

**Weierstraß-Institut
für Angewandte Analysis und Stochastik
Leibniz-Institut im Forschungsverbund Berlin e. V.**

Preprint

ISSN 2198-5855

Pricing American options by exercise rate optimization

Christian Bayer¹, Raúl F. Tempone^{2,3}, Sören Wolfers³

submitted: November 27, 2019

¹ Weierstrass Institute
Mohrenstr. 39
10117 Berlin
Germany
E-Mail: christian.bayer@wias-berlin.de

² RWTH Aachen
Departement for Mathematics
Kackertstraße 9
52072 Aachen
E-Mail: tempone@uq.rwth-aachen.de

³ KAUST
Computer, Electrical and Mathematical
Sciences & Engineering Division (CEMSE)
Thuwal 23955 – 6900
Saudi Arabia
E-Mail: raul.tempone@kaust.edu.sa
soeren.wolfers@kaust.edu.sa

No. 2651
Berlin 2019



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

Key words and phrases. Computational finance, American option pricing, stochastic optimization problem, Monte Carlo, multivariate approximation, rough volatility.

This work was supported by the KAUST Office of Sponsored Research (OSR, award URF/1/2584-01-01), the German Research Foundation (DFG, grant BA5484/1) and the Alexander von Humboldt Foundation. R. Tempone and S. Wolfers are members of the KAUST SRI Center for Uncertainty Quantification in Computational Science and Engineering.

Edited by
Weierstraß-Institut für Angewandte Analysis und Stochastik (WIAS)
Leibniz-Institut im Forschungsverbund Berlin e. V.
Mohrenstraße 39
10117 Berlin
Germany

Fax: +49 30 20372-303
E-Mail: preprint@wias-berlin.de
World Wide Web: <http://www.wias-berlin.de/>

Pricing American options by exercise rate optimization

Christian Bayer, Raúl F. Tempone, Sören Wolfers

Abstract

We present a novel method for the numerical pricing of American options based on Monte Carlo simulation and the 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 a given option is exercised. In contrast, our method determines the *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. The global optimum of this function can be approached gradually when starting from a constant exercise rate. Numerical experiments on vanilla put options in the multivariate Black–Scholes model and a 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 the exercise rates. Finally, we demonstrate the flexibility of our method through numerical experiments on max call options in the classical Black–Scholes model, and vanilla put options in both the Heston model and the non-Markovian rough Bergomi model.

1 Introduction

American options on $d \geq 1$ underlying assets $S_t = (S_{1,t}, \dots, S_{d,t})$ 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 $g: [0, T] \times \mathbb{R}_+^d \rightarrow \mathbb{R}_+$.

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 a risk-neutral measure \mathbb{Q} is determined solely by the current asset values. The value function $V: \mathbb{R}_+^d \rightarrow \mathbb{R}_+$ satisfies

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

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 $(S_t)_{0 \leq t \leq T}$ [24, Theorem 5.3]. In the remainder of this work, all expectations are taken with respect to the same risk-neutral measure \mathbb{Q} and denoted by \mathbb{E} .

Most state-of-the-art methods for American option pricing – including all variants of the Longstaff–Schwartz [27], PDE [1], binomial tree [16], and stochastic mesh [13] methods – exploit the dynamic programming principle to determine the value function using a backwards-iteration scheme. Further approaches are based on dual problems [30, 3], policy iteration [11], or (quasi-)analytic solutions [5, 26]. The computational cost of many methods grows exponentially

with respect to the number of dimensions, thus making them prohibitively expensive for options on many underlying assets. This phenomenon has been coined the curse of dimensionality [29, 7].

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

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

Here, the supremum is taken over Borel-measurable subsets of $E \subset [0, T] \times \mathbb{R}_+^d$, whose hitting times are given by $\tau_E := \inf\{t \geq 0 : (t, S_t) \in E\}$. To be precise, both Equation (1) and Equation (2) require some technical conditions on the processes $(Y_t)_{0 \leq t \leq T}$ and $(S_t)_{0 \leq t \leq T}$ [31, Corollary 2, Section 3.3.1]. Throughout this work, we assume that 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 [20] and developed in [2, 17, 23, 8, 18], but it has not yet found its way into the canon of numerical algorithms for American option pricing. In [20], separate exercise regions were determined for each exercise date of an American Asian option in a backwards iteration. The optimization at each step was performed in a brute force fashion, which explains why only two parameters were allowed in the parametrization of the exercise regions. In [17, 18], ad hoc parametrizations that exploit known behavior of the optimal exercise regions were used to optimize exercise regions as subsets of time-space without applying a backwards iteration.

In general, optimization of the exercise region faces two challenges. First, as mentioned in [18], it is not obvious how to parametrize the possible exercise regions in a multi-dimensional setting, or even in a one-dimensional setting that goes beyond vanilla options in the Black–Scholes model. Second, once a parametrization has been found, it is not obvious how to find the global optimum [17, 18]. Indeed, when the expectation in Equation (2) is replaced by an empirical average for the purpose of 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 [17].

To address these challenges, we introduce, in Section 2, a relaxation of the optimization problem in Equation (2) wherein the exercise regions $E \subset [0, T] \times \mathbb{R}_+^d$ are replaced by *exercise rates* $f: [0, T] \times \mathbb{R}_+^d \rightarrow \mathbb{R}_+$, which define *randomized exercise strategies* where options are exercised with an infinitesimal probability depending on the current time and asset values.¹ The space of exercise rates can easily be parametrized even in high dimensions using a finite-dimensional spaces of polynomials on $[0, T] \times \mathbb{R}_+^d$. The resulting optimization problem exhibits the same maximum as the original optimization problem over deterministic strategies but has the advantage of a differentiable objective function and a lower risk of getting stuck in local minima because of a richer search space. Indeed, by integrating analytically with respect to the exponential distribution that underlies the random exercise decision, we obtain an objective function that is smooth even when finitely many sample paths are used in the computations. We may then use gradient-based optimization routines to determine an optimal coefficient

¹We were informed after the initial submission of this manuscript that randomized stopping was previously studied from a theoretical perspective [21, 25]. These references do not contain discussions of numerical solution of the resulting stochastic optimization problem, however.

vector. Furthermore, we may start this optimization from an exercise rate that has a constant non-zero value across time and space and let the optimization routine gradually refine this neutral strategy towards an optimal one with marked variations in the exercise rate. This facilitates the search for a global optimum without requiring an informed initial guess that is already 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 discretization parameters. In particular, we provide heuristic bounds on the number of degrees of freedom in the exercise rate that are required for satisfactory randomized exercise strategies. These bounds are given in terms of the smoothness of the optimal exercise boundary as a manifold, not as a function of time.

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 $t \rightarrow T$. Despite the square-root singularity near the expiration time, the experiments presented in Section 3.1 show that low-degree polynomials suffice to capture the optimal exercise boundary well. In fact, we obtain a relative error of less than 0.1% with quadratic polynomials. 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.

Although we solve non-concave maximization problems, we are able to find global optima starting from a constant exercise rate. 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, 28]. Max call options pose a challenge to the direct determination of exercise regions because the optimal exercise regions are disconnected [12]. 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 for the generation of random sample paths only, its application in this scenario is straightforward. Finally, we consider the non-Markovian rough Bergomi model [6] in Section 3.5. 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

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 $f: \mathcal{T} \times \mathbb{R}_+^d \rightarrow \mathbb{R}_+$, the randomized exercise strategy with exercise rate f is given by early exercise at the time

$$\tau_f := \inf\{t \geq 0 : \int_0^t \lambda_u \, du \geq X\}, \quad (3)$$

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 dt .

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

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

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

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

and

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

Hence, we obtain

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

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

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

It is advisable to replace $\lambda U_t \, dt$ by $-dU_t$ in numerical implementations of this formula to avoid cancellations. The following proposition shows that, in theory, exercise rate optimization yields the correct option value. It is a special case of Theorem 2.2 in [21].

Proposition 2.2. We have

$$V(s_0) = \sup_{f: [0, T] \times \mathbb{R}_+^d \rightarrow \mathbb{R}_+} \psi(f). \quad (6)$$

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

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

into Equation (3) to obtain $\tau_{f_E} = \tau_E$. After replacing $+\infty$ with large numbers that diverge to $+\infty$ and applying Fatou's lemma, we may take the supremum over E to conclude from Equation (2) that $\sup_{f: [0, T] \times \mathbb{R}_+^d \rightarrow \mathbb{R}_+} \psi(f) \geq V(s_0)$.

Conversely, the law of total expectation shows, for any $f: [0, T] \times \mathbb{R}_+^d \rightarrow \mathbb{R}_+$, 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, $\psi(f) \leq V(s_0)$. \square

2.1 Numerical algorithm

To determine optimal exercise rates numerically, we

- (i) replace the time-continuous model of the stochastic process $(S_t)_{t \in \mathcal{T}}$ with a discretization with $N < \infty$ time steps, such as the Euler–Maruyama scheme;
- (ii) replace the expectation in Equation (5) with an average over $M < \infty$ fixed sample paths $(S_n^{(m)})_{1 \leq n \leq N, 1 \leq m \leq M}$;
- (iii) introduce a B -dimensional, $B < \infty$ parametrization $\mathbb{R}^B \ni \mathbf{c} \mapsto f_{\mathbf{c}}$ of the space of exercise rates;
- (iv) maximize the surrogate function

$$\begin{aligned} \bar{\psi}: \mathbb{R}^B &\rightarrow \mathbb{R} \\ \mathbf{c} &\mapsto \frac{1}{M} \sum_{m=1}^M \phi(f_{\mathbf{c}}, (S_t^{(m)})_{1 \leq n \leq N}).^2 \end{aligned}$$

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

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

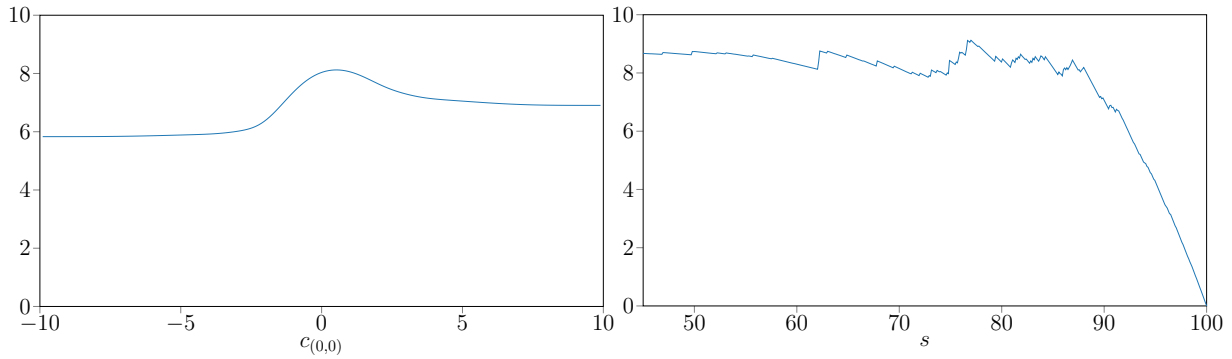
for any finite-dimensional linear space \mathcal{P} of functions on $\mathcal{T} \times \mathbb{R}^d$. After choosing a basis of \mathcal{P} , we obtain the desired parametrization $\mathbf{c} \mapsto f_{\mathbf{c}}$. Throughout the remainder of this manuscript, we work with spaces \mathcal{P}_k of polynomials of degree less than or equal to $k \geq 0$ in $d+1$ variables, and we use an orthonormal basis with respect to the inner product $\|f\|^2 := \frac{1}{NM} \sum_{n=1}^N \sum_{m=1}^M f(t_n, x_{n,m})$ induced by the time-space samples $(t_n, x_{n,m} := \log(S_n^m))_{1 \leq n \leq N, 1 \leq m \leq M}$.

Optimization Concerning step (iv), it is not clear that globally optimal coefficients, which may even lie at infinity, can be found numerically because $\bar{\psi}$ is not concave. However, in our numerical experiments, we found that the Quasi-Newton L-BFGS-B algorithm [14], as implemented in Python’s SciPy library³, performs well and does not get stuck in local maxima when started from a constant exercise rate.

²To evaluate ϕ , we use piecewise constant interpolation between the N nodes of the time-discretization scheme.

³<https://docs.scipy.org/doc/scipy/reference/optimize.minimize-lbfgsb.html>

The advantage of exercise rate optimization over exercise region optimization is illustrated by Figure 1. Even a simple gradient ascent algorithm could be used to maximize $\bar{\psi}$ in Figure 1a, where we show the dependence on the coefficient $c_{(0,0)}$ of the constant polynomial $p_{(0,0)} \equiv 1$ for a one-dimensional put option. For comparison, this is not possible for the function shown in Figure 1b, which arises from the optimization of deterministic exercise regions and requires the use of finite-difference stochastic-gradient algorithms.



(a) $\bar{\psi}: \mathbb{R} \rightarrow \mathbb{R}$ with \mathcal{P} the space of constant functions (b) $s \mapsto \frac{1}{M} \sum_{m=1}^M [Y_{\tau_{E_s} \wedge T}^{(m)}]$ 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 put option with $K = s_0 = 100$ and $T = 1$ in the Black–Scholes model with $r = 0.05$ and $\sigma = 0.3$. Both plots were generated using $M = 100$ sample paths with $N = 100$ time steps.

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

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

where

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

Figure 2 shows four snapshots of the search for an optimal exercise rate for max call options on two underlying securities.

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 .

For a fixed exercise rate and a fixed number of time steps, convergence with respect to the number of sample paths, M , occurs asymptotically at the Monte Carlo rate $M^{-1/2}$. Pre-asymptotically, the number of Monte Carlo samples has to be larger than a threshold depending on the dimension of the polynomial subspace to avoid overfitting, see the next paragraph.

For a fixed, smooth exercise rate, the expected payoff converges at the weak convergence rate of the discretization scheme with respect to the number of time steps (e.g., N^{-1} for the

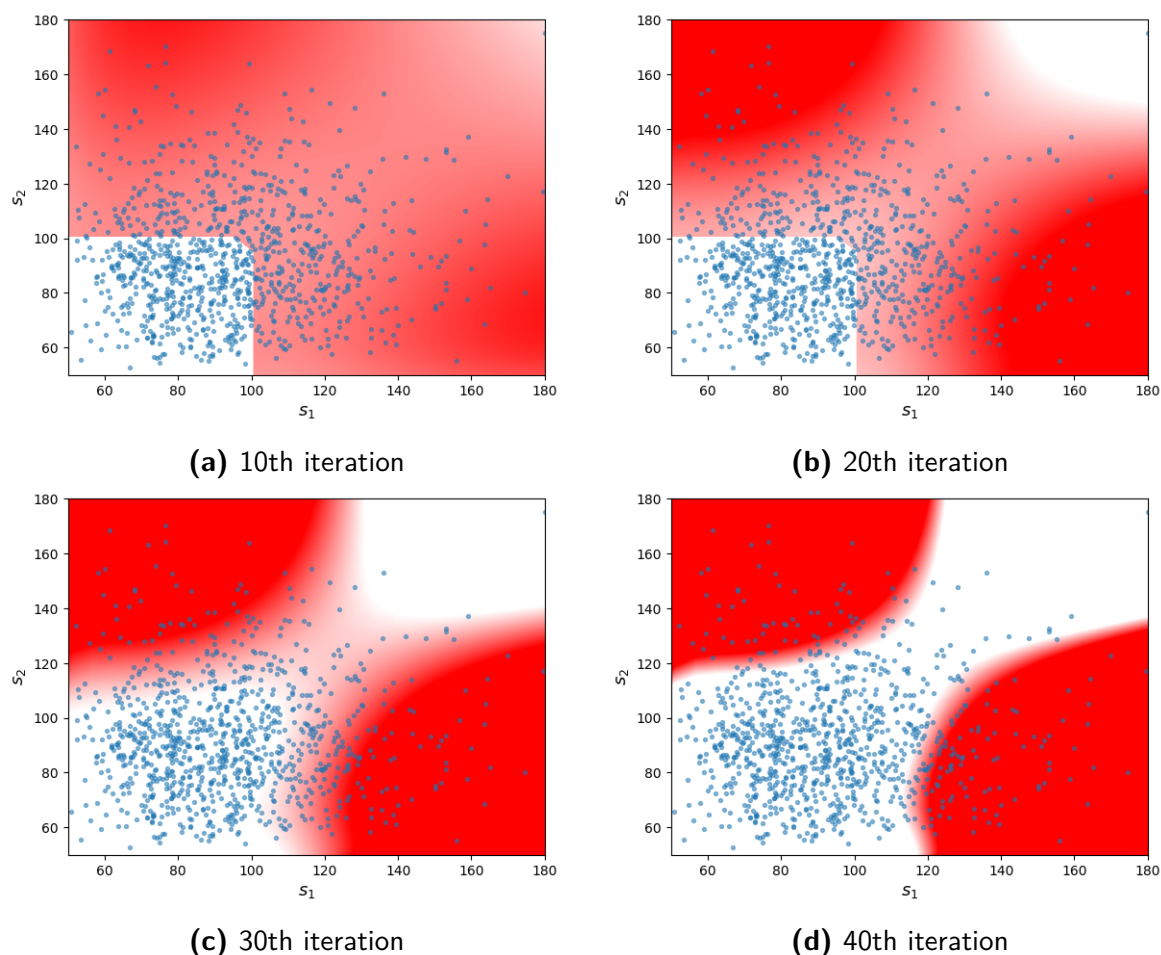


Figure 2: Four iterations of the exercise rate optimization for a max call option (all figures show a slice of the exercise rate at $t = T/2$). High color intensities represent high exercise rates. The white region in the bottom left contains the points with zero payoff, $\{g = 0\}$. Random sample values of the two underlying securities at $T/2$ are shown in blue.

Euler–Maruyama scheme). In the limit of increasingly steep exercise rates approaching the optimal deterministic exercise regions, the weak convergence rate is expected to deteriorate to $N^{-1/2}$. However, this effect does not become noticeable in our numerical experiments (see Section 3.1).

With everything else held fixed, we expect exponential or faster convergence with respect to ℓ , depending on what type of deterministic optimization routine is used. Figure 4 in Section 3 provides numerical evidence of exponential convergence using the L-BFGS-B algorithm.

To characterize the convergence of the optimal exercise rate with respect to k under the simplifying assumptions $M = \infty$ and $N = \infty$, 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 $L \rightarrow \infty$. Therefore, it suffices to study the approximability of the optimal exercise region E_* by polynomial superlevel sets, and the sensitivity of the expected payoff on the right-hand side of Equation (2) with respect to perturbations of the exercise region. Regarding the approximability of E_* , 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

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

for some $\Theta: B_* \rightarrow \mathbb{R}^{1+d}$ 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. Regarding the sensitivity of the expected payoff, [15] showed differentiability with respect to perturbations of the exercise region in spatial directions under the assumption that $(0, s_0) \notin E_*$ and that the payoff function 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.

A rigorous analysis of the interplay of the various discretizations will be the topic of future work; some numerical results are presented in Section 3.1 below.

Overfitting Choosing a subspace with a large number of degrees of freedom, $B \gg 1$, to improve the flexibility of the candidate exercise rates increases the cost of computations and the risk of *overfitting*. This means that the value of $\bar{\psi}(c^*)$ at the optimized coefficients c^* may overestimate the true value $\psi(f_{c^*})$ unless a correspondingly large number $M = M(B)$ of sample paths is used. Numerical experiments indicate that $M(B) \approx CB^2$ for some $C > 0$ but we were not able to prove such a formula. In practice, we can simply compute an unbiased estimate of $\psi(f_{c^*})$ 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 methods such as the Longstaff–Schwartz algorithm. Following statistical learning terminology, we refer to the biased and unbiased estimators of $\psi(f_{c^*})$ as *training* and *test* values, respectively. One way to avoid overfitting is to recompute the test value at each step of the optimization and to terminate as soon as the test value decreases. Note that, as in the case of the Longstaff–Schwartz algorithm, the test values are biased low, i.e., are Monte Carlo estimates of *lower bounds* of the option price.

Remark 1. *Note that time discretization does not correspond to restricting the exercise opportunities to a finite number of exercise times. In other words, the result of the exercise rate optimization based on a certain time-discretization of the above integrals does not correspond to the price of a Bermudan option. In fact, the randomization allows taking into account the possibility of exercising between time points.*

3 Numerical experiments

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

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 3a and 3b show that the prices found through exercise rate optimization with 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 $n \rightarrow \infty$. In particular, our maximization does not get stuck in local optima of $\bar{\psi}$. Furthermore, Figure 3a shows that test and training values converge at roughly the same speed, which means that we do not suffer from overfitting. This is not surprising, since the space of bivariate quadratic polynomials is only 6-dimensional. We restrict the 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 3b, we 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 about four significant digits, despite using only quadratic polynomials for the exercise boundary approximation. This confirms that singularities of the exercise boundary as a function of time do not pose a problem for our polynomial approximation scheme. For comparison, Figures 3c and 3d 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 closer inspection on a logarithmic scale reveals stagnation at a relative error of 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 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, for $n = 4$ and $k = 2$ the relative error between $\bar{\psi}(c_\ell)$ and the final value is already below 0.1% when $\ell = 20$ (Figure 4). For this reason, we limit the number of iterations below to 20.

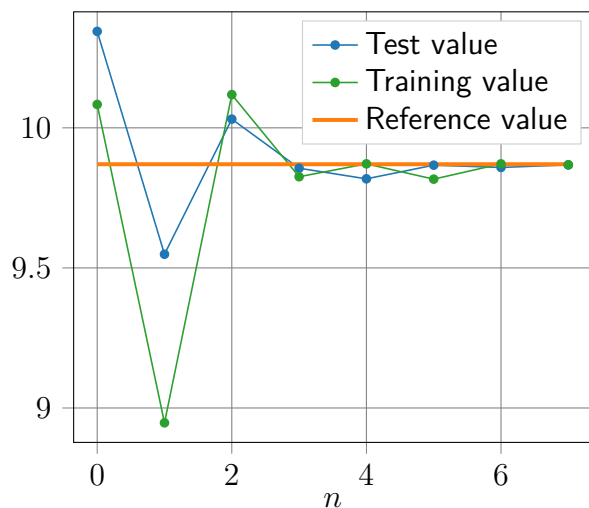
3.2 Comparison with Longstaff–Schwartz algorithm

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 $c_i := 1/d$, $1 \leq i \leq d$.

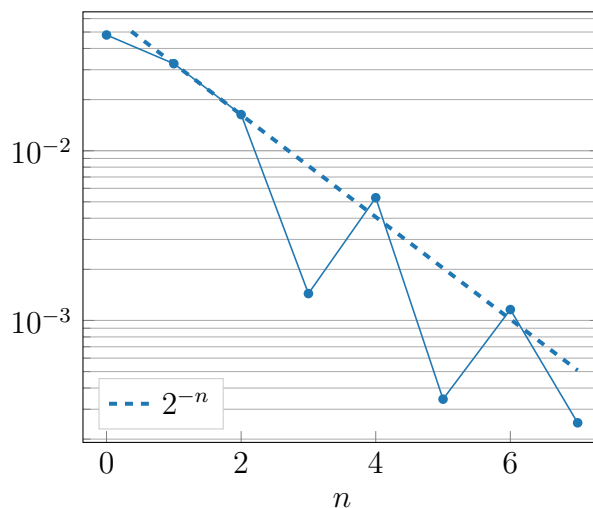
We compare our method to the Longstaff–Schwartz algorithm, as implemented in the freely available version 16 of the derivative pricing software Premia⁴. Like our method, the Longstaff–Schwartz algorithm requires specification of the number of sample paths, the number of time-steps used for their simulation, and the polynomial degree, 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 $M = 3.2 \times 10^6$ samples for both. Finally, we use a risk-free interest rate $r = 0.05$ and a diagonal volatility matrix $\Sigma_{ij} = 0.3^2 \delta_{ij}$, $1 \leq i, j \leq d$ 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,

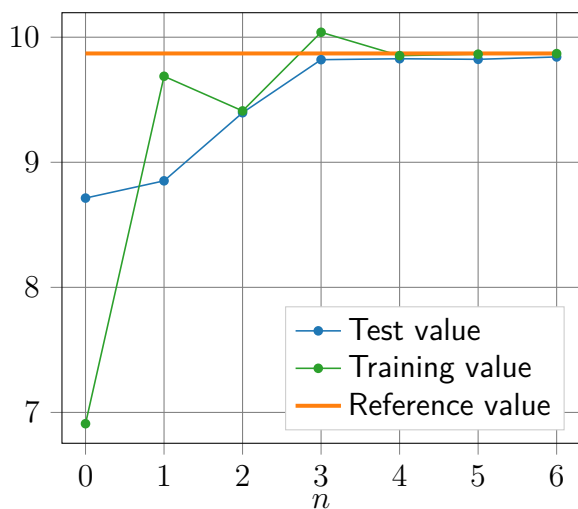
⁴<https://www.rocq.inria.fr/mathfi/Premia>



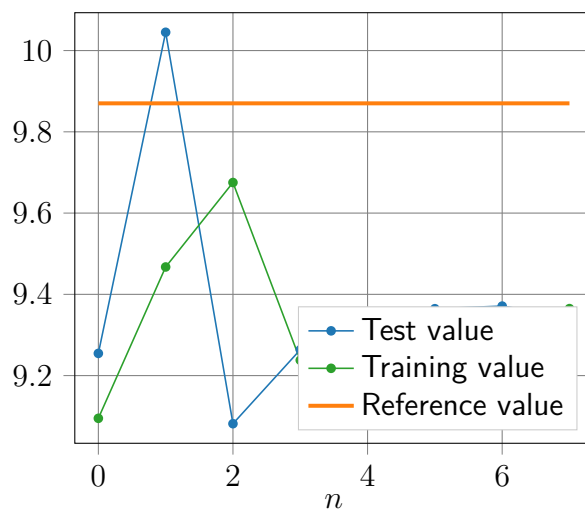
(a) $k = 2$



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



(c) $k = 1$



(d) $k = 0$

Figure 3: Exercise rate optimization with polynomial degree $0 \leq k \leq 2$, $M_n = 200 \times 4^n$ and $N_n = 2^n$, $0 \leq n \leq 7$ applied to a one-dimensional American put option in the Black–Scholes model with $\sigma = 0.3$, $r = 0.05$, $K = 100$, $s_0 = 100$, and $T = 1$.

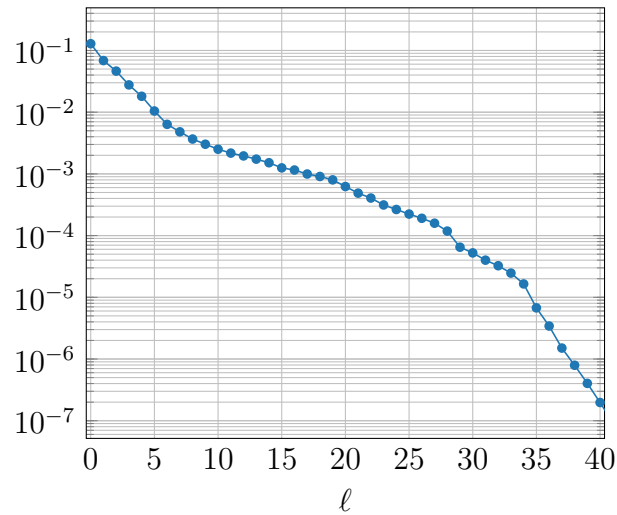


Figure 4: Convergence with respect to the number of function evaluations, ℓ , in the training step of exercise rate optimization for an American put option using the L-BFGS-B algorithm.

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

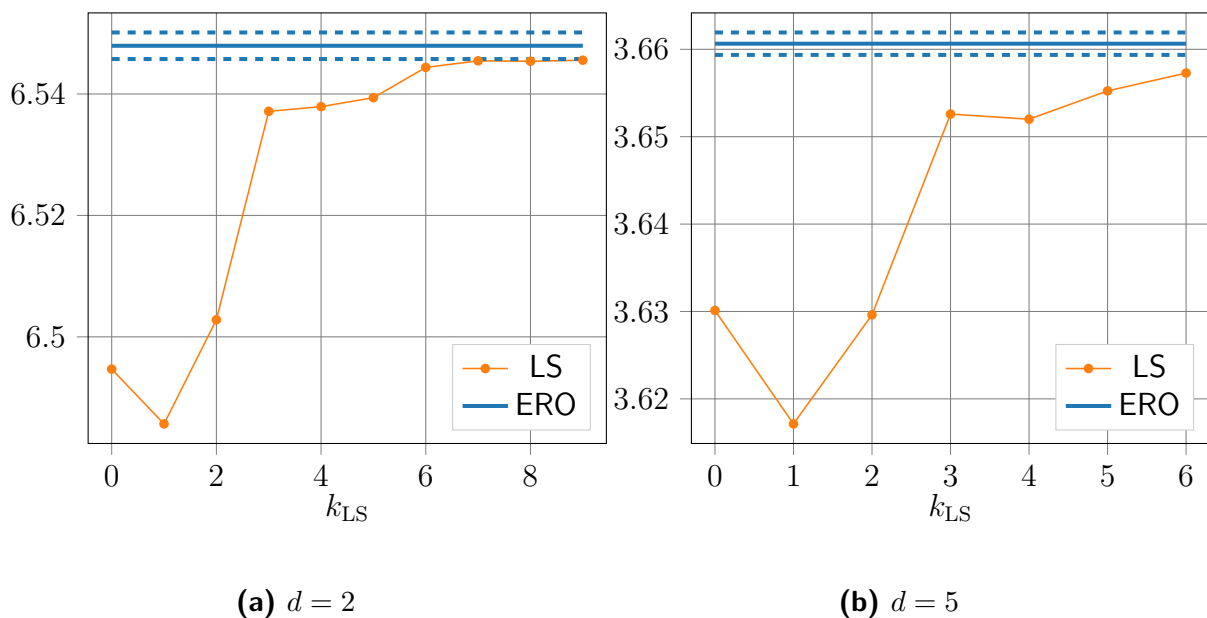


Figure 5: Convergence of the Longstaff–Schwartz algorithm (LS) for $\{2, 5\}$ -dimensional basket put options with increasing polynomial degree k_{LS} to reference values computed via exercise rate optimization (ERO) with polynomial degree $k_{\text{ERO}} = 2$ and 95% confidence bands (dashed).

Runtime comparison To obtain a fair runtime comparison, we created a Python package⁵ with straightforward implementations of both algorithms, which we ran on a 12 core Intel Xeon X5650 CPU.

For the same polynomial degree, exercise rate optimization is slower than the Longstaff–Schwartz algorithm. However, as we have seen above, the latter requires larger polynomial degrees for accurate results. Since the ratio between the dimensions of polynomial subspaces with degrees $k = 2$ and $k > 2$ grows with respect to the dimension of the domain, exercise rate optimization returns accurate results faster than the Longstaff–Schwartz algorithm in high-dimensional examples.

For example, for a basket put option as above with $d = 10$, the Longstaff–Schwartz algorithm returns 2.235 with $k = 2$ after 530 seconds and 2.237 with $k = 4$ after 7437 seconds. Exercise rate optimization, on the other hand, returns 2.240 with $k = 2$ after 2493 seconds. All these results were obtained with the same 3.2×10^6 Brownian motion samples.

3.3 Max call options

In this subsection, we consider max call options on two underlying assets, for which $g(s) := \max\{(s_1 - K)^+, (s_2 - K)^+\}$. These max call options present an interesting challenge for our method, since the optimal exercise region at any time before expiry has two connected components [12]. 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 $\delta = 0.1$ are taken from [3] and provided in Table 1 alongside 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 6. As expected, they are

		95% CI	1	k 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].

almost deterministic, which means that they exhibit steep slopes from values close to zero to values close to infinity. Since the specific values 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 a 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 modeling the disconnected regions is not possible with log-linear exercise rates available for $k = 1$, the hyperbolic conic sections available with $k = 2$ already provide satisfactory approximations.

⁵<https://pypi.org/project/pryce/>

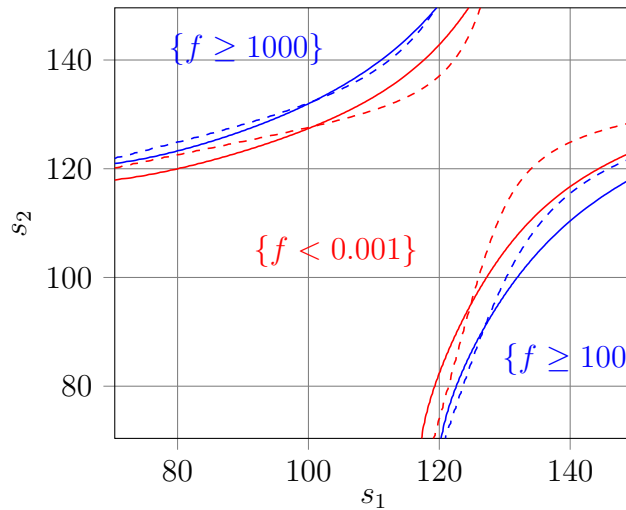


Figure 6: Level sets of optimal exercise rates for a max call option with $k = 2$ (dashed) and $k = 3$ (solid).

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 [22], which models the evolution of a single underlying asset X_t and its instantaneous variance v_t using the coupled system of stochastic differential equations

$$dX_t = \mu X_t dt + \sqrt{v_t} X_t dW_t^X, \quad (8)$$

$$dv_t = \kappa(\theta - v_t) dt + \xi \sqrt{v_t} dW_t^v, \quad (9)$$

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

Since our method requires Markovian markets, we must include the volatility and define $S_t := (X_t, v_t)$, $t \in \mathcal{T}$. This 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 replace μ with 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 $v_K(s_0)$ for 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` 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 7. The maximal relative difference between the two methods is 1% and occurs around $K^* = 130$. 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 for large enough K the initial point $(100, 0.15)$ lies within the optimal exercise region and the option is thus exercised immediately.

Figure 8 shows the numerically optimized exercise rates (with $k = 2$) at $t = 0.5$ for $K \in \{100, 110\}$.

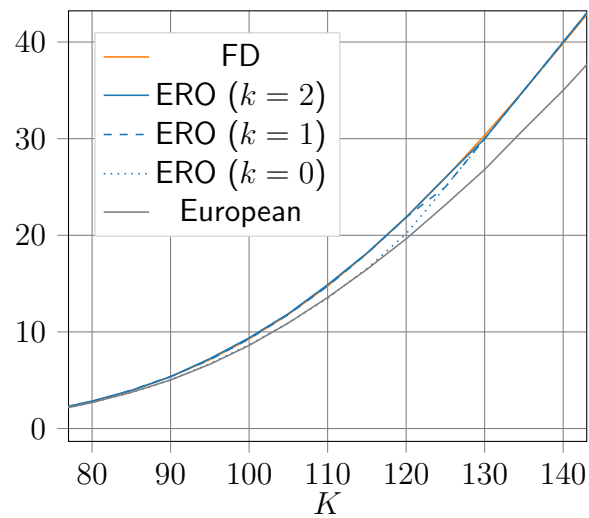


Figure 7: Dependence of the put option price on the strike in the Heston model; computed using exercise rate optimization (ERO) and a finite-difference method (FD).

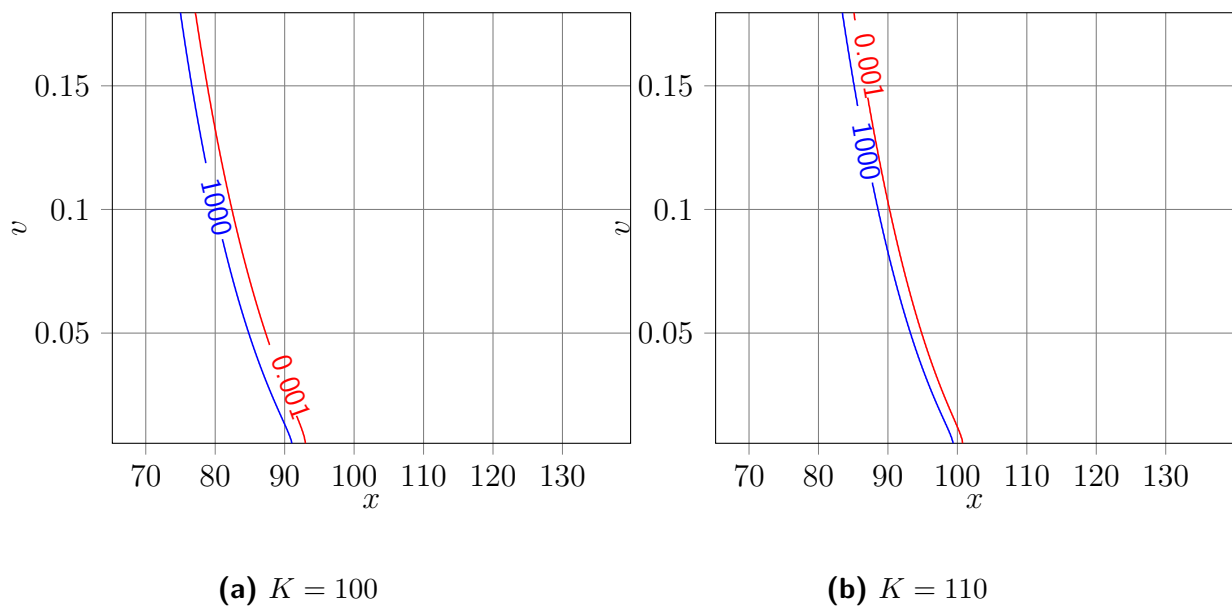


Figure 8: Level sets of optimal exercise rates at $t = 0.5$ for a put option in the Heston model.

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 $(v_t)_{t \in \mathcal{T}}$ (and the same parameter values as in the one-dimensional case) but different Wiener processes $(W^{X^i})_{t \in \mathcal{T}}, 1 \leq i \leq 10$ such

that the 11-dimensional Wiener process $(W_t^{X^1}, \dots, W_t^{X^{10}}, W_t^v)$ 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 9 shows estimates of the values of American basket put options (with coefficients $c \equiv 1/10$) that were obtained by exercise rate optimization for the corresponding 11-dimensional process $S_t := (X_t^1, \dots, X_t^{10}, v_t)$ using the same discretization parameters as before.

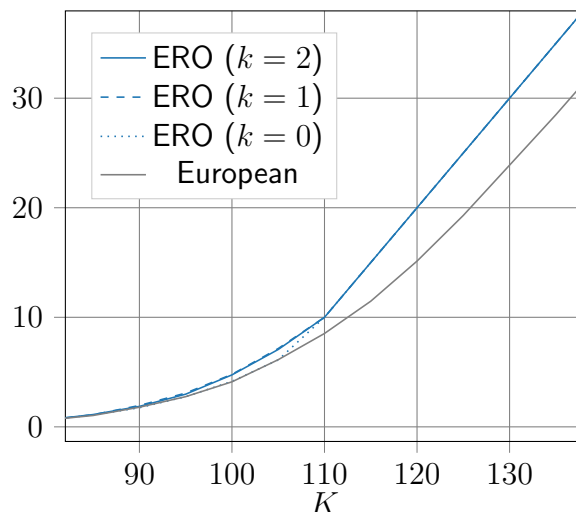


Figure 9: Dependence of basket put option price on the strike in the 10-dimensional Heston model.

3.5 Rough volatility

To illustrate the wide applicability of our method, we conclude this section with the non-Markovian *rough Bergomi* model, which was previously 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 $(v_t)_{t \in \mathcal{T}}$. Therefore, we consider the infinite-dimensional Markovian extension

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

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

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

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

for some $J < \infty$ and $0 < \Delta_1 < \dots < \Delta_J$. We 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

$$dX_t = rX_t dt + X_t \sqrt{v_t} dW_t^X, \quad X_0 = x_0, \quad (10)$$

$$v_t := v_0 \mathcal{E} \left(\eta \sqrt{2H} \int_0^t \frac{1}{(t-u)^{1/2-H}} dW_u^v \right), \quad (11)$$

where \mathcal{E} is the stochastic exponential in the Wick sense, $H = 0.07$, $r = 0.05$, $\eta = 1.9$, and W^X, W^v are Wiener processes with correlation $\rho = -0.9$. Since the asset price process X_t is a continuous local martingale, standard no arbitrage theory applies even though v_t is not a semi-martingale.

Table 2 shows the American option prices for $x_0 = 100$, $v_0 = 0.09$, $T = 1$, and different strikes, which we computed using the discretization parameters $M = 100\,000$, $N = 128$, $k = 2$, and $\Delta_j := j/8$, $1 \leq j \leq J$, $J \in \{0, 1, 3, 7\}$. 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, indicating that the exploitation of non-Markovian features does not yield significantly improved exercise strategies. This is not to say, however, that American option prices in non-Markovian and Markovian models are similar. The 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	
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	
J	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 options in the rough Bergomi model.

The numerically optimized exercise rates at $t = 0.5$ for $J = 0$ and $K \in \{100, 110\}$ are shown in Figure 10.

4 Conclusion

We have introduced a method of pricing American options by optimization of randomized exercise strategies, in which deterministic exercise regions are replaced by probabilistic exercise rates.

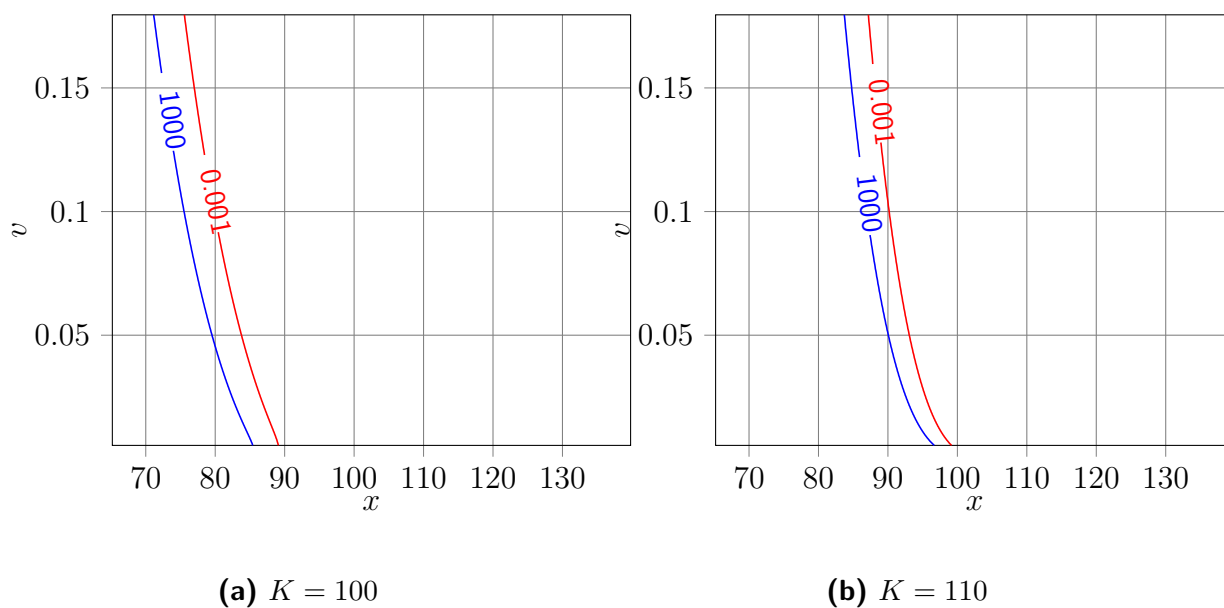


Figure 10: Level sets of optimal exercise rates at $t = 0.5$ for put options in the rough Bergomi model ($k = 2$, $J = 0$).

Since the objective function of the corresponding relaxed optimization problem is smooth, optimal exercise rates can be found using simple deterministic optimization routines. Our numerical experiments show that exercise rates based on quadratic polynomials are sufficient to obtain remarkably accurate price estimates and that the resulting non-concave objective functions 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 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 that 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 the number of samples that are required to determine a given number of degrees of freedom without significant overfitting would be of interest; similar but not directly transferable results were established in [9, 32].

To accelerate numerical implementations, multilevel Monte Carlo methods [19] could be used for evaluations of the expected payoff and its gradient.

It is an open question whether efficiently computable upper bounds on the option price [10] can be constructed using exercise rates as well.

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] Richard Bellman. *Adaptive control processes: A guided tour. (A RAND Corporation Research Study)*. Princeton University Press, 1961.
- [8] Denis Belomestny. On the rates of convergence of simulation-based optimization algorithms for optimal stopping problems. *Ann. Appl. Probab.*, 21(1):215–239, 02 2011.
- [9] Denis Belomestny et al. On the rates of convergence of simulation-based optimization algorithms for optimal stopping problems. *The Annals of Applied Probability*, 21(1):215–239, 2011.
- [10] Denis Belomestny et al. Solving optimal stopping problems via empirical dual optimization. *The Annals of Applied Probability*, 23(5):1988–2019, 2013.
- [11] Denis Belomestny and John Schoenmakers. *Advanced Simulation-Based Methods for Optimal Stopping and Control: With Applications in Finance*. Springer, 2018.
- [12] Mark Broadie and Jérôme Detemple. The valuation of American options on multiple assets. *Mathematical Finance*, 7(3):241–286, 1997.
- [13] Mark Broadie and Paul Glasserman. Pricing American-style securities using simulation. *Journal of Economic Dynamics and Control*, 21(8):1323 – 1352, 1997.
- [14] Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [15] 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.
- [16] John C. Cox, Stephen A. Ross, and Mark Rubinstein. Option pricing: A simplified approach. *Journal of financial Economics*, 7(3):229–263, 1979.
- [17] 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.
- [18] Simon Gemmrich. Multilevel Monte Carlo methods for American options. Master’s thesis, University of Oxford, 2012.

- [19] Michael B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.
- [20] Dwight Grant, Gautam Vora, and David Weeks. Path-dependent options: Extending the Monte Carlo simulation approach. *Management Science*, 43(11):1589–1602, 1997.
- [21] István Gyöngy and David Šiška. On randomized stopping. *Bernoulli*, 14(2):352–361, 2008.
- [22] 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.
- [23] 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.
- [24] Ioannis Karatzas and Steven E. Shreve. *Methods of mathematical finance*. Springer, 1998.
- [25] Nikolaj Vladimirovič Krylov. *Controlled diffusion processes*, volume 14. Springer Science & Business Media, 2008.
- [26] Rachel A. Kuske and Joseph B. Keller. Optimal exercise boundary for an American put option. *Applied Mathematical Finance*, 5(2):107–116, 1998.
- [27] 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.
- [28] Michael Ludkovski. Kriging metamodels and experimental design for Bermudan option pricing. *Journal of Computational Finance*, 22(1):37–77, 2018.
- [29] Christoph Reisinger and Gabriel Wittum. Efficient hierarchical approximation of high-dimensional option pricing problems. *SIAM Journal on Scientific Computing*, 29(1):440–458, 2007.
- [30] Leonard C. G. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.
- [31] Albert N. Shiryaev. *Optimal stopping rules*, volume 8. Springer Science & Business Media, 2007.
- [32] 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.