Interfaces with Other Disciplines

Recursive lower and dual upper bounds for Bermudan-style options☆

Alfredo Ibáñez [a,*], Carlos Velasco [b]

[a] *Universidad Pontificia Comillas, ICADE. c/ Alberto Aguilera 23, 28015 Madrid, Spain*
[b] *Department of Economics, Universidad Carlos III de Madrid. Calle Madrid, 126., Getafe, E-28903 Madrid, Spain*

A R T I C L E   I N F O

A B S T R A C T

Although Bermudan options are routinely priced by simulation and least-squares methods using lower and dual upper bounds, the latter are hardly optimized. In this paper, we optimize recursive upper bounds, which are more tractable than the original/nonrecursive ones, and derive two new results: (1) An upper bound based on (a martingale that depends on) stopping times is independent of the next-stage exercise decision and hence cannot be optimized. Instead, we optimize the recursive lower bound, and use its optimal recursive policy to evaluate the upper bound as well. (2) Less time-intensive upper bounds that are based on a continuation-value function only need this function in the continuation region, where this continuation value is less nonlinear and easier to fit (than in the entire support). In the numerical exercise, both upper bounds improve over state-of-the-art methods (including standard least-squares and pathwise optimization). Specifically, the very small gap between the lower and the upper bounds derived in (1) implies the recursive policy and the associated martingale are near optimal, so that these two specific lower/upper bounds are hard to improve, yet the upper bound is tighter than the lower bound.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Pricing Bermudan options in high dimensions requires Monte Carlo methods, and two simulation-based prices have been developed: lower and dual upper bounds. Specifically, Longstaff and Schwartz (2001) use a standard least-squares Monte Carlo (LSM) approach to compute lower bounds. Likewise, upper bounds are also based on least squares and simulation (Andersen & Broadie, 2004). Although both bounds are widely used, upper bounds are hardly optimized, which is important because simulation is time consuming, demanding a smart approach.

In this paper, we optimize (the more tractable) recursive upper bounds and provide two new results. Lower/upper bounds generated by simulation depend on an exercise policy, whereby the upper bound is derived from a martingale based on this policy. First, we show a recursive upper bound is independent of the next-stage

exercise decision and hence cannot be optimized. Therefore, we optimize the recursive lower bound, following Ibáñez and Velasco (2018) local LSM approach, and use its optimal recursive policy to evaluate the upper bound as well. We find these two bounds, which have a similar cost to the reciprocal bounds based on a policy estimated by the standard LSM method, are very tight. Second, we study separately an upper bound generated from a martingale based on continuation-value functions, a bound that is less time intensive yet more upper biased, and show how to reduce its bias as well.

In our first approach, we consider a given family of exercise policies—or stopping times. Ibáñez and Velasco maximize a recursive lower bound—or Bermudan price, L, with regard to this family at each exercise stage. An open question is which exercise strategy minimizes the upper bound, U. We show the exercise strategy that maximizes a recursive lower bound also minimizes not the recursive upper bound itself, but rather the gap between them, $U - L$. We provide a recursive expression for the gap (Theorem 1), and show a recursive upper bound U is independent of the next-stage exercise policy (Proposition 1). Therefore, minimizing the gap, $U - L$, is equivalent to maximizing the Bermudan price, L, recursively.

In the second approach, we consider a family of continuation-value functions. We show (i) a recursive upper bound is independent of the next-stage continuation-value function as well

(Proposition 1). (ii) By factorizing the two martingales that are based on either stopping times or continuation values, the latter martingale includes a third error term, which ensures the process is actually a martingale yet implies more biased upper bounds. The other two terms of the martingale are those of the standard factorization of the American option into an early-exercise premium and the European counterpart (Carr, Jarrow, & Myneni, 1992). This third term, however, depends only on the option continuation value in the waiting/continuation region.

The latter constraint is critical because Bermudan options are highly nonlinear near the exercise boundary but less so in the waiting region, and fitting a continuation-value function only in this region is easier. This new upper bound, based on a continuation value estimated only in the waiting region, is as accurate as an upper bound based on an exercise policy estimated by the LSM method (Andersen & Broadie, 2004), but in a fraction of the time. The new bound is especially accurate for at-/in-the-money options, which depend mostly on sample paths that cross the exercise region and that do not contribute to the martingale's third error term. This dual waiting-region constraint is the reciprocal constraint of using in-the-money paths, and hence, the exercise region, to estimate the continuation value in the LSM/primal method.

In the numerical exercise, we price up-and-out Bermudan max-options (Desai, Farias, & Moallemi, 2012). The up-and-out barrier makes this option very sensitive to suboptimal exercise, providing a good test. From the local LSM method (Ibáñez & Velasco, 2018), we derive the optimal recursive exercise policy and compute the two bounds associated with this policy: the lower bound improves upon the reciprocal bounds based on the standard LSM and pathwise optimization (Desai et al., 2012) by more than 100–200 cents; the upper bound yields a one-digit gap. This small gap implies the recursive policy and the associated martingale are near optimal and the two bounds are close to the true price. The local policy is so good that reducing the number of subsimulation paths by 20 (i.e., 5% of the original subsimulations) decreases the time effort by a factor of 10, yet the upper bound increases only by a few cents.

Notably, the upper bound based on the local-LSM policy only changes marginally with the number of subsimulations and is robust to all refinements, implying the upper bound is tighter and closer to the true price than the lower bound. With other methods (e.g., the standard LSM) that yield a nontrivial gap, this claim cannot be made. This result agrees with the two-period Bermudan upper bound, which is independent of the (one-period) exercise policy. A tighter upper bound implies a mid point between lower and upper bounds is lower biased.

The duality approach in option pricing (Haugh & Kogan, 2004; Rogers, 2001) has been extended in many ways. Chen and Glasserman (2007) and Rogers (2010) study optimal dual bounds; Belomestny, Schoenmakers, and Dickmann (2013) use a multilevel approach; Glasserman (2004) studies dual bounds based on regression methods; Christensen (2014) and Bhim and Kawai (2018) derive upper bounds using linear and semidefinite programming. Desai et al. (2012) use a pathwise-optimization approach that is less time consuming. In a novel extension, Brown, Smith, and Sun (2010) uses an information relaxation that nests the perfect-information assumption of the martingale approach, and also applies to other problems in operations research (Balseiro & Brown, 2019).

We tailor Haugh–Kogan and Andersen–Broadie results to our optimal recursive setting, which yields such tight bounds. Specifically, in the former case based on continuation values, we improve the upper bound by computing the continuation value only in the waiting region. In the latter case based on stopping times, the upper bound is both tight and efficient if we use fewer subsimulation paths, but a near-optimal exercise strategy as in the local LSM approach. In both ways, we bring the overall cost of the martingale approach in line with pathwise optimization or information relaxation. Moreover, the factorization of the dual martingales in terms of the components of the Bermudan process is mostly new. Our tight bounds are an useful benchmark for new methods that try to improve upper bounds in terms of accuracy or time effort.

Section 2 reviews the local LSM method and explains the exercise policies and continuation values needed later for dual bounding; Section 3 shows the independence of the recursive upper bound on the next-stage exercise policy; Section 4 shows an upper bound based on a continuation-value function only needs this function in the waiting region; Section 5 provides the complexity analysis and examples; Section 6 concludes. Proofs are left to the Appendix.

2. Bermudan options: a local Least-Squares setting

In this section, we explain the difference between the standard and the local LSM methods. Because the local approach yields an optimal recursive exercise policy, it is our method of choice to derive both the lower bound and the martingale associated with the upper bound. We next discuss precisely how we compute various exercise policies and continuation-value functions needed later for dual bounding. Lastly, we define the lower and the upper bound.

Consider a Bermudan option that can be exercised at $t \in \{1, 2, \ldots, T\}$, where $t = 0$ is today. We denote by $I_t \geq 0$ the intrinsic value (or option payoff) if the Bermudan option is exercised at t. Consider a vector of N stock prices S_t. Interest rates are stochastic, $R_t > 0$ is a bank-account process, $R_0 = 1$, and $R_{j,t} = R_t / R_j$. If the interest rate r is constant, $R_t = e^{rt\Delta t}$, $R_{t,t+1} = e^{r\Delta t}$, and Δt is the time between t and $t + 1$.

We introduce two binary auxiliary processes, Y and b; $b_0 = 1$ and

$$Y_t \in \{0, 1\} \quad \text{and} \quad b_t = b_{t-1} \times Y_t, \ t = 1, 2, \ldots, T. \tag{1}$$

Both processes are used in the case of barrier options (e.g., securities subject to default risk). In the numerical exercise, in which we study an up-and-out max-option, $Y_t = 1_{\{\max \{S_t\} < B\}}$ where B is the up-and-out barrier (and the no barrier case is equivalent to assume that $Y_t = 1$ for all t or $B \to \infty$). In this case, the intrinsic value is given by $I_t \times b_t$, $t \leq 1, 2, \ldots, T$.

From the Bellman principle, the continuation value $V^*(t, S_t)$ of a Bermudan option satisfies

$$V^*(t, S_t) = E_t^Q \left[\frac{1}{R_{t,t+1}} \times \max \{I_{t+1} \times b_{t+1}, V^*(t+1, S_{t+1})\} \right],$$
$$t = 0, 1, \ldots, T-1, \tag{2}$$

and $V^*(T, S_T) = 0$. $E_t^Q[\ ]$ is the expectation operator under the risk-neutral measure Q conditional to the information at time t. We refer to V^* as the "first-best" Bermudan price. Although the results below can be derived in terms of a bank account (in which $R_t = 1$, $t = 0, 1, \ldots, T$), we work in nominal terms; thus, all equations carry directly to the computer.

Remark. In Eq. (2), we assume $b_t = 1$, and hence $V^*(t, S_t)$ is the continuation value seen since t, that is, conditional on no previous barrier (otherwise, $V^*(t, S_t) = 0$ if $b_t = 0$).

2.1. The local LSM algorithm

Let $n_{\max} \geq 1$ be the number of iterations. We specify the final period $t = T$, and recursively solve the continuation value for $T - 1$, $T - 2$, until $t = 1$, where $\widetilde{V}_{T-1}^{LSM}$ is a standard LSM estimator of the continuation value at the initial stage $T - 1$. Specifically, y_t is the realized payoff at t stage (of the Bermudan option exercisable

between t and T), $\widetilde{V}_t^n$ is a local LSM estimator of the continuation value at t stage and n iteration, and $K(z, h)$ is a Gaussian kernel (which is evaluated at z, and h is the bandwith). The algorithm provides $\widetilde{V}_t^*$, which is the final local LSM estimator of the continuation value at t, and $\widetilde{V}_t^{co}$, which is the estimator of the continuation value in the continuation region at t ($t = 1, 2, \ldots, T - 1$).

We assume a specific set of regressors x_t, which depend on prices or state variables S_t (or earlier values if necessary), where $f_t(x_t)$ is an affine function to be estimated by least squares.

The local LSM algorithm: Consider a set of simulated paths, $\varpi \in \Omega$.

0. At maturity, set $t = T$. Define $y_{T+1} = 0$, $\widetilde{V}_T^* = 0$, and $\widetilde{V}_T^{co} = 0$.

1. Updating paths, $\varpi \in \Omega$

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

If a barrier exists, $Y_t = 0$ cancels (set to a 0 value) a path that hits the barrier.

Set $t = t - 1$.

2. The new Continuation value. Set $n = 1$ and $\widetilde{V}_t^0 = \widetilde{V}_{t+1}^*$. If $t = T - 1$, set $\widetilde{V}_{T-1}^0 = \widetilde{V}_{T-1}^{LSM}$.

2.1. Localizing the exercise boundary: a local regression

$$\widetilde{V}_t^n = \arg\min_{f_t \in F} \sum_{\varpi \in \Omega} \left(f_t(x_t) - R_{t,t+1}^{-1} \times y_{t+1} \right)^2 \times \mathcal{K}\left(\widetilde{V}_t^{n-1}(x_t) - I_t, h \right)$$

$$\times 1_{\{I_t > 0\}} 1_{\{Y_t = 1\}},$$

$$\widetilde{V}_t^n \leftarrow \frac{\widetilde{V}_t^n + \widetilde{V}_t^{n-1}}{2} \text{ (if necessary, to avoid potential loops)} \quad (3)$$

Set $n = n + 1$. Go back to step **2.1** until $n = n_{\max}$.
Set $\widetilde{V}_t^* = \widetilde{V}_t^{n_{\max}}$.

3. The new Continuation value in the waiting region

$$\widetilde{V}_t^{co} = \arg\min_{f_t \in F} \sum_{\varpi \in \Omega} \left(f_t(x_t) - R_{t,t+1}^{-1} \times y_{t+1} \right)^2 \times 1_{\{\widetilde{V}_t^* \geq I_t\}}. \quad (4)$$

Go back to step **1** until $t = 1$.
End of the local LSM algorithm

At $t = 1$, we have estimated all continuation values: $\widetilde{V}_t^*$ and $\widetilde{V}_t^{co}$, $t = 1, 2, \ldots, T - 1$. □

Four remarks. First, from the value-matching condition (i.e., $V^*(t, S) = I_t(S)$), the function $\mathcal{K}(\widetilde{V}_t^{n-1}(x_t) - I_t, h)$ is a Gaussian kernel that underweights paths that are far away from the optimal-exercise boundary at the t stage and $n - 1$ iteration. In Eq. (3), in the case of a barrier, if $1_{\{Y_t = 1\}} = 0$, this term excludes the paths that hit the barrier from the regression. Second, the standard LSM method is given by no iterations and no kernel ($n_{\max} = 1$ and $\mathcal{K}() = 1$), where only in-the-money paths are used (i.e., the term $1_{\{I_t > 0\}}$). Third, $\widetilde{V}_t^{co}$ is an estimator of the continuation value in the waiting region. At time t, we compute many samples of the discounted realized payoff, $R_{t,t+1}^{-1} \times y_{t+1}$. Then, the payoffs in the waiting region (i.e., if $\widetilde{V}_t^* \geq I_t$) are approximated using standard least-squares by the family F. (Without the binary function $1_{\{\widetilde{V}_t^* \geq I_t\}}$, $\widetilde{V}_t^{co}$ is estimated by a standard regression, only that the exercise policy is based on a local approach.)

Finally, fourth, step 2.1 in the local LSM algorithm is a Newton-type iteration (in a multi-dimensional setting), which converges in a few (from one to three) iterations because the intrinsic value—I_t is a linear function in the in-the-money region and the continuation value—V^* is a smooth and monotonic function, in the case of standard put/call payoffs. However, in the case of a barrier, which if hit cancels the option, V^* is not monotonic. Because of this lack of monotonicity, which makes harder to price the Bermudan option, we include the step $\widetilde{V}_t^n \leftarrow \frac{\widetilde{V}_t^n + \widetilde{V}_t^{n-1}}{2}$, to avoid potential loops (in the Newton iterations).

The stopping time—or exercise strategy

Let $\widetilde{\tau}(t) \in \{t, t+1, \ldots, T\}$, and hence $\widetilde{\tau}(t) \geq t$, be a stopping time indexed in t, for $t \in \{1, 2, \ldots, T\}$, and $\widetilde{\tau}(T) = T$. If $\widetilde{\tau}$ is not indexed in t, $\widetilde{\tau} = \widetilde{\tau}(1)$. First, from the continuation value obtained from the local LSM method ($\widetilde{V}^*$), the stopping time $\widetilde{\tau}(t)$ is recursively defined as follows:

$$\widetilde{\tau}(t) = t \quad \text{if} \quad I_t \geq \widetilde{V}_t^*; \quad \widetilde{\tau}(t) = \widetilde{\tau}(t+1) \quad \text{otherwise.} \quad (5)$$

The stopping time $\widetilde{\tau}$ is used to compute both the lower bound (V_0^{low}) and the martingale ($\widehat{M}/R$) associated with the upper bound, which are defined below. This stopping time $\widetilde{\tau}$ provides the exercise strategy that optimizes a recursive lower bound. Second, the other continuation-value function $\widetilde{V}^{co}$ is used to estimate a new martingale and second upper bound.

Henceforth, no other least-squares estimators are necessary. In addition, for simplicity, in the following lower/upper bounds, the intrinsic value is written as I_t (instead of $I_t \times b_t$).

2.2. Lower and dual upper bounds

We rewrite Eq. (2) for a Monte-Carlo setting. Let $\mathcal{T}$ be the set of stopping-times, $\tau \in \{1, 2, \ldots, T\}$. For a given $\widetilde{\tau} \in \mathcal{T}$, a lower bound V_0^{low} is defined as follows:

$$V_0^{low} := E_0^Q\left[\frac{I_{\widetilde{\tau}}}{R_{\widetilde{\tau}}} \right] \leq \sup_{\tau \in \mathcal{T}} E_0^Q\left[\frac{I_\tau}{R_\tau} \right] = E_0^Q\left[\frac{I_\tau}{R_\tau} \Big|_{\tau = \tau^*} \right] := V_0^*, \quad (6)$$

where $V_0^* = V^*(0, S_0)$ is the Bermudan price and τ^* is the associated first-best stopping time.

A dual upper bound is an estimator of the Bermudan price and allows us to build a mid point and to assess a lower bound. However, a dual upper bound depends on a martingale that is not specified. For a martingale $\frac{M_t}{R_t}$, $t \in \{0, 1, \ldots, T\}$, upper bounds V_0^{up} are based on the following result:

$$V_0^{up} := M_0 + E_0^Q\left[\max_{1 \leq t \leq T} \left\{ \frac{I_t - M_t}{R_t} \right\} \right] \geq M_0 + E_0^Q\left[\frac{I_\tau}{R_\tau} - \frac{M_\tau}{R_\tau} \Big|_{\tau = \tau^*} \right]$$
$$= V_0^*. \quad (7)$$

The last equality follows from the optional sampling theorem and the inequality follows from

$$\max_{1 \leq t \leq T} \left\{ \frac{I_t - M_t}{R_t} \right\} \geq \frac{I_\tau - M_\tau}{R_\tau} \Big|_{\tau = \tau^*}.$$

The upper bound is binding (i.e., $V_0^{up} = V_0^*$) for the process associated with the first-best Bermudan price, M^* (Andersen & Broadie, 2004; Rogers, 2001; our Proposition 1). To define M^* (in Section 3.1), we use the standard factorization of the American option into the early-exercise premium and the European counterpart.

V_0^{up} is independent of the initial value M_0 (see Appendix A). In particular, if the initial value of the process M is set (not to zero but) to $M_0 = V_0^{low}$, it follows, along with

$$V_0^{low} = M_0 \leq V_0^* \leq V_0^{up},$$

that the following expectation is a proper gap:

$$E_0^Q\left[\max_{1 \leq t \leq T} \left\{ \frac{I_t - M_t}{R_t} \right\} \right] = V_0^{up} - V_0^{low} \geq 0;$$

that is, the difference between the upper and the lower bound is nonnegative. Then, conditional on V_0^{low}, we approximate V_0^{up} by simulation in two ways that correspond to two different types of martingales.

3. Recursive upper bounds based on stopping times

Because optimizing the dual upper bound is not tractable, we study a recursive version. Ibáñez and Velasco (2018) maximize recursively the Bermudan price with respect to a family of stopping

times at each exercise period; we refer to this price as a recursive Bermudan price, which is the objective function of the primal problem. The dual problem is to minimize the recursive upper bound and to determine whether the solution to these two (recursive) primal and dual problems are linked. If we consider a family of stopping times that are specified in a recursive way, a martingale based on the exercise strategy that maximizes the Bermudan price also minimizes not the upper bound itself, but rather the gap between the lower and the upper bound.

We derive a simple recursive expression for this gap (Theorem 1), from which we prove all results. A recursive upper bound is independent of the next-stage stopping time (Proposition 1). Therefore, for a martingale based on stopping times, minimizing the gap is equivalent to maximizing the lower bound in a recursive way (Proposition 3). In Section 4, we show Theorem 1 and Proposition 1 also hold for an upper bound based on continuation values.

From Eq. (6), which is based on stopping times, we build a martingale by using the exercise policy in Eq. (5). We define $\widehat{Z}_t$ as follows:

$$\widehat{Z}_t = 1_{\{t < \tilde{\tau}(t)\}}\widehat{V}_t + 1_{\{t = \tilde{\tau}(t)\}}I_t, \quad t = 1, 2, \ldots, T. \tag{8}$$

The process $\widehat{Z}$ is similar to the value process of a Bermudan option (either to wait or to exercise) associated with $\tilde{\tau}$, in which $\widehat{V}_T = 0$ and

$$\widehat{V}_{t-1} = E_{t-1}^Q\left[\frac{\widehat{Z}_t}{R_{t-1,t}}\right], \quad t = 1, 2, \ldots, T. \tag{9}$$

We then define the process $\widehat{M}$ as follows; that is, $\widehat{M}_0 = \widehat{V}_0$ and

$$\widehat{M}_t = \widehat{M}_{t-1}R_{t-1,t} + \widehat{Z}_t - \widehat{V}_{t-1} \times R_{t-1,t}, \quad t = 1, 2, \ldots, T, \tag{10}$$

so that $\widehat{M}_t/R_t$ is a martingale (i.e., $E_{t-1}^Q[\widehat{M}_t/R_{t-1,t}] = \widehat{M}_{t-1}$), which follows from $\widehat{V}$ definition. In particular, $\widehat{M}_1 = \widehat{Z}_1$.

In addition, from Eq. (9) for $t = 1$, we also define the lower bound from $\widehat{V}_0$, namely, $V_0^{low} = \widehat{V}_0$ (given that $\widehat{V}_0 \leq V_0^*$). This lower bound $\widehat{V}_0$ is approximated by simulation.

3.1. Factorizing the martingale

Importantly, the process $\widehat{M}$ in (10) is explicitly defined by $\widehat{M}_0 = \widehat{V}_0$ and

$$\widehat{M}_t = \sum_{j=1}^{t-1}\left(I_j - \widehat{V}_j\right) \times 1_{\{j=\tilde{\tau}(j)\}} \times R_{j,t} + \left(1_{\{t<\tilde{\tau}(t)\}}\widehat{V}_t + 1_{\{t=\tilde{\tau}(t)\}}I_t\right), \tag{11}$$

which is equal to the sum of the early-exercise premium (reinvested in a bank account) plus the right to exercise at time t (Appendix A). For $t = T$, because $\widehat{V}_T = 0$, the risk-neutral expectation of the discounted value of Eq. (11)'s right-hand-side (rhs) implies the classical factorization of an American option into an early-exercise premium plus the European counterpart (if $\tilde{\tau} = \tau^*$). This factorization is related to the Doob-decomposition theorem, in which the Bermudan-option price process is the Snell envelope (e.g., Carr et al., 1992).

From Eq. (11), it follows for $t \leq \tilde{\tau}$ that

$$\widehat{M}_t = \widehat{V}_t, \text{ if } t < \tilde{\tau}; \text{ and } \widehat{M}_t = I_t, \text{ if } t = \tilde{\tau},$$

and therefore,

$$\max_{1 \leq t \leq T}\left\{I_t - \widehat{M}_t\right\} \geq I_{\tilde{\tau}} - \widehat{M}_{\tilde{\tau}} = 0. \tag{12}$$

The martingale associated with the optimal stopping-time family, τ^*, is denoted by M^*/R and is defined in a similar way as in Eq. (11), where $M_0^* = V_0^*$. The next result complements the literature (Rogers, 2001) on dual upper bounds for the optimal τ^*.

Proposition 1. *(i) An upper bound based on the optimal stopping time τ^* is binding, $V_0^{up} = V_0^*$. And (ii), for any path,*

$$\tau^* = \inf\left(\arg\max_{1 \leq t \leq T}\{I_t - M_t^*\}\right),$$

$$0 = \max_{1 \leq t \leq T}\{I_t - M_t^*\},$$

in which the "inf" is taken in the case of multiple solutions.

Proof. See Appendix A. □

Remark. Proposition 1 shows Eq. (12)'s inequality is binding and the maximum is equal to 0 path by path for the optimal martingale M^* (associated with τ^*). It follows that the term $\max_{1 \leq t \leq T}\{I_t - \widehat{M}_t\}$, as well as the (sample) gap between the lower and the upper bound, has little variance if the process $\widehat{M}$ is based on a good exercise policy $\tilde{\tau}$ (and if, in addition, $\widehat{V}$ in $\widehat{M}$ is estimated with little simulation error). Next, we define the recursive lower and upper bounds, and then try to minimize the recursive upper bound and a recursive gap.

3.2. Recursive lower and upper bounds

We define a new variable *GAP* at time s as follows:

$$GAP_s := \frac{\widehat{M}_s}{R_s} - \frac{\widehat{Z}_s}{R_s} + \max_{s \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}, \quad 1 \leq s \leq T, \tag{13}$$

where $\frac{\widehat{M}_s}{R_s} - \frac{\widehat{M}_t}{R_t}$, $s \leq t$, are Doob-martingale increments. In particular, because $\widehat{M}_1 = \widehat{Z}_1$,

$$GAP_1 := \max_{1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\},$$

and the upper bound is given by

$$V_0^{up} := \widehat{M}_0 + E_0^Q\left[\max_{1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}\right] = \widehat{M}_0 + E_0^Q[GAP_1],$$

where $\widehat{M}_0 = \widehat{V}_0$. The next result allows us to understand a recursive gap between lower and upper bounds (see Belomestny et al., 2013, for other recursive statements).

Theorem 1. *The process GAP defined in Eq. (13) for $1 \leq s \leq T$, with $GAP_{T+1} = 0$, satisfies that*

$$GAP_s = \frac{-\widehat{Z}_s}{R_s} + \max\left\{\frac{I_s}{R_s}, \frac{\widehat{V}_s}{R_s} + GAP_{s+1}\right\}, \tag{14}$$

where $GAP_T = 0$. Moreover, in Eq. (14)'s rhs, only $\widehat{Z}_s$ depends on the function $\tilde{\tau}(s)$.

Proof. See Appendix A. □

Example. Consider a Bermudan option with three exercise opportunities ($s = 1$ and $T = 3$), from Eq. (14),

$$GAP_1 + \frac{\widehat{Z}_1}{R_1} = \max\left\{\frac{I_1}{R_1}, \frac{\widehat{V}_1}{R_1} + GAP_2\right\}$$

$$= \max\left\{\frac{I_1}{R_1}, \frac{\widehat{V}_1}{R_1} - \frac{\widehat{Z}_2}{R_2} + \max\left\{\frac{I_2}{R_2}, \frac{\widehat{V}_2}{R_2} + \underbrace{GAP_3}_{=0}\right\}\right\},$$

which, from Eqs. (8) and (9), is independent of the stopping time $\tilde{\tau}(1)$, as in Theorem 1.

For tractability, we analyze the upper bound recursively. We consider the following lower and upper bounds, which correspond to a Bermudan option that can only be exercised from s to T, $1 \leq s \leq T - 1$. That is,

$$V_{s,0}^{low} := E_0^Q\left[\frac{\widehat{V}_{s-1}}{R_{s-1}}\right] \text{ and } V_{s,0}^{up} := E_0^Q\left[\frac{\widehat{V}_{s-1}}{R_{s-1}}\right] + E_0^Q[GAP_s]. \tag{15}$$

Consistent with our notation, we have $V_0^{low} = V_{1,0}^{low}$, $V_0^{up} = V_{1,0}^{up}$, and $V_{s,0}^{up} - V_{s,0}^{low} = E_0^Q[GAP_s]$.

Proposition 2. $V_{s,0}^{up}$ *is independent of* $\tilde{\tau}(s)$, *which is the next-stage exercise decision.*

Proof. From Eq. (15) and Theorem 1,

$$V_{s,0}^{up} := E_0^Q\left[\frac{\widehat{V}_{s-1}}{R_{s-1}}\right] + E_0^Q[GAP_s]$$

$$= E_0^Q\left[\frac{\widehat{V}_{s-1}}{R_{s-1}}\right] + E_0^Q\left[\frac{-\widehat{Z}_s}{R_s} + \max\left\{\frac{I_s}{R_s}, \frac{\widehat{V}_s}{R_s} + GAP_{s+1}\right\}\right]$$

$$= E_0^Q\left[\max\left\{\frac{I_s}{R_s}, \frac{\widehat{V}_s}{R_s} + GAP_{s+1}\right\}\right].$$

□

$V_{s,0}^{up}$ does not depend on $\tilde{\tau}(s)$ because the upper bound directly compares the intrinsic value and an estimated continuation value at time s (i.e., $\max\{I_s, \widehat{V}_s\}$). In particular, a two-period Bermudan upper bound is always binding.[1] For a two-period Bermudan, $T = 2$,

$$V_0^{up} = V_{1,0}^{up} = \widehat{M}_0 + E_0^Q[GAP_1]$$

$$= \underbrace{\widehat{V}_0 - E_0^Q\left[\widehat{Z}_1 \times \frac{1}{R_1}\right]}_{=\widehat{V}_0} + E_0^Q\left[\max\left\{\frac{I_1}{R_1}, \frac{\widehat{V}_1}{R_1} + \underbrace{GAP_2}_{=0}\right\}\right]$$

$$= E_0^Q\left[\max\left\{I_1, E_1^Q\left[I_2 \times \frac{1}{R_{1,2}}\right]\right\} \times \frac{1}{R_1}\right] = V_0^*,$$

where the last equality follows from V_0^* definition (i.e., the maximum between exercise and the European option at $t = 1$).

Following the example of a Bermudan option with three exercise opportunities, $t \in \{1, 2, 3\}$, consider a family of exercise strategies. The first-order conditions associated with maximizing the Bermudan price at $t = 0$ imply optimal exercise at $t \in \{1, 2\}$, but only for those paths that are alive for the exercise decision at $t = 2$ (Ibáñez & Velasco, 2018). Hence, if we consider all paths at $t = 2$, we can solve this problem recursively, which is tractable and close to the optimal one. By contrast, minimizing the upper bound depends only on the exercise decision at $t = 2$, not on $t = 1$. That is, Proposition 1 implies that we cannot minimize the upper bound $V_{s,0}^{up}$ but rather the gap, $E_0^Q[GAP_s]$, in a recursive way (where $\tilde{\tau}(t)$ is given for $t > s$).

3.3. An optimal recursive gap

Define

$$\tilde{\tau}^*(s) := \arg\max_{\tilde{\tau}(s) \in \mathcal{T} | \tilde{\tau}(s+1)} V_{s,0}^{low}, \qquad (16)$$

where $V_{s,0}^{low}$ is given in Eq. (15). The stopping time $\tilde{\tau}^*(s)$ means optimal exercise at time s, conditional on $\tilde{\tau}(s+1)$ and subject to a given set of stopping times $\mathcal{T}$ (in which now $\tau \in \{s, s+1, \ldots, T\}$). Namely, if $\tilde{\tau}^*(s) > s$, $\tilde{\tau}^*(s) = \tilde{\tau}(s+1)$ where $\tilde{\tau}(s+1)$ is computed in advance.

Proposition 3. *Consider a Bermudan option that can only be exercised from s to T, $1 \le s \le T - 1$; that is, $s \le \tilde{\tau}(s) \in \mathcal{T}$. Assume the stopping time $\tilde{\tau}(s+1)$ is given. Then $\tilde{\tau}^*(s)$, which is defined in Eq. (16), satisfies*

$$\tilde{\tau}^*(s) = \arg\min_{\tilde{\tau}(s) \in \mathcal{T} | \tilde{\tau}(s+1)} E_0^Q[GAP_s].$$

---

[1] Kaniel, Tompaidis, and Zemlianov (2008), Lemma 1, proved this result for a two-period Bermudan option, $t \in \{1, 2\}$; the upper bound does not depend on the $t = 1$ exercise decision, implying a two-period Bermudan upper bound is unbiased.

Further, if $\tilde{\tau}(s+1) = \tau^*(s+1)$ and $\tau^*(s) \in \mathcal{T}$, then $\tilde{\tau}^*(s) = \tau^*(s)$ and $V_{s,0}^{low} = V_{s,0}^{up}$.

In particular, for $s = 1$ (where $R_0 = 1$, $\widehat{V}_0 = V_0^{low}$ and $\widehat{M}_0 = \widehat{V}_0$),

$$\tilde{\tau}^*(1) := \arg\max_{\tilde{\tau}(1) \in \mathcal{T} | \tilde{\tau}(2)} V_0^{low} = \arg\min_{\tilde{\tau}(1) \in \mathcal{T} | \tilde{\tau}(2)} \left\{V_0^{up} - V_0^{low}\right\}.$$

Proof. See Appendix A. □

Minimizing $E_0^Q[GAP_s]$ corresponds to optimally exercising at time s conditional on $\tilde{\tau}(s+1)$ and is solved by the local LSM approach, which yields a continuation-value function $\tilde{V}_s^*$.

The lower/upper bound biases

From $E_0^Q[GAP_1] = (V_0^* - V_0^{low}) + (V_0^{up} - V_0^*)$, we obtain the bias associated with the lower bound,

$$0 \le V_0^* - V_0^{low}$$

$$= E_0^Q\left[\left(1_{\{1 = \tau^*(1)\}}I_1 + 1_{\{1 < \tau^*(1)\}}V_1^*\right) \times \frac{1}{R_1}\right]$$

$$- E_0^Q\left[\left(1_{\{1 = \tilde{\tau}(1)\}}I_1 + 1_{\{1 < \tilde{\tau}(1)\}}\widehat{V}_1\right) \times \frac{1}{R_1}\right],$$

and with the upper-bound (from $V_0^{up} = \widehat{M}_0 + E_0^Q[GAP_1]$ and Eq. (14) for GAP_1),

$$0 \le V_0^{up} - V_0^*$$

$$= E_0^Q\left[\max\left\{I_1, \widehat{V}_1 + GAP_2 \times \left(\frac{1}{R_1}\right)^{-1}\right\} \times \frac{1}{R_1}\right]$$

$$- E_0^Q\left[\max\{I_1, V_1^*\} \times \frac{1}{R_1}\right].$$

For instance, if $GAP_2 = 0$, because $\widehat{V}_1 \le V_1^*$ and $V_0^* \le V_0^{up}$, then $\widehat{V}_1 = V_1^*$ (if $I_1 < V_1^*$) and $V_0^{up} = V_0^*$ so that the upper bound is unbiased and independent of the ($t = 1$) next-period exercise decision, as in Proposition 1. Here, we have assumed $\widehat{V}_1$ is computed without simulation error.

3.4. Computing the martingale paths and the upper bound

From Eq. (10), one path of the process $\widehat{M}$ is approximated as follows, $\widehat{M}_0 = \widehat{V}_0$ and, for $t \in \{1, 2, \ldots, T\}$,

$$\widehat{M}_t = \widehat{M}_{t-1}R_{t-1,t} + \left(1_{\{t < \tilde{\tau}(t)\}}(\widehat{V}_t + \widehat{\xi}_t) + 1_{\{t = \tilde{\tau}(t)\}}I_t\right) - (\widehat{V}_{t-1} + \widehat{\xi}_{t-1})$$

$$\times R_{t-1,t}, \qquad (17)$$

where $\widehat{\xi}_t$ is a zero-mean approximation error (i.e., $E[\widehat{\xi}_t] = 0$), and hence $\widehat{M}_t/R_t$ is a martingale.

Based on Eq. (9), $\widehat{V}$ is computed separately from $\tilde{\tau}$ and

$$\widehat{V}_{t-1} = E_{t-1}^Q\left[R_{t-1} \times \frac{I_{\tilde{\tau}(t)}}{R_{\tilde{\tau}(t)}}\right].$$

Because the latter expectation is not analytical, $\widehat{V}$ is estimated by subsimulation. For every path, and for every $t \in \{1, 2, \ldots, T-1\}$, the value $\widehat{V}_t$ is approximated by a new subsimulation (from t to $\tilde{\tau}(t+1)$), where $\widehat{\xi}_t$ is the error, which introduces a second bias in the upper bound.

Then, we directly simulate the following gap,

$$\widehat{g} = E_0^Q\left[\max_{1 \le t \le T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}\right].$$

That is, we approximate $\widehat{V}_0$ and $\widehat{g}$ by two independent simulations and approximate the two bounds as follows:

$$V_0^{low} \approx \widehat{V}_0 + \xi_0^{low} \quad \text{and} \quad V_0^{up} \approx \widehat{V}_0 + \xi_0^{low} + \widehat{g} + \xi_0^g,$$

where ξ_0^{low} and ξ_0^g are the respective simulation errors. Although the upper bound is basically as in Andersen and Broadie (2004),

it is based on an optimal recursive exercise policy (i.e., $\widetilde{\tau}$ in Eq. (5))—and the martingale (17) associated with this optimal recursive policy.

In addition, from Eq. (11), the martingale (i.e., $\widehat{M}/R$) is explicitly given by

$$\widehat{M}_t = \sum_{j=1}^{t-1} \left(I_j - \left(\widehat{V}_j + \widehat{\xi}_j \right) \right)$$
$$\times 1_{\{j = \widetilde{\tau}(j)\}} \times R_{j,t} + \left(1_{\{t < \widetilde{\tau}(t)\}} \left(\widehat{V}_t + \widehat{\xi}_t \right) + 1_{\{t = \widetilde{\tau}(t)\}} I_t \right). \quad (18)$$

In the Appendix, we show how the computational cost of $\widehat{M}$ could be reduced.

4. Upper bounds based on continuation values

We compute an upper bound based on continuation-value functions. First, we show why an upper bound that is based on a good exercise strategy is less biased than a bound based on continuation values. In the latter case, the martingale has an additional third error term, which implies a larger bias (from Jensen inequality). Second, we show fitting the continuation value in the waiting region is sufficient. In this region, the continuation value is less nonlinear and easier to fit than in the entire support. We show the second result in two different ways.

Consider a new family of continuation-value functions $\widetilde{V}_t$, $t = 1, 2, \ldots, T - 1$, and $\widetilde{V}_T = 0$. Extending (the stopping-times based) Eq. (8), we define

$$\widehat{Z}_t = \max\{\widetilde{V}_t, I_t\}, \quad (19)$$

where $\widehat{V}$ is defined as in Eq. (9) with the new $\widehat{Z}_t$ (i.e., $\widehat{V}_{t-1} = E_{t-1}^Q[\frac{\max\{\widetilde{V}_t, I_t\}}{R_{t-1,t}}]$, $t = 1, 2, \ldots, T$).

We define new recursive lower and upper bounds akin to Eq. (15), but using the new process $\widehat{Z}$. Because Theorem 1 and Proposition 1 trivially hold for this new process $\widehat{Z}$ (which depends on the *max* function and $\widetilde{V}$), the new recursive upper bound is given by

$$V_{s,0}^{up} := E_0^Q\left[\frac{\widehat{V}_{s-1}}{R_{s-1}} \right] + E_0^Q[GAP_s] = E_0^Q\left[\max\left\{ \frac{I_s}{R_s}, \frac{\widehat{V}_s}{R_s} + GAP_{s+1} \right\} \right], \quad (20)$$

which does not depend on $\widetilde{V}_s$ at time s (i.e., $\widehat{V}_s$ depends on $\widetilde{V}_{s+1}$). Here, minimizing the gap is not well defined, because the lower bound based on $\widetilde{V}_s$,

$$V_{s,0}^{low} := E_0^Q\left[\frac{\widehat{V}_{s-1}}{R_{s-1}} \right] = E_0^Q\left[\max\left\{ I_s, \widetilde{V}_s \right\} \times \frac{1}{R_s} \right], \quad (21)$$

is not necessarily lower biased.

Hence, let us impose the best case $E_0^Q[GAP_s] = 0$ in (the second equality of) Eq. (20), and search for the best $\widetilde{V}$ guaranteeing this zero equality. If we assume $GAP_{s+1} = 0$,

$$E_0^Q\left[\max\left\{ I_s, \widetilde{V}_s \right\} \times \frac{1}{R_s} \right] = E_0^Q\left[\max\left\{ I_s, \widehat{V}_s \right\} \times \frac{1}{R_s} \right],$$

and the simple solution associated with the latter equation is that $\widetilde{V}_s = \widehat{V}_s$ subject to $\widehat{V}_s > I_s$. This solution implies a fitting of the function $\widetilde{V}_s$ only in the waiting region, in which $\widehat{V}_s > I_s$. This same (waiting-region constraint) result is derived next by factorizing the process $\widehat{M}$.

Remark. In the case of stopping times (i.e., $\widehat{Z}_t = 1_{\{t < \widetilde{\tau}(t)\}} \widehat{V}_t + 1_{\{t = \widetilde{\tau}(t)\}} I_t$), the last equation is given by

$$E_0^Q\left[\left(1_{\{s < \widetilde{\tau}(s)\}} \widehat{V}_s + 1_{\{s = \widetilde{\tau}(s)\}} I_s \right) \times \frac{1}{R_s} \right] = E_0^Q\left[\max\left\{ I_s, \widehat{V}_s \right\} \times \frac{1}{R_s} \right],$$

and the same $\widehat{V}_s$ appears on both sides of the equality, where the only possible difference is because of $\widetilde{\tau}(s)$.

4.1. Computing a martingale based on the continuation value

Extending Eq. (17), we define a martingale based on the continuation value ($\widetilde{V}$) in Eq. (19); that is, $\widehat{M}_0 = \widehat{V}_0$, and for $t \in \{1, 2, \ldots, T\}$,

$$\widehat{M}_t = \widehat{M}_{t-1} R_{t-1,t} + \max\left\{ \widetilde{V}_t, I_t \right\} - \left(\widehat{V}_{t-1} + \widehat{\xi}_{t-1} \right) \times R_{t-1,t}, \quad (22)$$

and

$$\widehat{V}_{t-1} = E_{t-1}^Q\left[R_{t-1,t}^{-1} \max\left\{ \widetilde{V}_t, I_t \right\} \right].$$

Now, for every path, the process $\widehat{V}_t$ is computed for $t \in \{1, 2, \ldots, T - 1\}$ by a one-period subsimulation (from t to $t + 1$), where $\widehat{\xi}$ is the one-period subsimulation error, $E[\widehat{\xi}_t] = 0$.

Then (see Appendix A)

$$\widehat{M}_t = \sum_{j=1}^{t-1} \left\{ I_j - \widetilde{V}_j \right\}^+ \times R_{j,t} + \sum_{j=1}^{t-1} \left(\widetilde{V}_j - \left(\widehat{V}_j + \widehat{\xi}_j \right) \right)$$
$$\times R_{j,t} + \max\left\{ \widetilde{V}_t, I_t \right\}, \quad (23)$$

where the second sum is a martingale error-correcting term (i.e., $\widetilde{V}_j - \widehat{V}_j$). However, $\widetilde{V}_j$ cancels if $\{I_j - \widetilde{V}_j\}^+ > 0$, from the first and second sums, implying the error $\widetilde{V}_j - \widehat{V}_j$ only matters in the waiting region, in which $I_j \leq \widetilde{V}_j$. Hence, we define $\widetilde{V} = \widetilde{V}^{co}$, where $\widetilde{V}^{co}$ is given in the local LSM algorithm (see Eq. (4)).

Two remarks. First, from the factorizations in Eqs. (18) and (23), the former implies a martingale that is close to the optimal martingale M^*/R, if $\widetilde{\tau}$ is close to τ^*. The latter factorization includes a second sum that depends on the error between $\widetilde{V}^{co}$ and the implicit $\widehat{V}$ in the waiting region, implying a more biased upper bound.

Second, in the examples below, we show a standard-LSM stopping time produces only slightly worse upper bounds than the local-LSM stopping time, assuming in both cases $\widetilde{V} = \widetilde{V}^{co}$. This finding implies estimating the continuation value in the waiting region (i.e., $\widetilde{V} = \widetilde{V}^{co}$) is the key insight for upper bounds based on a continuation-value function. That is, instead of localizing the estimation (of a continuation value) in the exercise boundary as in the local LSM method, we localize this estimation in the waiting region.

In addition, we can build a third martingale based on both $\widetilde{\tau}(t)$ and $\widetilde{V}_t$. However, this martingale yields a less tight upper bound than a martingale based exclusively on $\widetilde{V}_t$, and hence is relegated to the Appendix.

5. Complexity analysis and numerical example

For N state variables, T exercise dates, and M paths, the standard LSM method requires simulating O(NTM) sample points and computing T regressions. For n_{max} iterations, the local LSM requires O(NTM) sample points, O($TM \times n_{max}$) kernel evaluations, and $n_{max} \times T$ regressions. In practice, a few ($n_{max} = 1$ to 3) local regressions are sufficient, if the solution of stage t is used as the initial step at $t - 1$. The local approach increases the computational effort in a linear way, which is given by the number of iterations above one ($n_{max} > 1$). In addition, the lower bound, which is independent of the specific LSM method, requires up to $M_{low}T$ intrinsic values (where M_{low} is the number of simulated paths).

The upper bound is costly; it requires computing the Bermudan value ($\widehat{V}$) in every exercise stage and in every path, which implies $M_{up}T$ subsimulations, where M_{up} is the number of simulated paths. In the case of an upper bound based on stopping times (continuation values), the subsimulation is launched until a path is exercised and is bounded by T stages (is a one-period subsimulation).

Table 1
Lower and upper bounds.

S_0	binomial price V_0^*	Lower bound, V_0^{low}				Upper bound, V_0^{up}			
		LSM	local LSM iterations			τ-based		V^{co}-based	
			1st	2nd	3rd	3rd	LSM	3rd	LSM
					N = 2 assets (kernel = 0.5%)				
100	31.074	28.799 (.006)	30.869 (.008)	30.988 (.006)	31.016 (.006)	31.083 (.001)	31.278 (.028)	31.347 (.006)	31.331 (.007)
	$[V_{0,\mathrm{DFM}}^{low}, V_{0,\mathrm{DFM}}^{up}]$				*N = 4 assets (kernel= 1%)*				
90	[33.011, 34.989]	32.706 (.008)	34.612 (.004)	34.656 (.004)	34.667 (.004)	34.749 (.005)	34.934 (.015)	34.976 (.013)	34.962 (.010)
100	[41.541, 43.587]	40.328 (.008)	43.117 (.003)	43.138 (.004)	43.161 (.004)	43.251 (.004)	43.630 (.024)	43.558 (.011)	43.557 (.010)
110	[48.169, 49.909]	47.197 (.007)	49.398 (.004)	49.429 (.004)	49.430 (.004)	49.482 (.004)	49.998 (.028)	49.780 (.007)	49.851 (.007)
					N = 8 assets (kernel = 5%)				
90	[44.113, 45.847]	43.321 (.006)	45.460 (.004)	45.460 (.004)	45.460 (.004)	45.580 (.003)	46.743 (.015)	45.847 (.007)	45.830 (.008)
100	[50.252, 51.814]	49.523 (.007)	51.357 (.003)	51.360 (.003)	51.360 (.003)	51.433 (.003)	51.646 (.015)	51.668 (.003)	51.667 (.005)
110	[53.488, 54.890]	52.319 (.006)	54.525 (.002)	54.527 (.002)	54.527 (.002)	54.564 (.001)	54.898 (.018)	54.697 (.004)	54.744 (.003)
					N = 16 assets (kernel = 5%)				
90	[50.885, 52.316]	49.779 (.005)	51.916 (.003)	51.923 (.002)	51.925 (.002)	51.981 (.002)	52.252 (.015)	52.158 (.005)	52.184 (.006)
100	[53.638, 54.883]	52.574 (.002)	54.601 (.002)	54.603 (.002)	54.603 (.002)	54.633 (.002)	53.806 (.017)	54.718 (.002)	54.800 (.003)
110	[55.146, 56.201]	54.968 (.005)	55.994 (.003)	55.995 (.003)	55.995 (.003)	56.025 (.002)	56.200 (.018)	56.070 (.002)	56.125 (.003)

Hence, in addition to $M_{up}T$ intrinsic values, the upper bound requires up to $M_{up-st}T^2K_{st}$ ($M_{up-cv}TK_{cv}$) evaluations due to the martingale, where K_{st} and K_{cv} are the subsimulation paths. Because of the quadratic effort (T^2), M_{up-st} and K_{st} are kept small in the former case.

In the case of stopping times, the upper bound is equal to the lower bound plus the gap. Hence, we directly estimate the gap ($\widehat{g}$), which is less volatile (see Proposition 1) and requires a much smaller number of paths, $M_{up-st} \ll M_{low}$.

5.1. Numerical example: pricing up-and-out Bermudan max-options

We price up-and-out max-call Bermudan options. The up-and-out barrier feature makes call payoffs sensitive to suboptimal exercise. We define $I_t = \{\max\{S_t\} - K\}^+$, S are lognormal-distributed prices, K is the strike price, and $B > K$ is the barrier.[2] We define $Y_t = 1_{\{\max\{S_t\}<B\}}$, $t = 1, 2, \ldots, T$, in Eq. (1). Hence, $b_t = 0$ indicates the up-and-out barrier (B) has been hit ($b_j = 0$, $j = t, t+1, \ldots, T$). The Bermudan payoff is given by $I_t \times b_t$.

We follow the MC exercise in Table 1 in Desai et al. (DFM). This table has nine examples, corresponding to three numbers of stocks ($N = \{4, 8, 16\}$) and three initial stock prices ($S_0 = \{90, 100, 110\}$). The strike price $K = 100$ is common across all scenarios. The up-and-out barrier is $B = 170$. To derive the two exercise strategies associated with the local and standard LSM methods, we exclude those points that are out of the money or hit the barrier (i.e., if $\max\{S_t\} \le K$ or if $\max\{S_t\} \ge B$). In the Appendix, we emphasize a few points regarding the implementation of local LSM for this barrier problem.

We use the same basis of $N + 2$ variables as DFM, namely, a constant, every component of the price vector $S = (S^{(1)}, S^{(2)}, \ldots, S^{(N)})$, and $\{\max\{S_t\} - K\}^+$; that is,

$$x_t = \left(1, S_t, \{\max\{S_t\} - K\}^+\right),$$

and the same linear function, namely, $f_t(x_t) = \beta_t' \times x_t$, where $\beta_t \in \mathcal{R}^{N+2}$ are the parameters. For both bounds, we report the mean and standard error over 10 independent trials.

5.1.1. Lower and upper bounds based on stopping times

In our Table 1, we provide the lower bound produced by two regression methods: the standard and the local LSM method (first

to third iterations). From the exercise strategies associated with the standard LSM and the local third iteration, we also generate upper bounds.

Table 1. Prices of Bermudan up-and-out max-call options for $N = \{4, 8, 16\}$ uncorrelated stocks in a lognormal setting (where $r = 0.05$ is the riskfree rate, $\delta = 0$ is the dividend yield, and $\sigma = 0.20$ is volatility). $K = 100$ is the strike price, $B = 170$ is the barrier, $T = 3$ is maturity, and 54 exercise opportunities exist. The first column is the stock price and the second is the best lower and upper bound, $[V_{0,\mathrm{DFM}}^{low}, V_{0,\mathrm{DFM}}^{up}]$, reported by DFM (Desai et al., 2012). The third to sixth and seventh to tenth are the lower and upper bounds, respectively. The third column is the standard LSM method, and the fourth to sixth columns are the first three iterations of the local LSM method. The seventh and eighth are upper-bounds based on a stopping time $\widetilde{\tau}$, which are associated with the standard LSM (as Andersen–Broadie) and local LSM third-iteration exercise strategies, respectively. The last two columns are upper bounds based on a continuation value V^{co}, which is reestimated in the continuation region, which is associated with the LSM and third-iteration local LSM exercise strategies, respectively. As in DSM, for both LSM methods, we use 200,000 paths to recursively compute the continuation values and then 2 million paths to compute the Bermudan price. We report the mean and standard error (over 10 independent trials). For the gap of the upper bound based on $\widetilde{\tau}$, we use 3,000 external paths and 10,000 subsimulation paths. For the upper bound based on V^{co}, we use 10,000 external paths and 500 subsimulation paths. We also report the two-asset case, in which the true price is derived from the binomial method and linear extrapolation (to correct the erratic binomial prices).

The local-LSM lower bounds improve upon the reciprocal standard-LSM lower bounds by 100–280 cents (upon DFM by 85 to 160 cents). In the nine examples, the first iteration of the local method yields the most significant improvement. For four assets, this bound increases only by a couple of cents after the third iteration; for eight and 16 assets, the price converges in one iteration. In all cases, the local upper bound yields a one-digit gap.

In Table 2, we increase the number of paths that are used in the local regression to improve Table 1 numbers. We consider the hardest problem, $N = 4$ stocks. Improving the lower bound is difficult. We reduce the standard error and get smoother prices, which is intuitive in a least-squares setting. In Fig. 1, we show the lower bound's robustness to the kernel. More (less) than 1% of the paths that are used in the Table 1 kernel imply lower (slightly larger but erratic) prices. This 1% is our optimal kernel choice.

---

[2] See Ballotta and Bonfiglioli (2016) or Zeng and Kwok (2014) for European options in a rich jump setting.

Table 2
Increasing the number of paths in the regression.

| paths M | LSM | Lower bound, V_0^{low} | | | | | | Upper bound, V_0^{up} | |
| | | local LSM iterations | | | | | | τ-based | V^{co}-based |
		1st	2nd	3rd	4th	5th	10th	3rd	3rd
$5*10^4$	40.327 (.010)	43.106 (.004)	43.128 (.005)	43.147 (.004)	43.153 (.004)	43.160 (.004)	43.178 (.004)	43.251 (.004)	43.558 (.011)
10^5	40.329 (.008)	43.112 (.005)	43.135 (.004)	43.154 (.005)	43.161 (.005)	43.167 (.005)	43.183 (.004)	43.250 (.005)	43.558 (.011)
$2*10^5$	40.328 (.008)	43.117 (.003)	43.138 (.004)	43.161 (.004)	43.168 (.004)	43.175 (.004)	43.191 (.004)	43.251 (.004)	43.558 (.011)
$4*10^5$	40.325 (.006)	43.119 (.003)	43.143 (.003)	43.164 (.002)	43.172 (.002)	43.179 (.002)	43.193 (.002)	43.250 (.002)	43.559 (.011)
10^6	40.328 (.002)	43.118 (.002)	43.143 (.002)	43.163 (.002)	43.170 (.002)	43.177 (.002)	43.193 (.002)	43.250 (.003)	43.558 (.011)
$2*10^6$	40.322 (.003)	43.120 (.003)	43.144 (.002)	43.165 (.002)	43.172 (.002)	43.179 (.002)	43.194 (.002)	43.250 (.003)	43.559 (.011)
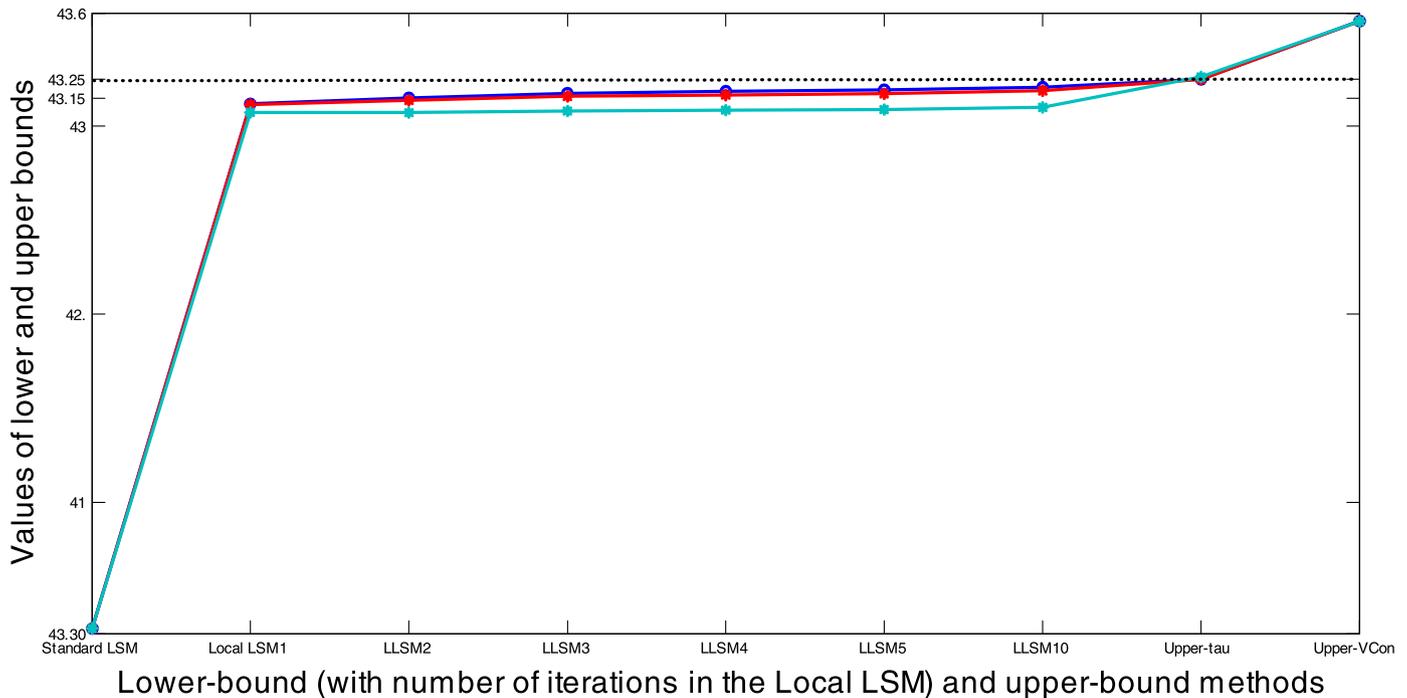
Fig. 1. Values of lower and upper bounds for different kernels and numbers of iterations in the local LSM method. For the kernels, the proportion p indicates the effective number of points used in the local regression ($p = 0.5\%$ in the dark blue line, $p = 1\%$ in the red line, and $p = 5\%$ in the light blue line). For example, Local LSM1 (LLSM3) indicates one (three) local regressions. We also show the lower-bound value of the standard LSM method and two upper bounds based on stopping times and a continuation value estimated in the waiting region. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 2. Prices of Bermudan up-and-out max-call options for $N = 4$ stocks and $S_0 = 100$, in the setup of Table 1. The first column (M) is the number of paths in the backward regressions to estimate the continuation values. The second column is the LSM method, and the third to eighth columns are the first five and the tenth iteration of the local-LSM lower bounds. The ninth and tenth are the upper bounds. We report the mean and standard error over 10 trials. Lower and upper bounds are as in Table 1.

Table 2 and Fig. 1 show the robustness of both upper bounds to the estimation of the continuation value by local least squares (i.e., number of simulation paths and kernel). Because the upper bound based on stopping times is also robust to the number of iterations, this upper bound is tighter (closer to the true price) than the lower bound.

5.1.2. Upper bounds

Table 3 shows the upper bound deteriorates little with the number of subsimulated paths. We can reduce the upper-bound cost without losing accuracy: By reducing the number of subsimulated paths from 10,000 to 500—a 5% (to 100—a 1%) of the original subsimulations, the gap rises by only 3 (15) cents.

Table 3. Gaps of lower and τ-based upper bounds for up-and-out Bermudan max-call options for $N = 4$ stocks, in the setup of Table 1. We directly compute the gap for different numbers of subsimulation paths (sub-paths): 100, 500, and 10,000. The upper bound is defined as the lower bound plus the gap (i.e., $V_0^{up} = V_0^{low} + \text{Gap}$). We report the mean and standard error over 10 trials. Table 1 represents the case using 10,000 subsimulation paths.

From Table 1, an upper bound based on continuation values is more biased. Yet fitting a continuation value (V^{co}) in the waiting region yields upper bounds that, especially if the option is at-/in-the-money, are as accurate as those based on stopping times and the standard LSM but in a fraction of the time. The improvement in this dual bound is mostly due to the waiting-region constraint and not the subsequent policy (e.g., if alternatively based on the standard LSM method); see Table 1, last two columns.[3]

---

[3] In addition (not reported here for brevity), for four assets, we can reduce this upper bound by another 10 bps if we use a quadratic (instead of a linear) function V^{co}, and by another few points if V^{co} is estimated from paths simulated exactly from S_0 (to compute $\tilde{V}^*$ for the $\tilde{\tau}$ stopping-time, it is convenient to simulate paths not from S_0 and $t = 0$ but from in-the-money values and $t < 0$).

Table 3
Gaps between lower and τ-based upper bounds.

S_0	Lower bound, V_0^{low} local LSM		Gap τ-based						Upper bound, V_0^{up}	
			100 sub-paths		500 sub-paths		10,000 sub-paths		10,000 sub-paths	
	LSM	3rd iter	LSM	3rd iter	LSM	3rd iter	LSM	3rd iter	LSM	3rd iter
90	32.706 (.008)	34.647 (.005)	2.405 (.015)	0.306 (.007)	2.257 (.015)	0.133 (.006)	2.228 (.015)	0.082 (.005)	34.934 (.015)	34.752 (.005)
100	40.328 (.008)	43.159 (.002)	3.483 (.025)	0.284 (.004)	3.329 (.026)	0.120 (.004)	3.302 (.024)	0.090 (.005)	43.630 (.024)	43.249 (.004)
110	46.197 (.007)	49.429 (.003)	3.965 (.029)	0.234 (.004)	3.833 (.027)	0.082 (.003)	3.801 (.028)	0.052 (.004)	49.998 (.028)	49.480 (.003)

Table 4a
Relative computing times for continuation-value parameters.

Method	Bandwidth	N		
		4	8	16
LSM	–	1	1	1
Local LSM 1st iter	Fixed	2.29	2.27	2.31
	Optimal	10.75	12.34	14.43
Local LSM 3rd iter	Fixed	4.87	4.89	4.98
	Optimal	13.34	14.96	16.59
LSM + LSM V^{co}	–	3.58	3.66	3.41
Local LSM 1st iter + LSM V^{co}	Fixed	4.87	4.93	4.73
	Optimal	13.33	15.00	16.84
Local LSM 3rd iter + LSM V^{co}	Fixed	7.45	7.55	7.39
	Optimal	15.92	17.62	19.00

Table 4b
Relative computing times for lower and upper bounds.

Bound	Method	Sub-paths	N		
			4	8	16
Lower V_0^{low}	LSM	–	1	1	1
	Local LSM 3rd iter	–	1.42	1.53	1.64
Upper V_0^{up} V^{co}-based	Local LSM 3rd iter	500	3.77	2.85	2.88
Upper V_0^{up} τ-based	Local LSM 3rd iter	10,000	113.40	125.40	197.39
		500	10.83	9.75	9.20
		100	6.73	5.69	4.77

In Tables 4a and 4b, we report the effort required to compute the optimal-recursive exercise policy along with V^{co} and the associated lower and upper bounds. We use a fixed and an optimized kernel, and use one and three local regressions. All times are relative to the standard LSM. First, to estimate the exercise policy and V^{co}, the cost increases between 1 and 19 times in the 30 entries of Table 4a. Second, to compute the lower bound, the total cost increases only between 42% and 64% (because the lower bound uses 2 million paths) in Table 4b. In the case of the upper bound based on a continuation value (stopping times), the cost increases by one order of magnitude (between 100 and 200 times). Yet, in the latter case of stopping times, reducing the number of subsimulation paths to 500, the increase is just 10 times.

Table 4a. Relative computing times for parameter estimation times normalized by column in the setup of Table 1. The LSM, local LSM, and LSM V^{co} methods use 200,000 paths to recursively compute the parameters associated with continuation values to define $\tilde{\tau}$ or V^{co} (Matlab R2017a 64-bit, HP Z620 Workstation, and Intel Xeon CPU E5-2620 0 @2.00 gigahertz).

Table 4b. Relative computing times for lower and upper bounds in the setup of Table 1. Lower V_0^{low} bounds (LSM and local LSM) use 200,000 paths to recursively compute the continuation values to define $\tilde{\tau}$ and another 2 million paths to compute the Bermudan price. Upper V_0^{up} V^{co}-based and Upper V_0^{up} τ-based use 30,000 and 3000 outer paths, respectively.

DSM report (see their Table 2) the cost of the upper bound is similar in these two methods, in their pathwise optimization and in the martingale-based continuation values. Pricing American options with stochastic parameters, Brown et al. (2010) also report both martingale-based upper bounds are time consuming compared to upper bounds based on information relaxation. Hence, our paper provides two extensions of martingale-based upper bounds. First, when using continuation values, we get a tighter upper bound by computing the continuation value only in the waiting region. Second, when using stopping times, the upper bound is both tight and efficient if we use fewer subsimulation paths, but a near-optimal exercise strategy as in the local LSM approach. With both methods, we bring the overall cost of the martingale approach in line with pathwise optimization or information relaxation.

6. Concluding remarks

In this paper, we show the exercise strategy that maximizes the Bermudan price/lower bound also minimizes the gap between the lower and the dual upper bound. We assume both bounds are specified recursively, and show the upper bound is independent of the next-stage policy. Upper bounds based on this optimal recursive exercise policy are very tight, as we show for up-and-out Bermudan max-options, and require few nested simulations. Upper bounds are tighter but more time intensive than lower bounds. In addition, a better upper bound based on continuation values, which is not as accurate but is more efficient than one based on stopping times, requires reestimating the continuation value *only* in the waiting region. From these results emerges the fact that although a tradeoff between tighter bounds and time effort is not straightforward in other methods, this tradeoff exists for optimal recursive lower/upper bounds.

Securities that provide flexibility of early exercise are ubiquitous in financial markets: from single-name American-equity options and Bermudan options to enter/cancel an interest-rate swap, to credit-risk models (Ayadi, Ben-Ameurb, & Fakhfakh, 2016). Specifically, for applications of lower/dual bounds in equity models with stochastic volatility, see Ibáñez and Velasco (2016) and Fabozzi, Paletta, and Tunaru (2017); for the appraisal of Bermudan swaptions prices, see Andersen and Andreasen (2001) and Svenstrup (2005); for term-structure applications, see Joshi and Tang (2014); and see Kogan and Mitra (2017) and Bender, Schweizer, and Zhuo (2017) for extensions to other economic problems. Lower/dual bounds applications are also common in operations research (Trigeorgis & Tsekrekos, 2018), such as when to launch a new product or halt a failing project, delayed-purchase options (Aydin, Birbil, & Topaloğlu, 2017), inventory problems (Brown et al., 2010), or energy real options (Nadarajaha, Margot, & Secomandi, 2017), which includes more references to the literature.

Appendix A. Proofs

V_0^{up} *does not depend on the martingale initial value* $\widehat{M}_0$.
Consider a different initial value $\widehat{x}_0 \neq \widehat{V}_0$. From Eq. (7),

$$V_0^{up} := \widehat{V}_0 + E_0^Q\left[\max_{1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}\right]$$

$$= \widehat{x}_0 + E_0^Q\left[\max_{1 \leq t \leq T}\left\{\frac{I_t - \left(\widehat{M}_{s,t} - (\widehat{V}_0 - \widehat{x}_0) \times R_t\right)}{R_t}\right\}\right],$$

and a more general martingale $\widehat{m}_t/R_t$ is given by $\widehat{m}_0 = \widehat{x}_0$ and $\widehat{m}_t = \widehat{M}_t - (\widehat{V}_0 - \widehat{x}_0) \times R_t$ for $t \geq 1$. And $\widehat{m}_t = \widehat{M}_t$ if $\widehat{x}_0 = \widehat{V}_0$. □

Proof of Eq. (11). From the same Eq. (11),

$$\widehat{M}_t - \widehat{M}_{t-1}R_{t-1,t} = \left(I_{t-1} - \widehat{V}_{t-1}\right) \times 1_{\{t-1=\widetilde{\tau}(t-1)\}}$$
$$\times R_{t-1,t} + \left(1_{\{t<\widetilde{\tau}(t)\}}\widehat{V}_t + 1_{\{t=\widetilde{\tau}(t)\}}I_t\right)$$
$$- \left(1_{\{t-1<\widetilde{\tau}(t-1)\}}\widehat{V}_{t-1} + 1_{\{t-1=\widetilde{\tau}(t-1)\}}I_{t-1}\right) \times R_{t-1,t}$$
$$= \left(1_{\{t<\widetilde{\tau}(t)\}}\widehat{V}_t + 1_{\{t=\widetilde{\tau}(t)\}}I_t\right) - \widehat{V}_{t-1} \times R_{t-1,t},$$

which is Eq. (10). □

Proof of Proposition 1. M^* is explicitly defined as $\widehat{M}$ in Eq. (11) for $\widetilde{\tau} = \tau^*$. The definition of τ^* (i.e., $\tau^*(t) = t$ if $I_t \geq V_t^*$; $\tau^*(t) > t$ otherwise) implies

$$(I_t - V_t^*) \times 1_{\{t=\tau^*(t)\}} = \{I_t - V_t^*\}^+ \geq 0.$$

From Eq. (11) for $\widetilde{\tau} = \tau^*$ and from the last equation, it follows that

$$M_t^* > I_t, \text{ if } t \neq \tau^*; \text{ and } M_t^* = I_t, \text{ if } t = \tau^*,$$

where τ^* means $\tau^*(1)$, and therefore,

$$\max_{1 \leq t \leq T}\{I_t - M_t^*\} = I_{\tau^*} - M_{\tau^*}^* = 0, \text{ and}$$

$$V_0^{up} := M_0^* + E_0^Q\left[\max_{1 \leq t \leq T}\left\{\frac{I_t - M_t^*}{R_t}\right\}\right] = V_0^* + E_0^Q[0] = V_0^*.$$

□

Proof of Theorem 1. From Eq. (13),

$$GAP_s = \frac{-\widehat{Z}_s}{R_s} + \max\left\{\frac{I_s}{R_s}, \frac{\widehat{M}_s}{R_s} + \max_{s+1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}\right\}$$

$$= \frac{-\widehat{Z}_s}{R_s} + \max\left\{\frac{I_s}{R_s}, \frac{\widehat{V}_s}{R_s} + \frac{\widehat{M}_s R_{s,s+1} + \widehat{Z}_{s+1} - \widehat{V}_s R_{s,s+1}}{R_{s+1}}\right.$$
$$\left. - \frac{\widehat{Z}_{s+1}}{R_{s+1}} + \max_{s+1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}\right\}$$

$$= \frac{-\widehat{Z}_s}{R_s} + \max\left\{\frac{I_s}{R_s}, \frac{\widehat{V}_s}{R_s} + \underbrace{\frac{\widehat{M}_{s+1}}{R_{s+1}} - \frac{\widehat{Z}_{s+1}}{R_{s+1}} + \max_{s+1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}}_{=GAP_{s+1}}\right\},$$

(24)

the last equality follows from the definition of $\widehat{M}$ in Eq. (10). Then, it is easy to show by induction that $GAP_s + \frac{\widehat{Z}_s}{R_s}$ does not depend on $\widetilde{\tau}(s)$ or $\widetilde{V}_s$, because GAP_{s+1} does not either.

Next, we prove $GAP_T = 0$. From Eq. (13),

$$GAP_T = \frac{I_T - \widehat{Z}_T}{R_T}.$$

And from Eq. (14) (because $\widehat{V}_T = 0$, $GAP_{T+1} = 0$, and $I \geq 0$), it follows that $GAP_T = (-\widehat{Z}_T + I_T)\frac{1}{R_T}$ as well. Then, $GAP_T = 0$ if $\widehat{Z}_T = I_T$, which is the case because $\widehat{V}_T = 0$. □

Proof of Proposition 3. $V_{s,0}^{up}$ does not depend on $\widetilde{\tau}(s)$, and hence, from Eq. (15),

$$\arg\min_{\widetilde{\tau}(s) \in \mathcal{T}|\widetilde{\tau}(s+1)} E_0^Q[GAP_s] = \arg\min_{\widetilde{\tau}(s) \in \mathcal{T}|\widetilde{\tau}(s+1)}\left\{V_{s,0}^{up} - V_{s,0}^{low}\right\}$$

$$= \arg\min_{\widetilde{\tau}(s) \in \mathcal{T}|\widetilde{\tau}(s+1)}\left\{-V_{s,0}^{low}\right\}$$

$$= \arg\max_{\widetilde{\tau}(s) \in \mathcal{T}|\widetilde{\tau}(s+1)} V_{s,0}^{low} = \widetilde{\tau}^*(s).$$

□

Proof of Eq. 23. We assume $\widehat{\widetilde{\xi}} = 0$ for simplicity. From the same Eq. (23),

$$\widehat{M}_t - \widehat{M}_{t-1}R_{t-1,t} = \left(\{I_{t-1} - \widetilde{V}_{t-1}\}^+ + \widetilde{V}_{t-1} - \widehat{V}_{t-1}\right)$$
$$\times R_{t-1,t} + \max\left\{\widetilde{V}_t, I_t\right\}$$
$$- \max\left\{\widetilde{V}_{t-1}, I_{t-1}\right\} \times R_{t-1,t}$$
$$= \max\left\{\widetilde{V}_t, I_t\right\} - \widehat{V}_{t-1} \times R_{t-1,t},$$

which is Eq. (22), because $\{I_{t-1} - \widetilde{V}_{t-1}\}^+ + \widetilde{V}_{t-1} = \max\{\widetilde{V}_{t-1}, I_{t-1}\}$. □

Reducing the upper-bound computational cost

Consider a path "ω" such that the stopping time satisfies that $t < \widetilde{\tau}(t)$ and $\widetilde{\tau}(t+1) = t+1$ (and let $\widehat{\widetilde{\xi}}_t = 0$). From Eq. (18) and $R_{j,t}/R_t = R_j^{-1}$,

$$\frac{\widehat{M}_t - I_t}{R_t} = \sum_{j=1}^{t-1}\left(I_j - \widehat{V}_j\right) \times 1_{\{j=\widetilde{\tau}(j)\}} \times \frac{1}{R_j} + \frac{\widehat{V}_t - I_t}{R_t}.$$

Likewise, noting the "$j = t$" term of the sum in $\widehat{M}_{t+1}$ is zero because $t < \widetilde{\tau}(t)$,

$$\frac{\widehat{M}_{t+1} - I_{t+1}}{R_{t+1}} = \sum_{j=1}^{t-1}\left(I_j - \widehat{V}_j\right) \times 1_{\{j=\widetilde{\tau}(j)\}} \times \frac{1}{R_j}.$$

Note that $t < \widetilde{\tau}(t)$ does not necessarily imply $\widehat{V}_t > I_t$ (i.e., $\widehat{V}_t$ is independent of the stopping time at time t); if it did, $\widetilde{\tau}(t) = \widetilde{\tau}(t+1)$ would be the optimal time – t recursive exercise policy. Hence, we have no guarantee that

$$\frac{-(\widehat{M}_t - I_t)}{R_t} < \frac{-(\widehat{M}_{t+1} - I_{t+1})}{R_{t+1}},$$

which would imply computing $\widehat{V}_t$ is not necessary if $t < \widetilde{\tau}(t)$. Similarly, given $\tau(t+1) = t+1$, computing $\widehat{V}_{t-j}$ is also not necessary for any previous period $t - j$ ($j \geq 0$) such that the path is in the continuation region, namely, $t - j < \widetilde{\tau}(t - j)$.

Hence, for any path ω, it does not necessarily follow that

$$\max_{1 \leq t \leq T: t=\widetilde{\tau}(t)}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\} = \max_{1 \leq t \leq T}\left\{\frac{I_t - \widehat{M}_t}{R_t}\right\}.$$

Using the lhs, however, reduces the number of periods in which to launch a subsimulation (especially for at-the-money/out-of-the-money options, as paths start in the waiting region), but introduces a negative bias that lowers the upper bound. If $\widetilde{\tau}$ is a good exercise policy, this bias may be negligible, especially compared to the time saved in subsimulations.[4]

A martingale based on both stopping times and continuation values, $\widetilde{\tau}(t)$ *and* $\widetilde{V}_t$

Extending Eqs. (8) and (19), we define,

$$\widehat{Z}_t^{(cvst)} = 1_{\{t<\widetilde{\tau}(t)\}}\widetilde{V}_t + 1_{\{t=\widetilde{\tau}(t)\}}I_t, \tag{25}$$

---

[4] Broadie and Cao (2008) and Joshi (2007) introduce a similar idea to reduce the cost of dual upper bounds.

which also asks for setting $\widetilde{V}_s = \widehat{V}_s$ only in the waiting region (where $\widehat{V}_s > I_s$ or $s < \widetilde{\tau}(s)$). Then $\widehat{M}$ is given by $\widehat{M}_0 = \widehat{V}_0$, and for $t = \{1, 2, \ldots, T\}$,

$$\widehat{M}_t = \widehat{M}_{t-1} R_{t-1,t} + \left(1_{\{t < \widetilde{\tau}(t)\}} \widetilde{V}_t + 1_{\{t = \widetilde{\tau}(t)\}} I_t\right) - \left(\widehat{V}_{t-1} + \widehat{\xi}_{t-1}\right)$$
$$\times R_{t-1,t},$$

and

$$\widehat{V}_{t-1} = E_{t-1}^Q \left[R_{t-1,t}^{-1} \times \left(1_{\{t < \widetilde{\tau}(t)\}} \widetilde{V}_t + 1_{\{t = \widetilde{\tau}(t)\}} I_t\right)\right].$$

We combine, in a one-period subsimulation, $\widetilde{\tau}$ from the local regression and $\widetilde{V}$ from least-squares (in the waiting region). This approach is intuitive if $\widetilde{\tau}$ is close to τ^*. It follows that

$$\widetilde{M}_t = \sum_{j=1}^{t-1} \left(I_j - \widetilde{V}_j\right) \times 1_{\{t = \widetilde{\tau}(t)\}} \times R_{j,t} + \sum_{j=1}^{t-1} \left(\widetilde{V}_j - \left(\widehat{V}_j + \widehat{\xi}_j\right)\right) \times R_{j,t}$$
$$+ \left(1_{\{t < \widetilde{\tau}(t)\}} \widetilde{V}_t + 1_{\{t = \widetilde{\tau}(t)\}} I_t\right). \tag{26}$$

However, a low-biased $\widetilde{V}_t$ implies a process based on $\max\{\widetilde{V}_t, I_t\}$ is closer to the optimal $Z_t^* = \max\{V_t^*, I_t\}$ (where V_t^* is defined in Eq. (2)) than the one based on $(1_{\{t < \widetilde{\tau}(t)\}} \widetilde{V}_t + 1_{\{t = \widetilde{\tau}(t)\}} I_t)$. That is, for any stopping time $\widetilde{\tau}(t)$ (including $\widetilde{\tau}(t) = \tau^*(t)$),

if $\widetilde{V}_t \leq V_t^*$, $\max\{V_t^*, I_t\} \geq \max\{\widetilde{V}_t, I_t\} \geq 1_{\{t < \widetilde{\tau}(t)\}} \widetilde{V}_t + 1_{\{t = \widetilde{\tau}(t)\}} I_t.$

We indeed find the latter martingale yields the most biased upper bound of the three.

Additional details on the local-regression example

First, for a local regression, having paths close to the unknown exercise boundary is critical; otherwise, we have no information to rely on. We start to simulate paths three months before the initial period $t = 0$, so rich price dispersion is present at the first exercise dates. For $N = \{4, 8\}$, we simulate paths from an in-the-money point (i.e., 120 for all assets). If the boundary is well above $K = 100$, and we simulate paths from 90 or 100 (and from $t = 0$), few paths overshoot the boundary at the first exercise dates. For $N = 16$ assets, we simulate from 100 because many paths will eventually hit the barrier. The simulated paths are the same for the standard LSM method. These changes improve the robustness of the local method. The local exercise strategy does not depend on moneyness.

Second, by using the continuation value estimated in the previous stage to define the kernel, one local regression produces very good prices. We iterate this local regression a couple of times to increase this price a few cents. Third, the up-and-out barrier implies the Bermudan price is not monotonic near the exercise boundary. To avoid potential cycles, we define the new continuation value as one half the local regressions of the present and previous periods (the last two iterations, in the case of more than one iteration).

Lastly, the optimal kernel uses approximately 1%–5% of the 200,000 simulated points that are closer to the exercise boundary. This percentage determines the value of the kernel bandwidth h in the local estimation of the continuation value for each period by an iterative procedure starting from a grid search to adapt to the dispersion of the exercise boundary. The effective number of local (to the exercise-boundary) points used in the local estimation is obtained as in (Fan & Gijbels, 1995, p. 374), so that a larger (lower) number of paths implies more biased (more erratic) prices due to a wider (narrower) kernel, that is, a larger (smaller) h. Fixing a unique h for all exercise periods after some limited number of trials can speed up the procedure at a limited cost in terms of price accuracy (e.g., $h \in [0.5, 1.5]$ produces lower bounds similar to the

optimal kernel). For eight and 16 stocks, many of those 200,000 paths eventually hit the barrier near expiry, which implies fewer available points for the local regression, requiring a less localized kernel of 5%.

References

Andersen, L., & Andreasen, J. (2001). Factor dependence on bermudan swaptions: fact or fiction? *Journal of Financial Economics, 62*, 3–37.

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

Ayadi, M. A., Ben-Ameurb, H., & Fakhfakh, T. (2016). A dynamic program for valuing corporate securities. *European Journal of Operational Research, 249*, 751–770.

Aydin, N., Birbil, Ş., & Topaloğlu, H. (2017). Delayed purchase options in single-leg revenue management. *Transportation Science, 51*, 1031–1386.

Ballotta, L., & Bonfiglioli, E. (2016). Multivariate asset models using Levy processes and applications. *The European Journal of Finance, 22*, 1320–1350.

Balseiro, S. R., & Brown, D. B. (2019). Approximations to stochastic dynamic programs via information relaxation duality. *Operations Research, 67*, 577–597.

Belomestny, D., Schoenmakers, J., & Dickmann, F. (2013). Multilevel dual approach for pricing american style derivatives. *Finance and Stochastics, 17*, 717–742.

Bender, C., Schweizer, N., & Zhuo, J. (2017). A primal-dual algorithm for BSDEs. *Mathematical Finance, 27*, 866–901.

Bhim, L., & Kawai, R. (2018). Smooth upper bounds for the price function of American style options. *International Journal of Theoretical and Applied Finance, 21*, 1–38.

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

Brown, D. B., Smith, J. E., & Sun, P. (2010). Information relaxations and duality in stochastic dynamic programs. *Operations Research, 58*, 785–801.

Carr, P., Jarrow, R., & Myneni, R. (1992). Alternative characterization of American put options. *Mathematical Finance, 2*, 87–105.

Chen, N., & Glasserman, P. (2007). Additive and multiplicative duals for American option pricing. *Finance and Stochastics, 11*, 153–179.

Christensen, S. (2014). A method for pricing American options using semi-infinite linear programming. *Mathematical Finance, 24*, 156–172.

Desai, V., Farias, V., & Moallemi, C. (2012). Pathwise optimization for optimal stopping problems. *Management Science, 58*, 2292–2308.

Fabozzi, F. J., Paletta, T., & Tunaru, R. (2017). An improved least squares monte carlo valuation method based on heteroscedasticity. *European Journal of Operational Research, 263*, 698–706.

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.

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

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

Ibáñez, A., & Velasco, C. (2016). One-factor based exercise strategies for american options in multi-factor models. Working Paper, U. Carlos III de Madrid and SSRN.

Ibáñez, A., & Velasco, C. (2018). The optimal method for pricing american options by simulation. *Mathematical Finance, 28*, 1143–1180.

Joshi, M., & Tang, R. (2014). Effective sub-simulation-free upper bounds for the Monte Carlo pricing of callable derivatives and various improvements to existing methodologies. *Journal of Economic Dynamics & Control, 40*, 25–45.

Joshi, M. S. (2007). A simple derivation of and improvements to Jamshidian's and rogers' upper bound methods for bermudan options. *Applied Mathematical Finance, 14*, 197–205.

Kaniel, R., Tompaidis, S., & Zemlianov, A. (2008). Efficient computation of hedging parameters for discretely exercisable options. *Operations Research, 56*, 811–826.

Kogan, L., & Mitra, I. (2017). Accuracy verification for numerical solutions of equilibrium models. Working Paper, MIT.

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

Nadarajaha, S., Margot, F., & Secomandi, N. (2017). Comparison of least squares Monte Carlo methods with applications to energy real options. *European Journal of Operational Research, 256*, 196–204.

Rogers, L. C. G. (2001). Monte carlo valuation of American options. *Mathematical Finance, 12*, 271–286.

Rogers, L. C. G. (2010). Dual valuation and hedging of Bermudan options. *SIAM Journal on Financial Mathematics, 1*, 604–608.

Svenstrup, M. (2005). On the suboptimality of single-factor exercise strategies for Bermudan swaptions. *Journal of Financial Economics, 78*, 651–684.

Trigeorgis, L., & Tsekrekos, A. E. (2018). Real options in operations research: a review. *European Journal of Operational Research, 270*, 1–24.

Zeng, P., & Kwok, Y. K. (2014). Pricing barrier and Bermudan style options under time-changed Levy processes: fast hilbert transform approach. *SIAM Journal on Scientific Computing, 36*, B450–B485.