S_t)_{t\in\mathcal{T}}), h \rangle = -\int_0^T Y_t \,\mathrm{d}\langle \nabla_f U_t, h\rangle + \langle \nabla_f U_T, h\rangle Y_T \quad \forall h \colon \mathcal{T} \times \mathbb{R}^d_+ \to \mathbb{R},$$

where

$$\langle \nabla_f U_t, h\rangle = -U_t \int_0^t h(u, S_u)\,\mathrm{d}u \quad \forall t \in \mathcal{T}.$$

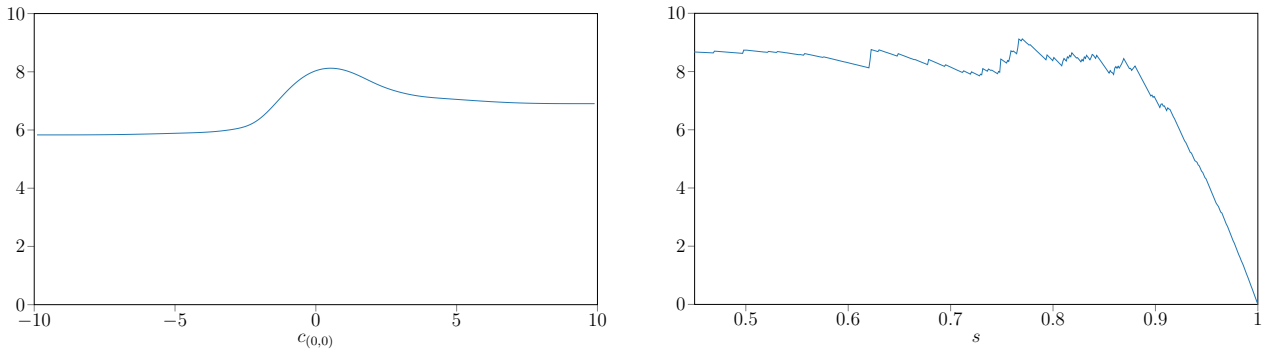

(a) $\overline{\psi} \colon \mathbb{R} \to \mathbb{R}$ with $\mathcal{P}$ the space of constant functions

(b) $s \mapsto \frac{1}{M} \sum_{m=1}^{M} [Y^{(m)}_{\tau_{E_s} \wedge T}]$ with $E_s := [0,T] \times [0,s]$

Figure 1: Functions to be maximized in a one-parameter optimization of a randomized exercise strategy (a) and a one-parameter optimization of a deterministic strategy (b) for a one-dimensional American at-the-money put option with $K = s_0 = 100$ and $T = 1$ in the Black–Scholes model with $r = 0.05$ and $\sigma = 0.3$. Each plot was generated using $M = 100$ sample paths with $N = 100$ time steps.

Accuracy To obtain accurate results, we must choose large enough values for the number of samples, M, the number of time steps, N, the number of iterations of the optimization routine, ℓ, and the polynomial degree, k.

[3] `https://docs.scipy.org/doc/scipy/reference/optimize.minimize-lbfgsb.html`

Convergence with respect to M occurs at the Monte Carlo rate $M^{-1/2}$. Convergence with respect to N is slightly more complicated. For a fixed, smooth, exercise rate, the surrogate $\overline{\psi}$ converges at the usual weak convergence rate of the discretization scheme with respect to the number of time steps (N^{-1} for the Euler–Maruyama scheme). However, in the limit of increasingly steep exercise rates approaching the optimal deterministic exercise regions, the weak convergence rate is expected to deteriorate to $N^{-1/2}$ though this effect seems to occur at a level that is overshadowed by the remaining Monte Carlo error in the experiments of Section 3.1.

Convergence with respect to ℓ, assuming we do have global convergence, is expected to occur at least at an exponential rate, depending on what type of deterministic optimization routine we use. Figure 3 below provides numerical evidence of such exponential convergence.

Finally, to characterize the convergence with respect to k, we note that for any polynomial $0 \neq p_k \in \mathcal{P}_k$ the randomized exercise strategies with exercise rates $f_L := \exp(Lp_k) \in F_k$ converge to a deterministic strategy with early exercise region $E_k := \{p_k \geq 0\}$ as $L \to \infty$. Therefore, it suffices to study (i) the approximability of the optimal exercise region E_* by polynomial superlevel sets, and (ii) the sensitivity of the expected payoff on the right hand side of Equation (2) with respect to perturbations of the exercise region. In the direction of (i), we observe that if E_* is a bounded C^m-submanifold, $m \geq 2$, of $(0,T) \times \{g > 0\}$, then there exists a sequence of polynomials p_k such that the boundaries $B_k := \partial E_k$ of the corresponding exercise regions $E_k := \{p_k \geq 0\}$ satisfy

$$B_k = \{(t,s) + \Theta(t,s) : (t,s) \in B_*\} \tag{7}$$

for some $\Theta\colon B_* \to \mathbb{R}^{1+d}$ such that

$$\sup_{(t,s)\in B_*} |\Theta(t,s)| < Ck^{-m}.$$

This follows from a combination of the multi-dimensional Jackson theorem [4] with a partition of unity and elementary geometry. Results in the direction of (ii) were presented in [11], where it was shown that the expected payoff is differentiable with respect to perturbations of the exercise region in spatial directions if $(0, s_0) \notin E_*$ and the payoff functions lies in some Hölder space $C^{1,\alpha}$, $\alpha > 0$. Unfortunately, this result is not quite general enough for our purposes, since we require bounds with respect to general, spatio-temporal perturbations of the domain (as in Equation (7)) and for payoff functions that are only Lipschitz. On the other hand, it would suffice for our purposes to have Hölder continuity instead of differentiability with respect to such perturbations. As discussed in the introduction, the optimal exercise boundaries of vanilla put options in the Black–Scholes model exhibit singularities near expiry when viewed as functions of time, which arise from the lack of differentiability of the payoff function. However, as argued above, the polynomial approximability of optimal exercise regions is determined by their smoothness as submanifolds of $\{g > 0\}$, not as functions of time until expiry. Indeed, our experiments indicate that the polynomial degree $k = 2$ is sufficient for most purposes, even in multivariate models not of Black–Scholes type.

Overfitting Choosing a subspace with a large number $B \gg 1$ of degrees of freedom to improve the flexibility of candidate exercise rates increases the cost of computations and the risk of *overfitting*. This means that the value of $\overline{\psi}(\boldsymbol{c}^*)$ at the optimized coefficients $\boldsymbol{c}^*$ may overestimate the true value $\psi(f_{\boldsymbol{c}^*})$ unless a correspondingly large number $M = M(B) \gg 1$ of sample paths is used. Since it is difficult to analyze the number of samples required to avoid overfitting, we simply compute another, unbiased, estimate of $\psi(f_{\boldsymbol{c}^*})$ using a new set of sample paths $(\tilde{S}_t^{(m)})_{t\in\mathcal{T}}$, $1 \leq m \leq M$. Similar techniques are used in classical regression-based method such as the Longstaff–Schwartz algorithm. Following statistical learning terminology, we refer to the biased and unbiased estimators of $\psi(f_{\boldsymbol{c}^*})$ as *training* and *test* values, respectively. In particular, the expectation of the test value is a lower bound of the correct option price according to Proposition 2.2.

3 Numerical experiments

Throughout this section, we use the L-BFGS-B algorithm with initial coefficients $\boldsymbol{c} \equiv 0$ to maximize $\overline{\psi}$.

3.1 Convergence with respect to discretization parameters

In this subsection, we study the convergence of our method with respect to the discretization parameters M, N, k and ℓ by pricing the vanilla put option from Figure 1 with strike $K = 100$ and expiry $T = 1$ in the Black–Scholes model with volatility $\sigma = 0.3$, risk-free interest rate $r = 0.05$, and spot price $s_0 = 100$. Using a binomial tree algorithm with 50 000 levels (i.e., 50 000 time steps and 50 000 spatial discretization nodes at $T = 1$), we obtain the reference value $v^* = 9.8701$. Figures 2a and 2b show that the prices found through exercise rate optimization based on polynomial degree $k = 2$ and $M_n := 200 \times 4^n$ sample paths with $N_n := 2^n$ time-steps converge towards this reference value as $n \to \infty$. In particular, our maximization does not get stuck in local optima of $\overline{\psi}$. Furthermore, Figure 2a shows that test and training value converge at roughly the same speed, meaning that we are not suffering from overfitting. This is not surprising, since the space of bivariate quadratic polynomials is only 6-dimensional. We restrict all following plots to the test value, which constitutes an unbiased estimate of the quality of a given exercise rate.

In the logarithmic scale of Figure 2b, we can see that our approximations converge to the reference value at roughly the rate $2^{-n} = \mathcal{O}(N_n^{-1} + M_n^{-1/2})$. We obtain an accuracy of roughly 4 significant digits despite using only quadratic polynomials for the exercise boundary approximation, which confirms that singularities of the exercise boundary as a function of time do not pose a problem to our polynomial approximation scheme. For comparison, Figures 2c and 2d show results for $k \in \{0, 1\}$, that is, for constant exercise rates and for exercise rates that depend only linearly on space and time, respectively. For $k = 1$, the results are astoundingly similar to the case $k = 2$, though an inspection on a logarithmic scale reveals stagnation at a relative error around 0.5%. For $k = 0$, our method stagnates around the value 9.35, which is roughly the price of a European option with the same parameters.

To study the effects of M, N and k, we performed the experiments in this and the following subsection with the tolerance of the L-BFGS-B optimization set to machine precision, which required between 70 and 200 function evaluations to achieve. However, an error comparable to that of the remaining discretization errors can already be achieved with significantly fewer evaluations. Indeed, Figure 3 shows that for $n = 4$ and $k = 2$ the relative error between $\overline{\psi}(\boldsymbol{c}_\ell)$ and the final value is below 0.1% already for $\ell = 20$. For this reason, we limit the number of iterations in Sections 3.4 and 3.5 to 20.

3.2 Comparison with Longstaff–Schwartz algorithm

In this subsection, we consider *basket put options* on linear combinations of $d \in \{2, 5\}$ underlying assets. The payoff function of such options is given by $g(s) := (K - c \cdot s)^+$ for $K > 0$ and $c \in \mathbb{R}^d$. In our experiments we use $K := 100$ and $c_i := 1/d$, $1 \le i \le d$.

We compare our method to the Longstaff–Schwartz algorithm as implemented in the freely available version 16 of the derivative pricing software Premia[4]. Just as our method, the Longstaff–Schwartz algorithm requires specification of the number M of sample paths, the number N of time-steps used for their simulation, and the polynomial degree k, which controls the accuracy of approximations of the value function. For simplicity, we restrict the simulations in this section to $N = 8$ time steps. To prevent our comparison being skewed by the fact that the two algorithms use different sample paths, we use the same large number of $M = 3 \times 10^6$ samples for both algorithms. Finally, we use a risk-free interest rate $r = 0.05$ and a diagonal volatility matrix $\Sigma_{ij} = 0.3^2 \delta_{ij}$, $1 \le i, j \le d$ in the underlying Black–Scholes model with $s_0 = (100, \dots, 100)$.

To emphasize the efficiency of exercise rate optimization with respect to the polynomial degree, we use it to compute reference values $v^* = 6.5479$ and $v^* = 3.6606$ with polynomial degree $k_{\mathrm{ERO}} = 2$ for $d = 2$ and $d = 5$, respectively. Figure 4 shows that the Longstaff–Schwartz algorithm converges to these values as $k_{\mathrm{LS}} \to \infty$ but only achieves a comparable performance for $k \approx 6$. We show 95% confidence bands around our reference value based on the empirical variance in the evaluation of our test value, which indicate that remaining difference between the two methods can be explained by the random sampling error.

[4] `https://www.rocq.inria.fr/mathfi/Premia`

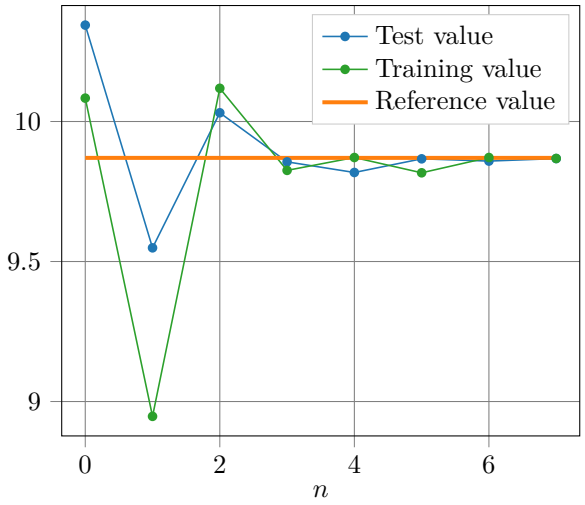

(a) $k = 2$

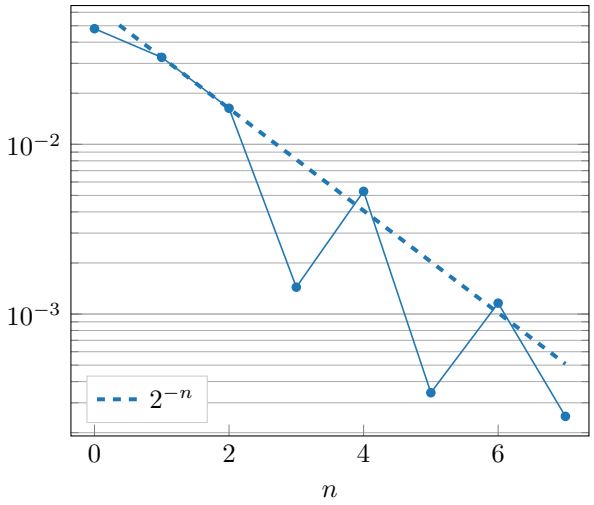

(b) Relative error of test value, $k = 2$

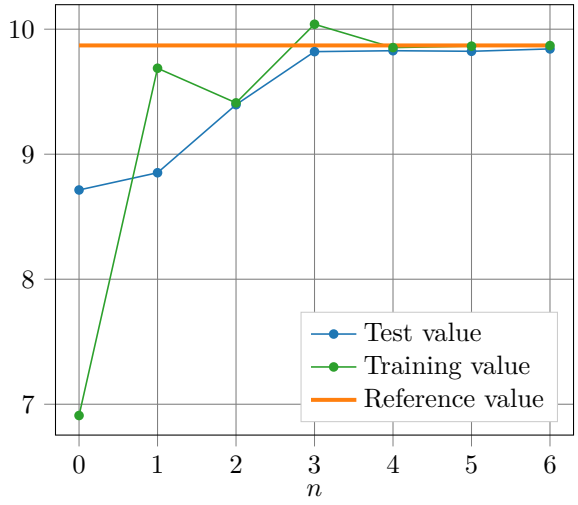

(c) $k = 1$

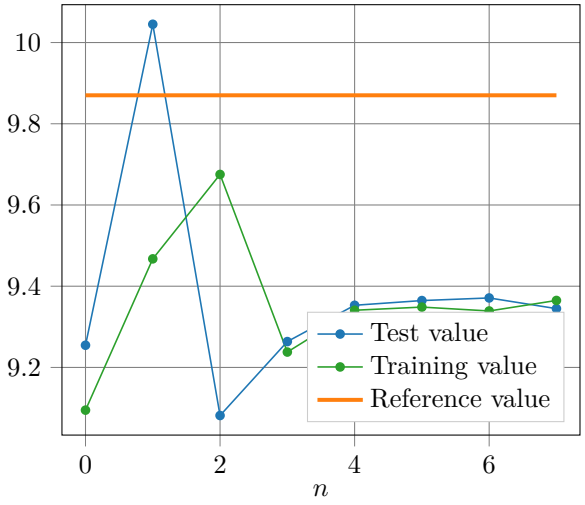

(d) $k = 0$

Figure 2: Exercise rate optimization with polynomial degree $0 \leq k \leq 2$ and $M_n = 200 \times 4^n$ and $N_n = 2^n$, $0 \leq n \leq 7$ applied to one-dimensional American put option in the Black–Scholes model with $\sigma = 0.3$, $r = 0.05$, $K = 100$, $s_0 = 100$, and $T = 1$.

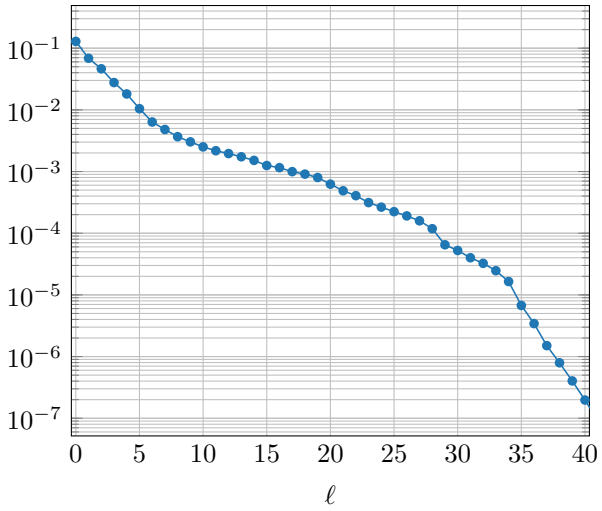

Figure 3: Convergence with respect to the number of function evaluations, ℓ, in the training step of exercise rate optimization for an American put option using the L-BFGS-B algorithm.

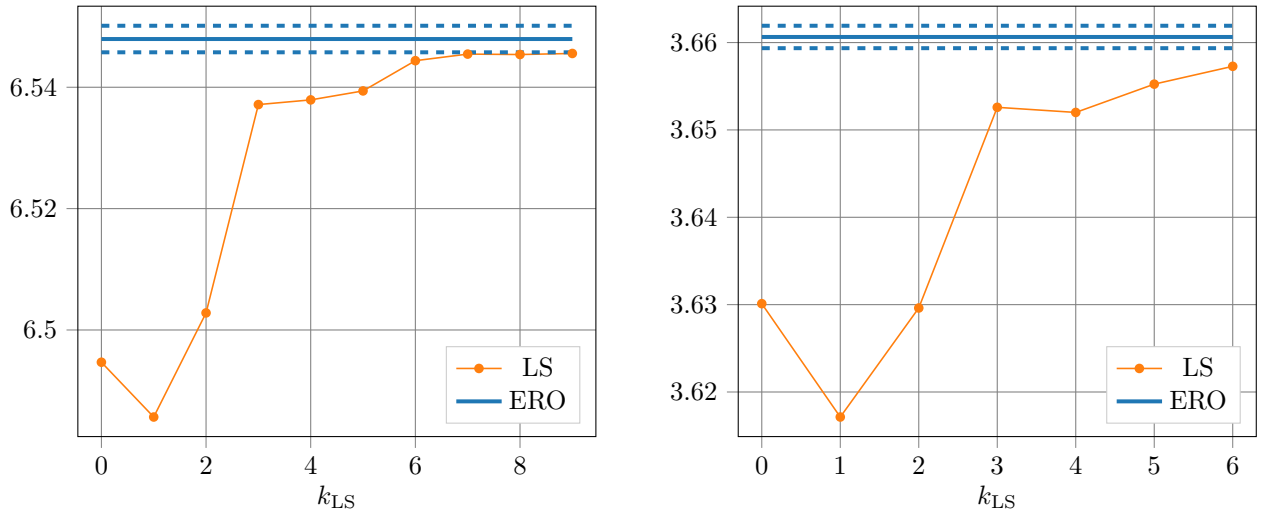

(a) $d = 2$ **(b)** $d = 5$

Figure 4: Convergence of Longstaff–Schwartz algorithm (LS) for $\{2, 5\}$-dimensional basket put options with increasing polynomial degree k_{LS} to reference values computed via exercise rate optimization (ERO) with polynomial degree $k_{\mathrm{ERO}} = 2$ and 95% confidence bands (dashed).

3.3 Max call options

In this subsection, we consider call options on two underlying assets of the form $g(s) := \max\{(s_1 - K)^+, (s_2 - K)^+\}$. These so-called max call options present an interesting challenge to our method, since the optimal exercise region at any time consists of 2 connected components of $\mathbb{R}^2_+$ [8]. Lower and upper bounds for the option prices in the Black–Scholes model with $r = 0.05$, $\Sigma_{ij} = 0.2^2\delta_{ij}$, $K = 100$, $N = 8$ and dividend $\delta = 0.1$ were provided in [3]. Table 1 lists these bounds along with the results of our method for $k \in \{1, 2, 3\}$ and $M = 1\,000\,000$. The optimized exercise rates with $k \in \{2, 3\}$ are shown in Figure 5. As expected, they are almost deterministic, which means that they exhibit steep jumps from values close to zero to values close to infinity. Since the specific values beyond

		95% CI	k 1	2	3
s_0	90	[8.053,8.082]	7.126	8.009	8.039
	100	[13.892,13.934]	12.311	13.821	13.865
	110	[21.316, 21.359]	19.133	21.220	21.256

Table 1: Prices of max call option. 95% confidence intervals (CI) taken from [3].

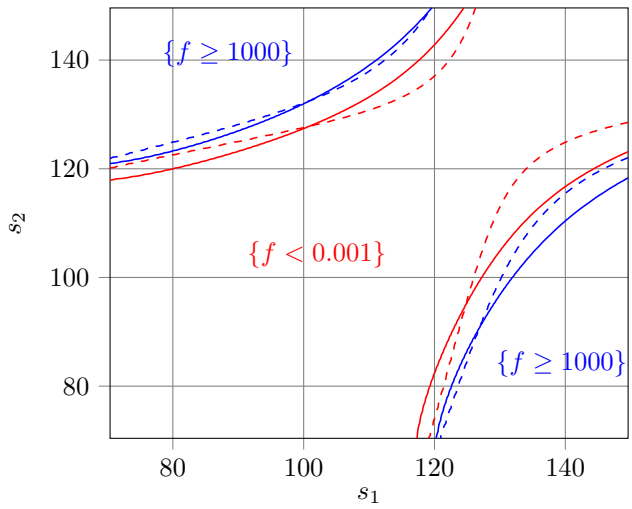

Figure 5: Level sets of optimal exercise rate for American max call option with $k = 2$ (dashed) and $k = 3$ (solid), respectively.

these jump are irrelevant, we restrict our plots to the level sets of exercise rate 0.001 and 1000. The results in this subsection were obtained using a maximal number of 20 optimization steps. Performing more steps would further reduce the distance between these level sets without noticeable difference in the resulting option price. As predicted by theory, there are two disjoint regions of high exercise rates. Furthermore, due to the symmetry of the underlying model and the payoff, the optimized exercise rate is almost axisymmetric even though we do not enforce this symmetry. While there is no hope to model disconnected regions with log-linear exercise rates available for $k = 1$, the hyperbolic conic sections available with $k = 2$ already provide satisfactory approximations.

3.4 Stochastic volatility

In this subsection, we apply our method to pricing in a stochastic volatility model.

For this purpose, we consider the basic Heston model as described in [17], which models the evolution of a single underlying asset X_t and volatility ν_t by a coupled system of stochastic differential equations,

$$\mathrm{d}X_t = \mu X_t \,\mathrm{d}t + \sqrt{\nu_t} X_t \,\mathrm{d}W_t^X, \tag{8}$$

$$\mathrm{d}\nu_t = \kappa(\theta - \nu_t)\,\mathrm{d}t + \xi\sqrt{\nu_t}\,\mathrm{d}W_t^\nu, \tag{9}$$

where $\mu > 0, \kappa > 0$, $\theta > 0$, $\xi > 0$ with $2\kappa\theta > \xi^2$ and W_t^X and W_t^ν are Wiener processes with correlation $-1 \le \rho \le 1$.

Since our method requires Markovian markets, we must include the volatility and define $S_t := (X_t, \nu_t)$, $t \in \mathcal{T}$. Intuitively, this simply means that knowledge of the current volatility is required to make optimal exercise decisions in stochastic volatility models.

To obtain a risk neutral measure, we simply replace μ by the risk-free rate $r = 0.05$ in Equation (8). We choose the remaining parameters $\kappa = 3$, $\theta = 0.05$, $\xi = 0.5$, $\rho = -0.5$ and compute estimates of the value $v_K(s_0)$ of a put

option with $s_0 = (100, 0.15)$ and 25 different values of the strike $K \in [90, 150]$. For this purpose, we use polynomials of degree $k \in \{0, 1, 2\}$ and $M = 100\,000$ samples with $N = 32$ time steps.

For comparison, we also show the results of the finite difference method `FD_Hout_Heston`, as implemented in Premia, with 32 time steps and a grid of 100×100 nodes in the discretization of the stock-volatility plane. The results are shown in Figure 6. The maximal relative difference between the two methods is 1% and occurs around

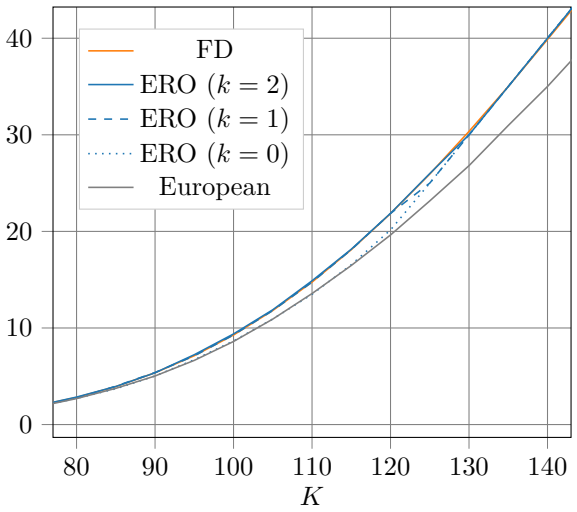

Figure 6: Dependence of American put option price on strike in the Heston model; computed using exercise rate optimization (ERO) and a finite-difference method (FD), respectively.

$K^* = 130$. Up to roundoff error, the prices computed by our method are equal to $K - 100$ for all $K \geq K^*$. This behavior is expected, since the initial point $(100, 0.15)$ lies in the optimal exercise region for large enough K, for which the put option is thus exercised immediately.

Figure 7 shows the numerically optimized exercise rates (with $k = 2$) at $t = 0.5$ for $K \in \{100, 110\}$.

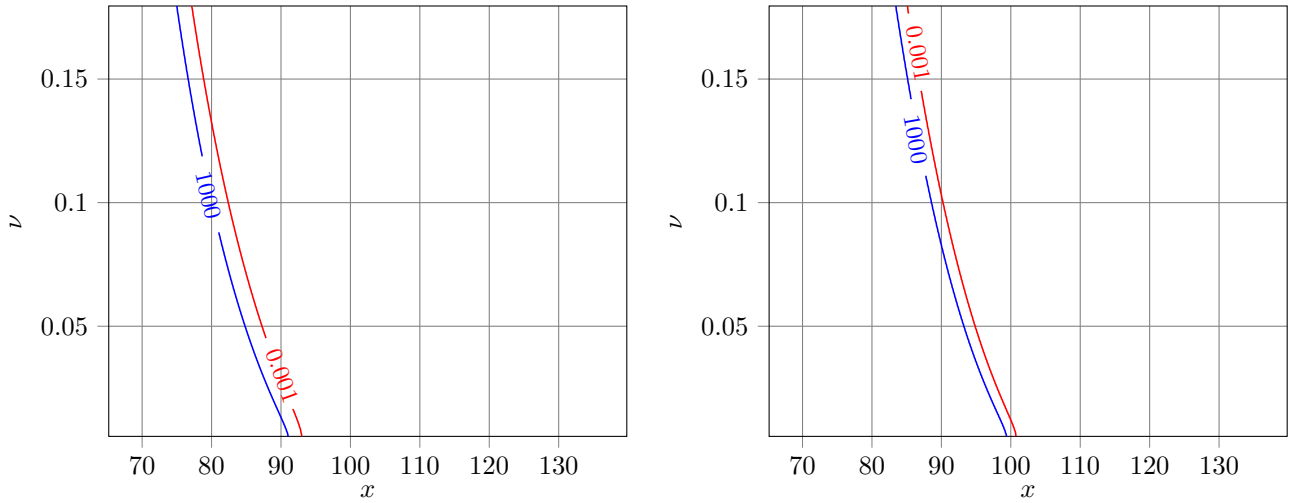

(a) $K = 100$ **(b)** $K = 110$

Figure 7: Level sets of optimal exercise rates at $t = 0.5$ for American put option in the Heston model.

Finally, we consider a 10-dimensional portfolio where each underlying $(X_t^i)_{t \in \mathcal{T}}, 1 \leq i \leq 10$ follows Equation (8)

with the same volatility process $(\nu_t)_{t\in\mathcal{T}}$ (and the same parameter values as in the one-dimensional case) but different Wiener processes $(W^{X^i})_{t\in\mathcal{T}}$, $1 \le i \le 10$ such that the 11-dimensional Wiener process $(W_t^{X^1}, \dots, W_t^{X^{10}}, W_t^{\nu})$ has the covariance matrix

$$\Sigma = \begin{pmatrix}
1. & 0.2 & 0.2 & 0.35 & 0.2 & 0.25 & 0.2 & 0.2 & 0.3 & 0.2 & -0.5 \\
0.2 & 1. & 0.2 & 0.2 & 0.2 & 0.125 & 0.45 & 0.2 & 0.2 & 0.45 & -0.5 \\
0.2 & 0.2 & 1. & 0.2 & 0.2 & 0.2 & 0.2 & 0.2 & 0.45 & 0.2 & -0.5 \\
0.35 & 0.2 & 0.2 & 1. & 0.2 & 0.2 & 0.2 & 0.2 & 0.425 & 0.2 & -0.5 \\
0.2 & 0.2 & 0.2 & 0.2 & 1. & 0.1 & 0.2 & 0.2 & 0.5 & 0.2 & -0.5 \\
0.25 & 0.125 & 0.2 & 0.2 & 0.1 & 1. & 0.2 & 0.2 & 0.35 & 0.2 & -0.5 \\
0.2 & 0.45 & 0.2 & 0.2 & 0.2 & 0.2 & 1. & 0.2 & 0.2 & 0.2 & -0.5 \\
0.2 & 0.2 & 0.2 & 0.2 & 0.2 & 0.2 & 0.2 & 1. & 0.2 & -0.1 & -0.5 \\
0.3 & 0.2 & 0.45 & 0.425 & 0.5 & 0.35 & 0.2 & 0.2 & 1. & 0.2 & -0.5 \\
0.2 & 0.45 & 0.2 & 0.2 & 0.2 & 0.2 & 0.2 & -0.1 & 0.2 & 1. & -0.5 \\
-0.5 & -0.5 & -0.5 & -0.5 & -0.5 & -0.5 & -0.5 & -0.5 & -0.5 & -0.5 & 1
\end{pmatrix}$$

Figure 8 shows estimates of the American basket put option (with coefficients $c \equiv 1/10$) value based on exercise rate optimization for the corresponding 11-dimensional process $S_t := (X_t^1, \dots, X_t^{10}, \nu_t)$, using the same discretization parameters as before.

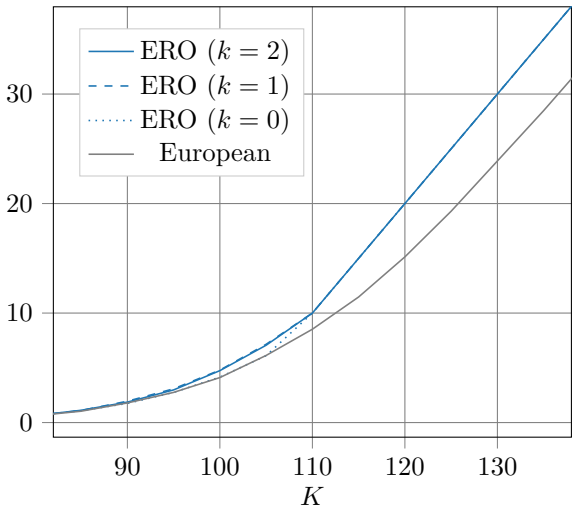

Figure 8: Dependence of American basket put option price on strike in the 10-dimensional Heston model.

3.5 Rough volatility

To illustrate the wide applicability of our method, we conclude this section with the non-Markovian *rough Bergomi* model, which has previously been applied to explain implied volatility smiles and other phenomena in the pricing of European options [6]. In non-Markovian models, Equation (2) does not hold because optimal exercise strategies may be based on the entire history of the path $(S_t)_{t\in\mathcal{T}}$ (which we again assume to include the underlying asset $(X_t)_{t\in\mathcal{T}}$ as well as the volatility $(\nu_t)_{t\in\mathcal{T}}$). Therefore, we consider the infinite-dimensional Markovian extension

$$\tilde{S}_t := (S_u)_{u\in[0,t]} \quad \forall t \in \mathcal{T}.$$

for which Equation (2) formally holds with subsets of $\mathcal{T} \times \mathbb{R}_+^d$ replaced by subsets of $\mathcal{T} \times \Gamma$, where $\Gamma := \bigcup_{t\in\mathcal{T}} \{s : [0,t] \to \mathbb{R}_+^d\}$.

For numerical purposes, we subsample realizations of S_t (with the convention that $S_t := S_0$ for $t < 0$) and define

$$\tilde{\boldsymbol{S}}_t := (S_t, S_{t-\Delta_1}, \dots, S_{t-\Delta_J}) \in \mathbb{R}^{d_{\text{eff}}} := \mathbb{R}^{2\times(1+J)} \quad \forall t \in \mathcal{T}$$

for some $J < \infty$ and $0 < \Delta_1 < \cdots < \Delta_J$. We may now apply the algorithm described in Section 2.1 to the resulting problem of finding exercise rates on the extended space $\mathcal{T} \times \mathbb{R}^{d_{\text{eff}}}$.

Following [6, Section 4], we generate samples from the risk-neutral measure induced by

$$X_t := x_0 \mathcal{E}\left(\int_0^t \sqrt{\nu_u}\, \mathrm{d}W_u^X\right) \tag{10}$$

$$\nu_t := \nu_0 \mathcal{E}\left(\eta\sqrt{2H} \int_0^t \frac{1}{(t-u)^{1/2-H}}\, \mathrm{d}W_u^\nu\right), \tag{11}$$

where $\mathcal{E}$ is the stochastic exponential, W^X and W^ν are Wiener processes with correlation $\rho = -0.9$, and $H = 0.07$, $\eta = 1.9$. Table 2 shows the American option prices for $x_0 = 100$, $\nu_0 = 0.09$, $T = 1$ and different strikes, which we computed using the discretization parameters $M = 100\,000$, $N = 128$, $k = 2$, and $\Delta_j := j/8$, $1 \le j \le J$, $J \in \{0, 1, 3, 7\}$. For comparison, we include the European prices computed by simple Monte Carlo simulation. The difference between our estimates for $J = 0$ and $J = 7$ is not consistently larger than the Monte Carlo sampling error, which indicates that exploitation of non-Markovian features does not yield significantly improved exercise strategies. This is not to say, however, that American option prices are similar in non-Markovian and Markovian models. Non-Markovianity of the samples of $(S_t)_{t\in\mathcal{T}}$ plays an important role in the evaluation of any given strategy, even when the strategy only depends on the spot values.

		K							
		70	80	90	100	110	120	130	140
Euro.		1.83	3.13	5.06	7.98	12.21	17.99	25.35	33.88
J	0	1.88	3.23	5.32	8.51	13.24	20	30	40
	1	1.88	3.23	5.31	8.50	13.22	20	30	40
	3	1.88	3.21	5.31	8.50	13.22	20	30	40
	7	1.88	3.22	5.30	8.50	13.23	20	30	40

Table 2: Prices of put option in rough Bergomi model.

The numerically optimized exercise rates at $t = 0.5$ for $J = 0$ and $K \in \{100, 110\}$ are shown in Figure 9.

4 Conclusion

We have introduced a novel method for the pricing of American options based on the optimization of randomized exercise strategies, which replace deterministic exercise regions by probabilistic exercise rates.

Since the objective function of the corresponding relaxed optimization problem is smooth, optimal exercise rates can be found using simple deterministic optimization routines. Furthermore, our numerical experiments have shown that exercise rates based on quadratic polynomials often suffice to obtain remarkably accurate price estimates and that the resulting non-concave objective function can be globally maximized using only a few iterations. Since the market model only appears in the simulation of sample paths, our method is quite flexible and easy to implement. We have demonstrated its practical applicability in uni- and multivariate Black–Scholes, Heston and rough Bergomi models.

In even higher-dimensional situations than those considered in this work, already the space of quadratic polynomials may be prohibitively large. In this case, the polynomial subspace $\mathcal{P}$ could be designed in an anisotropic way to exploit, for example, the fact that the exercise decision of basket put options with coefficients c is most sensitive to the coordinate $\tilde{s}_1 := c \cdot s$. For situations where large polynomial subspaces are unavoidable, a rigorous analysis of

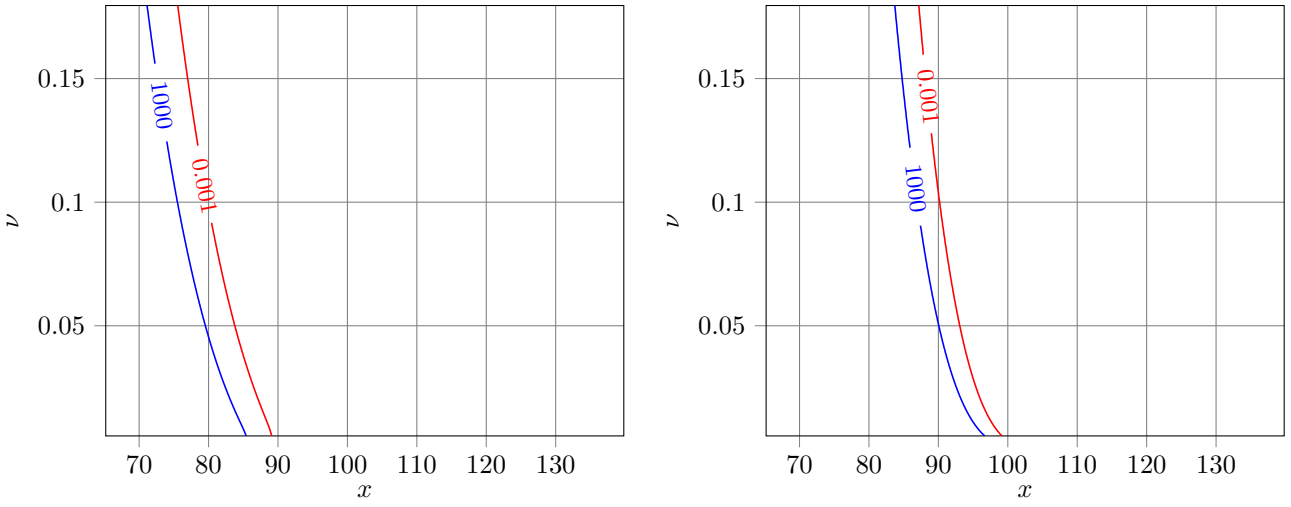

(a) $K = 100$ **(b)** $K = 110$

Figure 9: Level sets of optimal exercise rates at $t = 0.5$ for American put option in the rough Bergomi model ($k = 2$, $J = 0$)

the number of samples required to determine a given number of degrees of freedom without significant overfitting would be beneficial. Without such an analysis, it may be worthwhile to compute test values at each step during the optimization and terminate as soon as the test value decreases. To accelerate our method, multilevel Monte Carlo methods [15] could be used for function and gradient evaluations during the optimization.

Acknowledgments C. Bayer gratefully acknowledges support from the German Research Foundation (DFG, grant BA5484/1). R. Tempone and S. Wolfers are members of the KAUST SRI Center for Uncertainty Quantification at the Computer, Electrical and Mathematical Sciences & Engineering Division at King Abdullah University of Science and Technology (KAUST).

References

[1] Yves Achdou and Olivier Pironneau. *Computational methods for option pricing*. SIAM, 2005.

[2] Leif Andersen. A simple approach to the pricing of Bermudan swaptions in the multi-factor LIBOR market model. *Journal of Computational Finance*, 3:5–32, 1999.

[3] Leif Andersen and Mark Broadie. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, 50(9):1222–1234, 2004.

[4] Thomas Bagby, Len Bos, and Norman Levenberg. Multivariate simultaneous approximation. *Constructive Approximation*, 18(4):569, December 2002.

[5] Giovanni Barone-Adesi and Robert E. Whaley. Efficient analytic approximation of American option values. *The Journal of Finance*, 42(2):301–320, 1987.

[6] Christian Bayer, Peter Friz, and Jim Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6):887–904, 2016.

[7] Denis Belomestny and John Schoenmakers. *Advanced Simulation-Based Methods for Optimal Stopping and Control: With Applications in Finance*. Springer, 2018.

[8] Mark Broadie and Jérôme Detemple. The valuation of American options on multiple assets. *Mathematical Finance*, 7(3):241–286, 1997.

[9] Mark Broadie and Paul Glasserman. Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control*, 21(8):1323 – 1352, 1997.

[10] Peter Carr. Randomization of the American put. *The Review of Financial Studies*, 11:597–626, 1998.

[11] Cristina Costantini, Emmanuel Gobet, and Nicole El Karoui. Boundary sensitivities for diffusion processes in time dependent domains. *Appl. Math. Optim.*, 54(2):159–187, 2006.

[12] John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of financial Economics*, 7(3):229–263, 1979.

[13] Diego Garcıa. 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.

[14] Simon Gemmrich. Multilevel Monte Carlo methods for American options. Master's thesis, University of Oxford, 2012.

[15] Michael B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.

[16] Dwight Grant, Gautam Vora, and David Weeks. Path-dependent options: Extending the Monte Carlo simulation approach. *Management Science*, 43(11):1589–1602, 1997.

[17] 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.

[18] Alfredo Ibanez and Fernando Zapatero. Monte Carlo valuation of American options through computation of the optimal exercise frontier. *Journal of Financial and Quantitative Analysis*, 39(2):253–275, 2004.

[19] Ioannis Karatzas and Steven E. Shreve. *Methods of mathematical finance*. Springer, 1998.

[20] Rachel A. Kuske and Joseph B. Keller. Optimal exercise boundary for an American put option. *Applied Mathematical Finance*, 5(2):107–116, 1998.

[21] Francis A. Longstaff and Eduardo S. Schwartz. Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies*, 14(1):113–147, 2001.

[22] Michael Ludkovski. Kriging metamodels and experimental design for Bermudan option pricing. *Journal of Computational Finance*, 22(1):37–77, 2018.

[23] Leonard C. G. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.

[24] Albert N. Shiryaev. *Optimal stopping rules*, volume 8. Springer Science & Business Media, 2007.

[25] 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.