WILEY MATHEMATICAL FINANCE An International Journal of Mathematics, Statistics and Financial Economics

ORIGINAL ARTICLE

The optimal method for pricing Bermudan options by simulation

Alfredo Ibáñez[1] | Carlos Velasco[2]

[1]Department of Business Administration, ITAM, Río Hondo 1, Ciudad de México 01080, México

[2]Department of Economics, Universidad Carlos III de Madrid, C/Madrid, 126, E-28903 Getafe (Madrid), Spain

Correspondence
Alfredo Ibáñez, Department of Business Administration, ITAM, Río Hondo 1, Ciudad de México 01080, México.
Email: alfre.ibanez@gmail.com

Abstract

Least-squares methods enable us to price Bermudan-style options by Monte Carlo simulation. They are based on estimating the option continuation value by least-squares. We show that the Bermudan price is maximized when this continuation value is estimated near the exercise boundary, which is equivalent to implicitly estimating the optimal exercise boundary by using the value-matching condition. Localization is the key difference with respect to global regression methods, but is fundamental for optimal exercise decisions and requires estimation of the continuation value by iterating local least-squares (because we estimate and localize the exercise boundary at the same time). In the numerical example, in agreement with this optimality, the new prices or lower bounds (i) improve upon the prices reported by other methods and (ii) are very close to the associated dual upper bounds. We also study the method's convergence.

KEYWORDS

American and Bermudan options, local least-squares, optimal stopping-times, optimization, simulation

1 | INTRODUCTION

Pricing Bermudan options by simulation has drawn the interest of practitioners and academics alike, because many securities contain early-exercise features and depend on several stochastic factors. A Bermudan option is a classical example of an optimal stopping-time problem. Simulation is a flexible method that does not suffer from the curse of dimensionality. Carriére (1996), Tsitsiklis and Van Roy (1999), and Longstaff and Schwartz (2001) develop a least-squares Monte Carlo (LSM) approach, which is widely used in practice.

Pricing Bermudan options by simulation requires providing an exercise policy (and evaluating this policy by a standard simulation). We derive the right optimization function for this problem, the price of a Bermudan option for a given parametric set of exercise boundaries, and the associated first-order conditions (FOCs), from which we obtain the optimal exercise policy (for this family). Our approach is similar to Merton's (1973) pricing of perpetual American options. In a multifactor setting, however, the optimal exercise boundary functional form is unknown and, hence, is approximated by a family (whereas it is constant in Merton's perpetual case). In both problems, the optimal exercise rule is obtained from the FOCs.[1]

We maximize the Bermudan price at each exercise period with regard to a given family of exercise boundaries, which depend on a set of basis functions (the regressors) of the state variables. The FOCs imply value-matching errors, the difference between continuation and intrinsic values, are orthogonal to the regressors precisely when they are localized at the exercise boundary. These FOCs characterize the optimal boundary (for a given family) because value-matching errors are zero at the optimal exercise boundary by definition. The FOCs, however, depend on the initial state; therefore, not all points of the optimal boundary are (statistically) equally important, implying a simple basis is sufficient to provide a good local approximation. This paper focuses on the FOCs corresponding to a linear approximation (with polynomial basis); the regressors become more complex for nonlinear functions (e.g., neural networks).

Further, the FOCs are given in two different ways: by parameterizing, in the intrinsic value, either the continuation value or the exercise boundary. In the former case, the parametric continuation value replaces the intrinsic value, whereas, in the latter case, the intrinsic value depends explicitly on the parametric exercise boundary. The former case estimates the continuation value along the exercise boundary; that is, it characterizes the exercise boundary in an implicit way (by using the value-matching condition). The latter case explicitly estimates this boundary. That is, in the value-matching error, which is given by the difference between continuation and intrinsic values, the continuation value is the dependent variable and the intrinsic value contains one of the two parametric regression functions.

The FOCs of the Bermudan price-maximizing problem lead to an iterative procedure, because we estimate and localize the exercise boundary at the same time, like a root-finding problem. That is, we run local least-squares regressions by using the discounted realized payoffs of the simulated paths that are close to the exercise boundary identified in the previous iteration. When parameterizing the continuation value (the exercise boundary), a kernel down-weights large value-matching errors (the distance to the boundary), focusing the estimation on points around the implicit (explicit) exercise boundary. Once solved, the FOCs do provide an exercise rule, which also allows us to update all sample-path realized payoffs and proceed to the previous exercise period in a recursive way (like the Longstaff–Schwartz method).

We call this algorithm the "local LSM" method. Localizing the continuation value at the exercise boundary is the key difference with respect to the LSM method, but is fundamental to produce optimal exercise decisions. By taking the recursive solution of the previous period in every initial iterative step, a couple of iterations are sufficient for this localization if we use the continuation value. In computation terms, the extra time effort over the LSM method depends on the number of iterations (if > 1), because local and global regressions are alike.

We simulate paths starting before time 0 from the same in-the-money point. The latter feature implies a richer price dispersion (generating points near the exercise boundary at the first exercise dates), which is important for a local regression. In this way, the local LSM method saves time for pricing portfolios of Bermudans with different strike prices, because the initial simulation point and the derived exercise rules focusing on the boundary do not depend on the option moneyness.

In the numerical example, we price Bermudan options on the maximum of five stocks. This max-call option is a challenging security, having separated exercise regions and multiple (five) boundaries, and is a benchmark problem, for which lower and upper bounds are present (Broadie & Cao, 2008). We use a second-order polynomial basis of functions for the regressors.[2] The local LSM prices are always above the best available lower bounds, up to 20 to 30 basis points/bps in some cases with just one or two iterations (see Tables 6.1 to 6.3). Both algorithms, using either the explicit or the implicit boundary, produce similar good prices. We compute the associated dual upper bound (Andersen & Broadie, 2004), and the gap between lower and upper bounds is reduced to a few bps. This simultaneous near optimality of both bounds, which are based on the local exercise strategy derived in this paper, is addressed in a second paper, which provides additional numerical results (Ibáñez & Velasco, 2016).

We stress that improving upon the prices provided here by the local LSM method by enlarging the dimension of the basis from a quadratic fit (and just using the LSM method) is not easy. The variance of the local LSM prices is less than or equal to the variance in the LSM method case (which can be theoretically justified). Although we optimize the kernel's bandwidth in these simulations, a simple kernel, which is held constant across iterations and exercise periods, produces similar good results. These two features are relevant in practice.

We also study the convergence of the local LSM method. We assume an infinite-dimension linear approximation, a kernel that localizes the exercise boundary, one set of simulated paths, and standard technical assumptions on function smoothness and stopping-time identification. First, we provide rates of convergence of the exercise-boundary estimates, which depend on the number of paths, basis dimension, and kernel bandwidth.[3] Second, we adapt Stentoft's (2004b) results for the LSM method[4] to prove the convergence of the local LSM Bermudan prices.

In recent work, Belomestny (2011a, 2011b) and Belomestny, Dickmann, and Nagapetyan (2015) also use local regression and study the error and complexity of simulation-based algorithms. Zanger (2016) analyzes the speed of convergence of the LSM method using both linear and nonlinear approximations. Although we provide some estimates of the convergence rate for the exercise boundary, a thorough analysis of the error and complexity of our local LSM algorithm, along the lines of these papers, is left for future research.

The organization of the rest of the paper is as follows. Section 2 studies the cost of suboptimal exercise for Bermudan options. For the sake of simplicity, we consider a two-factor exchange or Margrabe option. Section 3 derives and analyzes the optimal FOCs. Section 4 studies the implementation and convergence of the local LSM method. Section 5 considers a multifactor setting and general payoffs. Section 6 provides numerical results of a Bermudan max-call option on five securities. Section 7 concludes. Proofs are left to the Appendix.

2 | THE COST OF SUBOPTIMAL EXERCISE

We address the maximization of Bermudan option prices by first deriving the cost of suboptimal exercise. For simplicity, we consider a two-factor problem, which contains all ingredients of a multifactor setting. We consider multiple state variables and general payoffs in Section 5. Let $(\Omega, \mathcal{F}, \mathcal{P}, \{\mathcal{F}_t\}_{0 \leq t \leq T})$ be an augmented filtered probability space (Duffie, 2001), where $S(t)$ is the vector of state variables adapted to the filtration $\mathcal{F}_t$, and Q denotes the unique risk-neutral measure that is taken as given.

A Bermudan-style option can be exercised a finite number of periods $t_j > t_0$, where $t_j \in \{t_1, t_2, \ldots, t_J\}$, t_0 is the initial time, and $t_J = T$ is maturity. Consider a spread or Margrabe option, with intrinsic value

$$I(S_1, S_2) = \{S_2 - S_1 - K\}^+,$$

where $S(t) = (S_1(t), S_2(t))$ are two random prices and K is the strike price. We denote by $C(t, S_1, S_2)$ the true Bermudan spread option price (i.e., the true continuation value) at time t, and $f_j(S_1)$ denotes the true exercise boundary at time $t_j > t_0$. That is, $f_j(S_1)$ verifies the value-matching condition $C(t_j, S_1, f_j(S_1)) = I(S_1, f_j(S_1))$, and it is optimal to exercise if $S_2(t_j) \geq f_j(S_1(t_j))$. In particular, if S_1 (S_2) is constant, the spread option is a standard call (put) option.

Denote by $\widetilde{f}_j(S_1)$ a given exercise boundary. That is, we exercise if $S_2(t_j) \geq \widetilde{f}_j(S_1(t_j))$. Therefore, $\widetilde{f}_j(S_1) < f_j(S_1)$ ($\widetilde{f}_j(S_1) > f_j(S_1)$) implies "suboptimally" accelerating (delaying) exercise compared to the optimal stopping time of the Bermudan spread option. We assume that $\widetilde{f}$ implies the same number of exercise boundaries as f, which is one for the spread option. Let us denote by $\widetilde{f}_{j:J} = \{\widetilde{f}_j, \widetilde{f}_{j+1}, \ldots, \widetilde{f}_J\}$ (similar for $f_{j:J}$) the exercise boundaries for $t_j, t_{j+1}, \ldots, t_J$. Let $P(t_1, S, \widetilde{f}_{2:J})$ denote the price (i.e., the continuation value) of the Bermudan option at time t_1 for a given policy $\widetilde{f}_{2:J}$, and $P(t_0, S, \widetilde{f}_{1:J})$ the initial price for a policy $\widetilde{f}_{1:J} = (\widetilde{f}_1, \widetilde{f}_{2:J})$. Consistent with the previous definitions, it holds that

$$
\begin{aligned}
P\left(t_0, S(t_0), \left(\widetilde{f}_1, \widetilde{f}_{2:J}\right)\right) = E_0^Q \Big[D_1 \Big((S_2(t_1) - S_1(t_1) - K) \times 1_{\{S_2(t_1) \geq \widetilde{f}_1(S_1(t_1))\}} \\
+ P(t_1, S(t_1), \widetilde{f}_{2:J}) \times 1_{\{S_2(t_1) < \widetilde{f}_1(S_1(t_1))\}} \Big) \Big],
\end{aligned}
$$

where D_1 is the inverse of a bank account between t_0 and t_1 ($\mathcal{F}_1$-adapted), E_0^Q denotes the Q expectation conditional on $\mathcal{F}_0$, and Q is given by $Q(S_2(t_1), S_1(t_1)|S(t_0))$. To further simplify notation, prices are written in discounted terms (or zero interest rates, $D_1 = 1$).

The cost of suboptimal exercise

The cost of suboptimal exercise is nonnegative and given by

$$
C(t_0, S) - P\left(t_0, S, \widetilde{f}_{1:J}\right) \geq 0.
$$

We denote by $\Delta\widetilde{f}_1(S_1(t_1)) = \widetilde{f}_1(S_1(t_1)) - f_1(S_1(t_1))$ the exercise boundary error. First, assume $\Delta\widetilde{f}(S_1) \leq 0$, which implies suboptimally accelerating exercise at t_1. The associated cost is given by

$$
\begin{aligned}
E_0^Q \Bigg[& \int_{\widetilde{f}_1(S_1(t_1))}^{f_1(S_1(t_1))} (C(t_1, S_1(t_1), S_2(t_1)) - (S_2(t_1) - S_1(t_1) - K)) \times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \\
& + \int_0^{\widetilde{f}(S_1(t_1))} \left(C(t_1, S_1(t_1), S_2(t_1)) - P\left(t_1, (S_1(t_1), S_2(t_1)), \widetilde{f}_{2:J}\right)\right) \\
& \times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \Bigg],
\end{aligned}
\tag{2.1}
$$

where, by the definition of conditional probability,

$$
Q(S_1(t_1), S_2(t_1)|S(t_0)) = Q(S_2(t_1)|S_1(t_1), S(t_0)) \times Q(S_1(t_1)|S(t_0)),
$$

and the expectation $E_0^Q[.]$ depends only on $Q(S_1(t_1)|S(t_0))$.[5] Equation (2.1) includes two costs. For any value of S_1, the first integral is the cost associated with exercise at t_1, when S_2 is in the suboptimal exercise region (i.e., $S_2 \in [\widetilde{f}(S_1), f(S_1)]$). The second integral is the continuation cost associated with suboptimal exercise in the future $(t_2, \ldots, t_J)$, when S_2 is outside of the exercise region ($S_2 \in [0, \widetilde{f}(S_1))$).

Second, for policies that suboptimally delay exercise at t_1 $(\Delta \widetilde{f}_1(S_1) \geq 0)$, the cost is given by

$$
\begin{aligned}
E_0^Q \Bigg[&\int_{f_1(S_1(t_1))}^{\widetilde{f}_1(S_1(t_1))} ((S_2(t_1) - S_1(t_1) - K) - C(t_1, S_1(t_1), S_2(t_1))) \times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \\
&+ \int_0^{\widetilde{f}(S_1(t_1))} (C(t_1, S_1(t_1), S_2(t_1)) - P(t_1, (S_1(t_1), S_2(t_1)), \widetilde{f}_{2:J})) \\
&\times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \Bigg],
\end{aligned}
\tag{2.2}
$$

which also includes two costs associated with the rights to exercise at and after t_1, respectively.

The first integral in equations (2.1) and (2.2) is the same, irrespective of the sign of $\Delta \widetilde{f}_1(S_1)$, and is nonnegative. Hence, equations (2.1) and (2.2) are the same. Let us define by $G(\widetilde{f}_{1:J}) = G(S_1(t_1), \widetilde{f}_{1:J})$ the argument within the two expectations in (2.1) and (2.2), that is,

$$
\begin{aligned}
G\left(\widetilde{f}_{1:J}\right) = &\int_{f_1(S_1(t_1))}^{\widetilde{f}_1(S_1(t_1))} ((S_2(t_1) - S_1(t_1) - K) - C(t_1, S_1(t_1), S_2(t_1))) \\
&\times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \\
&+ \int_0^{\widetilde{f}(S_1(t_1))} \left(C(t_1, S_1(t_1), S_2(t_1)) - P\left(t_1, (S_1(t_1), S_2(t_1)), \widetilde{f}_{2:J}\right) \right) \\
&\times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1).
\end{aligned}
$$

The following proposition gives the cost of suboptimal exercise for Bermudan options.[6]

Proposition 2.1. *Consider a Bermudan spread option that can be exercised at times $t_1, t_2, \ldots, t_J$. Let $f_1(S_1)$ $(\widetilde{f}_1(S_1))$ be the true (a given) exercise boundary at t_1, respectively. And let $C(t_1, S_1, S_2)$ $(P(t_1, (S_1, S_2), \widetilde{f}_{2:J}))$ be the Bermudan's price or continuation value for the true (a given $\widetilde{f}_{2:J}$) policy from t_2 to t_J. Then the cost of suboptimal exercise is given by*

$$
C(t_0, S) - P\left(t_0, S, \widetilde{f}_{1:J}\right) = E_0^Q \left[G\left(\widetilde{f}_{1:J}\right) \right].
\tag{2.3}
$$

If $S_1(t_1)$ is constant, for a call option, $E_0^Q[G(\widetilde{f}_{1:J})] = G(\widetilde{f}_{1:J})$. And if $S_2(t_1)$ is constant, for a put option, we rewrite f_1 and $\widetilde{f}_1$ as functions of $S_2(t_1)$ (instead of $S_1(t_1)$).

Proof. It follows from equations (2.1) and (2.2), which, as discussed above, are the same. $\qquad\square$

This result is valid for a general exercise boundary function $\widetilde{f}$. For instance, $\widetilde{f}$ can depend on a linear basis (polynomials) as well as a nonlinear basis (exponential or neural networks).

3 | MAXIMIZING THE BERMUDAN OPTION PRICE

The function $E_0^Q[G(\widetilde{f}_{1:J})]$ gives the cost of suboptimal exercise for a policy $\widetilde{f}_{1:J}$. Because minimizing the cost of suboptimal exercise is equivalent to maximizing the Bermudan price, this cost is the right objective function to determine the optimal exercise policy.

Equation (2.3) can be rewritten as follows:

$$P\left(t_0, S, \widetilde{f}_{1:J}\right) = C(t_0, S) - E_0^Q\left[G\left(\widetilde{f}_{1:J}\right)\right].$$

Then, for any given family of exercise boundaries F, it holds that

$$\max_{\widetilde{f}_{1:J}\in F} P\left(t_0, S, \widetilde{f}_{1:J}\right) = C(t_0, S) - \min_{\widetilde{f}_{1:J}\in F} E_0^Q\left[G\left(\widetilde{f}_{1:J}\right)\right].$$

Hence, we minimize $E_0^Q[G(\widetilde{f}_{1:J})]$ subject to $\widetilde{f}_{1:J} \in F$.

Optimizing among exercise boundaries is Merton's (1973) approach for pricing perpetual American options (where this price is given in closed-form solution). For a given risk-neutral measure Q, the result does not depend on the objective probability measure. The richer F is, the lower the cost. That is, consider two families F_a and F_b, where $F_a \subseteq F_b$,

$$\min_{\widetilde{f}_{1:J}\in F_a} E_0^Q\left[G\left(\widetilde{f}_{1:J}\right)\right] \geq \min_{\widetilde{f}_{1:J}\in F_b} E_0^Q\left[G\left(\widetilde{f}_{1:J}\right)\right] \geq 0.$$

And F can be chosen according to the dimensionality of the problem in practice.

For tractability, this minimization can be solved in a recursive way, where $\widetilde{f}_{1:J} = (\widetilde{f}_1, \widetilde{f}_{2:J})$. First solve the exercise boundaries in $\widetilde{f}_{2:J} = \{\widetilde{f}_2, \widetilde{f}_3, \ldots, \widetilde{f}_J\}$, upon which the price function $P(t_1, (S_1, S_2), \widetilde{f}_{2:J})$ depends; then solve $\widetilde{f}_1(S_1)$ at t_1. This way we do not include the variable $\widetilde{f}_{2:J}$ in the function G below. We discuss this recursive approach and consider stochastic interest rates at the end of this section.

Our exposition depends only on two consecutive exercise periods, which we fix at t_0 and t_1. Accordingly, we eliminate the dependence on time t_1 below. That is, $f(S_1), \widetilde{f}(S_1), C(S_1, S_2)$, and $P((S_1, S_2), \widetilde{f}_{2:J})$ denote $f_1(S_1), \widetilde{f}_1(S_1), C(t_1, S_1, S_2)$, and $P(t_1, (S_1, S_2), \widetilde{f}_{2:J})$, respectively. This notation shortens all the expressions.

3.1 | The first-order conditions: Small (and highly probable) value-matching errors are orthogonal to the regressors

Let us denote by $\widetilde{f} = \widetilde{f}(S_1, b)$ a parameterized exercise boundary (at t_1), where b is a vector of parameters. For example, $\widetilde{f}(S_1, b)$ can be an n-order polynomial in S_1, and $b \in \mathcal{R}^{n+1}$ are the polynomial coefficients.[7] Assume the exercise boundary $\widetilde{f}_{2:J}^*$ is taken as given.

Then consider

$$\min_b E_0^Q \left[G(\widetilde{f}(S_1, b)) \right],$$ (3.1)

with FOCs given by the product of two terms,

$$E_0^Q \left[G_b \left(\widetilde{f}(S_1, b^*) \right) \right] = E_0^Q \left[\frac{dG\left(\widetilde{f}^*\right)}{d\widetilde{f}} \times \frac{\partial \widetilde{f}(S_1, b^*)}{\partial b} \right] = 0,$$ (3.2)

where we assume $\widetilde{f}^* = \widetilde{f}(S_1, b^*)$ is the optimal solution of equation (3.1).

Equation (3.2) is given by

$$E_0^Q \left[\left(C\left(S_1, \widetilde{f}^*(S_1)\right) - \left(\widetilde{f}^*(S_1) - S_1 - K \right) \right) \times \frac{\partial \widetilde{f}(S_1, b^*)}{\partial b} \times Q\left(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0) \right) \right] = \Sigma,$$ (3.3)

where the expectation depends only on $Q(S_1(t_1)|S(t_0))$, and

$$\Sigma = E_0^Q \left[\left(C\left(S_1, \widetilde{f}^*(S_1)\right) - P\left(\left(S_1, \widetilde{f}^*(S_1)\right), \widetilde{f}_{2:J}^* \right) \right) \times \frac{\partial \widetilde{f}(S_1, b^*)}{\partial b} \times Q\left(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0) \right) \right].$$

The vector Σ is the (recursive) error in the FOCs, given by the difference between the true and estimated continuation values at t_1, C, and P, respectively, where $C \geq P$ implies $\Sigma \geq 0$. The value of Σ depends on the error in the policy $\widetilde{f}_{2:J}^*$ (if $\widetilde{f}_{2:J}^* = f_{2:J}^*$, $C = P$, and $\Sigma = 0$), which is taken as given, but not on the error in $\widetilde{f}^*(S_1)$, which we want to determine.

From equation (3.3), the optimal (though not the "first best" if $\widetilde{f}^* \neq f$) exercise boundary is obtained when the value-matching errors, localized in the most significant part of the exercise boundary, are orthogonal to the regressors (up to an error Σ). That is, the value-matching errors are weighted by $Q(S_2(t_1) = \widetilde{f}^*(S_1(t_1)|S_1(t_1), S(t_0))$, which has two implications. First, these errors are precisely zero at the optimal boundary, $S_2(t_1) = f(S_1(t_1))$. Second, because Q depends on $S(t_0)$, not all states of the exercise boundary $(S_1(t_1), S_2 = \widetilde{f}^*(S_1(t_1)))$ are equally likely conditional on $S(t_0)$. Hence, at time t_1, equation (3.3) ensures that small and highly probable value-matching errors are orthogonal to the regressors.[8]

Proposition 3.1. *Consider the minimization of the cost of suboptimal exercise in equation (3.1). The FOCs are given by equation (3.3). In addition, if $\widetilde{f}(S_1, b)$ is linear in b, we have a "beta" type (implicitly defined) solution. Let us denote by $p(S_1) = \frac{\partial \widetilde{f}(S_1, b^*)}{\partial b}$ the vector of regressors associated with b, and $q(S_1) = Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0))$. Then,*

$$b^* = E_0^Q \left[p(S_1) \times p(S_1)' \times q(S_1) \right]^{-1} \times E_0^Q \left[\left(C(S_1, \widetilde{f}^*(S_1)) + S_1 + K \right) \right.$$

$$\left. \times p(S_1) \times q(S_1) - \Sigma \right].$$ (3.4)

Proof. It follows from equation (3.3) and because $\widetilde{f}(S_1, b)$ is linear in b. ☐

Remark 3.2. *b^* appears in both sides of equation (3.4) (as $\widetilde{f}^*(S_1) = \widetilde{f}(S_1, b^*)$), and hence this equation is solved iteratively. Propositions 2.1 and 3.1 (the cost of suboptimal exercise and the associated*

FOCs, respectively) are quite general; they depend only on knowing the number of exercise bound-
aries, which is one for the spread option. This assumption is easy to check and simplifies the exposi-
tion, but is not necessary, as shown in Section 5 (for a max-call option). We do not need to assume
the exercise boundary $\widetilde{f}^*(S_1)$ is explicitly defined. We can use a continuation value function, and
then $\widetilde{f}^*(S_1)$ is implicitly defined from the value-matching condition. This result is explained below
in Section 3.3 (in this case, some extrapolation errors can happen, but they are easily controlled
for).

Because the regression function $\widetilde{f}(S_1, b^*)$ solves the orthogonality conditions only for $S_2 = \widetilde{f}^*(S_1)$,
in the space $S_1 \times S_2$, we call this method a local least-squares method in Section 4. For other
payoffs, $\{S_2 - S_1 - K\}^+$ is replaced by the corresponding intrinsic value, $I(S_1, S_2)$, distinguish-
ing between put and call payoffs, which have complementary exercise regions. We summarize these
results.

Corollary 3.3. *Consider a parametric family of exercise boundaries, F. For example, F can depend*
on a linear basis (e.g., polynomials) or a nonlinear basis (neural networks). Assume (i) the exer-
cise boundary $\widetilde{f}^*_{2:J}$ is taken as given and (ii) the FOCs are sufficient for optimality. Then the policy
$\widetilde{f}^*(S_1) = \widetilde{f}(S_1, b^*)$, which solves equation (3.3), is the optimal exercise policy (for the family F).

Proof. $\widetilde{f}(S_1, b^*)$ minimizes the cost of suboptimal exercise, and equivalently (from equation (2.3))
maximizes the Bermudan price. □

Corollary 3.3 points to a recursive algorithm, as the boundary $\widetilde{f}^*_{2:J}$ is taken as given. An algorithm
that is based on equation (3.3) will be optimal in the sense that it maximizes Bermudan prices for a
given family F. Yet the exercise boundary $\widetilde{f}^*_{1:J}$ yields a price that is a lower bound of the true price,
when a second independent simulation evaluates this policy (as we do in our numerical exercise). In
Section 4, we show that an algorithm based on the sample equivalent of equation (3.3) converges (i.e.,
yields the true price in the limit) under appropriate conditions.

3.2 | The exercise boundary errors (adjusted by convexity) are orthogonal to the regressors

Let us provide more intuition on the FOCs, which say that "small" or "local" value-matching errors
are orthogonal to the regressors. Assuming the continuation value C is smooth, a second-order Taylor
approximation (at $S_2 = f(S_1)$) implies, with error $O(\Delta \widetilde{f}(S_1)^3)$,

$$C\left(S_1, \widetilde{f}(S_1)\right) \approx C(S_1, f(S_1)) + C_2 \Delta \widetilde{f}(S_1) + \frac{1}{2} C_{22} \Delta \widetilde{f}(S_1)^2,$$

where C_2 and C_{22} denote the partial derivatives with regard to S_2.[9]
Therefore, from $C(S_1, f(S_1)) = f(S_1) - S_1 - K$ and $\widetilde{f}(S_1) = f(S_1) + \Delta \widetilde{f}(S_1)$,

$$C(S_1, \widetilde{f}(S_1)) - (\widetilde{f}(S_1) - S_1 - K) \approx -(1 - C_2) \Delta \widetilde{f}(S_1) + \frac{1}{2} C_{22} \Delta \widetilde{f}(S_1)^2.$$

Equation (3.3) simplifies to

$$E_0^Q \left[\left((1 - C_2) \Delta \widetilde{f}^*(S_1) - \frac{1}{2} C_{22} \Delta \widetilde{f}^*(S_1)^2 \right) \times \frac{\partial \widetilde{f}(S_1, b^*)}{\partial b} \times Q(S_2 = \widetilde{f}^*(S_1) | S_1, S(t_0)) \right] \approx \Sigma,$$

and, equivalently (if $C_2 < 1$),[10]

$$E_0^Q\left[\left(\Delta\widetilde{f}^*(S_1) - \frac{1}{2}\frac{C_{22}}{1-C_2}\Delta\widetilde{f}^*(S_1)^2\right)(1-C_2)\times\frac{\partial\widetilde{f}(S_1,b^*)}{\partial b}\right.$$

$$\left.\times Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0))\right] \approx \Sigma. \tag{3.5}$$

Redefine the regressors as $(1-C_2)\times\frac{\partial\widetilde{f}(S_1,b^*)}{\partial b}$. A second result follows. The errors of the optimal exercise boundary, $\Delta\widetilde{f}^*$, are orthogonal to the regressors (adjusted by convexity, $C_{22} > 0$). In brief, the FOCs estimate the most significant part of the optimal exercise boundary (adjusted by convexity), which, intuitively, should provide the best solution/stopping time.

For a constant regressor (i.e., $\frac{\partial\widetilde{f}(S_1,b^*)}{\partial b_{1:1}} = 1$), we further have

$$E_0^Q\left[\left(\Delta\widetilde{f}^*(S_1) - \frac{1}{2}\frac{C_{22}}{1-C_2}\Delta\widetilde{f}^*(S_1)^2\right)\times\left(1-C_2\right)\times Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0))\right] \approx \Sigma_{1:1}.$$

The minus convexity adjustment implies that the optimal boundary $\widetilde{f}^*$ will be upper biased (because $\Sigma_{1:1} \geq 0$); that is, the average of $\Delta\widetilde{f}^*(S_1)$ is nonnegative. The strategy associated with $\widetilde{f}^*$ is slightly biased toward delaying exercise (compared to the first-best f), which, from convexity, is less costly than advancing exercise.

3.3 | Two models for the exercise boundary

We can explicitly or implicitly parameterize the exercise boundary, $\widetilde{f}(S_1, b)$. That is, in the value-matching error (i.e., $C - I$), the intrinsic value I depends on the exercise boundary, and C is approximated by an estimator.

To explicitly use the exercise boundary
Consider, for example, a second-order polynomial function in S_1,

$$\widetilde{f}(S_1, b) = \sum_{j=0}^{2} b_j S_1^j, \tag{3.6}$$

where $b = (b_0, b_1, b_2)'$, implying

$$\frac{\partial\widetilde{f}(S_1,b)}{\partial b} = \left(\frac{\partial\widetilde{f}(S_1,b)}{\partial b_0}, \frac{\partial\widetilde{f}(S_1,b)}{\partial b_1}, \frac{\partial\widetilde{f}(S_1,b)}{\partial b_2}\right)' = \left(1, S_1, S_1^2\right)',$$

and we substitute these regressors in equation (3.3).

To use a continuation value function
Similar to the LSM method, we can use a new continuation value function $\widetilde{C}(S_1, S_2, b)$, which is used to make exercise decisions. Note that $\widetilde{C}(S_1, S_2, b)$ and the above $P((S_1, S_2), \widetilde{f}_{2:J})$ are two different

functions (e.g., $\widetilde{C}(b) > C$ implies suboptimally delaying exercise at t_1). Then the exercise boundary is implicitly defined from the value-matching condition,

$$\widetilde{C}(S_1, \widetilde{f}(S_1, b), b) = \widetilde{f}(S_1, b) - S_1 - K. \tag{3.7}$$

The regressors (i.e., the term "$\partial \widetilde{f} / \partial b$") are defined from the implicit function theorem; that is,

$$\frac{\partial \widetilde{C}}{\partial \widetilde{f}} \times \frac{\partial \widetilde{f}}{\partial b} + \frac{\partial \widetilde{C}}{\partial b} = \frac{\partial \widetilde{f}}{\partial b} \Rightarrow \frac{\partial \widetilde{f}}{\partial b} = \left(1 - \frac{\partial \widetilde{C}}{\partial \widetilde{f}}\right)^{-1} \times \frac{\partial \widetilde{C}}{\partial b}.$$

The FOCs change as follows,[11] where $\widetilde{C}_2 = \frac{\partial \widetilde{C}}{\partial \widetilde{f}}$:

$$E_0^Q \left[\frac{C(S_1, \widetilde{f}^*(S_1)) - \widetilde{C}(S_1, \widetilde{f}^*(S_1), b^*)}{\widetilde{C}_2(S_1, \widetilde{f}^*(S_1), b^*) - 1} \times \frac{\partial \widetilde{C}(S_1, \widetilde{f}^*(S_1), b^*)}{\partial b} \right.$$

$$\left. \times Q(\widetilde{C}(S_1, S_2, b^*) = S_2 - S_1 - K | S_1, S(t_0)) \right] = 0. \tag{3.8}$$

Note that $\widetilde{f}^*(S_1) = \widetilde{f}(S_1, b^*)$ is never explicitly computed in this case, but is implicitly defined from equation (3.7).

For example, assume $\widetilde{C}(S_1, S_2, b)$ is also a second-order polynomial in S; that is,

$$\widetilde{C}(S_1, S_2, b) = b_{00} + \sum_{j=1}^{2} b_{0j} S_j + \sum_{j=1}^{2} \sum_{n=j}^{2} b_{jn} S_j S_n, \tag{3.9}$$

where $b = \{b_{00}, b_{01}, b_{02}, b_{11}, b_{12}, b_{22}\}'$. The partial derivatives are given by

$$\widetilde{C}_2 = \frac{\partial \widetilde{C}}{\partial \widetilde{f}} = b_{02} + b_{12} S_1 + 2b_{02} S_2 \text{ and}$$

$$\frac{\partial \widetilde{C}}{\partial b} = \left(\frac{\partial \widetilde{C}}{\partial b_{00}}, \frac{\partial \widetilde{C}}{\partial b_{01}}, \frac{\partial \widetilde{C}}{\partial b_{02}}, \frac{\partial \widetilde{C}}{\partial b_{11}}, \frac{\partial \widetilde{C}}{\partial b_{12}}, \frac{\partial \widetilde{C}}{\partial b_{22}}\right)' = \left(1, S_1, S_2, S_1^2, S_1 S_2, S_2^2\right)'.$$

We make two remarks. (i) Equation (3.5) simplifies to

$$E_0^Q \left[\left(\Delta f^*(S_1) - \frac{1}{2} \frac{C_{22}}{1 - C_2} \Delta f^*(S_1)^2\right) \frac{\partial \widetilde{C}(S_1, f^*(S_1), b^*)}{\partial b} Q(S_1, \widetilde{C}(S_1, S_2, b^*)) \right.$$

$$\left. = I(S) | S_1, S(t_0)) \right] \approx \Sigma,$$

where $I(S) = S_2 - S_1 - K$, and provides the same intuition. (ii) From equation (3.7), if $\widetilde{C}(S_1, S_2, b)$ is quadratic in S, the pair "S_1 and $S_2 = \widetilde{f}(S_1, b)$" is a hyperbole. Equations (3.7) and (3.9) contain the quadratic equation (3.6), and they should provide a better fit and, hence, a lower cost.

A continuation value function that implies the same exercise boundary

For a given exercise boundary, $\widetilde{f}(S_1)$, a continuation value function that implies the same exercise policy always exists. That is, for any given $0 < \delta < 1$, if we define

$$\widetilde{C}(S_1, S_2, b) = (S_2 - S_1 - K) - \delta\left(S_2 - \widetilde{f}(S_1, b)\right),$$

then

$$\text{if } S_2 = \widetilde{f}(S_1, b), \widetilde{C}(S_1, S_2, b) = S_2 - S_1 - K,$$

$$\text{and if } S_2 \lessgtr \widetilde{f}(S_1, b), \widetilde{C}(S_1, S_2, b) \gtrless S_2 - S_1 - K.$$

Therefore, "$S_2 = \widetilde{f}(S_1, b)$" is also the implicit exercise boundary in $\widetilde{C}$. Now the partial derivatives are given by

$$\widetilde{C}_2 = 1 - \delta \quad \text{and} \quad \frac{\partial \widetilde{C}}{\partial b} = \frac{\partial \widetilde{f}(S_1, b)}{\partial b} \delta,$$

which implies the FOCs in equations (3.3) and (3.8) are the same.

3.4 | The dependent variable $\widehat{C}$

We denote by $\widehat{C}$ an estimator of the Bermudan price C. As in any least-squares problem, it is convenient that the dependent variable, $\widehat{C} = \widehat{C}(t_1, S_1, S_2)$, is an efficient and unbiased estimator (e.g., if feasible, one can include control variate techniques to estimate $\widehat{C}$).

Discounting realized payoffs
Like the LSM method, we define the following, where D_j is the discount factor between t_0 and t_j:

$$\widehat{C}(t_1, S_1, S_2) = \frac{D_j}{D_1} \{ S_2(\tau_j) - S_1(\tau_j) - K \}^+, \text{ where } S(t_i) \sim Q(S(t_i)|S(t_{i-1})), \tag{3.10}$$

and the stopping time $\tau_j \in \{t_2, t_3, \ldots, t_J\}$ is defined recursively as follows. If we estimate the exercise boundary, $\widetilde{f}^*$,

$$\tau_j = t_j, \quad \text{if } S_2(t_j) \geq \widetilde{f}^*(t_j, S_1(t_j)); \quad \text{and} \quad \tau_j = \tau_{j+1}, \text{ otherwise.} \tag{3.11}$$

And if we estimate the continuation value, $\widetilde{C}^*$, τ_j is reciprocally defined as

$$\tau_j = t_j, \quad \text{if } \{ S_2(t_j) - S_1(t_j) - K \}^+ \geq \widetilde{C}^*(t_j, S_1(t_j), S_2(t_j)); \quad \text{and} \quad \tau_j = \tau_{j+1}, \text{ otherwise.} \tag{3.12}$$

Discounting a continuation value versus discounting realized payoffs
An alternative to the previous definition of $\widehat{C}$, which uses the discounted realized payoff and the stopping time $\tau_j \in \{t_2, t_3, \ldots, t_J\}$, would be to discount an estimated continuation value at time t_2, $\widetilde{C}(t_2, S(t_2), b)$. Then

$$\widehat{C}(t_1, S_1, S_2) = \frac{D_2}{D_1} \max \left\{ \widetilde{C}(t_2, S_1(t_2), S_2(t_2), b), \{ S_2(t_2) - S_1(t_2) - K \}^+ \right\}, \tag{3.13}$$

where $S(t_2) \sim Q(S(t_2)|S(t_1))$.

The literature stresses that this recursion yields (upper) biased estimators. Note that in our local approach, we focus on the boundary $\widetilde{f}(t_2, S_1, b^*)$ or a localized continuation value $\widetilde{C}(t_2, S_1, S_2 \approx \widetilde{f}(S_1, b^*), b^*)$. Therefore, the recursion based on (3.13) is less intuitive in our setting.

3.5 | A recursive approach and stochastic interest rates

We discuss the conditions under which the boundary $f(t_j, S_1)$, $t_j \in \{t_1, t_2, \ldots, t_{J-1}\}$, can be derived in a recursive way. We also extend Proposition 3.1 to the case of stochastic interest rates, which is important for pricing Bermudan swaptions and hybrid securities. We need to introduce more notation and write out the corresponding FOCs.

Let r_t be the instantaneous short interest rate, and let $D_j = e^{-\int_{t_0}^{t_j} r_u du}$ denote the inverse of a bank account associated with the Q risk-neutral measure (Duffie, 2001). Assume r follows a one-factor Markovian model. In this case, the exercise boundary depends on the short rate too, which is denoted by $f_j = f_j(S_1(t_j), r_j)$. Now the term of the cost of suboptimal exercise that depends on $\widetilde{f}_j = \widetilde{f}_j(S_1(t_j), r_j)$ is given by the following (see Appendix B):

$$
E_0^Q \left[\left(\prod_{i=1}^{j-1} 1_{\{0 \leq S_2 \leq \widetilde{f}_i(S_1, r)\}} \right) \times D_{j-1} \right.
$$

$$
\times E_{j-1}^Q \left[\int_{\widetilde{f}_j}^{f_j} \frac{D_j}{D_{j-1}} (C(t_j, S, r_j) - (S_2 - S_1 - K)) \times Q(S_2(t_j) | (S_1(t_j), r_j), (S(t_{j-1}), r_{j-1})) dS_2(t_j) \right.
$$

$$
\left. \left. + \int_0^{\widetilde{f}_j} \frac{D_j}{D_{j-1}} \left(C(t_j, S, r_j) - P(t_j, S, r, \widetilde{f}_{j+1:J}) \right) \times Q(S_2(t_j) | (S_1(t_j), r_j), (S(t_{j-1}), r_{j-1})) dS_2(t_j) \right] \right].
$$

For $j = 1$, we obtain equation (2.1) in Section 2.1 (where $D_0 = 1$, $\Pi_{i=1}^0 = 1$). The associated FOCs are given by

$$
E_0^Q \left[\prod_{i=1}^{j-1} 1_{\{0 \leq S_2 \leq \widetilde{f}_i(S_1, r)\}} \times D_j \left(C \left(t_j, S_1, \widetilde{f}_j^*(S_1, r), r \right) - \left(\widetilde{f}_j^*(S_1, r) - S_1 - K \right) \right) \right.
$$

$$
\left. \times \frac{\partial \widetilde{f}_j(S_1, r, b^*)}{\partial b} \times Q \left(S_2 = \widetilde{f}_j^*(S_1, r) | (S_1, r), (S(t_0), r(t_0)) \right) \right] = \Sigma,
$$

where this expectation and $\widetilde{f}_j^*$ depend on the previous boundaries $\widetilde{f}_i(S_1)$, $1 \leq i \leq j - 1$.

This condition holds only for those paths such that the option is not exercised before t_j (i.e., $\Pi_{i=1}^{j-1} 1_{\{0 \leq S_2 \leq \widetilde{f}_i(S_1, r)\}} = 1$). If we omit this term, the FOCs associated with $\widetilde{f}_j^*$ are equivalent to a Bermudan option that can be exercised only from t_j to t_J. In practice, we consider this recursive approach (similar to equation (3.3)) because it is intuitive and tractable:

$$
E_0^Q \left[D_j \left(C \left(t_j, S_1, \widetilde{f}_j^*(S_1, r), r \right) - \left(\widetilde{f}_j^*(S_1, r) - S_1 - K \right) \right) \right.
$$

$$
\left. \times \frac{\partial \widetilde{f}_j(S_1, r, b^*)}{\partial b} \times Q \left(S_2 = \widetilde{f}_j^*(S_1, r) | (S_1, r), (S(t_0), r(t_0)) \right) \right] = \Sigma.
$$

The discount factor D_j weights the value-matching errors. If r is constant, the discount factor cancels. In the case of fixed-income securities, using a different numeraire and pricing measure can be more convenient. In the rest of the paper, we consider the initial FOCs in equation (3.3).

3.6 │ The initial simulation of paths

The orthogonality conditions in equation (3.3) hold only for

$$\left(S_1(t_1), \widetilde{f}^*(S_1(t_1)) \right) \in S_1 \times S_2,$$

conditional on $S(t_0)$, which defines the exercise boundary. For simplicity, assume S_1 is constant; that is, the spread option is a call option. Optimally, all samples of $S_2(t_1)$ are just $f(S_1)$, which is the optimal exercise price and independent of $S_2(t_0)$. For example, because $f(S_1)$ is unknown, but $f(S_1) > S_1 + K$, for estimating f we sample better from $Q(S_2(t_1)|S_1 + K)$ than from $Q(S_2(t_1)|S_2(t_0))$, where the price $S_2(t_0)$ could be deep away-from-the-money. This condition is relevant and necessary for American-style options, which are approximated by a Bermudan option, because little price dispersion exists in the first exercise dates (and few paths go through the exercise boundary).[12]

Hence, in all numerical examples, we start (and recommend) simulating from an in-the-money point. Once we have computed the exercise boundaries, we use a second simulation to price the Bermudan option, now sampling from $Q(S(t_1)|S(t_0))$.

4 │ THE LOCAL LSM ALGORITHM

The results derived in Section 3 are novel optimality properties of Bermudan options. In this section, we implement these results by using least-squares and simulation, where the recursive updating of the path's realized payoffs is similar to the LSM method.

To solve equation (3.3) (or, equivalently, (3.4)), the exercise boundary $\widetilde{f}^*$ is unknown and must be determined at the same time. Therefore, we proceed iteratively. Our algorithm solves, in every iteration, for the optimal $\widetilde{f}_i(S_1) = \widetilde{f}(t_1, S_1, b_i)$ with the previous $\widetilde{f}_{i-1}(S_1)$ fixed, where i is the iteration number. The conditional density $Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0))$ is replaced by a kernel evaluating the distance $S_2 - \widetilde{f}^*(S_1)$ with a bandwidth converging to zero with the number of paths, so the exercise boundary is identified asymptotically (implicitly or explicitly). Below, the objective function is written as a least-squares problem, because it yields the same FOCs. We call this method the local least-squares MC (i.e., local LSM) method.

4.1 │ The explicit exercise boundary

Consider the local least-squares problem:

$$\min_{\widetilde{f}_i} E_0^Q \left[\left(\hat{C}\left(S_1, S_2 \right) + S_1 + K - \widetilde{f}_i(S_1) \right)^2 \times \ker_Q \left(S_2 \approx \widetilde{f}_{i-1}(S_1)|S_1, S(t_0) \right) \right], \qquad (4.1)$$

where $\hat{C}$ (see equations (3.10) and (3.11)) is an estimator of C and $\ker_Q$ is a kernel that localizes S_2 in the previous estimator of $\widetilde{f}_{i-1}$. This problem implies a local least-squares regression, where $\hat{C} + S_1 + K$ is the dependent variable and $\widetilde{f}_i(S_1)$ is the regression function. In this iterative process, the stop criterion for $\widetilde{f}^*$ is given by $\widetilde{f}_i - \widetilde{f}_{i-1} \approx 0$ (whereas in practice, we use a small fixed number of iterations).

The FOCs are given by

$$E_0^Q \left[\left(\hat{C}(S_1, S_2) - (\widetilde{f}_i(S_1) - S_1 - K) \right) \times \frac{\partial \widetilde{f}_i(S_1)}{\partial b} \times \ker_Q(S_2 \approx \widetilde{f}_{i-1}(S_1)|S_1, S(t_0)) \right] = 0, \quad (4.2)$$

where we take $\widetilde{f}_{i-1}$ as fixed. An initial (zero) iteration of the period t_{J-1} is a global regression, $\ker_Q = 1$, where we estimate a continuation value (which is used in the kernel of the next iteration, or, alternatively, we take $\widetilde{f}_0(S_1) = S_1 + K$). The rest of the iterations are local regressions (to estimate $\widetilde{f}_i(S_1), i \geq 1$). For any other time $t_j < t_{J-1}$, the initial boundary is better taken from the solution of the previous step t_{j+1}. Note how the expectation in equation (4.2) depends on S_2, in contrast to equation (3.3), because we use a smoothing kernel.

Given a sample of M prices $(S_1^{(m)}, S_2^{(m)})$ distributed according to $Q(S_1, S_2 | S(t_0))$, the sample equivalent of equation (4.2) is given by

$$\frac{1}{M} \sum_{m=1}^{M} \left(\widehat{C}\left(S_1^{(m)}, S_2^{(m)}\right) - \left(\widetilde{f}_i\left(S_1^{(m)}\right) - S_1^{(m)} - K\right) \right) \times \frac{\partial \widetilde{f}_i(S_1^{(m)})}{\partial b} \times \mathcal{K}_h\left(S_2^{(m)} - \widetilde{f}_{i-1}\left(S_1^{(m)}\right)\right) = 0,$$

where $\mathcal{K}_h(x) = \mathcal{K}(x/h)/h$ focuses on paths for which $x \approx 0$ (that is, $S_2 \approx \widetilde{f}_{i-1}(S_1)$). $\mathcal{K}$ is a kernel function integrating to 1 and h is a bandwidth that converges to zero with M. We explain the kernel and its bandwidth in Section 6.

- **The local LSM algorithm** Let $ni \geq 1$ be a fixed number of iterations. Because it is a recursive approach, we specify the final period $t = T$, and recursively solve the exercise boundary for $T - 1$, $T - 2$, until $t = 1$. The intrinsic value is denoted by $I_t = \{S_{2,t} - S_{1,t} - K\}^+, t = \{1, 2, \ldots, T\}$. And $\widetilde{f}_t^n$ is the exercise boundary at time t and iteration n.

Consider a set of simulated paths, $\varpi \in \Omega$.

0. Set $t = T$. Define $y_{T+1} = 0$ and $\widetilde{f}_T^* = S_{1,T} + K$.

1. UPDATING PATHS, $\varpi \in \Omega$

$$y_t = \begin{cases} I_t, & \text{if } S_{2,t} \geq \widetilde{f}_t^* \\ e^{-r\Delta t} \times y_{t+1}, & \text{otherwise.} \end{cases}$$

Set $t = t - 1$.

2. The new EXERCISE BOUNDARY

Set $n = 1$ and $\widetilde{f}_t^0 = \widetilde{f}_{t+1}^*$.

2.1. LOCALIZING THE EXERCISE BOUNDARY

Provide a kernel $\mathcal{K}_h$. Then,

$$\widetilde{f}_t^n = \arg\min_{f_t \in F} \sum_{\varpi \in \Omega} \left(f_t(x_t) - (S_{1,t} + K) - \frac{y_{t+1}}{e^{r\Delta t}} \right)^2 \times \mathcal{K}_h\left(\widetilde{f}_t^{n-1}(x_t) - S_{2,t}\right).$$

Set $n = n + 1$. Go back to step **2.1** until $n = ni$.
Set $\widetilde{f}_t^* = \widetilde{f}_t^{ni}$. Go back to step **1** until $t = 1$.
End of the local LSM algorithm
Note that at $t = 1$, we have already estimated all exercise boundaries; $\widetilde{f}_{T-1}^*, \widetilde{f}_{T-2}^*, \ldots, \widetilde{f}_1^*$.

Remark 4.1.

(1) $\frac{y_{t+1}}{e^{r\Delta t}} + (S_{1,t} + K)$ *is the dependent variable in the regression, where, in our case, the boundary* $f(x) \in F$ *is a simple polynomial and x are the regressors.* y_t *is the realized payoff between t and* T *(following the* $\widetilde{f}^*_{t:T}$ *stopping-time rule) for any path,* $\varpi \in \Omega$.

(2) *With the function*

$$W_t = \mathcal{K}_h \left(\widetilde{f}^{n-1}_t(x_t) - S_{2,t} \right),$$

the local least-squares estimator is easily computed by ordinary least squares, by redefining x_t *and* $\frac{y_{t+1}}{e^{r\Delta t}} + (S_{1,t} + K)$ *as* $x_t \times \sqrt{W_t}$ *and* $(\frac{y_{t+1}}{e^{r\Delta t}} + (S_{1,t} + K)) \times \sqrt{W_t}$, *respectively.*

(3) *We use a Gaussian kernel,* $\mathcal{K}_h(\widetilde{f}^{n-1}_t(x_t) - S_{2,t})$, *where h is the bandwidth. Because we explicitly estimate the exercise boundary, no equivalent LSM method exists.*

4.2 | The continuation value

If we use a continuation value (where $\widehat{C}$ is given by equations (3.10) and (3.12)), instead of equations (4.1) and (4.2), we have, respectively,

$$\min_{\widetilde{C}_i} E_0^Q \left[\frac{\left(\widehat{C}(S_1, S_2) - \widetilde{C}_i(S_1, S_2) \right)^2}{\partial \widetilde{C}_{i-1}/\partial S_2^* - 1} \times \ker_Q \left(I(S_1, S_2) \approx \widetilde{C}_{i-1}(S_1, S_2) | S_1, S(t_0) \right) \right],$$

where $I(S_1, S_2) = \{S_2 - S_1 - K\}^+$ is the intrinsic value, and

$$E_0^Q \left[\frac{\widehat{C}(S_1, S_2) - \widetilde{C}_i(S_1, S_2)}{\partial \widetilde{C}_{i-1}(S_1, S_2, b_{i-1})/\partial S_2 - 1} \times \frac{\partial \widetilde{C}_i(S_1, S_2, b_i)}{\partial b} \times \ker_Q(I \approx \widetilde{C}_{i-1} | S_1, S(t_0)) \right] = 0. \quad (4.3)$$

This equation implies is a generalized local regression, but in practice, the regression is not sensitive to the term $\partial \widetilde{C}_{i-1}/\partial S_2$ in the denominator (e.g., if close to constant).

- **The local LSM algorithm** Let $ni \geq 1$ be a fixed number of iterations. Because it is a recursive approach, we specify the final period $t = T$ and recursively solve the continuation value for $T - 1$, $T - 2$, until $t = 1$. The intrinsic value is denoted by $I_t = \{S_{2,t} - S_{1,t} - K\}^+, t = \{1, 2, \ldots, T\}$. And $\widetilde{V}^n_t$ is the continuation value at time t and iteration n.

Consider a set of simulated paths, $\varpi \in \Omega$.

0. Set $t = T$. Define $y_{T+1} = 0$ and $\widetilde{V}^*_T = 0$.
1. UPDATING PATHS, $\varpi \in \Omega$

$$y_t = \begin{cases} I_t, & \text{if } I_t \geq \widetilde{V}^*_t \\ e^{-r\Delta t} \times y_{t+1}, & \text{otherwise.} \end{cases}$$

Set $t = t - 1$.
2. The new CONTINUATION VALUE

Set $n = 1$ and $\widetilde{V}^0_t = \widetilde{V}^*_{t+1}$. If $t = T - 1$, set $\widetilde{V}^0_{T-1} = \widetilde{V}^{LSM}_{T-1}$ from the LSM method.

2.1. LOCALIZING THE CONTINUATION VALUE

Provide a kernel $\mathcal{K}_h$. Then,

$$\widetilde{V}_t^n = \arg\min_{v_t \in F} \sum_{\varpi \in \Omega} \left(v_t(x_t) - \frac{y_{t+1}}{e^{r\Delta t}} \right)^2 \times \mathcal{K}_h \left(\widetilde{V}_t^{n-1}(x_t) - I_t \right).$$

Set $n = n + 1$. Go back to step **2.1** until $n = ni$.
Set $\widetilde{V}_t^* = \widetilde{V}_t^{ni}$. Go back to step **1** until $t = 1$.
End of the local LSM algorithm
Note that at $t = 1$, we have already estimated all continuation values; $\widetilde{V}_{T-1}^*, \widetilde{V}_{T-2}^*, \ldots, \widetilde{V}_1^*$.

Remark 4.2.

(1) $\frac{y_{t+1}}{e^{r\Delta t}}$ *is the dependent variable in the regression, where the continuation value $v(x) \in F$ is a simple polynomial and x are the regressors.*

(2) *With the function*

$$W_t = \mathcal{K}_h \left(\widetilde{V}_t^{n-1}(x_t) - I_t \right),$$

the local least squares estimator is easily computed by ordinary least squares by redefining x_t and y_{t+1} as $x_t \times \sqrt{W_t}$ and $y_{t+1} \times \sqrt{W_t}$, respectively.

(3) *We also use a Gaussian kernel, $\mathcal{K}_h(\widetilde{V}_t^{n-1}(x_t) - I_t)$, where h is the bandwidth. In the standard LSM method, $ni = 1$ and $\mathcal{K} = 1$ (or, in practice, $\mathcal{K} = 1_{\{S_2 \geq S_1 + K\}}$ for in-the-money paths).*

4.3 | Convergence

We now analyze the convergence properties of the local LSM algorithm in a general setup, which includes the spread option, considering the case of the exercise boundary function estimation $\widetilde{f}$ and then the continuation value estimation $\widetilde{C}$ (which only defines the boundary implicitly). In the former case, we assume the intrinsic value is linear in one price and the boundary function is smooth in the other prices, so we can explicitly solve for this frontier. For the latter case, we assume standard smoothness conditions on the value function (see Appendix C for details). In this section, we concentrate on the former case of the exercise boundary.

First, we provide new results on the rate of convergence of local series estimates of the exercise boundary or continuation value function at the exercise boundary (i.e., Theorem 4.3) or iterated versions of it (Corollary 4.4).[13] This result extends Newey (1997) and previous global results on the approximation of the continuation value function. Second, using these exercise boundaries, we give sufficient conditions for the convergence of the corresponding option price estimates to the true price as the number of paths M increases (i.e., Theorem 4.5) following Stentoft's (2004b) approach, which depends on an identification condition of the corresponding stopping time.

Theorem 4.3 shows the first ($i = 1$) iteration (or regression) parameter estimates $\hat{b}_1$ of the boundary approximation are consistent at any period t_j, $j = 1, \ldots, J$ (as well as the corresponding exercise-boundary-series estimate). These estimates are based on an initial ($i = 0$) global-consistent estimate $\hat{C}(S)$ of the continuation value at each period (provided by, e.g., the LSM method), with a convergence rate ξ_M depending on the number of paths M. This rate has to be sufficiently fast relative to the

increasing order of the polynomial H used in the series approximation and to the shrinking bandwidth h localizing the boundary region.

Theorem 4.3. *Under Assumptions C.1–C.4 in Appendix C, and assuming that for some estimate $\hat{C}(S)$ independent of our sample and some sequence $\xi_M \rightarrow 0$ as $M \rightarrow \infty$, it holds that*

$$\int \left[\hat{C}(S) - C(S)\right]^2 Q(S)dS = O_p(\xi_M), \quad H^4 h^{-3}\xi_M \rightarrow 0. \tag{4.4}$$

As $M \rightarrow \infty$, we have that

$$\int \left[\tilde{f}\left(S_1, \hat{b}_1\right) - \tilde{f}(S_1)\right]^2 Q_1(S_1)dS_1 = O_p\left(H(Mh^2)^{-1} + h^{-1}H^{-2\alpha} + h^4H^2 + H^4h^{-3}\xi_M\right) = o_p(1).$$

Proof. See Appendix C. □

The assumptions in Appendix C establish smoothness of the functions C, f, and Q; the availability of M independent paths; and restrictions on the asymptotic rates of h and H with respect to M. The conditions on the joint distribution of S strengthen those in Stentoft (2004a) to guarantee that conditioning on the boundary function is well defined and this function is smooth enough. Theorem 4.3 proves the convergence of $\hat{b}_1$ and $\tilde{f}(S_1, \hat{b}_1)$ under the additional restrictions (4.4) on a preliminary global estimate $\hat{C}(S)$ of the continuation value $C(S)$ (which can be obtained from the standard LSM algorithm). Both choices of H and h show a bias and variability trade-off, because a larger H implies a better sieve approximation to $\tilde{f}$ at the cost of estimating more parameters, whereas a smaller h imposes a stronger local conditioning to implicitly identify $\tilde{f}$ at the cost of reducing the effective number of paths used in the regression.

Iterating the previous local estimate of f is also possible, though $\tilde{f}(S_1, \hat{b}_1)$ might converge more slowly than an initial series estimates of C calculated without further smoothing as in LSM. Nevertheless, as described in the next corollary, proving the convergence holds for the second local iteration ($i = 2$) is possible if the initial estimates converge fast enough and the target functions are smooth.

Corollary 4.4. *Under the conditions of Theorem 4.3 and*

$$H^4 h^{-3}\left\{H(Mh^2)^{-1} + h^{-1}H^{-2\alpha} + h^4H^2 + H^4h^{-3}\xi_M\right\} \rightarrow 0$$

as $M \rightarrow \infty$ for some $\alpha > 1$, we have that $\int[\tilde{f}(S_1, \hat{b}_2) - \tilde{f}(S_1)]^2 Q_1(S_1)dS_1 = o_p(1)$.

Proof. See Appendix C. □

Now, using Theorem 4.3 and the methods of proposition 1 in Stentoft (2004b), it is immediate to show that the option price estimate given by the local LSM using these exercise boundary approximations converges to the true price under Assumption C.1(ii) in Appendix C.

Theorem 4.5. *Under the conditions of Theorem 4.3, $\hat{C}(t_0, S)$ is a consistent estimate of $C(t_0, S)$.*

Proof. Immediate from Theorem 4.3 and proposition 1 of Stentoft (2004b). □

4.4 | The continuation value versus the explicit exercise boundary

Using the continuation value offers two advantages. Consider the option on the maximum of N securities. First, the max-option depends on N exercise boundaries, but the continuation value is a single function. Second, this method easily converges (in practice), because an initial value is not necessary and the LSM provides consistent estimates of the continuation value. On the other hand, the exercise

boundary depends on one dimension less than the continuation value (e.g., for a Black–Scholes call/put, the boundary is a single point) and extrapolation is never done, which avoids extrapolation errors. For more complex payoffs, such as the max-option, the N exercise boundaries are obtained by regression too. Therefore, to explicitly estimate the boundary, we need to know more properties of the particular Bermudan option, but then it is a standard local least-squares problem.

In a horse race, the best approach is the one that yields a larger Bermudan price for a similar computational cost. This race can be translated to a richer family of exercise boundaries. The two methods, however, produce two different families of boundary functions that are not directly comparable. For the spread option, a quadratic continuation value function implies a hyperbolic exercise boundary, which is richer than a quadratic boundary, but the continuation value uses six parameters and the explicit boundary just three. In practice, we prefer the continuation value because it converges in only a few iterations (yet one can combine both methods, because a new iteration is independent of the previous one).[14]

4.5 | Complexity analysis

We briefly outline the computational cost of our algorithm. One advantage of the local LSM algorithm is that it is directly comparable to the LSM method, which has been exhaustively analyzed (e.g., Belomestny et al., 2015). The basic difference is that the local LSM method uses a kernel regression (localized in the exercise boundary). The cost of a local regression is similar to a standard regression; it requires computation of an extra weight for every path and exercise period (besides a small fixed cost of tuning the kernel). To localize the exercise boundary, we follow an iterative procedure. The number of iterations (per exercise opportunity) equals the number of regressions. By using the solution of the previous period t as the initial step at $t - 1$, a couple of iterations already yield the best prices. Additional iterations provide marginal gains, which can go either way because of the random error.

For N state variables, J exercise opportunities, M paths, and ni iterations, the LSM method requires $O(NJM)$ sample points and J regressions. The local LSM also requires $O(NJM)$ sample points, $O(JM \times ni)$ kernel evaluations, and $ni \times J$ regressions. In practice, $ni = 1$ to 3 local regressions are sufficient.

Further, because the local LSM focus on the exercise boundary, the local exercise rules (i.e., stopping times) do not depend on the option moneyness, which saves computational time in the case of portfolios of derivatives. Most LSM implementations start to simulate paths from the initial value of the state variables, which depends on moneyness. In brief, the local LSM is not more complex than the standard LSM method in practice, and is based on the FOCs.

5 | MULTIPLE STATE VARIABLES AND GENERAL PAYOFFS

The results derived for a two-factor spread option extend to a general setting, such as a multifactor model with separate exercise regions, that is, multiple exercise boundaries. We show this extension with a rather difficult payoff, the option on the maximum of N securities. Although modeling and estimating the continuation value looks easier than modeling and estimating multiple exercise boundaries, the latter case is straightforward too.

5.1 | The continuation value

Consider N securities, $S = \{s_1, s_2, \dots, s_N\}$. Denote $S_1 = \{s_1, s_2, \dots, s_{N-1}\}$, $S_2 = s_N$, $\overline{k}(S_1) \geq 0$ as the number of optimal exercise boundaries, and $f_k(S_1)$ and $\widetilde{f}_k(S_1)$ for $k = \{1, 2, \dots, \overline{k}(S_1)\}$ as the

implicit exercise boundaries. The intrinsic value is $I(S) = \{\max\{S\} - K\}^+$. Then the cost of suboptimal exercise (see the first term of equation (2.1)) is now given by

$$\min_b \int_0^\infty \dots \int_0^\infty \sum_{k=1}^{\overline{k}(S_1)} \left(\int_{\widetilde{f}_k(S_1,b)}^{f_k(S_1)} (C(S_1, S_2) - I(S_1, S_2)) \times Q(S_2|S_1, S(t_0)) \times dS_2 \right) \times Q(S_1|S(t_0)) \times dS_1,$$

where, for the $N-1$ most exterior integrals, $dS_1 = ds_1 \times ds_2 \times \dots \times ds_{N-1}$. The FOC are given by

$$E_0^Q \left[\sum_{k=1}^{\overline{k}(S_1)} \frac{C\left(S_1, \widetilde{f}_k^*(S_1)\right) - I\left(S_1, \widetilde{f}_k^*(S_1)\right)}{\widetilde{C}_2\left(S_1, \widetilde{f}_k^*(S_1), b^*\right) - I_2\left(S_1, \widetilde{f}_k^*(S_1)\right)} \times \frac{\partial \widetilde{C}\left(S_1, \widetilde{f}_k^*(S_1), b^*\right)}{\partial b} \right.$$

$$\left. \times Q(S_2 = \widetilde{f}_k^*(S_1)|S_1, S(t_0)) \right] = 0,$$

which are implemented as equation (4.3) above. The multiple integral is rewritten as an expectation, and to save notation, we write $S_2 = \widetilde{f}_k^*(S_1)$ if $I(S_1, \widetilde{f}_k^*(S_1)) = \widetilde{C}(S_1, \widetilde{f}_k^*(S_1), b^*)$. Let us stress that specifying or even knowing the number of exercises boundaries $\overline{k}(S_1)$ is not necessary, because the local regression of continuation values directly weights the value-matching errors.[15]

5.2 | The explicit exercise boundary

If we directly model the exercise boundary, it is convenient to rewrite the exercise boundaries as a function of the whole vector of prices S (instead of $S_1 = \{s_1, s_2, \dots, s_{N-1}\}$). Define the integer "$n^o(S) = n$" if $s_n = \max\{s_1, s_2, \dots, s_N\}$. We write $n^o = n^o(S)$, and $S_{/n} \in \mathcal{R}^{N-1}$ denotes the vector S without s_n. We assume the option is exercised at the largest price, s_{n^o}. Now N exercise boundaries exist, one for every security, which are denoted by $F_n(S_{/n}, b_n)$ and $\widetilde{F}_n(S_{/n}, b_n)$, $n = 1, 2, \dots, N$ (cap letters are used to avoid confusion with f). Then the cost is given by

$$\min_b \int_0^\infty \dots \int_0^\infty \left(\int_{\widetilde{F}_{n^o}(S_{/n^o})}^{F_{n^o}(S_{/n^o})} \left(C(S_{/n^o}, s_{n^o}) - I\left(S_{/n^o}, s_{n^o}\right) \right) \times Q(s_{n^o}|S_{/n^o}, S(t_0)) \times ds_{n^o} \right)$$

$$\times Q(S_{/n^o}|S(t_0)) \times dS_{/n^o},$$

where the expression within the parentheses is given by

$$\sum_{n=1}^N 1_{\{n^o = n\}} \times \int_{\widetilde{F}_n(S_{/n})}^{F_n(S_{/n})} (C(S_{/n}, s_n) - I(S_{/n}, s_n)) \times Q(s_n|S_{/n}, S(t_0)) \times ds_n.$$

The FOCs are given by

$$\int_0^\infty \dots \int_0^\infty \left[C\left(S_{/n^o}, F_{n^o}^*(S_{/n^o})\right) - I\left(S_{/n^o}, F_{n^o}^*(S_{/n^o})\right) \right] \times \frac{\partial F_{n^o}(S_{/n^o}, b_{n^o}^*)}{\partial b}$$

$$\times Q\left(S_{/n^o}, s_{n^o} = \widetilde{F}_{n^o}^*(S_{/n^o})|S(t_0)\right) \times dS_{/n^o} = 0.$$

Implementing these FOCs is easy. We have to estimate N different local least-squares regressions. Consider M simulation points, $S^{(m)}$, $m = \{1, 2, \ldots, M\}$. Separate these M points into N subsets, where $S_{/n^o}^{(m)}$ denotes a point ($S^{(m)}$ without the $s_{n^o}^{(m)}$ component) of the subset $n^o = \{1, 2, \ldots, N\}$. That is, for $n = 1, 2, \ldots, N$,

$$\int_0^\infty \cdots \int_0^\infty \left[C\left(S_{/n}, F_n^*(S_{/n})\right) - I\left(S_{/n}, F_n^*(S_{/n})\right) \right] \times \frac{\partial F_n^*\left(S_{/n}, b_n^*\right)}{\partial b}$$

$$\times Q\left(S_{/n}, s_n = \widetilde{F}_n(S_{/n}) | S(t_0)\right) \times dS_{/n} \times 1_{\{n^o = n\}} = 0,$$

where $I(S_{/n}, F_n^*(S_{/n})) = \{F_n^*(S_{/n}) - K\}^+$.

6 | NUMERICAL EXERCISE: MAX-OPTIONS ON FIVE STOCKS

We study a call option on the maximum of five securities in a lognormal setting. We distinguish between symmetric (same dividends, correlations, and volatilities) and asymmetric securities. If the five securities are symmetric, better explanatory variables are given by the ordered prices. From the prices $S = \{s_1, s_2, \ldots, s_N\}$, we obtain a new vector $X = \{s_{(1)}, s_{(2)}, \ldots, s_{(N)}\}$, where $s_{(n)} \geq s_{(n-1)}$ and $s_{(n)}, s_{(n-1)} \in S$. We consider three examples: symmetric securities with regressors X or S, and asymmetric securities with regressors $\{S, s_{(N)}\}$. The maximum price, $s_{(N)}$, turns out to be a good regressor. In the three cases, we consider a second-order polynomial as the regression function and a Gaussian kernel.

To avoid any upper bias, we use a second (or out-of-sample) independent simulation to compute Bermudan prices in all tables. In the first (in-sample) estimation of the boundaries, we use a large number of paths (4 million, including antithetic variables). In Table 6.4, we show how these results depend on a lower number of in-sample paths. In Tables 6.1 to 6.4, the kernel is defined such that p indicates the proportion of in-the-money paths effectively used in the local LSM estimation. In Table 7.1A, we show the results of a simpler kernel choice.

Symmetric securities and ordered prices, X
The literature has reported similar prices in most LSM implementations, and these prices, or lower-bound estimators, are close to the upper bound estimators (e.g., Andersen & Broadie, 2004; Haugh & Kogan, 2004). Consider the at-the-money option, the five stock prices, and strike price as equal to 100. Broadie and Cao (2008) report a tight confidence interval, lower- and upper-bounds, $[26.125, 26.152]$ for a Bermudan option with 3 years to maturity and nine exercise dates.

In Table 6.1, for the LSM, we get a price of 26.092 (0.0036). For the local LSM, we get a price very close to 26.15 (0.0036) from the first iteration and for all p, which is just 1 bps below our upper-bound estimator. If we explicitly estimate the boundary, we obtain a close price of 26.14. Because the symmetry implies the five boundaries are the same, estimating the five boundaries subject to this constraint, as is done in our implementation, is more efficient.

TABLE 6.1 Local LSM prices for 5-asset symmetric Bermudan max-call options: Ordered basis

S_0	p	$\hat{L}_0$ (s.e.)	$\hat{L}_1$	$\hat{L}_2$	$\hat{L}_3$	$\hat{L}_5$	$\hat{L}_{B,10}$	$\hat{U}_5$	$\hat{U}_0$ (s.e.)
80	**0.01**	8.991	9.009	9.012	9.012	9.012	9.012	9.018	9.025
$\overline{BC}$[9.014]	**0.03**	(.0016)	9.011	9.013	9.013	9.013	9.013	9.019	(.0117)
BC[9.008]	**0.05**		9.012	9.013	9.013	9.014	9.013	9.018	
	0.10		9.012	9.013	9.013	9.014	9.012	9.018	
90	**0.01**	16.599	16.621	16.637	16.637	16.638	16.635	16.642	16.659
$\overline{BC}$[16.644]	**0.03**	(0026)	16.636	16.637	16.637	16.639	16.634	16.642	(.0116)
BC[16.627]	**0.05**		16.637	16.637	16.637	16.638	16.632	16.645	
IZ^3[16.6307]	**0.10**		16.636	16.635	16.636	16.637	16.627	16.649	
100	**0.01**	26.092	26.146	26.150	26.152	26.152	26.146	26.159	26.180
$\overline{BC}$[26.152]	**0.03**	(.0036)	26.149	26.150	26.151	26.153	26.144	26.164	(.0099)
BC[26.125]	**0.05**		26.149	26.149	26.150	26.152	26.140	26.164	
IZ^3[26.1470]	**0.10**		26.147	26.146	26.148	26.148	26.131	26.168	
110	**0.01**	36.704	36.773	36.776	36.777	36.778	36.764	36.782	36.804
$\overline{BC}$[36.781]	**0.03**	(.0046)	36.776	36.776	36.778	36.777	36.762	36.791	(.0095)
BC[36.722]	**0.05**		36.775	36.774	36.776	36.776	36.757	36.791	
IZ^3[36.7669]	**0.10**		36.774	36.771	36.773	36.773	36.746	36.793	
120	**0.01**	47.860	47.931	47.939	47.939	47.940	47.933	47.950	47.974
$\overline{BC}$[47.988]	**0.03**	(.0056)	47.935	47.939	47.941	47.942	47.916	47.956	(.0090)
BC[47.862]	**0.05**		47.937	47.939	47.940	47.940	47.924	47.962	
	0.10		47.936	47.936	47.938	47.937	47.910	47.962	

Notes: The local LSM price, $\hat{L}_i$, with i iterations using the continuation value ($i = 0$ is LSM). $\hat{L}_B$ is based on fitting the exercise boundary starting from $\hat{L}_5$. Option parameters (lognormal setting) are volatility $\sigma = 20\%$, dividend-yield $q = 10\%$, interest rate $r = 5\%$, strike price $K = 100$, maturity $T = 3$, correlation $\rho = 0$, exercise dates $n = 9$, and paths $M = 4,000,000$. Quadratic-ordered basis. We start to simulate prices 6 months in advance. BC ($\overline{BC}$) are lower- (upper-) bound Bermudan prices from Broadie and Cao (2008). IZ^3 is a lower bound based on a cubic frontier from Ibáñez and Zapatero (2004). Standard errors are in parentheses. p is the proportion of the effective number of local points to M used by the Gaussian kernel. $\hat{U}_5$ and $\hat{U}_0$ are the upper bounds associated with the exercise strategy implicit in the lower bounds $\hat{L}_5$ and $\hat{L}_0$, respectively (where these upper bounds are equal to the lower bound plus a gap computed with 3,000 paths and 10,000 sub simulation paths).

Symmetric securities and non-ordered prices, S

This choice is an example of a poor basis of state variables. In Table 6.2, the local LSM yields a price approximately 0.60% larger than the LSM. For instance, for the at-the-money option, the LSM and the local LSM yield a price of 25.93 and 26.11, respectively. The results are robust and converge after three iterations.

Asymmetric securities

Using the ordered vector X is suboptimal now, because loss of information occurs. For a deep in-the-money security, the optimal exercise decision depends also on whether it is a high-volatility/dividend security. Broadie and Cao (2008) report a less tight confidence interval of [37.75, 37.99] for the same at-the-money option. Our local LSM price is near 37.96 (0.005) by estimating the continuation value, which is a notable improvement and a few bps below than our upper-bound estimator of 38.00 (0.005). The prices are maximized for a small kernel with $p = 0.01$ (a large kernel with $p = 0.10$ for the exercise boundary). Similar results hold for the other in- and out-of-the-money cases.

In brief, the results produced by the local LSM improve over the LSM method are a few (one-digit) bps below the associated upper bound, and require one or two iterations (i.e., number of regressions).

TABLE 6.2 Local LSM prices for 5-asset symmetric Bermudan max-call option: Nonordered basis

S_0	p	$\hat{L}_0$ (s.e.)	$\hat{L}_1$	$\hat{L}_2$	$\hat{L}_3$	$\hat{L}_5$	$\hat{L}_{B,10}$	$\hat{U}_5$	$\hat{U}_0$ (s.e.)
80	**0.01**	8.890	8.967	8.982	8.989	8.991	8.983	9.017	9.042
$\overline{BC}$[9.014]	**0.03**	(.0016)	8.970	8.987	8.992	8.995	8.989	9.018	(.0119)
BC[9.008]	**0.05**		8.969	8.988	8.994	8.997	8.990	9.018	
	0.10		8.967	8.987	8.994	8.998	8.990	9.016	
90	**0.01**	16.461	16.564	16.594	16.602	16.603	16.592	16.649	16.676
$\overline{BC}$[16.644]	**0.03**	(.0026)	16.570	16.599	16.606	16.610	16.602	16.647	(.0130)
BC[16.627]	**0.05**		16.570	16.598	16.606	16.609	16.601	16.648	
IZ³[16.6307]	**0.10**		16.568	16.595	16.605	16.608	16.597	16.647	
100	**0.01**	25.933	26.059	26.090	26.097	26.108	26.098	26.169	26.208
$\overline{BC}$[26.152]	**0.03**	(.0036)	26.063	26.097	26.105	26.110	26.107	26.167	(.0135)
BC[26.125]	**0.05**		26.063	26.098	26.107	26.112	26.107	26.168	
IZ³[26.1470]	**0.10**		26.061	26.094	26.104	26.112	26.100	26.169	
110	**0.01**	36.527	36.667	36.696	36.705	36.712	36.710	36.802	36.851
$\overline{BC}$[36.781]	**0.03**	(.0046)	36.673	36.706	36.714	36.717	36.722	36.804	(.0142)
BC[36.722]	**0.05**		36.674	36.706	36.715	36.718	36.723	36.803	
IZ³[36.7669]	**0.10**		36.671	36.703	36.713	36.718	36.713	36.803	
120	**0.01**	47.672	47.827	47.856	47.865	47.867	47.866	47.965	48.028
$\overline{BC}$[47.988]	**0.03**	(.0056)	47.826	47.859	47.870	47.871	47.870	47.968	(.0200)
BC[47.862]	**0.05**		47.826	47.861	47.871	47.872	47.878	47.967	
	0.10		47.823	47.857	47.868	47.872	47.872	47.966	

Notes: The local LSM price, $\hat{L}_i$, with i iterations using the continuation value ($i = 0$ is LSM). $\hat{L}_B$ is based on fitting the exercise boundary starting from $\hat{L}_5$. Option parameters are $\sigma = 20\%$, $q = 10\%$, $r = 5\%$, $K = 100$, $T = 3$, $\rho = 0$, $n = 9$, and $M = 4,000,000$. Quadratic nonordered basis. We start to simulate prices 6 months in advance. BC ($\overline{BC}$) are lower- (upper-) bound Bermudan prices from Broadie and Cao (2008). IZ³ is a lower bound based on a cubic frontier from Ibáñez and Zapatero (2004). Standard errors are in parentheses. p is the proportion of the effective number of local points to M used by the Gaussian kernel. $\hat{U}_5$ and $\hat{U}_0$ are the upper bounds associated with the exercise strategy implicit in the lower bounds $\hat{L}_5$ and $\hat{L}_0$, respectively (where these upper bounds are equal to the lower bound plus a gap computed with 3,000 paths and 10,000 sub simulation paths).

These results are robust to the kernel p and moneyness. The largest improvement happens in the first two iterations, where additional iterations provide smaller gains, which may go either way because of the associated random error.

Number of paths

In Table 6.4, we show how the results depend on the number of paths used in the in-sample regression. Tables 6.1 to 6.3 use 4 million paths. The first block of Table 6.4 shows that when we have a good basis (as in Table 6.1), the loss is less than 2 cents for all cases, from 100,000 paths on. From the second block of Table 6.4, for the hardest asymmetric case, a small trade-off occurs between better prices and the number of paths for the local LSM method, which is intuitive in any regression model.

6.1 | The kernel and the bandwidth

Our local regressions are estimated by weighted least squares, where a Gaussian kernel function, $\ker_Q$, gives more weight to the data points close to the exercise boundary. The kernel depends on a bandwidth parameter h, which normalizes the boundary errors. The smaller the bandwidth h, the narrower the localization around the boundary.

TABLE 6.3 Local LSM prices for 5-asset asymmetric Bermudan max-call option: Nonordered basis and the maximum price

S_0	p	$\hat{L}_0$(s.e.)	$\hat{L}_1$	$\hat{L}_2$	$\hat{L}_3$	$\hat{L}_5$	$\hat{L}_{B,10}$	$\hat{U}_5$	$\hat{U}_{B,10}$	$\hat{U}_0$ (s.e.)
80	**0.01**	18.626	18.828	18.832	18.836	18.836	18.800	18.860	18.865	18.894
$\overline{BC}$[18.866]	**0.03**	(.0045)	18.835	18.813	18.814	18.820	18.820	18.862	18.866	(.003)
BC[18.744]	**0.05**		18.829	18.798	18.805	18.807	18.825	18.863	18.861	
	0.10		18.814	18.778	18.784	18.786	18.825	18.865	18.861	
90	**0.01**	27.318	27.616	27.619	27.622	27.624	27.578	27.655	27.677	27.702
$\overline{BC}$[27.659]	**0.03**	(.0053)	27.619	27.593	27.598	27.602	27.605	27.656	27.674	(.005)
BC[27.480]	**0.05**		27.612	27.574	27.582	27.586	27.610	27.657	27.660	
	0.10		27.593	27.545	27.554	27.557	27.609	27.660	27.665	
100	**0.01**	37.554	37.925	37.941	37.961	37.958	37.920	38.004	38.008	38.051
$\overline{BC}$[37.988]	**0.03**	(.0051)	37.929	37.908	37.933	37.932	37.944	38.006	38.003	(.005)
BC[37.746]	**0.05**		37.923	37.887	37.914	37.913	37.945	38.006	38.004	
	0.10		37.901	37.854	37.881	37.878	37.940	38.001	38.003	
110	**0.01**	49.020	49.409	49.436	49.443	49.444	49.367	49.510	49.513	49.587
$\overline{BC}$[49.492]	**0.03**	(.0065)	49.421	49.403	49.417	49.421	49.436	49.513	49.510	(.008)
BC[49.175]	**0.05**		49.421	49.382	49.398	49.405	49.439	49.516	49.512	
	0.10		49.401	49.348	49.363	49.369	49.434	49.517	49.510	
120	**0.01**	61.180	61.541	61.592	61.596	61.600	61.537	61.701	61.696	61.772
$\overline{BC}$[61.686]	**0.03**	(.0071)	61.567	61.570	61.579	61.586	61.610	61.704	61.697	(.014)
BC[61.294]	**0.05**		61.578	61.551	61.564	61.570	61.613	61.702	61.696	
	0.10		61.561	61.520	61.533	61.541	61.607	61.704	61.694	

Notes: The local LSM price, $\hat{L}_i$, with i iterations using the continuation value ($i = 0$ is LSM). $\hat{L}_B$ is based on fitting the exercise boundary starting from $\hat{L}_5$. Option parameters are as in Table 6.1, except for volatilities, $\sigma_i = 8\%, 16\%, 24\%, 32\%$ and 40% for $i = 1, \dots, 5$, respectively. Quadratic nonordered basis and the maximum price. We start to simulate prices 6 months in advance. BC ($\overline{BC}$) are lower- (upper-) bound Bermudan prices from Broadie and Cao (2008). Standard errors are in parentheses. p is the proportion of the effective number of local points to M used by the Gaussian kernel. $\hat{U}_5$, $\hat{U}_{B,10}$, and $\hat{U}_0$ are the upper bounds associated with the exercise strategy implicit in the lower bounds $\hat{L}_5$, $\hat{L}_{B,10}$, and $\hat{L}_0$, respectively (where these upper bounds are equal to the lower bound plus a gap computed with 3,000 paths and 10,000 sub simulation paths).

The two kernels, $\ker_Q(\hat{C}(S^{(m)}) \approx I(S^{(m)}))$ and $\ker_Q(S_2^{(m)} \approx \widetilde{f}(S_1^{(m)})))$, are implemented by

$$\frac{1}{h}\phi\left(\frac{\hat{C}(S^{(m)}) - I(S^{(m)})}{h}\right) \quad \text{and} \quad \frac{1}{h}\phi\left(\frac{S_2^{(m)} - \widetilde{f}\left(S_1^{(m)}\right)}{h}\right),$$

respectively, where

$$\phi(z) = (2\pi)^{-1/2} \exp\left(-\frac{1}{2}z^2\right).$$

The bandwidth, h, is chosen to maximize price estimates. Optimal values typically represent between 1% to a 10% (i.e., $p = \{0.01, 0.03, 0.05, 0.10\}$) of the overall sample information as computed in local polynomial nonparametric estimation (Fan & Gijbels, 1995, p. 364).[16] p is equivalent (for a given h) to the ratio of the effective number of local points to the number of simulated paths available to estimate

TABLE 6.4 Number of paths and convergence: Symmetric and asymmetric options ($p = 0.05$)

	10	10	10	10	10
IN	100,000	200,000	400,000	1,000,000	2,000,000
OUT	4,000,000	4,000,000	4,000,000	4,000,000	4,000,000
Symmetric Option and Ordered Basis					
0	26.015	26.005	26.008	26.007	26.092
1	26.012	26.078	26.129	26.144	26.149
2	26.118	26.122	26.126	26.125	26.149
3	26.128	26.132	26.134	26.135	26.150
4	26.128	26.127	26.131	26.130	26.151
5	26.126	26.129	26.129	26.135	26.152
s.e.	(0.0028)	(0.0028)	(0.0019)	(0.0027)	(0.0031)
Asymmetric Option and Nonordered Basis					
0	37.556	37.547	37.564	37.547	37.554
1	37.740	37.863	37.866	37.929	37.923
2	37.861	37.881	37.888	37.882	37.908
3	37.876	37.893	37.902	37.903	37.933
4	37.877	37.891	37.906	37.904	37.934
5	37.883	37.898	37.898	37.913	37.932
s.e.	(0.0061)	(0.0069)	(0.0056)	(0.0040)	(0.0051)

Notes: Convergence of the local LSM Bermudan price as a function of the number of paths to estimate the continuation value, for both symmetric and asymmetric securities, with a quadratic ordered and nonordered basis (parameters are those of Tables 6.1 and 6.3, respectively). The first row is the number of simulation experiments, the second is the number of paths to estimate the continuation value, the third is the number of paths to price the option, and the rest are the Bermudan prices associated with each iteration of the local LSM (zero is LSM), where s.e. is the standard error.

the unconditional expectation of the dependent variable. A smaller p indicates a narrower bandwidth and a more localized regression.

Kernel robustness

Tables 6.1 to 6.3 already indicate the optimal kernel is robust (to different values of p), though h can change from period to period. In Table 7.1A and 7.1B, we use a simple kernel, where h depends on the standard deviation of the errors, the number of in-the-money paths, M_*, and a control parameter, δ; that is,

$$h = \delta \times \text{Var}[\hat{C}(S_m) - I(S_m)]^{1/2} \times M_*^{-0.2}.$$

δ is constant and, once tuned, is the same for all iterations and periods. This simple kernel is very robust; the losses are a few cents compared to the optimal kernel. We show only the results for the asymmetric case, for both the continuation value and the explicit boundary. For low δ, both methods are less robust because very few (in-the-money) sample paths are used in the estimation. Otherwise, explicitly estimating the five exercise boundaries yields the best prices, namely, 37.95 (for $0.75 \leq \delta \leq 2$).

7 | CONCLUDING REMARKS

LSM methods enable us to price Bermudan options by simulation, and are based on estimating the option continuation value by least squares (and proceeding in a recursive way). We show the

TABLE 7.1A Alternative h choice (continuation value): Asymmetric case ($S_0 = 100$)

Iter \ δ	0.15	0.2	0.3	0.6	0.8	1	1.2	1.5	2	2.5
0	37.520	37.520	37.520	37.520	37.520	37.520	37.520	37.520	37.520	37.520
1	37.799	37.802	37.803	37.804	37.804	37.802	37.799	37.792	37.779	37.761
2	37.844	37.861	37.867	37.870	37.867	37.864	37.862	37.854	37.836	37.812
3	37.581	37.879	37.871	37.887	37.888	37.887	37.882	37.873	37.853	37.827
4	37.881	37.779	37.826	37.895	37.896	37.891	37.890	37.882	37.859	37.830
5	37.875	37.873	37.884	37.890	37.898	37.894	37.893	37.884	37.862	37.831

TABLE 7.1B Alternative h choice (exercise boundary): Asymmetric case ($S_0 = 100$)

Iter \ δ	0.3	0.4	0.5	0.6	0.75	1	1.25	1.5	2	3
1	37.257	37.301	37.461	37.542	37.742	37.887	37.919	37.936	37.928	37.869
2	37.670	37.850	37.933	37.945	37.948	37.951	37.951	37.948	37.936	37.872
3	37.935	37.941	37.946	37.947	37.949	37.954	37.953	37.939	37.937	37.870
4	37.899	37.191	37.925	37.740	37.906	37.959	37.956	37.951	37.935	37.861
5	37.501	37.943	37.952	37.953	37.952	37.819	37.953	37.951	37.933	37.851
6	37.693	37.917	37.251	37.956	37.957	37.959	37.951	37.949	37.930	37.844
7	37.696	37.932	37.929	37.953	37.957	37.906	37.954	37.949	37.927	37.839
8	37.900	37.928	37.155	37.946	37.950	37.958	37.951	37.948	37.926	37.835
9	37.828	37.865	37.304	36.578	37.955	37.935	37.952	37.946	37.924	37.831
10	37.931	37.080	37.944	36.913	37.954	37.954	37.950	37.946	37.923	37.830

Notes: Price of the at-the-money asymmetric Bermudan max-option for different iterations and control parameter δ using a quadratic nonordered basis (parameters are those of Table 6.3). The bandwidth h depends on the standard deviation of the errors, the number of in-the-money paths, M_*, and a control parameter, δ; that is, $h = \delta \times \mathrm{Var}[\hat{C}(S_m) - I(S_m)]^{1/2} \times M_*^{-0.2}$. The parameter δ is constant through all the iterations and all the exercise opportunities. We show both cases: lower bounds based on parametrizing the continuation value and the exercise boundary, respectively.

Bermudan option price is maximized when this continuation value is estimated near the exercise boundary, which is equivalent to implicitly estimating the optimal exercise boundary by using the value-matching condition. Localizing is the main difference with respect to (global) regression methods, but is the key for optimal exercise decisions, and requires iteration of a few local least-squares regressions per exercise date.

We call this new algorithm the local LSM method (which can be implemented in a second way, by explicitly estimating the exercise boundary) and also study its convergence. The new prices, or lower bounds, improve over the LSM method and are close to the associated dual upper bounds (e.g., for max-call options). The small gap between the lower and the dual upper bounds implies the near optimality of the local LSM strategy, a property that is further addressed in a second paper (Ibáñez & Velasco, 2016).

A direct by-product of the local LSM is the (optimal) exercise boundary. In real options, this boundary or "trigger" function is necessary to determine the option value. Grenadier and Malenko (2010) use a variation of the standard LSM method to approximate this trigger function (for an American option). The local LSM method focuses precisely on this boundary.

Finally, let us point out that an American option (under diffusion processes) is characterized by the value-matching and smooth-pasting conditions (Merton, 1973). The equivalent Bermudan option, for a given family of exercise boundaries, depends only on value-matching errors being orthogonal to

the regressors at the exercise boundary, which also holds for processes with jumps (whereas smooth-pasting does not hold, because exercise is discrete).

ACKNOWLEDGMENTS

ENDNOTES

[1] This type of idea has appeared in the previous literature, most notably in the construction of lower-bound estimators (Broadie & Detemple, 1996), where optimization is over constant exercise thresholds. We, however, optimize over more general classes of functions (e.g., polynomials) and combine optimization and simulation.

[2] For N state variables and an H-degree polynomial, the number of coefficients to estimate by least-squares is of the order $O(N^H)$. Hence, most LSM applications use a quadratic function, $H = 2$, as we do here.

[3] This result extends Newey (1997) to local conditional expectation series estimates given a restriction, which is imposed iteratively starting from unrestricted estimates.

[4] Tsitsiklis and Van Roy (2001), Clément, Lamberton, and Protter (2002), Stentoft (2004b), Egloff (2005), and Glasserman and Yu (2004) study convergence properties. Moreno and Navas (2003), Stentoft (2004a), and Areal, Rodrigues, and Armada (2004) study numerical issues. Dutt and Welke (2008), Rasmussen (2005), and Wang and Caflisch (2010) study further refinements. For other MC methods, see Glasserman (2004), Ibáñez (2003), Ibáñez and Zapatero (2004), and references therein. The LSM is used from pricing Bermudan swaptions (Trolle & Schwartz, 2009) to optimal portfolios (Brandt, Goyal, Santa-Clara, & Stroud, 2005; Longstaff, 2001), credit risk (Jarrow, Li, Liu, & Wu, 2010), energy derivatives (Cartea & Williams, 2008), real estate (Longstaff, 2005), real options (Gamba, 2003), executive stock options (León & Vaello-Sebastià, 2009), control problems (Belomestny, Kolodko, & Schoenmakers, 2010), or nonlinear PDEs (Bouchard & Touzi, 2004), among other applications.

[5] For clarity, equation (2.1) can also be written as a double integral, namely,

$$
\int_0^\infty \left(\int_{\widetilde{f}_1(S_1(t_1))}^{f_1(S_1(t_1))} (C(t_1, S_1(t_1), S_2(t_1)) - (S_2(t_1) - S_1(t_1) - K)) \times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \right.
$$

$$
+ \int_0^{\widetilde{f}_1(S_1(t_1))} (C(t_1, S_1(t_1), S_2(t_1)) - P(t_1, (S_1(t_1), S_2(t_1)), \widetilde{f}_{2:J})) \times Q(S_2(t_1)|S_1(t_1), S(t_0)) \times dS_2(t_1) \Bigg)
$$

$$
\times Q(S_1(t_1)|S(t_0)) \times dS_1(t_1).
$$

[6] See Ibáñez and Paraskevopoulos (2010) for a second-order Taylor approximation of the cost of suboptimal exercise, for the case of continuously exercisable American put/call options in a diffusion setting.

[7] Because $S_1 \in \mathcal{R}$, we can divide the real line in several intervals, and then define a different polynomial exercise boundary, $\widetilde{f}$, for each interval. For large dimensions, this simple idea suffers from the curse of dimensionality.

[8] For a call option (i.e., S_1 and $\widetilde{f}$ are constants and $S_1 + K$ is the strike price), equation (3.3) simplifies to

$$C(S_1, \widetilde{f}^*(S_1)) - (\widetilde{f}^*(S_1) - (S_1 + K)) = \Sigma/Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0)).$$

The value-matching error equals $\Sigma/Q(S_2 = \widetilde{f}^*(S_1))$ at $\widetilde{f}^*(S_1)$, the estimated boundary. If $\Sigma = 0$, $\widetilde{f}^* = f$. $Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0)) = Q(S_2|S_1, S(t_0))|_{S_2 = \widetilde{f}^*(S_1)}$ is a conditional probability evaluated at $S_2 = \widetilde{f}^*(S_1)$.

[9] Intuitively, C is smooth because it is given by an expectation, which smooths nonlinear payoffs, and $O(\Delta \widetilde{f}(S_1)^3)$ is small because otherwise the value-matching errors in the FOCs are large.

[10] Because $C_2 = 1$ for the continuously exercisable American option, $C_2 < 1$ for the equivalent Bermudan case.

[11] We assume $\widetilde{C}_2(S_1, \widetilde{f}(S_1), b) < 1$, which is consistent with our previous definition of the suboptimal exercise region. That is, for accelerating (delaying) exercise, $S_2 \in [\widetilde{f}(S_1), f(S_1)]$ ($S_2 \in [f(S_1), \widetilde{f}(S_1)]$).

[12] Broadie and Cao (2008), Rasmussen (2005), and Wang and Caflisch (2010) start to simulate paths before the initial time t_0 to improve the efficiency of regression methods, for example, if $S(t_0)$ is deep out-of-the-money or to simultaneously obtain the price and the Greeks of the Bermudan option.

[13] A series estimates is a least-squares estimation where the regressors are H approximation functions (e.g., a series of H polynomials); see Newey (1997).

[14] Some extrapolation errors appear when we use a continuation value function. Although the whole function $\widetilde{C}(S_1, S_2)$ is used to make exercise decisions, a local LSM focuses on the exercise boundary, namely, $\widetilde{C}(S_1, f^*(S_1))$. Therefore, we are extrapolating the value $\widetilde{C}(S_1, S_2)$ from $\widetilde{C}(S_1, f^*(S_1))$, which implies some raw exercise errors can appear for S_2 far away from $f^*(S_1)$. For example, for deep in-the-money paths, and convex payoffs, a quadratic function will overshoot the linear intrinsic value. This overshooting can be controlled by checking if the partial derivatives are larger than (lower than minus) one for call- (put-) type payoffs. Indeed proving the convergence of a model that estimates the boundary is easier because we do not have to control for extrapolation errors. These errors may occur with any basis of functions and regression method.

[15] Consider the option on the maximum of two assets, $N = 2$. From Ibáñez and Zapatero (2004), $\bar{k}(S_1) = \{0, 1, 2\}$. $\bar{k}(S_1) = 0$ if S_1 is very deep in-the-money, being optimal to exercise to the largest of the two prices; $\bar{k}(S_1) = 1$ if S_1 is deep out-of-the-money, being optimal to exercise for S_2 large enough; $\bar{k}(S_1) = 2$ if S_1 is at or in-the-money, being optimal to exercise to the largest of S_1 and S_2 if (approximately) $S_1 \neq S_2$.

[16] In standard nonparametric estimation by local polynomials, the regressions use local information in a neighborhood of a particular value of the regressors, which, unlike in our case, is explicitly known.

[17] If we model the continuation value, $\widetilde{f}_{bb} \neq 0$, then,

$$G_{bb}\left(\widetilde{f}(S_1, b^*)\right) = \frac{d[g(\widetilde{f}^*(S_1))]}{d\widetilde{f}} \times \frac{\widetilde{C}_b \widetilde{C}_b'}{\left(\widetilde{C}_2 - 1\right)^2} + g(\widetilde{f}^*(S_1)) \times \frac{\widetilde{C}_b\left(\widetilde{C}_{2b} + \widetilde{C}_{22b}\right)}{\left(\widetilde{C}_2 - 1\right)^2},$$

where $\widetilde{C}_{bb} = 0$. But a similar interpretation follows if $\widetilde{f}^*$ is close to f, because $g(f(S_1)) = 0$.

REFERENCES

Andersen, L., & Broadie, M. (2004). A primal-dual simulation algorithm for pricing multi-dimensional American options. *Management Science*, *50*, 1222–1234.

Areal, N., Rodrigues, A., & Armada, M. R. (2008). Improvements to the least squares Monte Carlo option valuation method. *Review of Derivatives Research*, *11*, 119–151.

Belomestny, D. (2011a). On the rates of convergence of simulation-based optimization algorithms for optimal stopping problems. *The Annals of Applied Probability*, *21*, 215–239.

Belomestny, D. (2011b). Pricing Bermudan options using nonparametric regression: Optimal rates of convergence for lower estimates. *Finance and Stochastics*, *15*, 655–683.

Belomestny, D., Dickmann, F., & Nagapetyan, T. (2015). Pricing American options via multi-level approximation methods. *SIAM Journal of Financial Mathematics*, *6*, 448–466.

Belomestny, D., Kolodko, A., & Schoenmakers, J. (2010). Regression methods for stochastic control problems and their convergence analysis. *SIAM Journal on Control and Optimization*, *48*, 3562–3588.

Bouchard, B., & Touzi, N. (2004). Discrete time approximation and Monte Carlo simulation for backward stochastic differential equations. *Stochastic Processes and Their Applications*, *111*, 175–206.

Brandt, M., Goyal, A., Santa-Clara, P., & Stroud, J. (2005). A simulation approach to dynamic portfolio choice with an application to learning about return predictability. *Review of Financial Studies*, *28*, 831–873.

Broadie, M., & Cao, M. (2008). Improved lower and upper bound algorithms for pricing American options by simulation. *Quantitative Finance*, *8*, 845–861.

Broadie, M., & Detemple, J. (1996). American option valuation: New bounds, approximations, and a comparison of existing methods. *Review of Financial Studies*, *9*, 1211–1250.

Carriére, J. F. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics*, *19*, 19–30.

Cartea, A., & Williams, T. (2008). UK gas markets: The market price of risk and applications to multiple interruptible supply contracts. *Energy Economics*, *30*, 829–846.

Clément, E., Lamberton, D., & Protter, P. (2002). An analysis of a least squares regression method for American option pricing. *Finance and Stochastics*, *6*, 449–471.

Duffie, D. (2001). *Dynamic asset pricing theory* (3rd ed.). Princeton, NJ: Princeton University Press.

Dutt, S. K., & Welke, G. M. (2008). Just-in-time Monte Carlo for path-dependent American options. *Journal of Derivatives*, *15*, 29–47.

Egloff, D. (2005). Monte Carlo algorithms for optimal stopping and statistical learning. *Annals of Applied Probability*, *15*, 1396–1432.

Fan, J., & Gijbels, I. (1995). Data-driven bandwidth selection in local polynomial fitting: Variable bandwidth and spatial adaptation. *Journal of the Royal Statistical Society, Series B*, *57*, 371–394.

Gamba, A. (2003). Real options valuation: A Monte Carlo approach, in real options. In G. A. Sick & S. C. Myers (Eds.), *Handbooks in finance*. North Holland: Elsevier.

Glasserman, P. (2004). *Monte Carlo methods in financial engineering*. New York: Springer.

Glasserman, P., & Yu, B. (2004). Number of paths versus number of basis functions in American option pricing. *Annals of Applied Probability*, *14*, 2090–2119.

Grenadier, S., & Malenko, A. (2010). A Bayesian approach to real options: The case of distinguishing between temporary and permanent shocks. *Journal of Finance*, *65*, 1949–1986.

Haugh, M., & Kogan, L. (2004). Pricing American options: A duality approach. *Operations Research*, *52*, 258–270.

Ibáñez, A. (2003). Valuation by simulation of contingent claims with multiple exercise opportunities. *Mathematical Finance*, *14*, 223–248.

Ibáñez, A., & Paraskevopoulos, I. (2010). The sensitivity of American options to suboptimal exercise strategies. *Journal of Financial and Quantitative Analysis*, *45*, 1563–1590.

Ibáñez, A., & Velasco, C. (2016). *Recursive lower- and dual upper-bounds for Bermudan-style options.* Working paper at SSRN.

Ibáñez, A., & Zapatero, F. (2004). Valuation by simulation of American options through computation of the optimal exercise frontier. *Journal of Financial and Quantitative Analysis*, *39*, 253–275.

Jarrow, R., Li, H., Liu, S., & Wu, C. (2010). Reduced-form valuation of callable corporate bonds: Theory and evidence. *Journal of Financial Economics*, *95*, 227–248.

León, A., & Vaello-Sebastià, A. (2009). American GARCH employee stock option valuation. *Journal of Banking & Finance*, *33*, 1129–1143.

Longstaff, F. (2001). Optimal portfolio choice and the valuation of illiquid securities. *Review of Financial Studies*, *14*, 407–431.

Longstaff, F. (2005). Borrower credit and the valuation of mortgage-backed securities. *Real Estate Economics*, *33*, 619–661.

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, *14*, 113–147.

Merton, R. (1973). The theory of rational option pricing. *Bell Journal of Economics and Management Science*, *4*, 141–183.

Moreno, M., & Navas, J. F. (2003). On the robustness of least-squares Monte Carlo (LSM) for pricing American derivatives. *Review of Derivatives Research*, *6*, 107–128.

Newey, W. K. (1997). Convergence rates and asymptotic normality for series estimators. *Journal of Econometrics*, *79*, 147–168.

Rassmussen, N. S. (2005). Control variates for Monte Carlo valuation of American options. *Journal of Computational Finance*, *9*, 2–12.

Stentoft, L. (2004a). Assessing the least squares Monte-Carlo approach to American option valuation. *Review of Derivatives Research*, *7*, 129–168.

Stentoft, L. (2004b). Convergence of the least squares Monte Carlo approach to American option valuation. *Management Science*, *50*, 1193–1203.

Trolle, A., & Schwartz, E. S. (2009). A general stochastic volatility model for the pricing of interest rate derivatives. *Review of Financial Studies*, *22*, 2007–2057.

Tsitsiklis, J., & Roy, B. V. (1999). Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Transactions on Automatic Control*, *44*, 1840–1851.

Tsitsiklis, J. N., & Roy, B. V. (2001). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, *12*, 694–703.

Wang, Y., & Caflisch, R. (2010). Pricing and hedging American-style options: A simple simulation-based approach. *Journal of Computational Finance*, *13*, 95–125.

Zanger, D. Z. (2016). Convergence of a least-squares Monte Carlo algorithm for American option pricing with dependent sample data. *Mathematical Finance*. Advance online publication. https://doi.org/10.1111/mafi.12125.

APPENDIX A: SECOND-ORDER CONDITIONS

We show the conditions under which the local LSM is a local minimum (i.e., the convexity of the cost function). From equations (3.1) and (3.2), the matrix of second derivatives is given by

$$E_0^Q \left[G_{bb} \left(\widetilde{f}(S_1, b^*) \right) \right] = E_0^Q \left[G_{\widetilde{f}\widetilde{f}} \left(\widetilde{f}^* \right) \times \widetilde{f}_b(b^*) \times \widetilde{f}_b(b^*)' + G_{\widetilde{f}} \left(\widetilde{f}^* \right) \times \widetilde{f}_{bb}(b^*) \right].$$

Let us assume $\widetilde{f}(S_1, b)$ is linear in b; that is, $\widetilde{f}_{bb} = 0$. It follows that

$$G_{bb} \left(\widetilde{f}(S_1, b^*) \right) = \frac{d \left[\left(\left(\widetilde{f}^*(S_1) - S_1 - K \right) - C(S_1, \widetilde{f}^*(S_1)) \right) \times Q(S_2 = \widetilde{f}^*(S_1) | S_1, S(t_0)) \right]}{d \widetilde{f}}$$

$$\times \left(\frac{\partial \widetilde{f}(S_1, b^*)}{\partial b} \right) \times \left(\frac{\partial \widetilde{f}(S_1, b^*)}{\partial b} \right)',$$

where $\widetilde{f}_b(b^*) \times \widetilde{f}_b(b^*)' = (\frac{\partial \widetilde{f}(S_1, b^*)}{\partial b}) \times (\frac{\partial \widetilde{f}(S_1, b^*)}{\partial b})'$ denotes the regressors cross-product matrix. Similar to a nonlinear least-squares problem, the matrix G_{bb} depends on an additional term, which is denoted by $\frac{d[g(\widetilde{f}^*(S_1))]}{d\widetilde{f}}$. If $\frac{d[g(\widetilde{f}^*(S_1))]}{d\widetilde{f}} \geq 0$ for all S_1, $E_0^Q[G_{bb}]$ is a standard second-order cross-moment. In this case, $E_0^Q[G_{bb}]$ is a variance–covariance matrix, which is semidefinitive positive.

That is,

$$\frac{d[g(\widetilde{f}^*(S_1))]}{d\widetilde{f}} = \left(1 - C_2(S_1, \widetilde{f}^*(S_1)) \right) \times Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0))$$

$$+ \left(\left(\left(\widetilde{f}^*(S_1) - S_1 - K \right) - C(S_1, \widetilde{f}^*(S_1)) \right) \times \frac{d \left[Q(S_2 = \widetilde{f}^*(S_1)|S_1, S(t_0)) \right]}{d\widetilde{f}}.$$

The first term is positive (as $C_2 < 1$). The second term is close to zero on average (which follows from the FOCs), and is small if the local LSM exercise boundary is close to the true optimal boundary ($\widetilde{f}^*$ and f, respectively).

In particular, if a first-order approximation of $g(\widetilde{f}^*(S_1))$ at the point $f(S_1)$ is accurate (e.g., $\widetilde{f}^* = f$ or $(\widetilde{f}^* - f)^2$ is negligible),

$$g(\widetilde{f}^*(S_1)) \approx g(f(S_1)) + \frac{d[g(f(S_1))]}{d\widetilde{f}} \times \left(\widetilde{f}^*(S_1) - f(S_1) \right)$$

$$= \left(1 - C_2(S_1, f(S_1)) \right) \times Q(S_2 = f(S_1)|S_1, S(t_0)) \times \left(\widetilde{f}^*(S_1) - f(S_1) \right),$$

and then $g(\widetilde{f}^*(S_1))$ is increasing in $\widetilde{f}$ and $E_0^Q[G_{bb}]$ is semidefinitive positive.[17]

APPENDIX B: A RECURSIVE APPROACH AND STOCHASTIC INTEREST RATES

Let r_t be the short interest rate, where $D_j = e^{-\int_{t_0}^{t_j} r_u du}$ is the inverse of a bank account associated with Q. The cost of suboptimal exercise is given by

$$\min_{\widetilde{f}_{1:J}} E_0^Q \left[G \left(\widetilde{f}_{1:J} \right) \right] = \min_{\widetilde{f}_1, \widetilde{f}_{2:J}} C(t_0, S) - P(t_0, S, (\widetilde{f}_1, \widetilde{f}_{2:J}))$$

$$= \min_{\widetilde{f}_1, \widetilde{f}_{2:J}} E_0^Q \left[\int_{\widetilde{f}_1(S_1)}^{f_1(S_1)} D_1 \left(C(t_1, S) - (S_2 - S_1 - K) \right) \times Q(S_2|S_1, S(t_0)) \times dS_2 \right]$$

$$+ E_0^Q \left[\int_0^{\widetilde{f}_1(S_1)} D_1 \left(C(t_1, S) - P \left(t_1, S, \widetilde{f}_{2:J} \right) \right) \times Q(S_2|S_1, S(t_0)) \times dS_2 \right].$$

Note that only the second term depends on $\widetilde{f}_2$; that is,

$$
E_0^Q \left[\int_0^{\widetilde{f}_1(S_1)} D_1 \left(C(t_1, S) - P\left(t_1, S, \widetilde{f}_{2:J}\right) \right) \times Q(S_2(t_1) | S_1(t_1), S(t_0)) \times d S_2(t_1) \right]
$$

$$
= E_0^Q \left[\int_0^{\infty} 1_{\{0 \le S_2 \le \widetilde{f}_1(S_1)\}} D_1 \left(C(t_1, S) - P\left(t_1, S, \widetilde{f}_{2:J}\right) \right) \times Q(S_2(t_1) | S_1(t_1), S(t_0)) \times d S_2(t_1) \right]
$$

$$
= E_0^Q \left[1_{\{0 \le S_2 \le \widetilde{f}_1(S_1)\}} \times D_1 \times \left(C(t_1, S_1, S_2)) - P\left(t_1, (S_1, S_2), \widetilde{f}_{2:J}\right) \right) \right]
$$

$$
= E_0^Q \left[1_{\{0 \le S_2 \le \widetilde{f}_1(S_1)\}} \times D_1 \right.
$$

$$
\times E_1^Q \left[\int_{\widetilde{f}_2}^{f_2} \frac{D_2}{D_1} \left(C(t_2, S) - \left(S_2 - S_1 - K\right) \right) \times Q(S_2(t_2) | S_1(t_2), S(t_1)) \times d S_2(t_2) \right.
$$

$$
\left. \left. + \int_0^{\widetilde{f}_2} \frac{D_2}{D_1} \left(C(t_2, S) - P(t_2, S, \widetilde{f}_{3:J}) \right) \times Q(S_2(t_2) | S_1(t_2), S(t_1)) \times d S_2(t_2) \right] \right].
$$

It follows that for any period t_j, the cost that depends on $\widetilde{f}_j(S_1(t_j))$ is given by

$$
E_0^Q \left[\left(\prod_{i=1}^{j-1} 1_{\{0 \le S_2 \le \widetilde{f}_i(S_1)\}} \right) \times D_{j-1} \right.
$$

$$
\times E_{j-1}^Q \left[\frac{D_j}{D_{j-1}} \int_{\widetilde{f}_j}^{f_j} \left(C(t_j, S) - \left(S_2 - S_1 - K\right) \right) \times Q(S_2(t_j) | S_1(t_j), S(t_{j-1})) \times d S_2(t_j) \right.
$$

$$
\left. \left. + \frac{D_j}{D_{j-1}} \int_0^{\widetilde{f}_j} \left(C(t_j, S) - P(t_j, S, \widetilde{f}_{j+1:J}) \right) \times Q(S_2(t_j) | S_1(t_j), S(t_{j-1})) \times d S_2(t_j) \right] \right].
$$

By minimizing with regard to $\widetilde{f}_j(S_1(t_j))$, the FOCs are given by

$$
E_0^Q \left[\prod_{i=1}^{j-1} 1_{\{0 \le S_2 \le \widetilde{f}_i(S_1)\}} \times D_{j-1} \right.
$$

$$
\times E_{j-1}^Q \left[\frac{D_j}{D_{j-1}} \left(P\left(t_j, \left(S_1, \widetilde{f}_j^*(S_1)\right), \widetilde{f}_{j+1:J}\right) - \left(\widetilde{f}_j^*(S_1) - S_1 - K\right) \right) \frac{\partial \widetilde{f}_j(S_1, b^*)}{\partial b} \right.
$$

$$
\left. \left. Q(S_2(t_j) = \widetilde{f}_j^*(S_1) | S_1(t_j), S(t_{j-1})) \right] \right]
$$

$$= E_0^Q \left[\prod_{i=1}^{j-1} 1_{\{0 \leq S_2 \leq \widetilde{f}_i(S_1)\}} \times D_j \times \left(P\left(t_j, \left(S_1, \widetilde{f}_j^*(S_1)\right), \widetilde{f}_{j+1:J}\right) - \left(\widetilde{f}_j^*(S_1) - S_1 - K\right)\right)\right.$$

$$\left. \times \frac{\partial \widetilde{f}_j(S_1, b^*)}{\partial b} \times Q\left(S_2(t_j) = \widetilde{f}_j^*(S_1) | S_1(t_j), S(t_0)\right)\right] = 0.$$

If we enforce $1_{\{0 \leq S_2 \leq \widetilde{f}_i(S_1)\}} = 1$ for $i = \{1, 2, \ldots, j-1\}$, the FOCs associated with $\widetilde{f}_j^*$ are equivalent to the case of a Bermudan option that can be exercised only from t_j to t_J. And if interest rates are stochastic (assume r follows a one-factor Markovian model), the exercise boundary depends also on r; that is, $f_j(S_1(t_j), r_j)$. The FOCs are given by

$$E_0^Q \left[D_j \left(P\left(t_j, \left(S_1, \widetilde{f}_j^*(S_1, r)\right), r, \widetilde{f}_{j+1:J}\right) - \left(\widetilde{f}_j^*(S_1, r) - S_1 - K\right)\right)\right.$$

$$\left. \times \frac{\partial \widetilde{f}_j(S_1, r, b^*)}{\partial b} \times Q\left(S_2 = \widetilde{f}_j^*(S_1, r) | (S_1, r), (S(t_0), r(t_0))\right)\right] = 0.$$

APPENDIX C: CONVERGENCE OF LOCAL LSM ESTIMATION

In this appendix, we analyze the convergence properties of our new local LSM. Convergence of LSM algorithms has received much attention in the literature. Glasserman and Yu (2004) considered the case in which the basis functions are polynomials and the underlying process is either Brownian or geometric Brownian motion. Stentoft (2004b), however, assumes that the factors are bounded, possibly after rescaling. Both find upper bounds on the rate of increment of the number of approximating functions with the number of paths. Egloff (2005) evaluates the convergence within a generalized statistical learning problem, whereas Tsitsiklis and Van Roy (2001) and Clément, Lamberton, and Protter (2002) consider a similar approach but holding the number of basis fixed, and Belomestny (2011a) and Belomestny et al. (2015) consider local polynomial kernel regression instead of global series approximation.

Belomestny (2011b) obtains optimal convergence rates for the estimate of the continuation-value function, and studies the optimal relation between the number of paths in the first and second simulations to estimate the optimal policy and the option price, respectively. Zanger (2016) provides error estimates when a single set of simulated sample paths is used for the estimation of the continuation-value function, including linear and nonlinear nonconvex approximations.

Our case differs from all previous analyses in that we add to the nonparametric series estimation of the boundary or continuation value a further kernel smoothing that locates the exercise boundary by using information on previous estimates. The bandwidth controlling this additional smoothing is simultaneously converging to zero with the increasing number of basis and paths, allowing consistent estimation of the exercise boundary either directly or by inverting the value-matching condition. We also study in detail the effect of starting from preliminary continuation value estimates obtained by LSM, which are consistent for any combination of factors.

First, we adapt the results of Newey (1997) on consistency of series estimators, to analyze the convergence rate of the exercise boundary estimate, which is based on the same localization principle as the local continuation value (see Theorem 4.3). Second, we use Stentoft's (2004b) arguments, which depend on a regularity condition to identify the corresponding stopping time—cf. Assumption C.1(ii) below—to justify the convergence of option price estimates using our exercise-boundary estimates (see Theorem 4.5).

Exercise Boundary. We first analyze a local series regression to estimate the parameters b in a parameterization $\tilde{f}(S_1; b) = p^H(S_1)'b$ of the exercise boundary $\tilde{f}(S_1)$, cf. (3.6), with $p^H(S_1) = (p_1^H(S_1), \ldots, p_H^H(S_1))'$ being a vector of H approximating functions (given, e.g., by a set of orthonormal polynomials), so that $\partial \tilde{f}(S_1, b)/\partial b = p^H(S_1)$. For each path we construct an estimate of $\tilde{f}(S_1)$ based on $\hat{C} = \hat{C}(S_1, S_2)$, an unbiased estimate of $C(S_1, S_2)$, and on the value-matching condition, $C(S_1, S_2) = I(S_1, \tilde{f}(S_1))$. We assume the domain of S can be divided into a finite number of regions such that in each one, the intrinsic value is linear in the component S_2,

$$I(S_1, S_2) = c_1(S_1) + c_2 S_2, \tag{C.1}$$

where the function c_1 and the scalar c_2 are known and might differ from region to region with $S_1 \in \mathbb{R}^{N-1}$. Because $I(S_1, \tilde{f}(S_1))$ is linear in $\tilde{f}$, (C.1) implies $Y = (\hat{C} - c_1(S_1))/c_2$ is an unbiased estimate of $\tilde{f}(S_1)$ when $S_2 = \tilde{f}(S_1)$, and motivates the local regression of

$$Y^{(m)} = \left(\hat{C}^{(m)} - c_1\left(S_1^{(m)} \right) \right)/c_2 \quad \text{on } p^H\left(S_1^{(m)} \right), \quad m = 1, \ldots, M,$$

when $S_2 \approx \tilde{f}(S_1)$ to approximate $\tilde{f}(S_1)$. Note $E[Y|S_1, S_2 = \tilde{f}(S_1)] = \tilde{f}(S_1)$ as $E[\hat{C}|S_1, S_2] = C(S_1, S_2)$.

We focus on estimation of a particular exercise boundary, and we omit in the notation reference to it and to the specific time period in the backward induction of LSM. If S is redefined in an appropriate way, reducing the number of regions to just one for particular problems would be possible, for example, by replacing S in C and I by $S^* = (S_{(1)}, S_{(2)}) = (S_{\min}, S_{\max})$ in a two-asset max-option. In case of symmetric regions and exercise boundaries, these restrictions could be incorporated in the least-squares fit, increasing efficiency of estimates.

Consider the local-series estimate at iteration $i = 0, 1, 2, \ldots,$

$$\hat{b}_i = \left(\mathbf{P}' \hat{\mathcal{K}}_h^{(i-1)} \mathbf{P} \right)^{-1} \mathbf{P}' \hat{\mathcal{K}}_h^{(i-1)} \mathbf{Y} = \left(\sum_{m=1}^M \mathcal{K}_h\left(S_2^{(m)} - \tilde{f}\left(S_1^{(m)}, \hat{b}_{i-1} \right) \right) p^H\left(S_1^{(m)} \right) p^H\left(S_1^{(m)} \right)' \right)^{-1}$$

$$\times \sum_{m=1}^M \mathcal{K}_h\left(S_2^{(m)} - \tilde{f}\left(S_1^{(m)}, \hat{b}_{i-1} \right) \right) p^H\left(S_1^{(m)} \right) Y^{(m)},$$

where $\mathcal{K}_h(\cdot) = h^{-1}\mathcal{K}(h^{-1}\cdot)$ and $\mathcal{K}$ is a positive, symmetric, continuously differentiable with square integrable derivative $\dot{\mathcal{K}}$, kernel function with $\inf_{S_1} \mathcal{K}(S_1) \geq \epsilon > 0$ and $\int_{S_1} \mathcal{K}(S_1) dS_1 = 1$, where $S \in \mathcal{S} = S_1 \times \cdots \times S_N$, all compact intervals and $\hat{\mathcal{K}}_h^{(i-1)} = diag\{\mathcal{K}_h(S_2^{(m)} - \tilde{f}(S_1^{(m)}, \hat{b}_{i-1}))\}_{m=1}^M$, with $\mathbf{Y} = (Y^{(1)}, \ldots, Y^{(M)})'$ and $\mathbf{P} = (p^H(S_1^{(1)}), \ldots, p^H(S_1^{(M)}))'$ stacking all paths information. The bandwidth parameter h tends toward zero with M, simultaneously with an increasing H, to approximate $\tilde{f}$ as sample information gets larger.

For the initial iteration, $i = 0$, we can use any previous estimate of the boundary function $\tilde{f}$ based on estimates $\hat{b}_0$, or, starting from the LSM algorithm, we can use the value-matching condition to define

$$\hat{b}_1 = \left(\sum_{m=1}^M \mathcal{K}_h\left(\hat{C}\left(S^{(m)} \right) - I\left(S^{(m)} \right) \right) p^H\left(S_1^{(m)} \right) p^H\left(S_1^{(m)} \right)' \right)^{-1}$$

$$\times \sum_{m=1}^M \mathcal{K}_h\left(\hat{C}\left(S^{(m)} \right) - I\left(S^{(m)} \right) \right) p^H\left(S_1^{(m)} \right) Y^{(m)},$$

where $C(S^{(m)}) \approx I(S^{(m)})$ defines an equivalent neighborhood to $S_2^{(m)} \approx \tilde{f}(S_1^{(m)})$. Typically, $\hat{C}(S^{(m)}) = C(S^{(m)}, \hat{\delta}^n) = p^n(S^{(m)})' \hat{\delta}^n$ is an (order n series) estimate as in the LSM algorithm.

For concreteness, we work with shifted orthonormal Legendre polynomials, with support in $(0, 1)$, resulting after normalization of factors, as in Stentoft (2004b). Then $\sup_{S_1 \in S_1} \|p^H(S_1)\| \leq \zeta_0$ $(H) \sim H$. The following conditions follow assumptions C.1 to C.3 in Stentoft (2004b), who shows they are sufficient for Newey's (1997) original assumptions when using a series estimate based on shifted Legendre polynomials. We add smoothness conditions on both C and $\tilde{f}$, as we obtain $\tilde{f}$ as a restriction of C around a neighborhood defined by $\tilde{f}$ itself.

Assumption C.1.

(i) *The simulated paths, $S^{(m)}$, $m = 1, \ldots, M$, are independent.*

(ii) $Pr(I(t_j, S) = C(t_j, S)) = 0$, $j = 1, \ldots, J$.

Assumption C.2. *The support S of S is a Cartesian product of compact connected intervals on which S has a probability density Q that is bounded away from zero, $Q(S) \geq \epsilon > 0$, and is twice continuously differentiable, and the conditional density of S_2 given S_1 is also bounded away from zero, $Q(S_2|S_1) \geq \epsilon > 0$.*

Assumption C.3. $C(S_1, S_2)$ *and* $\tilde{f}(S_1)$ *are continuously differentiable of order $s \geq 2N$ on the support of S.*

Assumption C.4. *As $M \to \infty$, for $\alpha = s/N$ and assuming $\zeta_0(H) \sim H$,*

$$\frac{H^3}{Mh} + \frac{1}{hH^{2\alpha}} + h^2 H + \frac{H}{Mh^2} \to 0.$$

If $\lambda_{\min}$ denotes the smallest eigenvalue of a matrix, Assumption C.2 implies the following regularity condition:

$$\lambda_{\min} \left(E \left[\mathcal{K}_h \left(S_2 - \tilde{f} \left(S_1 \right) \right) p^H(S_1) p^H(S_1)' \right] \right) \geq \epsilon^2 > 0 \tag{C.2}$$

uniformly in H and h, using that $\int \mathcal{K}_h(S_2 - \tilde{f}(S_1)) dS_2 = 1$, which is needed to control the behavior of the series least-squares estimates. Note also that Assumption C.2 implies the left-hand side of (C.2) converges for $h \to 0$ to

$$\lambda_{\min} \left(\int p^H(S_1) p^H \left(S_1 \right)' Q \left(S_1, \tilde{f}(S_1) \right) dS_1 \right) = \lambda_{\min} \left(\tilde{Q}_1 E \left[p^H \left(S_1 \right) p^H(S_1)' | S_2 = \tilde{f} \left(S_1 \right) \right] \right),$$

where $\tilde{Q}_1 = \int Q(S_1, \tilde{f}(S_1)) dS_1$, so (C.2) is a conditional version of the original condition stated by Newey (1997) without $\mathcal{K}_h(S_2 - \tilde{f}(S_1))$. Assumption C.3 guarantees a fast-converging series representation of the exercise boundary; that is, for $\tilde{f}(S_1) = E[Y|S_1, S_2 = \tilde{f}(S_1)]$, a $b^{(H)}$ exists such that $|\tilde{f} - p^{H'} b^{(H)}|_0 = O(H^{-\alpha})$ as $H \to \infty$, for $\alpha = s/N > 1$, where $|g|_0 = \sup_{S_1 \in S_1} |g(S_1)|$. Assumption C.4 imposes some important restrictions on the order H of the polynomial basis and the bandwidth h of the kernel $\mathcal{K}_h$. These conditions result basically from reducing the effective number of paths from M to Mh in Newey's conditions.

Continuation Value. In this case, the regressions are local versions of the LSM algorithm, fitting directly the continuation value. The exercise boundary is now only defined implicitly by the value-matching condition. We need to generalize our setup and assumptions to include in p^H all elements

in S, not just S_1, which are used to approximate the continuation value $C(S_1, S_2)$, which is the new target in the estimation. Then the identification condition (C.2) has to be replaced by

$$\lambda_{\min} \left(E \left[\mathcal{K}_h \left(S_2 - \tilde{f} \left(S_1 \right) \right) p^H(S) p^H(S)' \right] \right) \geq \varepsilon > 0 \ \forall H, h. \tag{C.3}$$

Note that (C.3) requires that $\lambda_{\min}(\int p^H(S_1, \tilde{f}(S_1)) p^H(S_1, \tilde{f}(S_1))' Q(S_1, \tilde{f}(S_1)) dS_1) \geq \varepsilon > 0$, $\forall H$, implying $\tilde{f}(S_1)$ is not in the span of functions given by $p^H(S_1)$, for example. polynomials of finite order. If continuation values are estimated to make exercise decisions, the consistency of price estimates requires some control on the extrapolation to other regions, because the approximation of the continuation value is not uniform in S, but only consistent around the exercise boundary. This result can be achieved by using a fixed smoothing in the asymptotics, the local estimation giving more weight to paths closer to the exercise boundary, but maintaining uniform convergence for all possible paths. An alternative route is to ensure that the derivatives of the continuation-value function estimates are in the $(0, 1)$ interval in case of simple parameterizations such as a quadratic fitting, which has at most one critical point in the relevant range. Then similar results to those in Theorem 4.3 and Corollary 4.4 would be obtained.

Proof of Theorem 4.3. Consider first the unfeasible estimate with the true $\tilde{f}$ inside $\mathcal{K}_h$,

$$\tilde{b}_1 = (\mathbf{P}' \mathcal{K}_h \mathbf{P})^{-1} \mathbf{P}' \mathcal{K}_h \mathbf{Y},$$

where $\mathcal{K}_h = diag\{\mathcal{K}_h^{(m)}\}_{m=1}^M$ and $\mathcal{K}_h^{(m)} = \mathcal{K}_h(S_2^{(m)} - \tilde{f}(S_1^{(m)}))$. We assume without loss of generality that $E[\mathcal{K}_h(S_2 - \tilde{f}(S_1)) p^H(S_1) p^H(S_1)'] = \mathbf{I}_H$, with $\mathbf{I}_H$ as the H-th order identity matrix, because its smallest eigenvalue is bounded away from zero uniformly in H and h. Define $\tilde{\mathbf{R}} = \mathbf{P}' \mathcal{K}_h \mathbf{P} / M$; then

$$E \left[\left\| \tilde{\mathbf{R}} - \mathbf{I}_H \right\|^2 \right] = \sum_{k=1}^H \sum_{j=1}^H E \left[\left\{ \frac{1}{M} \sum_{m=1}^M p_k^H \left(S_1^{(m)} \right) p_j^H \left(S_1^{(m)} \right) \mathcal{K}_h^{(m)} - I_{jk} \right\}^2 \right]$$

$$\leq \frac{1}{M} \sum_{k=1}^H \sum_{j=1}^H E \left[p_k^H \left(S_1 \right)^2 p_j^H (S_1)^2 \mathcal{K}_h^2 \right]$$

$$\leq \frac{1}{Mh} E \left[\sum_{k=1}^H p_k^H (S_1)^2 \sum_{j=1}^H p_j^H (S_1)^2 \mathcal{K}_h \right]$$

$$\leq \frac{1}{Mh} \zeta_0 (H)^2 E \left[\sum_{j=1}^H p_j^H (S_1)^2 \mathcal{K}_h \right]$$

$$\leq \frac{1}{Mh} \zeta_0 (H)^2 H \sim \frac{H^3}{Mh}$$

so that $\|\tilde{\mathbf{R}} - \mathbf{I}_H\| = O_p(H^{1/2}(Mh)^{-1/2} \zeta_0(H)) = O_p(H^{3/2}(Mh)^{-1/2}) = o_p(1)$, by Assumption C.4, so the smallest eigenvalue of $\tilde{\mathbf{R}}$ converges to 1.

Define $1_M = I(\underline{\lambda}(\tilde{\mathbf{R}}) > \frac{1}{2}) = 1 + o_p(1)$. Then, given that the elements of $E[\varepsilon \varepsilon' | S]$ are bounded for $\varepsilon = \mathbf{Y} - \mathbf{G}$, $\mathbf{G} = (g(S^{(1)}), \dots, g(S^{(M)}))'$, $g(S) = E[Y | S]$,

$$E \left[1_M \left\| \tilde{\mathbf{R}}^{-1/2} \mathbf{P}' \mathcal{K}_h \varepsilon / M \right\|^2 | S \right] = 1_M E \left[\varepsilon' \mathcal{K}_h \mathbf{P} \left(\mathbf{P}' \mathcal{K}_h \mathbf{P} \right)^{-1} \mathbf{P}' \mathcal{K}_h \varepsilon | S \right] / M$$

$$= 1_M tr \left\{ \mathcal{K}_h \mathbf{P} \left(\mathbf{P}' \mathcal{K}_h \mathbf{P} \right)^{-1} \mathbf{P}' \mathcal{K}_h E \left[\varepsilon \varepsilon' | S \right] \right\} / M$$

$$\leq C_0 h^{-1} 1_M tr \left\{ \left(\mathbf{P}' \mathcal{K}_h \mathbf{P} \right)^{-1} \mathbf{P}' \mathcal{K}_h \mathbf{P} \right\} / M$$

$$\leq C_0 H (Mh)^{-1},$$

so that $\| \tilde{\mathbf{R}}^{-1/2} \mathbf{P}' \mathcal{K}_h \varepsilon / M \| = O_p(H^{1/2}(Mh)^{-1/2})$, and it follows that

$$1_M \left\| \tilde{\mathbf{R}}^{-1} \mathbf{P}' \mathcal{K}_h \varepsilon / M \right\| \leq 1_M \left\{ \left[\left(\varepsilon' \mathcal{K}_h \mathbf{P} / M \right) \tilde{\mathbf{R}}^{-1/2} \tilde{\mathbf{R}}^{-1/2} \left(\mathbf{P}' \mathcal{K}_h \varepsilon / M \right) \right] \right\}^{1/2}$$

$$= O_p(1) 1_M \left\| \tilde{\mathbf{R}}^{-1/2} \mathbf{P}' \mathcal{K}_h \varepsilon / M \right\| = O_p \left(H^{1/2} (Mh)^{-1/2} \right).$$

Now, writing $\tilde{\mathbf{F}} = (\tilde{f}(S_1^{(1)}), \dots, \tilde{f}(S_1^{(M)}))'$, with the matrix $1_M \mathbf{P}(\mathbf{P}' \mathcal{K}_h \mathbf{P})^{-1} \mathbf{P}' \mathcal{K}_h$ idempotent, $1_M \| \tilde{\mathbf{R}}^{-1} \mathbf{P}' \mathcal{K}_h (\tilde{\mathbf{F}} - \mathbf{P}b^{(H)}) / M \|$ is bounded by

$$O_p(1) 1_M \left(\left(\tilde{\mathbf{F}} - \mathbf{P}b^{(H)} \right)' \mathcal{K}_h \mathbf{P} \left(\mathbf{P}' \mathcal{K}_h \mathbf{P} \right)^{-1} \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{P}b^{(H)} \right) / M \right)^{1/2}$$

$$= O_p \left(h^{-1/2} \right) 1_M \left(\left(\tilde{\mathbf{F}} - \mathbf{P}b^{(H)} \right)' \left(\tilde{\mathbf{F}} - \mathbf{P}b^{(H)} \right) / M \right)^{1/2}$$

$$= O_p(h^{-1/2} H^{-\alpha}).$$

Then $1_M \| \tilde{\mathbf{R}}^{-1} \mathbf{P}' \mathcal{K}_h (\tilde{\mathbf{F}} - \mathbf{G}) / M \| = O_p(1) 1_M \| \mathbf{P}' \mathcal{K}_h (\tilde{\mathbf{F}} - \mathbf{G}) / M \|$, and

$$\left\| \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M \right\| \leq \left\| \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M - E \left\{ \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M \right\} \right\| + \left\| E \left\{ \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M \right\} \right\|,$$

where

$$\left\| E \left[\mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M \right] \right\| = O \left(h^2 \zeta_0 (H) \right) = O \left(h^2 H \right). \tag{C.4}$$

Bound (C.4) follows because $\tilde{f}(S_1) = E[g(S) | S_2 = \tilde{f}(S_1)]$ and $g(S) = E[Y|S]$, so that $E[\mathbf{P}' \mathcal{K}_h(\tilde{\mathbf{F}} - \mathbf{G})/M]$ is

$$\int \mathcal{K}_h \left(S_2 - \tilde{f}(S_1) \right) \left\{ g \left(S_1, \tilde{f}(S_1) \right) - g \left(S_1, S_2 \right) \right\} p^H(S_1) Q \left(S_2 | S_1 \right) Q(S_1) dS,$$

and we use that

$$\left| \int \mathcal{K}_h \left(S_2 - \tilde{f}(S_1) \right) \left\{ g \left(S_1, \tilde{f}(S_1) \right) - g \left(S_1, S_2 \right) \right\} Q \left(S_2 | S_1 \right) dS_2 \right|$$

$$\leq C_0 \int \left(S_2 - \tilde{f}(S_1) \right)^2 \mathcal{K}_h \left(S_2 - \tilde{f}(S_1) \right) dS_2 = O \left(h^2 \right),$$

uniformly in S_1, by a second-order Taylor expansion of $g(S_1, S_2) Q(S_2 | S_1)$ around $S_2 = \tilde{f}(S_1)$, and $\int \| p^H(S_1) \| Q(S_1) dS_1 = O(H)$.

On the other hand,

$$E \left\| \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M - E \left\{ \mathbf{P}' \mathcal{K}_h \left(\tilde{\mathbf{F}} - \mathbf{G} \right) / M \right\} \right\|^2$$

$$\leq \frac{1}{M} E\left[p^H(S_1)' p^H(S_1)\, \mathcal{K}_h\left(S_2 - \tilde{f}(S_1)\right)^2 \left\{\tilde{f}(S_1) - g(S)\right\}^2\right]$$

$$\leq \frac{\zeta_0(H)^2}{M} E\left[\mathcal{K}_h\left(S_2 - \tilde{f}(S_1)\right)^2 \left\{\tilde{f}(S_1) - g(S)\right\}^2\right]$$

$$\leq \frac{\zeta_0(H)^2 h}{M} = O\left(\frac{H^2 h}{M}\right),$$

so that $\|\mathbf{P}'\mathcal{K}_h(\tilde{\mathbf{F}} - \mathbf{G})/M - E\{\mathbf{P}'\mathcal{K}_h(\tilde{\mathbf{F}} - \mathbf{G})/M\}\| = O_p(H(h/M)^{1/2})$, and together with (C.4), they imply

$$\left\|\mathbf{P}'\mathcal{K}_h\left(\tilde{\mathbf{F}} - \mathbf{G}\right)/M\right\| = O_p\left(h^2 H + H(h/M)^{1/2}\right).$$

Then, collecting terms, it follows that

$$1_M\left\|\tilde{b}_1 - b^{(H)}\right\| \leq 1_M\left\|\tilde{\mathbf{R}}^{-1}\mathbf{P}'\mathcal{K}_h\varepsilon/M\right\| + 1_M\left\|\tilde{\mathbf{R}}^{-1}\mathbf{P}'\mathcal{K}_h\left(\tilde{\mathbf{F}} - \mathbf{P}b^{(H)}\right)/M\right\|$$

$$+1_M\left\|\tilde{\mathbf{R}}^{-1}\mathbf{P}'\mathcal{K}_h\left(\tilde{\mathbf{F}} - \mathbf{G}\right)/M\right\|$$

$$= O_p\left(H^{1/2}(Mh)^{-1/2} + h^{-1/2}H^{-\alpha} + h^2 H + \zeta_0(H)(h/M)^{1/2}\right)$$

$$= O_p\left(H^{1/2}(Mh)^{-1/2} + h^{-1/2}H^{-\alpha} + h^2 H\right) = o_p(1)$$

by Assumption C.4.

Now we wish to bound $\|\tilde{b}_1 - \hat{b}_1\|$. We first study $\|\tilde{\mathbf{R}} - \hat{\mathbf{R}}\|$. Then for points $C_*^{(m)}$ between $C(S^{(m)})$ and $\hat{C}(S^{(m)})$, we apply the mean value theorem and Cauchy–Schwarz inequality,

$$\left\|\tilde{\mathbf{R}} - \hat{\mathbf{R}}\right\| \leq \zeta_0(H)^2 \frac{1}{M}\sum_{m=1}^{M}\left|\mathcal{K}_h\left(C\left(S^{(m)}\right) - I\left(S^{(m)}\right)\right) - \mathcal{K}_h\left(\hat{C}\left(S^{(m)}\right) - I\left(S^{(m)}\right)\right)\right|$$

$$= O_p\left(h^{-3/2}\zeta_0(H)^2\right)\left(\frac{1}{Mh}\sum_{m=1}^{M}\left|\dot{\mathcal{K}}\left(C_*^{(m)}h^{-1}\right)\right|^2 \frac{1}{M}\sum_{m=1}^{M}\left|C\left(S^{(m)}\right) - \hat{C}\left(S^{(m)}\right)\right|^2\right)^{1/2}$$

$$= O_p\left(h^{-3/2}H^2\xi_M^{1/2}\right) = o_p(1), \tag{C.5}$$

by (4.4), noting that $\frac{1}{M}\sum_{m=1}^{M}|C(S^{(m)}) - \hat{C}(S^{(m)})|^2 = \int_S |C(S^{(m)}) - \hat{C}(S^{(m)})|^2 dQ_M(S) = O_p(1)\int_S |C(S^{(m)}) - \hat{C}(S^{(m)})|^2 Q(S)dS = O_p(\xi_M)$, where dQ_M is the empirical cdf of the paths $S^{(m)}$ and $\dot{\mathcal{K}}$ is the first derivative of $\mathcal{K}$, so that $1_M^* = I(\underline{\lambda}(\tilde{\mathbf{R}}) > \frac{1}{2})I(\underline{\lambda}(\hat{\mathbf{R}}) > \frac{1}{2}) = 1 + o_p(1)$.

Next, $E[1_M^*\|\hat{\mathbf{R}}^{-1/2}\mathbf{P}'(\mathcal{K}_h - \hat{\mathcal{K}}_h)\varepsilon/M\|^2|S]$ is

$$1_M^* E\left[\varepsilon'\left(\mathcal{K}_h - \hat{\mathcal{K}}_h\right)\mathbf{P}\left(\mathbf{P}'\hat{\mathcal{K}}_h\mathbf{P}\right)^{-1}\mathbf{P}'\left(\mathcal{K}_h - \hat{\mathcal{K}}_h\right)\varepsilon|S\right]/M$$

$$= 1_M^* tr\left\{\left(\mathcal{K}_h - \hat{\mathcal{K}}_h\right)\mathbf{P}\left(\mathbf{P}'\hat{\mathcal{K}}_h\mathbf{P}\right)^{-1}\mathbf{P}'\left(\mathcal{K}_h - \hat{\mathcal{K}}_h\right)E\left[\varepsilon\varepsilon'|S\right]\right\}/M$$

$$\leq C_0 h^{-2} 1_M^* tr\left\{\left(\mathbf{P}'\hat{\mathcal{K}}_h\mathbf{P}\right)^{-1}\mathbf{P}'\mathbf{P}\right\}/M$$

$$\leq C_0 h^{-2} H/M,$$

so that $1^*_M \|\hat{\mathbf{R}}^{-1/2} \mathbf{P}'(\mathcal{K}_h - \hat{\mathcal{K}}_h)\varepsilon/M\| = O_p(h^{-1}(H/M)^{1/2})$.

Finally

$$\left\| \mathbf{P}'\left(\mathcal{K}_h - \hat{\mathcal{K}}_h\right)\mathbf{G}/M \right\| \leq \zeta_0(H) \frac{1}{M} \sum_{m=1}^{M} \left| \left(\mathcal{K}_h\left(S^{(m)}\right) - \hat{\mathcal{K}}_h\left(S^{(m)}\right)\right)\right|$$

$$= O_p\left(H h^{-3/2} \xi_M^{1/2}\right),$$

reasoning as in the proof of (C.5). Then $\|\hat{b}_1 - \tilde{b}_1\| = O_p(H^2 h^{-3/2} \xi_M^{1/2} + h^{-1}(H/M)^{1/2}) = o_p(1)$ by condition (4.4) and Assumption C.4.

Therefore,

$$1^*_M \int \left[p^H(S_1)\left(\tilde{b}_1 - b^{(H)}\right) + p^H(S_1)\left(\hat{b}_1 - \tilde{b}_1\right) + p^H(S_1)'b^{(H)} - \tilde{f}(S_1)\right]^2 Q(S_1) dS_1$$

$$\leq 1^*_M \left\| \tilde{b}_1 - b^{(H)} \right\|^2 + 1^*_M \left\| \hat{b}_1 - \tilde{b}_1 \right\|^2 + 1^*_M \int \left[p^H(S_1)'b^{(H)} - \tilde{f}(S_1)\right]^2 Q(S_1) dS_1$$

$$= O_p(H(Mh)^{-1} + h^{-1}H^{-2\alpha} + \left(h^2 H\right)^2) + O_p\left(H^4 h^{-3} \xi_M + H(Mh^2)^{-1}\right)$$

$$= o_p(1), \tag{C.6}$$

under Assumption C.4, and we conclude $\int [\tilde{f}(S_1, \hat{b}_1) - \tilde{f}(S_1)]^2 Q(S_1) dS_1 = o_p(1)$ as $1^*_M \to 1$ with probability approaching one, and the proof is completed. □

Proof of Corollary 4.4. The proof would proceed similarly, but to compare the feasible and unfeasible estimates, $\tilde{b}_1$ and $\hat{b}_1$, respectively, evaluate (C.6) with ξ_M replaced by the bound given in Theorem 4.3 for the L^2 error of $\tilde{f}$, which itself depends on ξ_M of the original global estimate of C. □