

Efficient Computation of Hedging Parameters for Discretely Exercisable Options

Ron Kaniel

Fuqua School of Business, Duke University, Durham, North Carolina 27708, ron.kaniel@duke.edu

Stathis Tompaidis

McCombs School of Business, University of Texas at Austin, Austin, Texas 78712, stathis.tompaidis@mcombs.utexas.edu

Alexander Zemlianov

Lehman Brothers, New York, New York 10019, alex.zemlianov@lehman.com

We propose an algorithm to calculate confidence intervals for the values of hedging parameters of discretely exercisable options using Monte Carlo simulation. The algorithm is based on a combination of the duality formulation of the optimal stopping problem for pricing discretely exercisable options and Monte Carlo estimation of hedging parameters for European options. We show that the width of the confidence interval for a hedging parameter decreases, with an increase in the computer budget, asymptotically at the same rate as the width of the confidence interval for the price of the option. The method can handle arbitrary payoff functions, general diffusion processes, and a large number of random factors. We also present a fast, heuristic, alternative method and use our method to evaluate its accuracy.

Subject classifications: finance: derivative hedging.

Area of review: Financial Engineering.

History: Received December 2003; revisions received July 2006, May 2007; accepted May 2007.

1. Introduction

The idea behind no-arbitrage option pricing is that, in a complete market, one can almost surely replicate the payoff of an option using a suitably chosen portfolio of instruments. The construction of this replicating portfolio is based on the computation of the hedging parameters, or sensitivities, of option prices with respect to parameters of the underlying process.¹ Indeed, the first derivative of the option price with respect to the initial asset price, Δ , corresponds to the amount of the underlying asset held in the replicating portfolio, whereas the second derivative, Γ , corresponds to the characteristic time interval between rebalancings. Reliable estimation of option prices and hedging parameters, or option price sensitivities, has become very important with the ever-expanding range of applications of options from, for example, problems in supply chain management, to problems in energy finance and real estate.

In this paper, we develop an algorithm that uses Monte Carlo simulation to estimate option price sensitivities for options with multiple exercise dates and a potentially large number of underlying assets. The advantage of Monte Carlo simulation and the reason it is the method of choice for problems with many assets is that, by its nature, Monte Carlo simulation does not suffer from an exponential increase in effort for a linear increase in the number of underlying assets, a common problem with finite difference discretizations of partial differential equations and high-dimensional lattice algorithms. In addition, Monte Carlo simulation offers an estimate of its own accuracy and is

relatively easy to perform in parallel, leading to significant increases in computational speed.

In the literature, Monte Carlo simulation has been used to calculate sensitivities of the option price for European options, i.e., options with a single exercise date; see Glynn (1989), Broadie and Glasserman (1996), Fournié et al. (1999), and Glasserman (2003) for an overview. We show how the simulation-based algorithms for calculating sensitivities of European options, and in particular the likelihood-ratio algorithm proposed by Broadie and Glasserman (1996), can be used to compute confidence intervals for the values of sensitivities of discretely exercisable options with multiple exercise dates. To this end, we combine the likelihood-ratio algorithm for estimating sensitivities of European options with an algorithm based on a dual representation of discretely exercisable option prices, originally proposed by Davis and Karatzas (1994) and further developed by Rogers (2002), Haugh and Kogan (2004), and Andersen and Broadie (2004). The duality-based algorithm provides confidence intervals for the price of a discretely exercisable option. Our algorithm uses these intervals in a multistage algorithm, where Monte Carlo simulation is employed at every stage. We term the algorithm the *LRD algorithm*, standing for the combination of the likelihood ratio and duality algorithms.

The intuition behind the LRD algorithm is that a discretely exercisable option is equivalent to a European option that expires on the first exercise date of the discretely exercised option. The value of the payoff of this European

option is equal to the price of a discretely exercisable option, starting on the first exercise date of the original option and having one exercise date less. To estimate the sensitivities of the original, discretely exercisable option, we apply the likelihood ratio algorithm on the corresponding European option. We use the duality-based Monte Carlo algorithm to approximate the value function of the corresponding European option which delivers a new, discretely exercisable function. Judicious use of the confidence intervals for the option price then leads to confidence intervals for the option sensitivities.

In addition to proposing the LRD algorithm, we provide estimates of the asymptotic width of the confidence intervals of the option price sensitivities in terms of the simulation parameters and the available computer budget. We demonstrate that, choosing the simulation parameters optimally, the width of the confidence interval for the option price sensitivities decreases with an increase in the computer budget at the same rate as the width of the confidence interval for the option prices themselves, computed by the duality-based algorithm.²

We also compare the LRD algorithm to alternative algorithms for calculating option price sensitivities, also entirely based on Monte Carlo simulation. A comparison with an algorithm based on the duality algorithm for estimating confidence intervals for option prices and finite differences for approximating option price sensitivities shows that the LRD algorithm is superior in terms of both accuracy and speed. We also propose a heuristic algorithm that does not produce confidence intervals but can be faster than the LRD algorithm. With the aid of the LRD algorithm, we evaluate the heuristic alternative in numerical examples, and establish its accuracy.

The remainder of this paper is organized as follows: In §2, we describe the asset market, and review the literature on the calculation of sensitivities for the case of European options and on the confidence intervals for the price of discretely exercisable options. In §3, we describe the LRD algorithm that combines the two algorithms of §2 to provide confidence intervals for the values of sensitivities of discretely exercisable options and obtain the simulation parameters that provide the tightest confidence interval for a given computer budget. In §4, we propose two alternative algorithms and compare them to the LRD algorithm. Section 5 provides numerical examples, and §6 concludes. The proofs of the propositions are presented in an online appendix, which is available as part of the online version that can be found at <http://or.journal.informs.org/>.

2. Model Setup and Literature Review

We consider a market with n tradable risky assets with prices S and one riskless asset with deterministic instantaneous return rate $r(t)$. We assume that there are n

independent sources of uncertainty. The risk-neutral dynamics for the price of the risky and the riskless assets are given by

$$\begin{aligned} dS_i(t) &= r(t)S_i(t)dt + \sigma(S, t)S_i(t)dW_i, \\ S_i(0) &= x_i, \quad i = 1, \dots, n, \\ dB(t) &= r(t)B(t)dt, \quad B(0) = 1, \end{aligned} \quad (1)$$

where we assume that all the components W_i of the Wiener process W are independent.³ We also assume that the volatility matrix σ is invertible.

For the market model given in Equation (1), we consider an option with payoff function $h(S)$. For the case of a European option, with a single exercise date, T , the price of the option is given by

$$Q_0(x) = \mathbb{E}_0 \left[\frac{h(S(T))}{B(T)} \right],$$

where \mathbb{E}_0 is the expectation under the risk-neutral measure conditional on $S(0) = x$, and $B(T) = \exp(\int_0^T r(t) dt)$.

When there are multiple, discrete exercise opportunities, we use the dynamics of the stock prices in Equation (1) to determine a vector-valued discrete-time Markov process $S(t) = (S_1(t), \dots, S_n(t))$ on \mathbb{R}_+^n with fixed initial state x , taking values at times $0 = t_0 < t_1 < t_2 < \dots < t_d = T$. Assuming that the option can be exercised at times t_1, t_2, \dots, t_d , the problem of pricing this option reduces to solving the optimal stopping problem:

$$\text{Primal: } Q_0 = \sup_{\tau} \mathbb{E}_0 \left[\frac{h(S(\tau))}{B(\tau)} \right], \quad (2)$$

where τ is a stopping time taking values in the finite set $\tau \in \{t_1, t_2, \dots, t_d\}$, and $h(S(t))$ is the payoff from exercise at time t . The quantity $h(S(\tau))/B(\tau)$ is the exercise value measured in time 0 dollars.

Below we review results from the literature on how to calculate option price sensitivities for European options, and on how to obtain a confidence interval for the option price for options with multiple, discrete exercise dates.

2.1. Sensitivities of European Options: The Likelihood-Ratio Algorithm

Broadie and Glasserman (1996) present an algorithm based on Monte Carlo simulation for estimating sensitivities of European options. The algorithm, called the likelihood-ratio algorithm, is based on the observation that, for a European option, the price of the option Q_0 satisfies

$$\begin{aligned} \Delta_i(x) &= \frac{dQ_0}{dx_i} \Big|_{S(0)=x} = \frac{d}{dx_i} \mathbb{E}_0 [e^{-\int_0^T r(t) dt} h(S(T))] \\ &= \frac{d}{dx_i} \left[\int e^{-\int_0^T r(t) dt} h(S(T)) g(S(T), x) dS(T) \right] \\ &= \int e^{-\int_0^T r(t) dt} h(S(T)) \frac{\partial g(S(T), x)}{\partial x_i} dS(T) \end{aligned}$$

$$\begin{aligned}
 &= \int e^{-\int_0^T r(t) dt} h(S(T)) \frac{\partial \ln g(S(T), x)}{\partial x_i} \\
 &\quad \cdot g(S(T), x) dS(T) \\
 &= \mathbb{E}_0 \left[e^{-\int_0^T r(t) dt} h(S(T)) \frac{\partial \ln g(S(T), x)}{\partial x_i} \right], \quad (3)
 \end{aligned}$$

where all expectations are conditional on $S(0) = x$, $g(S(T), x)$ is the transition density of reaching $S(T)$ given $S(0) = x$, and we are assuming that g satisfies standard regularity conditions that allow the interchange of the order of differentiation and integration. Equation (3) indicates that it is possible to express the sensitivity of the option price with respect to the initial condition as a weighted expectation of the payoff function.

For the simple case of geometric Brownian motion with constant coefficients, the i th component of the first derivative of the option price with respect to the initial stock price for asset i , Δ_i , is given by

$$\Delta_i(x) = \mathbb{E}_0 \left[e^{-\int_0^T r(t) dt} h(S(T)) \frac{W_i(T)}{x_i \sigma_i T} \right], \quad (4)$$

where σ_i is the diagonal element of the i th row of the volatility matrix. The ij th component of the second derivative of the option price with respect to the initial prices of assets i and j , Γ_{ij} , is given by

$$\begin{aligned}
 \Gamma_{ij} = \mathbb{E}_0 \left[e^{-\int_0^T r(t) dt} \frac{h(S(T))}{x_i x_j \sigma_i \sigma_j T} \left\{ \frac{W_i(T) W_j(T)}{T} \right. \right. \\
 \left. \left. - \delta_{ij} (1 + \sigma_i W_i(T)) \right\} \right], \quad (5)
 \end{aligned}$$

where δ_{ij} is equal to one if $i = j$, and zero otherwise.

From Equations (4) and (5), we have that the weights that allow the sensitivities of the European option price to be expressed as an expectation of the payoff are given by

$$\begin{aligned}
 \Delta_i(x) &= \mathbb{E}_0 [e^{-\int_0^T r(t) dt} h(S(T)) (\xi_\Delta)_i], \quad (\xi_\Delta)_i = \frac{W_i(T)}{x_i \sigma_i T}, \\
 \Gamma_{ij}(x) &= \mathbb{E}_0 [e^{-\int_0^T r(t) dt} h(S(T)) (\xi_\Gamma)_{ij}], \\
 (\xi_\Gamma)_{ij} &= \frac{1}{x_i x_j \sigma_i \sigma_j T} \left\{ \frac{W_i(T) W_j(T)}{T} - \delta_{ij} (1 + \sigma_i W_i(T)) \right\}. \quad (6)
 \end{aligned}$$

We note that the algorithm fails when the time to expiration tends to zero, $T \rightarrow 0$, because the variance of the weights tends to infinity.⁴

2.2. Primal-Dual Algorithm for Computing Confidence Intervals for Prices of Discretely Exercisable Options

Davis and Karatzas (1994), and later Rogers (2002) and Haugh and Kogan (2004), introduced a methodology for obtaining confidence intervals for the price of a discretely

exercisable option based on the dual representation of the option pricing problem described in Equation (2).⁵ Although Rogers (2002) was able to apply the method by judiciously choosing functions that approximate the true option price in a case-by-case basis, Haugh and Kogan (2004) and Andersen and Broadie (2004) described an algorithm that is independent of the option payoff. We present below an overview of the algorithm as described in Andersen and Broadie (2004).

As described in Equation (2), the price of a discretely exercisable option at time 0 is the solution to an optimal stopping problem. More generally, the option price at time $t_k < T$ is given by

$$Q_{t_k} = \max \left(h_{t_k}, \mathbb{E}_{t_k} \left[\frac{B_{t_k}}{B_{t_{k+1}}} Q_{t_{k+1}} \right] \right), \quad (7)$$

where \mathbb{E}_{t_k} is the expectation, conditional on information up to and including time t_k , and we use the notation $h_{t_k} = h(S(t_k))$.

Equation (7) states that the option price at time t_k , Q_{t_k} , is the maximum of the immediate exercise value $h_{t_k}(S_{t_k})$ and the expected present value of not exercising and following an optimal exercise policy in the future. The process Q_{t_k}/B_{t_k} is the smallest supermartingale that dominates h_{t_k}/B_{t_k} at the possible exercise times; i.e., the Snell envelope of h_{t_k}/B_{t_k} . The terminal condition is $Q_T = h_T$.

To estimate the option price at time 0, Q_0 , we have that for any specific exercise strategy τ , the value achieved by following this exercise strategy is dominated by the value achieved by the optimal strategy τ^* :

$$\mathbb{E}_0 \left[\frac{h_\tau}{B_\tau} \right] \leq \mathbb{E}_0 \left[\frac{h_{\tau^*}}{B_{\tau^*}} \right] = Q_0,$$

i.e., any algorithm that gives a stopping rule can be used to compute a lower bound on the price Q_0 .

To find an upper bound for Q_0 , consider an arbitrary adapted martingale π_{t_k} . Then, we have

$$\begin{aligned}
 Q_0 &= \sup_\tau \mathbb{E}_0 \left[\frac{h_\tau}{B_\tau} + \pi_\tau - \pi_\tau \right] = \pi_0 + \sup_\tau \mathbb{E}_0 \left[\frac{h_\tau}{B_\tau} - \pi_\tau \right] \\
 &\leq \pi_0 + \mathbb{E}_0 \left[\max_k \left(\frac{h_{t_k}}{B_{t_k}} - \pi_{t_k} \right) \right], \quad (8)
 \end{aligned}$$

where the second equality follows from the martingale property of π_{t_k} and the inequality follows from the optional sampling theorem. Because π_{t_k} is an arbitrary adapted martingale, we can achieve an upper bound on the option price by taking any adapted martingale. Davis and Karatzas (1994), Rogers (2002), Haugh and Kogan (2004), and Andersen and Broadie (2004) show that taking the infimum over all adapted martingales in (8) results in equality; i.e.,

$$Q_0 = \inf_\pi \left(\pi_0 + \mathbb{E}_0 \left[\max_k \left(\frac{h_{t_k}}{B_{t_k}} - \pi_{t_k} \right) \right] \right). \quad (9)$$

This result implies that the “duality gap;” i.e., the difference between the solution of the dual problem (9) and the primal problem (2), is zero.⁶

Andersen and Broadie (2004) suggest that to get tight lower and upper bounds, we need to find an exercise policy that is close to the optimal policy τ^* . Assuming that we have chosen some exercise policy τ , we define an adapted exercise indicator process $\mathbf{1}_{t_k}$ that equals one if exercise should take place at time t_k , and zero otherwise. For all $0 \leq t \leq T$, we define the t_k -indexed stopping times τ_k as

$$\tau_k = \inf \{ u \in \tau \cap [t_k, T] : \mathbf{1}_u = 1 \}.$$

Thus, τ_k denotes the first instance, either at time t_k or later, at which the option should be exercised according to a given strategy. With this definition, the process L_{t_k} defined as

$$\frac{L_{t_k}}{B_{t_k}} = \mathbb{E}_{t_k} \left[\frac{h_{\tau_k}}{B_{\tau_k}} \right]$$

provides a lower bound to the option price process.

To calculate the upper bound, we make use of the lower-bound process L_{t_k} . We define a process π_{t_k} by $\pi_0 = L_0$, and for $0 \leq k \leq d - 1$,

$$\pi_{t_{k+1}} = \pi_{t_k} + \frac{L_{t_{k+1}}}{B_{t_{k+1}}} - \frac{L_{t_k}}{B_{t_k}} - \mathbf{1}_{t_k} \mathbb{E}_{t_k} \left[\frac{L_{t_{k+1}}}{B_{t_{k+1}}} - \frac{L_{t_k}}{B_{t_k}} \right]. \quad (10)$$

We note that, conditional on no exercise at time t_k , i.e., when $\mathbf{1}_{t_k} = 0$, we have that

$$\mathbb{E}_{t_k} \left[\frac{L_{t_{k+1}}}{B_{t_{k+1}}} \right] = \frac{L_{t_k}}{B_{t_k}};$$

thus, the process π_{t_k} is a martingale.

We have

$$L_0 \leq Q_0 \leq U_0 = L_0 + \mathbb{E}_0 \left[\max_k \left(\frac{h_{t_k}}{B_{t_k}} - \pi_{t_k} \right) \right] := L_0 + \delta_0. \quad (11)$$

As described above, the method produces upper and lower bounds for the option prices. Because Monte Carlo simulation is used for estimating these bounds, the end result is a confidence interval for the option price.

The algorithm proposed by Andersen and Broadie (2004) can be summarized as follows:⁷

ALGORITHM 1. Calculation of a confidence interval for the price of discretely exercisable options.

Step 1. Find an approximate exercise strategy τ . Several algorithms for finding such strategies are available. One popular algorithm, proposed by Longstaff and Schwartz (2001) (see also Carriere 1996 and Tsitsiklis and Van Roy 1999), uses Monte Carlo simulation and projection of the option price function on a set of approximating functions.

Step 2. Once an approximate exercise strategy is selected, perform Monte Carlo simulation for the stock and bond prices with N_L simulation paths, and estimate a confidence interval of the lower bound for the option price.

Step 3. Generate an additional N_δ simulation paths for the stock price to be used for estimating the upper bound. Along these paths, at each exercise date $k = 1, \dots, d$, check whether the option is exercised or not.

Step 3a. If the option is exercised on date k , set $L_{t_k}/B_{t_k} = h_{t_k}/B_{t_k}$ and generate N_i paths for the stock and bond prices to estimate $\mathbb{E}_{t_k} [L_{t_{k+1}}/B_{t_{k+1}}] = \mathbb{E}_{t_k} [h_{\tau_{k+1}}/B_{\tau_{k+1}}]$, where τ_{k+1} is the first stopping time at, or after, time t_{k+1} .

Step 3b. If the option is not exercised on date k , generate N_i simulation paths for the stock and bond prices and estimate $L_{t_k}/B_{t_k} = \mathbb{E}_{t_k} [h_{\tau_k}/B_{\tau_k}]$.

Step 4. Find the maximum value in Equation (11) along each of the N_δ paths and estimate a confidence interval for δ_0 .

Figure 1 illustrates the procedure.

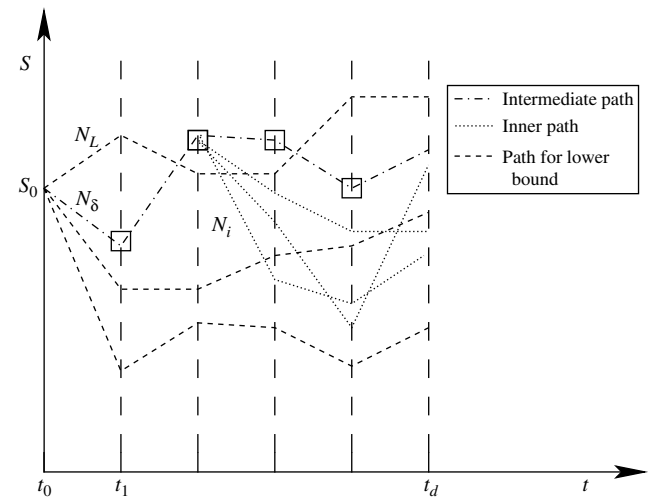
Based on the algorithm, a confidence interval for the option price can be calculated by taking the estimated values for the upper and lower bounds and subtracting and adding an appropriate number of standard errors, computed from the variance of the upper and lower bounds from the simulation.

The approximate computer time required by Steps 2 through 4 is

$$N_L T_{\text{path}} + (d - 1) N_\delta N_i T_{\text{path}},$$

where T_{path} is the time required to generate one path of the underlying process, and d the number of exercise opportunities. The term $N_L T_{\text{path}}$ corresponds to the time necessary for generating the N_L paths for estimating the lower bound. The term $(d - 1) N_\delta N_i T_{\text{path}}$ approximates the time required to generate N_i paths on each of the $d - 1$ exercise opportunities along each of the N_δ paths used to estimate the upper bound for the option price.

Figure 1. Graphical representation of Algorithm 1.



Notes. Starting at time t_0 , N_L paths are generated for the calculation of the lower bound of the option price, and N_δ paths for the calculation of the upper bound on the option price. For each time t_1, \dots, t_{d-1} (starting at each square), an additional N_i inner paths are generated along the N_δ paths to estimate the martingales necessary for calculating the upper bounds on the price.

3. Sensitivities for Discretely Exercisable Options

In this section, we combine the formulas for the calculation of sensitivities of European options with the algorithm that computes bounds for the price of discretely exercisable options to obtain an algorithm that computes confidence intervals for the sensitivities. Assuming that the option is not exercised at time t_0 , discretely exercisable options can be thought of as European options that expire on the first exercise date. These European options have a payoff with value equal to the price of a new option that can be exercised on either that date, or any subsequent exercise date.⁸

The price of the option is given by

$$Q_0 = \mathbb{E}_0 \left[\frac{Q_{t_1}}{B_{t_1}} \right],$$

which is the same as the price of a European option with the payoff function $h(S_{t_1}) = Q_{t_1}(S_{t_1})/B_{t_1}$, suggesting that the sensitivities Δ and Γ can be computed via the formulas given for the European case. Hence, we obtain

$$\Delta_i = \mathbb{E}_0 \left[\left(\xi_{\Delta} \right)_i \frac{Q_{t_1}(S_{t_1})}{B_{t_1}} \right], \quad \Gamma_{ij} = \mathbb{E}_0 \left[\left(\xi_{\Gamma} \right)_{ij} \frac{Q_{t_1}(S_{t_1})}{B_{t_1}} \right],$$

where the weights ξ_{Δ} , ξ_{Γ} are given in Equation (6).

The formulas above can be used to calculate Δ and Γ by Monte Carlo simulation. However, the formulas are based on the assumption that the values for the option price at time t_1 , Q_{t_1} are known. Knowing the option price is equivalent to knowing the optimal exercise policy for the option, and is not achievable with the methods we have presented. Instead, as discussed in §2.2, tight confidence intervals for the option prices can be obtained via the primal-dual algorithm. We use these confidence intervals for the option price to calculate the confidence intervals for the option price sensitivities. We denote the upper and lower bounds for the price of the option at time t_k as U_{t_k} and L_{t_k} , respectively. For time t_1 , we have

$$L_{t_1} \leq Q_{t_1} \leq U_{t_1}.$$

Then, the upper and lower bound for the i th component of Δ , Δ_i is given by

$$\begin{aligned} & \mathbb{E}_0 \left[\mathbf{1}_{\xi_{\Delta_i} < 0} \xi_{\Delta_i} \frac{U_{t_1}}{B_{t_1}} + \mathbf{1}_{\xi_{\Delta_i} \geq 0} \xi_{\Delta_i} \frac{L_{t_1}}{B_{t_1}} \right] \\ & := \Delta_i^l \leq \Delta_i \leq \Delta_i^u := \mathbb{E}_0 \left[\mathbf{1}_{\xi_{\Delta_i} \geq 0} \xi_{\Delta_i} \frac{U_{t_1}}{B_{t_1}} + \mathbf{1}_{\xi_{\Delta_i} < 0} \xi_{\Delta_i} \frac{L_{t_1}}{B_{t_1}} \right], \quad (12) \end{aligned}$$

where $\mathbf{1}_A$ is equal to one if A is true and zero otherwise. A similar expression can be used to bound the components of Γ . The gap between the upper and lower estimate of the

sensitivities is directly related to the gap between the upper and lower estimate of the prices because

$$\begin{aligned} \Delta_i^u - \Delta_i^l &= \mathbb{E}_0 \left[\mathbf{1}_{\xi_{\Delta_i} \geq 0} \xi_{\Delta_i} \frac{U_{t_1}}{B_{t_1}} + \mathbf{1}_{\xi_{\Delta_i} < 0} \xi_{\Delta_i} \frac{L_{t_1}}{B_{t_1}} \right] \\ & \quad - \mathbb{E}_0 \left[\mathbf{1}_{\xi_{\Delta_i} < 0} \xi_{\Delta_i} \frac{U_{t_1}}{B_{t_1}} + \mathbf{1}_{\xi_{\Delta_i} \geq 0} \xi_{\Delta_i} \frac{L_{t_1}}{B_{t_1}} \right] \\ &= \mathbb{E}_0 \left[\mathbf{1}_{\xi_{\Delta_i} > 0} \xi_{\Delta_i} \left(\frac{U_{t_1}}{B_{t_1}} - \frac{L_{t_1}}{B_{t_1}} \right) \right. \\ & \quad \left. + \mathbf{1}_{\xi_{\Delta_i} < 0} \xi_{\Delta_i} \left(\frac{L_{t_1}}{B_{t_1}} - \frac{U_{t_1}}{B_{t_1}} \right) \right] \\ &= \mathbb{E}_0 \left[|\xi_{\Delta_i}| \left(\frac{U_{t_1}}{B_{t_1}} - \frac{L_{t_1}}{B_{t_1}} \right) \right]. \end{aligned}$$

3.1. Confidence Interval for Sensitivities of Discretely Exercisable Options

The algorithm for estimating a confidence interval for the option sensitivities extends Algorithm 1, and is described below.

ALGORITHM 2. Calculation of a confidence interval for the sensitivities of discretely exercisable options.

Step 1. Find an approximate exercise strategy.

Step 2. Simulate N_o Monte Carlo paths for the stock and bond prices between time t_0 and time t_1 . This stage will be referred to as the “outer loop” of the simulation. For each path i out of N_o , get the approximate bounds for the price at time t_1 , $\bar{L}_{t_1}^i$ and $\bar{U}_{t_1}^i$, using Steps 2–4 of Algorithm 1. The simulations at this stage involve generating N_L sample paths for the estimation of the lower bound of the price and N_{δ} intermediate paths, and N_i inner paths for each intermediate path at each time t_2, \dots, t_{d-1} , for the estimation of the upper bound of the price.

Step 3. Average the results from Step 2 according to Equation (12), and compute approximate upper and lower bounds, as well as the standard deviation of the upper and lower bounds.

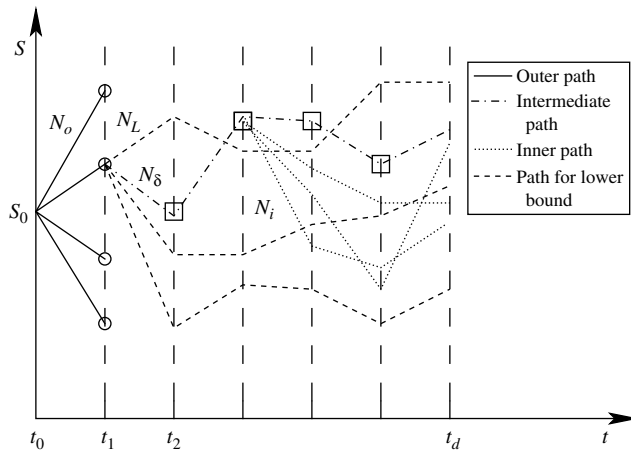
Figure 2 illustrates the algorithm.

The computer time required for the algorithm is approximately given by

$$N_o [N_L T_{\text{path}} + (d-1)N_{\delta} (N_i T_{\text{path}} + T_{\text{path}}) + T_{0 \rightarrow 1}],$$

where $T_{0 \rightarrow 1}$ is the computer time necessary to generate one of the outer loop paths. The terms $N_L T_{\text{path}}$ and $(d-1)N_{\delta} (N_i T_{\text{path}} + T_{\text{path}})$ correspond to the time spent calculating the approximate bounds L_{t_1} and U_{t_1} , respectively. Because these bounds are calculated for each of the N_o paths that end at t_1 , the total computational time is given by $N_o [N_L T_{\text{path}} + (d-1)N_{\delta} (N_i T_{\text{path}} + T_{\text{path}})]$. The term $N_o T_{0 \rightarrow 1}$ corresponds to the time spent generating the outer loop paths.

Figure 2. Graphical representation of the LRD algorithm.



Notes. N_o outer paths are generated from time t_0 to time t_1 . Starting at the end of each of the N_o paths (from each circle), N_L paths are generated for the calculation of the lower bound of the option price, and N_δ intermediate paths for the calculation of the upper bound on the option price. For each time t_2, \dots, t_{d-1} (starting at each square), additional N_i inner paths are generated to estimate the martingales necessary for calculating the upper bounds on the price.

We have the following proposition:

PROPOSITION 1. *With at least q^2 probability, the value of a sensitivity of a discretely exercisable option, G , is in the interval*

$$\left[\frac{1}{N_o} \sum_{j=1}^{N_o} \bar{G}_j^l - \frac{z_q \sigma(\bar{G}^l)}{\sqrt{N_o}}, \frac{1}{N_o} \sum_{j=1}^{N_o} \bar{G}_j^u + \frac{z_q \sigma(\bar{G}^u)}{\sqrt{N_o}} \right],$$

where \bar{G}_j^u, \bar{G}_j^l are the computed upper and lower bounds for G , $\sigma(\bar{G}^l), \sigma(\bar{G}^u)$ are the standard deviations for the computed upper and lower bounds, and j is the index for the outer loop paths and where the barred quantities are inner loop sample averages.

The value z_q is the z -value for the standard normal distribution, corresponding to a centered confidence interval with probability q ; e.g., for $q = 95\%$, $z_q = 1.96$. The confidence level can be made as close to 100% as desired by using the appropriate multiple of the standard deviations $\sigma(\bar{G}^l), \sigma(\bar{G}^u)$.

The proof is provided in the online appendix.

3.2. Optimization of the LRD Algorithm

The width of the confidence interval depends on two types of biases. The first type can be referred to as the “policy bias.” This bias is independent of the choice of the sampling parameters, and is attributed to the fact that the exercise policy used to calculate the confidence interval for the option price may not be optimal. As a consequence, in general, there is a nonzero “duality gap,” which results in a finite distance between the lower and upper bounds for the

price, and hence in a nonzero difference between the lower and upper bounds for sensitivities even when samples of infinite size are used at each stage of the simulation.

The second type of bias is the “sampling bias,” which is due to the finite number of paths used at each stage of the simulation. Although it might be hard to control the policy bias, it is possible to reduce the sampling bias by choosing optimally the sampling parameters N_o, N_L, N_δ , and N_i . This optimization is important, particularly because the LRD algorithm involves a triply nested Monte Carlo simulation, and can be, potentially, time-consuming. In calculations we do not report, we have estimated the asymptotic width of the confidence interval in terms of the parameters N_o, N_L, N_δ , and N_i . Our calculations show that, for a fixed computer budget, increasing the number of the outer and inner paths— N_o and N_i , respectively—has the biggest impact on the width. On the other hand, the number of paths, N_L , used for calculating the lower bound, is of secondary importance. The role of intermediate paths, N_δ , is ambivalent: it appears that, for a fixed computer budget, an increase in N_δ results in wider confidence intervals. This paradoxical phenomenon is due to the fact that increasing the number of intermediate paths necessitates a decrease in the number of inner and outer loop paths to keep the computer budget fixed. This calculation suggests that the number of intermediate paths should be rather small.

Based on this intuition, we set the number of intermediate paths to one, effectively reducing the complexity of the algorithm to a doubly nested simulation. For this choice, we were able to estimate the asymptotic width of the confidence interval for the option price sensitivities.

PROPOSITION 2. *When the number of intermediate paths is set to $N_\delta = 1$, there exist nonnegative constants, c_1, c_2 , independent of the sampling parameters N_o, N_L, N_i , such that, with at least q^2 -probability, asymptotically in N_o, N_L, N_i , the width \bar{D} of the confidence interval for the difference between the upper and lower bounds for the sensitivity of a discretely exercisable option satisfies*

$$\bar{D} - D_\infty \leq \frac{c_1}{\sqrt{N_o}} + \frac{c_2}{\sqrt{N_i}} + o\left(\frac{1}{\sqrt{N_o}}, \frac{1}{\sqrt{N_i}}\right),$$

where D_∞ is the policy bias in the limit of using an infinite number of simulation paths.

The proof of the proposition is given in the online appendix. Although it is surprising that we are able to compute asymptotic bounds for the confidence interval even though one of the simulation parameters is set to one, the intuition behind the proof is that sampling the state space with N_o outer paths between times t_0 and t_1 , as well as with N_i inner paths for each of the intermediate paths starting at time t_1 , compensates for the lack of intermediate sampling. The proof of the proposition uses probabilistic estimates for the different random variables involved, and

Lindeberg’s theorem for the limiting distribution of the sum of independent, but not necessarily identically distributed random numbers, to bound the sampling error.

For a fixed computer budget, \mathcal{B} , we choose the simulation parameters N_o, N_i, N_L so that the width of the confidence interval for an option sensitivity of interest is minimized. This objective is formulated as:

$$\text{choose } N_o, N_i, N_L \text{ to minimize: } \frac{c_1}{\sqrt{N_o}} + \frac{c_2}{\sqrt{N_i}}$$

under the constraint:

$$N_o \{ N_L T_{\text{path}} + (d-1) \cdot (N_i T_{\text{path}} + T_{0 \rightarrow 1}) + T_{0 \rightarrow 1} \} \leq \mathcal{B}. \quad (13)$$

The solution can be found using Lagrange multipliers. It is given by

$$N_o = \frac{c_1}{c_2} \sqrt{\frac{\mathcal{B}}{(d-1)T_{\text{path}}}}, \quad N_i = \frac{c_2}{c_1} \sqrt{\frac{\mathcal{B}}{(d-1)T_{\text{path}}}}, \quad (14)$$

from which we observe that, asymptotically,

$$\frac{N_o}{N_i} = \frac{c_1^2}{c_2^2} = \lambda;$$

i.e., the ratio of the number of outer loop paths to inner loop paths is independent of the computer budget \mathcal{B} . The optimal choice balances the sampling bias due to the outer loop simulations with the sampling bias due to the inner loop simulations.⁹

COROLLARY 1. For a fixed computer budget \mathcal{B} , choosing the simulation parameters according to Equation (14) results in a confidence interval with width of the order $O(\mathcal{B}^{-1/4})$.

To numerically implement the algorithm, we start with a small fraction of the total computer budget and choose different ratios of the sampling parameters N_o/N_i , keeping the computer budget fixed. For each value of the ratio, we run trial simulations, keeping track of the width of the confidence interval. We choose the ratio that minimizes the width of the confidence interval, and then scale the values of N_o, N_i to the extent permitted by the total computer budget. The procedure is further discussed in §5.4.

We note that for the optimal choice of the simulation parameters, the asymptotic order of convergence of the LRD algorithm, with respect to the computer budget, is the same as the asymptotic order of convergence of the duality-based algorithm for computing confidence intervals for option prices.¹⁰ Intuitively, this similarity is due to the fact that the algorithm for estimating a confidence interval for option prices relies on a doubly nested simulation, which cannot be further reduced, whereas setting $N_\delta = 1$ reduces the computational burden of estimating the confidence interval for option price sensitivities to a doubly nested simulation as well. We point out that even though the asymptotic width of the confidence interval with the computer budget is of the same order with respect to the computer budget, the confidence intervals for prices are typically smaller due to the larger variance of the weights involved in the calculation of the sensitivities.

4. Alternative Algorithms

4.1. An Algorithm Based on Monte Carlo Simulation and Finite Differences

An alternative algorithm for estimating a confidence interval for the sensitivities of option prices is to rely only on the confidence intervals obtained using the primal-dual algorithm for the option price. We illustrate the idea for the case of one dimension.

To estimate the first derivative of the option price at a point x_0 , assuming that for x in a small neighborhood of x_0 the approximate option price \hat{Q} satisfies $|\hat{Q}(x) - Q(x)| < \epsilon$, we can approximate the first derivative of Q at x_0 by the centered difference $(\hat{Q}(x_0 + h) - \hat{Q}(x_0 - h))/2h$. We have

$$\begin{aligned} & \left| \frac{\hat{Q}(x_0 + h) - \hat{Q}(x_0 - h)}{2h} - Q'(x_0) \right| \\ &= \left| \frac{[\hat{Q}(x_0 + h) - Q(x_0 - h)] - [\hat{Q}(x_0 - h) - Q(x_0 - h)]}{2h} \right. \\ & \quad \left. + \frac{Q(x_0 + h) - Q(x_0 - h)}{2h} - Q'(x_0) \right| \\ &\leq \frac{\epsilon}{h} + \left| \frac{Q(x_0 + h) - Q(x_0 - h)}{2h} - Q'(x_0) \right|. \end{aligned}$$

Assuming that in a neighborhood of x_0 , Q has continuous, bounded third derivatives, we have, from Taylor’s theorem,

$$\begin{aligned} Q(x_0 + h) &= Q(x_0) + hQ'(x_0) + \frac{h^2}{2}Q''(x_0) + \frac{h^3}{6}Q'''(y_+), \\ Q(x_0 - h) &= Q(x_0) - hQ'(x_0) + \frac{h^2}{2}Q''(x_0) + \frac{h^3}{6}Q'''(y_-), \end{aligned}$$

for $y_+ \in [x_0, x_0 + h]$, $y_- \in [x_0 - h, x_0]$. If $|Q'''(x)| \leq \alpha$ for $x \in [x_0 - h, x_0 + h]$, we have

$$\left| \frac{\hat{Q}(x_0 + h) - \hat{Q}(x_0 - h)}{2h} - Q'(x_0) \right| \leq \frac{\epsilon}{h} + \frac{\alpha h^2}{6}.$$

We can minimize the error by choosing $h = (\epsilon/2\alpha)^{1/3}$, achieving accuracy of

$$\begin{aligned} & \left| \frac{\hat{Q}(x_0 + h) - \hat{Q}(x_0 - h)}{2h} - Q'(x_0) \right| \\ &\leq \epsilon^{2/3} \left(\left(\frac{1}{2\alpha} \right)^{1/3} + \frac{\alpha^{1/3}}{2^{2/3}} \right). \end{aligned} \quad (15)$$

If confidence intervals are used for the lower and upper bounds, then the algorithm leads to a confidence interval for the value of the first derivative.

ALGORITHM 3. Alternative calculation of confidence intervals for the sensitivities of discretely exercisable options.

Step 1. Find approximate exercise strategies $\tau(S(0))$, conditional on starting at initial conditions $S(0) = x - h$, $S(0) = x + h$.

Step 2. Use Steps 2–4 of Algorithm 1 to obtain confidence intervals for the upper and lower bounds for the option price, $L_0(x - h)$, $L_0(x + h)$, $U_0(x - h)$, $U_0(x + h)$.

Step 3. Use Equation (15) to obtain a confidence interval for the first derivative of the option price.

Although Algorithm 3 is easier to implement than the LRD algorithm, we can show that it is asymptotically less efficient, both in terms of speed and in terms of accuracy. For the price bounds from Algorithm 1, we have that, asymptotically, the difference between the upper and lower bounds δ_0 defined in Equation (11) can be expressed as

$$\delta_0 = \text{policy bias} + o\left(\frac{1}{\sqrt{N_L}}\right) + o\left(\frac{1}{\sqrt{N_\delta}}\right) + o\left(\frac{1}{\sqrt{N_i}}\right).$$

The computer budget \mathcal{B} for Algorithm 1 is equal to

$$\mathcal{B} = N_L T_{\text{path}} + (d - 1)N_\delta(N_i T_{\text{path}} + T_{\text{path}}).$$

Minimizing the width of the confidence interval for a fixed budget \mathcal{B} , we have:

choose N_δ, N_L to minimize: δ_0

under the constraint:

$$N_L T_{\text{path}} + (d - 1)N_\delta \cdot (N_i T_{\text{path}} + T_{\text{path}}) \leq \mathcal{B}. \quad (16)$$

The solution is given by

$$yN_\delta = O(N_i) = O\left(\sqrt{\frac{\mathcal{B}}{T_{\text{path}}(d - 1)}}\right),$$

$$N_L = O\left(\left(\frac{\mathcal{B}}{T_{\text{path}}}\right)^{5/6} \frac{1}{(d - 1)^{1/6}}\right).$$

Assuming the policy bias is zero gives

$$\delta_0 \approx O\left(\left(\frac{T_{\text{path}}}{\mathcal{B}}\right)^{1/4}\right).$$

The difference $|\hat{Q}(x) - Q(x)|$ can be approximated by δ_0 . The width of the confidence interval achieved for the first-order price sensitivity, Δ , is obtained by Equation (15) and is described by the following corollary.

COROLLARY 2. *If the policy bias is equal to zero, for a fixed computer budget \mathcal{B} , Algorithm 3 achieves an asymptotic order of convergence of the order $O(\mathcal{B}^{-1/6})$.*

Similarly, if the policy bias D_∞ is not zero, the width of the confidence interval is described by the following corollary.

COROLLARY 3. *If the policy bias D_∞ is greater than zero, for a fixed computer budget \mathcal{B} , Algorithm 3 achieves an asymptotic order of convergence of the order $O(D_\infty + \mathcal{B}^{-1/6})^{2/3}$.*

We point out that, based on the discussion in the previous section, the LRD algorithm dominates the alternative. In particular, in the case of zero policy bias, the asymptotic width of the confidence interval, computed using the LRD algorithm, is $O(\mathcal{B}^{-1/4})$. For nonzero policy bias, the asymptotic width of the confidence interval computed using the LRD algorithm is $D_\infty + O(\mathcal{B}^{-1/4})$. If we require that the two algorithms achieve the same width for the confidence interval for the first sensitivity Δ , Algorithm 3 requires the time required by the LRD algorithm, raised to the power 3/2. Thus, the LRD algorithm dominates Algorithm 3 both in terms of accuracy and in terms of computer time. We point out that in the case of higher-order derivatives, the LRD algorithm has an even larger advantage over Algorithm 3.

4.2. A Fast, Heuristic Alternative Algorithm

To speed up the calculation of sensitivities, we propose a heuristic algorithm together with some justification for its accuracy. Although the algorithm does not achieve a confidence interval for the sensitivities, it requires significantly less time, and we can use the confidence intervals produced by the LRD algorithm to numerically study the accuracy of the heuristic algorithm.

ALGORITHM 4. Heuristic estimation of option price sensitivities.

Step 1. Find an approximate exercise strategy.

Step 2. Simulate N Monte Carlo paths for the stock and bond prices between time t_0 and time t_1 . For each path i , calculate, at time t_1 , ξ_Δ^i , ξ_Γ^i , as well as the approximation to the option price A^i , according to the approximate exercise strategy from Step 1.

Step 3. To estimate the option price sensitivities at time t_0 , average the results from Step 2 according to Equation (6), using ξ_Δ^i , ξ_Γ^i and the approximate prices A^i from Step 2.¹¹

We note that Step 2 requires the calculation of approximate option prices A^i for N different initial values for the stock price. This calculation may require additional nested Monte Carlo simulations. However, if one uses the algorithm proposed by Longstaff and Schwartz (2001), these additional nested simulations are unnecessary because an approximate lower bound for the option price can be obtained using the regression coefficients of the basis functions at time t_1 . An alternative algorithm would be to choose an approximate exercise policy and then treat the discretely exercisable option as an Asian-style option, with the payoff given by the random time when the path crosses the approximate exercise boundary.^{12,13} This algorithm would have the advantage of not requiring an approximation to the option price at time t_1 for all states, but is otherwise similar to Algorithm 4.

Algorithm 4 is based on the following argument: if we let $\Delta^{\text{approximate}}$ denote the approximate value of Δ calculated using the algorithm, V the true option price at time t_1 , and θ

the difference between V and the approximation price A , $\theta = V - A$, then, from Equation (6), we have that

$$\begin{aligned} & |\Delta(x) - \Delta(x)^{\text{approximate}}| \\ &= |\mathbb{E}_0[e^{-\int_0^{t_1} r(t) dt} (V(S(t_1)) - A(S(t_1))) \xi_\Delta]| \\ &= |\mathbb{E}_0[e^{-\int_0^{t_1} r(t) dt} \theta(S(t_1)) \xi_\Delta]| \\ &\leq (\mathbb{E}_0[(e^{-\int_0^{t_1} r(t) dt} \theta(S(t_1)))^2])^{1/2} (\mathbb{E}_0[(\xi_\Delta)^2])^{1/2}, \end{aligned}$$

where the last line follows from the Schwartz inequality.

Estimating the expectation $\mathbb{E}_0[(e^{-\int_0^{t_1} r(t) dt} \theta(S(t_1)))^2]$ requires an algorithm similar to the LRD algorithm.¹⁴

One important advantage of having developed the LRD algorithm is that we can use the confidence intervals obtained by it as a measure of accuracy for Algorithm 4. For example, for a particular option of interest, one could calculate accurate confidence intervals using the LRD algorithm, and then compare the results with those obtained from Algorithm 4. If the accuracy of Algorithm 4 is satisfactory, then it can be used for the calculation of option sensitivities in similar situations. In §5.5, we undertake such an approach and numerically study the performance of Algorithm 4.

5. Numerical Results

In this section, we present numerical results of the behavior of the LRD algorithm and the heuristic Algorithm 4. Our examples are similar to the ones described in Andersen and Broadie (2004), but in contrast to Andersen and Broadie (2004), we focus on the sensitivities of option prices rather than the prices themselves.

5.1. Effect of the Number of Assets

The interest in using Monte Carlo simulation is mostly driven by the possibility of working with options that depend on many underlying assets because, in this case, alternative algorithms based on partial differential equations and multidimensional lattices suffer from an exponential increase in effort for a linear increase in the number of underlying assets. To study the performance of the LRD algorithm for a large number of underlying assets, we consider an example of an option written on the maximum of n risky assets. The option payoff is given by

$$h(S_1, \dots, S_n) = (\max(S_1, \dots, S_n) - K)^+,$$

and is called a *max-call option*. We consider the case where the assets S_1, \dots, S_n follow the dynamics

$$\frac{dS_i}{S_i} = (r - \delta)dt + \sigma dW_i, \quad i = 1, \dots, n,$$

where the W_i are standard, uncorrelated Brownian motions under the risk-neutral measure and the initial conditions for the asset prices are taken to be equal; i.e., $S_i(0) = S_0$,

$i = 1, \dots, n$. We concentrate on a set of parameters used in the paper by Andersen and Broadie (2004), where the interest rate is $r = 5\%$ per year, the dividend rate is $\delta = 10\%$ per year, the annualized volatility is $\sigma = 20\%$, and the strike price is $K = \$100$. The option can be exercised at times $T/d, 2T/d, \dots, T$, where $d = 9$ and $T = 3$ years.

To calculate confidence intervals for the prices and the sensitivities, we construct an exercise strategy based on the algorithm by Longstaff and Schwartz (2001), using the following basis functions: a constant function, the largest and second-largest asset prices, monomials of degree one, two, and three in the asset prices, the price of a European max-call option on the largest of n assets, and the square and cube of this price.

In Table 1, we present our numerical results for the cases of 2, 3, 5, and 10 assets, and for initial conditions that are out of the money (all initial asset prices equal to \$90), at the money (all initial asset prices equal to \$100), and in the money (all initial asset prices equal to \$110). We implemented the algorithm in a parallel cluster of CPUs, access to which was provided by the Texas Advanced Computing Center, and verified that the algorithm was scalable; i.e., that computational speed increased approximately linearly with an increase in the number of processors. The computer time for all cases was set to the equivalent of 10 CPU hours for a single processor 3.3 GHz PC computer. The calculations were performed using between 16 and 32 1 GHz processors, for a total time between one and two hours.

The results indicate that even for a large number of assets, one can achieve tight confidence intervals for both the option prices and the sensitivities of the option prices.¹⁵ In percentage terms, the tightness of the confidence intervals does not appear to depend on the moneyness of the option. We point out that the results in Table 1 indicate that the approximate exercise strategies are relatively accurate, given the tightness of the confidence intervals for the sensitivities.

The tightness of the bounds deteriorates with the order of the sensitivity. This deterioration is expected, and is a direct consequence of the fact that the variance of the weights used to estimate the sensitivities increases with the order of the sensitivity.

5.2. Effect of the Number of Exercise Opportunities

To investigate the effect of the number of exercise opportunities on the accuracy of the LRD algorithm, we consider the same example as in the previous section and vary the number of exercise opportunities, keeping the time to expiration fixed at three years. Because the magnitude of the weights ξ_Δ, ξ_Γ from Equation (6) increases as the time between exercise opportunities decreases, and because we keep the computer budget fixed, we expect a deterioration of the sensitivity estimates as the number of exercise opportunities increases.

Table 1. Bounds for sensitivities of high-dimensional options.

| S_0 | L_0 | U_0 | $(\Delta_1)_{\text{lower}}$ | $(\Delta_1)_{\text{upper}}$ | $(\Gamma_{11})_{\text{lower}}$ | $(\Gamma_{11})_{\text{upper}}$ | $(\Gamma_{12})_{\text{lower}}$ | $(\Gamma_{12})_{\text{upper}}$ |
|----------|-------------------|-------------------|-----------------------------|-----------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|
| $n = 2$ | | | | | | | | |
| 90 | 8.062 (0.003) | 8.079 (0.003) | 0.246 (0.002) | 0.248 (0.002) | 0.0117 (0.0003) | 0.0120 (0.0003) | -0.0028 (0.0002) | -0.0026 (0.0002) |
| 100 | 13.890 (0.003) | 13.913 (0.003) | 0.331 (0.002) | 0.334 (0.002) | 0.0125 (0.0003) | 0.0130 (0.0003) | -0.0055 (0.0002) | -0.0052 (0.0002) |
| 110 | 21.327 (0.004) | 21.362 (0.004) | 0.404 (0.003) | 0.412 (0.003) | 0.0127 (0.0004) | 0.0135 (0.0004) | -0.0077 (0.0003) | -0.0071 (0.0003) |
| $n = 3$ | | | | | | | | |
| 90 | 11.266 (0.003) | 11.287 (0.003) | 0.213 (0.002) | 0.215 (0.002) | 0.0101 (0.0003) | 0.0103 (0.0003) | -0.0021 (0.0001) | -0.0019 (0.0001) |
| 100 | 18.667 (0.004) | 18.705 (0.004) | 0.273 (0.002) | 0.277 (0.002) | 0.0103 (0.0003) | 0.0108 (0.0003) | -0.0031 (0.0002) | -0.0028 (0.0002) |
| 110 | 27.528 (0.004) | 27.586 (0.004) | 0.314 (0.003) | 0.319 (0.003) | 0.0101 (0.0003) | 0.0107 (0.0003) | -0.0041 (0.0002) | -0.0037 (0.0002) |
| $n = 5$ | | | | | | | | |
| 90 | 16.614 (0.004) | 16.649 (0.004) | 0.173 (0.001) | 0.177 (0.001) | 0.0084 (0.0002) | 0.0088 (0.0002) | -0.0013 (0.0001) | -0.0011 (0.0001) |
| 100 | 26.107 (0.004) | 26.171 (0.004) | 0.200 (0.002) | 0.205 (0.002) | 0.0079 (0.0002) | 0.0085 (0.0002) | -0.0016 (0.0001) | -0.0013 (0.0001) |
| 110 | 36.711 (0.005) | 36.805 (0.005) | 0.218 (0.002) | 0.224 (0.002) | 0.0076 (0.0003) | 0.0083 (0.0003) | -0.0017 (0.0001) | -0.0013 (0.0001) |
| $n = 10$ | | | | | | | | |
| 90 | 26.218 (0.004) | 26.279 (0.004) | 0.115 (0.001) | 0.121 (0.001) | 0.0059 (0.0002) | 0.0066 (0.0002) | -0.0008 (0.00005) | -0.0003 (0.00005) |
| 100 | 38.283 (0.005) | 38.367 (0.005) | 0.121 (0.002) | 0.129 (0.002) | 0.0053 (0.0002) | 0.0061 (0.0002) | -0.0008 (0.00005) | -0.0003 (0.00005) |
| 110 | 50.815 (0.005) | 50.921 (0.005) | 0.122 (0.002) | 0.130 (0.002) | 0.0052 (0.0003) | 0.0060 (0.0003) | -0.0008 (0.00006) | -0.0003 (0.00006) |

Notes. Bounds for prices and sensitivities of Bermudan max-call options with $n = 2, 3, 5, 10$ assets. The payoff of the option is $(\max(S_1(t), \dots, S_n(t)) - K)^+$. The parameters are $K = 100$, $r = 5\%$, $\delta = 10\%$, $T = 3$, $\sigma = 20\%$. The initial price vector is $S(0) = (S_0, \dots, S_0)$, with $S_0 = 90, 100$, or 110 . There are nine exercise opportunities at $t_i = iT/d$, $i = 1, \dots, 9$, $d = 9$. The lower and upper bounds for the option price (L_0 , U_0 , respectively), as well as the lower and upper bounds for the first derivative, Δ , of the option price with respect to the initial asset price, and lower and upper bounds for the second derivative, Γ , are provided. Because the problem is symmetric with respect to the n assets, we only provide the first component of Δ , Δ_1 , and one diagonal term Γ_{11} and one off-diagonal term Γ_{12} for the second derivative Γ . Standard errors are given in parentheses below the estimated values. The 3.3 GHz-PC-equivalent computational time for all the calculations of the sensitivities was set to 10 hours. The calculations were performed using between 16 and 32 1 GHz processors in parallel, for a total time between one and two hours.

We point out that when there are only two exercise opportunities, the LRD algorithm automatically produces tight estimates of the upper bound of the option price sensitivities. Indeed, for the option payoff we consider, the upper bound for the option price, calculated from Algorithm 1, is independent of the exercise policy. This result is stated in the following proposition.

LEMMA 1. *For the case of a twice-exercisable option at times t_1, t_2 , with $0 < t_1 < t_2 = T$, the upper bound for the option price, obtained by Algorithm 1, described in §3.2.2, is independent of the choice of the exercise strategy, and is an unbiased estimate of the true price of the discretely exercisable option.*

The online appendix provides the proof of Lemma 1 in detail.

The intuition behind Lemma 1 is that, with only two exercise opportunities, the only nontrivial decision is made on the next-to-last exercise date. Because the optimal decision at that time is based on the comparison between the immediate exercise value and the discounted continuation value, and because the nested Monte Carlo simulation used in calculating the upper bound provides an unbiased estimate of the discounted continuation value, the estimate of the upper bound is an unbiased estimate of the true option price. Because the upper bound for the option price is an unbiased estimate of the true option price, the upper bound for the sensitivity of the option price is also close to the

Table 2. Dependence on the number of exercise opportunities.

| d | L_0 | U_0 | $(\Delta_1)_{\text{lower}}$ | $(\Delta_1)_{\text{upper}}$ | $(\Gamma_{11})_{\text{lower}}$ | $(\Gamma_{11})_{\text{upper}}$ | $(\Gamma_{12})_{\text{lower}}$ | $(\Gamma_{12})_{\text{upper}}$ |
|----------|-------------------|-------------------|-----------------------------|-----------------------------|--------------------------------|--------------------------------|--------------------------------|--------------------------------|
| $n = 2$ | | | | | | | | |
| 5 | 13.63 (0.03) | 13.64 (0.03) | 0.328 (0.001) | 0.337 (0.001) | 0.0127 (0.0001) | 0.0135 (0.0001) | -0.0050 (0.00008) | -0.0044 (0.00008) |
| 10 | 13.92 (0.03) | 13.94 (0.03) | 0.335 (0.002) | 0.338 (0.002) | 0.0128 (0.0003) | 0.0132 (0.0003) | -0.0047 (0.0002) | -0.0044 (0.0002) |
| 15 | 14.03 (0.03) | 14.07 (0.03) | 0.339 (0.003) | 0.343 (0.003) | 0.0128 (0.0005) | 0.0134 (0.0005) | -0.0054 (0.0004) | -0.0050 (0.0004) |
| 20 | 14.03 (0.03) | 14.07 (0.03) | 0.338 (0.004) | 0.344 (0.004) | 0.0120 (0.0009) | 0.0129 (0.0009) | -0.0059 (0.0006) | -0.0053 (0.0005) |
| 30 | 14.09 (0.03) | 14.15 (0.03) | 0.336 (0.007) | 0.348 (0.007) | 0.0134 (0.0016) | 0.0156 (0.0016) | -0.0030 (0.0011) | -0.0015 (0.0011) |
| $n = 3$ | | | | | | | | |
| 5 | 18.41 (0.04) | 18.44 (0.04) | 0.268 (0.001) | 0.276 (0.001) | 0.0106 (0.0002) | 0.0114 (0.0002) | -0.0032 (0.0001) | -0.0027 (0.0001) |
| 10 | 18.75 (0.04) | 18.79 (0.04) | 0.268 (0.002) | 0.271 (0.002) | 0.0106 (0.0003) | 0.0111 (0.0003) | -0.0034 (0.0002) | -0.0031 (0.0002) |
| 15 | 18.858 (0.04) | 18.913 (0.04) | 0.274 (0.002) | 0.280 (0.002) | 0.0103 (0.0004) | 0.0110 (0.0004) | -0.0034 (0.0001) | -0.0029 (0.0001) |
| 20 | 18.913 (0.04) | 18.978 (0.04) | 0.274 (0.004) | 0.283 (0.004) | 0.0111 (0.0008) | 0.0124 (0.0009) | -0.0031 (0.0004) | -0.0022 (0.0004) |
| 30 | 18.976 (0.04) | 19.060 (0.04) | 0.263 (0.007) | 0.277 (0.007) | 0.0060 (0.0016) | 0.0087 (0.0016) | -0.0032 (0.0008) | -0.0014 (0.0008) |
| $n = 5$ | | | | | | | | |
| 5 | 25.739 (0.002) | 25.791 (0.002) | 0.199 (0.001) | 0.206 (0.001) | 0.0082 (0.0001) | 0.0088 (0.0001) | -0.0017 (0.00003) | -0.0013 (0.00003) |
| 10 | 26.161 (0.002) | 26.229 (0.002) | 0.203 (0.002) | 0.208 (0.002) | 0.0076 (0.0003) | 0.0082 (0.0003) | -0.0015 (0.0001) | -0.0011 (0.0001) |
| 15 | 26.308 (0.002) | 26.388 (0.002) | 0.197 (0.003) | 0.205 (0.003) | 0.0071 (0.0005) | 0.0083 (0.0005) | -0.0017 (0.0002) | -0.0010 (0.0002) |
| 20 | 26.378 (0.004) | 26.472 (0.004) | 0.200 (0.005) | 0.212 (0.005) | 0.0079 (0.0009) | 0.0099 (0.0009) | -0.0024 (0.0003) | -0.0011 (0.0003) |
| 30 | 26.453 (0.004) | 26.561 (0.005) | 0.189 (0.007) | 0.210 (0.007) | 0.0071 (0.002) | 0.0111 (0.002) | -0.0040 (0.0006) | -0.0013 (0.0006) |
| $n = 10$ | | | | | | | | |
| 5 | 37.871 (0.002) | 37.947 (0.002) | 0.116 (0.001) | 0.131 (0.001) | 0.0051 (0.0002) | 0.0059 (0.0002) | -0.00100 (0.00004) | -0.00035 (0.00004) |
| 10 | 38.352 (0.002) | 38.446 (0.002) | 0.119 (0.002) | 0.128 (0.002) | 0.0049 (0.0003) | 0.0059 (0.0003) | -0.00093 (0.00006) | -0.00031 (0.00006) |
| 15 | 38.524 (0.002) | 38.630 (0.002) | 0.116 (0.003) | 0.129 (0.003) | 0.0049 (0.0006) | 0.0067 (0.0006) | -0.0014 (0.0001) | -0.0003 (0.0001) |
| 20 | 38.621 (0.005) | 38.729 (0.005) | 0.117 (0.005) | 0.135 (0.005) | 0.0055 (0.0009) | 0.0084 (0.0009) | -0.0019 (0.0002) | 0.0000 (0.0002) |
| 30 | 38.711 (0.005) | 38.835 (0.005) | 0.102 (0.008) | 0.132 (0.008) | 0.0033 (0.002) | 0.0092 (0.002) | -0.0024 (0.0004) | 0.0015 (0.0004) |

Notes. Bounds for prices and sensitivities of Bermudan max-call options with $n = 5, 10$ assets, and number of exercise opportunities between $d = 5, 10, 15, 20, 30$. The payoff, parameters and notations are the same as in Table 1. Standard errors are given in parentheses below the estimated values. The 3.3 GHz-PC-equivalent computational time for all the calculations of the sensitivities was set to 10 hours. The calculations were performed using between 16 and 32 1 GHz processors in parallel, for a total time between one and two hours.

true sensitivity (although it is not an unbiased estimate). This is due to the fact that to obtain an upper bound for Δ and Γ , we multiply either the upper or the lower bound for the option price by a weight function, depending on the sign of the weight function. For the max-call payoff

we consider, positive values of the weight function correspond to situations where the option is in the money, and negative values to situations when the option is out of the money. Because out-of-the-money options are relatively cheap, the biggest contribution to the upper bound of the

option price sensitivities arises from cases when the weight is positive. Because in those cases the weight multiplies the upper bound for the option price, which is an unbiased estimate of the true price, the upper bound for the sensitivity of the option price is affected relatively little by the exercise strategy.

The confidence intervals for the option price sensitivities, with a varying number of exercise opportunities, are given in Table 2. The confidence intervals remain relatively tight, even as the number of exercise opportunities increases. The computer budget was the same in all cases, set to the equivalent of 10 CPU hours for a single processor 3.3 GHz PC computer. As in the previous subsection, the calculations were performed using between 16 and 32 1 GHz processors, for a total time between one and two hours.

5.3. Correlated Assets

In Table 3, we study the behavior of the LRD algorithm for the case of an option on the maximum of two correlated assets, for different values of the correlation. We are interested in whether the confidence intervals for the various sensitivities vary with the correlation between the assets. The model for the asset prices is given by

$$\frac{dS_1}{S_1} = (r - \delta)dt + \sigma dW_1, \quad \frac{dS_2}{S_2} = (r - \delta)dt + \sigma dW_2,$$

where W_1, W_2 are standard correlated Brownian motions, with correlation ρ . The initial values of the two assets are set to the same value $S_1(0) = S_2(0) = S$.

For this case, we need to slightly modify the weights given in Equation (6) to account for the correlation. The new weights are given by

$$\begin{aligned} \xi_{\Delta_1} &= \frac{W_1(t_1) - W_2(t_1)}{S\sigma t_1(1 - \rho^2)}, \\ \xi_{\Delta_2} &= \frac{W_2(t_1) - W_1(t_1)}{S\sigma t_1(1 - \rho^2)}, \\ \xi_{\Gamma_{11}} &= \frac{(W_1(t_1) - W_2(t_1)\rho)^2 + t_1(\rho^2 - 1)(1 + W_1(t_1)\sigma - W_2(t_1)\rho\sigma)}{t_1^2 S^2 (\rho^2 - 1)^2 \sigma^2}, \\ \xi_{\Gamma_{22}} &= \frac{(W_2(t_1) - W_1(t_1)\rho)^2 + t_1(\rho^2 - 1)(1 + W_2(t_1)\sigma - W_1(t_1)\rho\sigma)}{t_1^2 S^2 (\rho^2 - 1)^2 \sigma^2}, \\ \xi_{\Gamma_{12}} &= \xi_{\Gamma_{21}} \\ &= \frac{W_1(t_1)W_2(t_1)(1 + \rho^2) + t_1\rho(1 - \rho^2) - (W_1^2(t_1) + W_2^2(t_1))\rho}{t_1^2 S^2 (\rho^2 - 1)^2 \sigma^2}, \end{aligned} \tag{17}$$

where t_1 is the first exercise date of the discretely exercisable option.

In Table 3, we vary the correlation between the two assets from -75% to 75% . For each case, the computer budget is set to the equivalent of one CPU hour of a single processor 3.3 GHz computer. We note that although the width of the confidence intervals for the option price sensitivities remain relatively small as the correlation changes, the widths increase as the correlation tends to $\pm 100\%$. This increase is due to the increased variance of the correlated weights, as can be observed from Equation (17).

Table 3. Correlated assets.

| Correlation (%) | L_0 | U_0 | $(\Delta_1)_{lower}$ | $(\Delta_1)_{upper}$ | $(\Gamma_{11})_{lower}$ | $(\Gamma_{11})_{upper}$ | $(\Gamma_{12})_{lower}$ | $(\Gamma_{12})_{upper}$ |
|-----------------|-------------------|-------------------|----------------------|----------------------|-------------------------|-------------------------|-------------------------|-------------------------|
| -75 | 15.440 (0.003) | 15.470 (0.003) | 0.384 (0.005) | 0.390 (0.005) | 0.0126 (0.0009) | 0.0136 (0.0009) | -0.0031 (0.0008) | -0.0022 (0.0008) |
| -50 | 15.014 (0.003) | 15.040 (0.003) | 0.375 (0.004) | 0.379 (0.004) | 0.0127 (0.0006) | 0.0134 (0.0006) | -0.0036 (0.0005) | -0.0030 (0.0005) |
| -25 | 14.508 (0.003) | 14.532 (0.003) | 0.354 (0.003) | 0.358 (0.003) | 0.0125 (0.0005) | 0.0131 (0.0005) | -0.0038 (0.0003) | -0.0034 (0.0003) |
| 0 | 13.890 (0.003) | 13.913 (0.003) | 0.331 (0.002) | 0.334 (0.002) | 0.0125 (0.0003) | 0.0130 (0.0003) | -0.0055 (0.0002) | -0.0052 (0.0002) |
| 25 | 13.088 (0.003) | 13.110 (0.003) | 0.310 (0.003) | 0.314 (0.003) | 0.0134 (0.0005) | 0.0139 (0.0005) | -0.0061 (0.0003) | -0.0057 (0.0003) |
| 50 | 12.141 (0.003) | 12.162 (0.003) | 0.288 (0.003) | 0.292 (0.003) | 0.0146 (0.0005) | 0.0153 (0.0005) | -0.0076 (0.0004) | -0.0071 (0.0004) |
| 75 | 10.885 (0.003) | 10.909 (0.003) | 0.267 (0.004) | 0.272 (0.004) | 0.0179 (0.0008) | 0.0189 (0.0008) | -0.0109 (0.0007) | -0.0101 (0.0007) |

Notes. Bounds for prices and sensitivities of Bermudan max-call options on two correlated assets. The payoff of the option is $(\max(S_1(t), S_2(t)) - K)^+$. The parameters are $K = 100, r = 5\%, \delta = 10\%, T = 3, \sigma = 20\%$. The initial prices for each asset are set to 100. There are nine exercise opportunities that are equally spaced at $t_i = iT/d, i = 1, \dots, 9, d = 9$. The correlation between the asset varies between -75% and 75% . The lower and upper bounds for the option price are provided (L_0, U_0 , respectively), as well as the lower and upper bounds for the first derivative, Δ , of the option price with respect to the initial asset price, and lower and upper bounds for the second derivative, Γ , with respect to the initial asset price. Because the problem is symmetric with respect to the two assets, we only provide the first component of Δ, Δ_1 , and one diagonal term (the second derivative with respect to the initial asset price of the first asset, Γ_{11}) for Γ . Standard errors are given in parentheses below the estimated values. The 3.3 GHz-PC-equivalent computational time for the calculations of the sensitivities was set to one hour.

5.4. Numerical Performance of the LRD Algorithm

5.4.1. Determining the Optimal Simulation Parameters. In §3.3.2, we described a framework for choosing the optimal simulation parameters based on the asymptotic scaling of the confidence intervals. One parameter that needs to be determined numerically is the ratio of the number of outer simulation paths N_o to the number of inner simulation paths N_i , denoted by λ . Panel A of Table 4 presents the dependence of the width of the confidence interval for different choices of this ratio, for the case of a max-call option on five identical assets. In each case, we have fixed the total computer budget to less than one minute, and calculated the width of the confidence interval for the Δ of the option. We then used interpolation to find the optimal value of λ . In particular, we fitted the width of the confidence intervals to the least-squares quadratic polynomial of the logarithm of the ratio λ . The minimum of that polynomial corresponds to the value $\lambda = 75$.

Although the optimal value of λ depends on the characteristics of the options; e.g., number of assets, moneyness, etc., in computations that we do not report we verified that even choices that are not optimal do not lead to significantly larger confidence intervals. This observation allows us to scale the values of the simulation parameters with the value of the computer budget without fine-tuning the value of λ . We have also verified that increasing the computer budget does not significantly change the estimate of the optimal value of λ . Once the optimal value is determined for a small computer budget, we fix the ratio of the number of outer simulation paths N_o to the number of inner simulation paths N_i to that value, and scale the values of N_o, N_i to exhaust the computer budget.

5.4.2. Scaling of the Width of the Confidence Intervals. In §3.2, we showed that for an optimal choice of the simulation parameters, the width of the confidence interval for a sensitivity is inversely proportional to the fourth power of the computer budget. To determine whether in our numerical implementation the width of the confidence interval follows this scaling relationship, we compute the values of option Δ s for different computer budgets, keeping the ratio of the number of outer simulation paths to inner simulation paths, λ , fixed.

In Table 4, Panel B, we present the widths of the confidence intervals for Δ , for the case of a max-call option on five identical assets. From the table, we can estimate the speed at which the width of the confidence interval decreases with an increase in the computer budget by regressing the logarithm of the width of the confidence interval on the logarithm of the computer budget. The strength of the relationship, measured by R^2 , indicates the extent to which asymptotic estimates are valid, whereas the value of the coefficient of the logarithm of the computer budget indicates the extent to which we achieve the theoretical value of -0.25 . The results are encouraging

Table 4. Numerical performance.

| λ | $(\Delta_1)_{\text{lower}}$ | $(\Delta_1)_{\text{upper}}$ | Std. dev. lower | Std. dev. upper | CI width |
|-----------------|-----------------------------|-----------------------------|-----------------|-----------------|----------|
| Panel A | | | | | |
| 4 | 0.190 | 0.198 | 0.024 | 0.024 | 0.102 |
| 25 | 0.193 | 0.208 | 0.016 | 0.016 | 0.078 |
| 100 | 0.191 | 0.215 | 0.011 | 0.011 | 0.067 |
| 400 | 0.182 | 0.227 | 0.008 | 0.008 | 0.076 |
| 1,000 | 0.171 | 0.241 | 0.007 | 0.007 | 0.097 |
| 2,500 | 0.156 | 0.255 | 0.005 | 0.005 | 0.119 |
| Computer budget | $(\Delta_1)_{\text{lower}}$ | $(\Delta_1)_{\text{upper}}$ | Std. dev. lower | Std. dev. upper | CI width |
| Panel B | | | | | |
| 1 | 0.1990 | 0.2071 | 0.0034 | 0.0034 | 0.0214 |
| 3 | 0.2013 | 0.2076 | 0.0026 | 0.0026 | 0.0165 |
| 5 | 0.2035 | 0.2092 | 0.0023 | 0.0023 | 0.0147 |
| 10 | 0.1998 | 0.2049 | 0.0019 | 0.0019 | 0.0125 |

Notes. Numerical performance of the LRD algorithm for estimating confidence intervals for the sensitivities of Bermudan max-call options with $n = 5$ assets and $d = 9$ exercise opportunities. The payoff and parameters are the same as in Table 1. In Panel A, we vary the ratio of outer to inner simulation paths, $\lambda = N_o/N_i$. We report the lower and upper estimates Δ , the standard deviations of the estimates, and the confidence interval that corresponds to $z = 1.96$. The time allocated to each different value of the ratio is under one minute on a 3.3 GHz PC computer. The value of the ratio that minimizes the width of the confidence interval is approximately equal to $\lambda = 75$. In Panel B, we fix the ratio to $\lambda = 75$, and vary the total computer budget available for estimating the option price sensitivities. We provide the estimate of the lower and upper bounds for Δ , the standard deviation of the bounds, and the confidence interval that corresponds to $z = 1.96$. The computer budget is expressed in terms of CPU hours for a 3.3 GHz PC computer.

because the estimate of R^2 is 99.98%, and the estimate of the coefficient of the logarithm of the computer budget is -0.232 ± 0.002 . The slight loss of efficiency, indicated by the slight deviation from the theoretical value of -0.25 , can be attributed to two reasons: (a) to an inaccurate choice of the optimal value of λ ; and (b) to efficiency losses in the parallelization of the program due to increased communication needs as the number of processors and the workload per processor increases. For this particular algorithm, communication needs are not significant because only a few values need to be communicated between the different processors.¹⁶

5.5. Numerical Study of the Accuracy of Algorithm 4

To determine the accuracy of Algorithm 4, we perform a numerical study similar to that of §5.1. For the same option payoff and the same parameter values, we estimate the sensitivities Δ and Γ of the option using Algorithm 4, and compare the results to the confidence intervals obtained by the LRD algorithm. The results are presented in Table 5.

Table 5. Estimates of sensitivities using a heuristic method.

| n | Time | Δ_1 | Δ_1 LRD interval | Γ_{11} | Γ_{11} LRD interval | Γ_{12} | Γ_{12} LRD interval |
|-----|------|------------------|-------------------------|--------------------|----------------------------|---------------------|----------------------------|
| 2 | 0.3 | 0.332 (0.005) | 0.327 to 0.338 | 0.0131 (0.0008) | 0.0120 to 0.0135 | -0.0046 (0.0005) | -0.0059 to -0.0048 |
| | 1 | 0.333 (0.003) | | 0.0131 (0.0004) | | -0.0049 (0.0003) | |
| | 3 | 0.334 (0.002) | | 0.0129 (0.0002) | | -0.0048 (0.0002) | |
| 3 | 0.3 | 0.268 (0.007) | 0.269 to 0.281 | 0.0105 (0.0009) | 0.0097 to 0.0114 | -0.0031 (0.0007) | -0.0034 to -0.0025 |
| | 1 | 0.266 (0.004) | | 0.0098 (0.0005) | | -0.0033 (0.0004) | |
| | 3 | 0.272 (0.002) | | 0.0107 (0.0003) | | -0.0030 (0.0002) | |
| 5 | 0.3 | 0.216 (0.010) | 0.196 to 0.209 | 0.0098 (0.0013) | 0.0074 to 0.0090 | -0.0012 (0.0009) | -0.0018 to -0.0011 |
| | 1 | 0.208 (0.006) | | 0.0087 (0.0008) | | -0.0013 (0.0005) | |
| | 3 | 0.205 (0.003) | | 0.0087 (0.0004) | | -0.0011 (0.0003) | |
| 10 | 0.3 | 0.127 (0.010) | 0.117 to 0.133 | 0.0047 (0.0009) | 0.0049 to 0.0065 | -0.0004 (0.0009) | -0.0009 to -0.0002 |
| | 1 | 0.120 (0.005) | | 0.0052 (0.0007) | | -0.0011 (0.0005) | |
| | 3 | 0.128 (0.003) | | 0.0056 (0.0003) | | -0.0004 (0.0003) | |

Notes. This table presents estimates of option sensitivities using the heuristic method of §4.2. The option payoff is that of a Bermudan max-call option, with $n = 2, 3, 5, 10$ assets. The payoff of the option is $(\max(S_1(t), \dots, S_n(t)) - K)^+$. The parameters are $K = 100$, $r = 5\%$, $\delta = 10\%$, $T = 3$, $\sigma = 20\%$. The initial price vector is $S(0) = (S_0, \dots, S_0)$, with $S_0 = 100$. There are nine exercise opportunities that are equally spaced at $t_i = iT/d$, $i = 1, \dots, 9$, $d = 9$. The sensitivities are given in columns $\Delta_1, \Gamma_{11}, \Gamma_{12}$. Time denotes the 3.3 GHz PC computer time equivalent, expressed in number of hours, used in each case. The probabilistic intervals for $\Delta_1, \Gamma_{11}, \Gamma_{12}$, are obtained from the estimates in Table 1 by subtracting 1.96 standard errors from the estimate of the lower bound and adding 1.96 standard errors to the estimate of the upper bound. Standard errors for $\Delta_1, \Gamma_{11}, \Gamma_{12}$, are given in parentheses below the estimated values.

From Table 5, we note that the values estimated by Algorithm 4 in all cases fall within the confidence intervals. In addition, the time required for Algorithm 4 is only a fraction of that required for the LRD algorithm. One needs to use caution, however. For example, the standard errors produced by Algorithm 4 are misleading. Indeed, the limit to which the sensitivities computed using Algorithm 4 converge is not necessarily the correct sensitivity value. It is important to note that having the confidence intervals computed using the LRD algorithm is what enables us to determine whether Algorithm 4 is accurate. After validating the accuracy of Algorithm 4 in several cases, it can be used to quickly estimate sensitivity values in similar situations.

6. Conclusions

We have presented a new algorithm, called the LRD algorithm, for obtaining confidence intervals for the values of sensitivities of the price of a discretely exercisable option with respect to the initial price of the underlying assets. The

algorithm combines a Monte Carlo algorithm for computing sensitivities for European options, and a Monte Carlo algorithm that obtains a confidence interval for the price of discretely exercisable options based on the dual representation of the optimal stopping problem associated with the calculation of the option price. The only input necessary for the algorithm is an approximate exercise policy, which can also be obtained using Monte Carlo simulation. We have shown that the accuracy and speed of the LRD algorithm is asymptotically of the same order as the accuracy and speed for the duality-based algorithm, proposed by Andersen and Broadie (2004) for estimating confidence intervals for the price of the option. We have also shown that the LRD algorithm is superior, both in terms of speed and accuracy, to an alternative algorithm based on a finite-difference approximation of the option price sensitivities.

The LRD algorithm can be applied in cases with complicated price dynamics, and, largely, arbitrary payoff functions. It can be used for options with a large number of underlying assets. In contrast to algorithms using finite

difference discretization of partial differential equations and high-dimensional lattice algorithms, it does not suffer from the curse of dimensionality.

In addition, the LRD algorithm can be used to evaluate the accuracy of alternative algorithms. Based on such an evaluation, we have described a simpler, faster, heuristic, alternative algorithm. Although the alternative algorithm does not provide confidence intervals, using the LRD algorithm as a measure of accuracy, we have found the alternative to be accurate in numerical simulations.

In future work, we plan to extend the algorithm to estimate confidence intervals for the values of sensitivities of the value function for other stochastic control problems with optimal stopping.

7. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Endnotes

1. These sensitivities are frequently referred to as “Greeks.”
2. This result is similar to the case of European options, where the width of the confidence interval for the price of an option decreases at the same rate as the width of the confidence interval for option price sensitivities computed using the likelihood-ratio algorithm.
3. The assumption of independence is made without loss of generality. When the components W_i are not independent, one can first perform a transformation to construct a process with independent components. We present a numerical example with correlated assets in §5.3.
4. An alternative algorithm for calculating sensitivities of options as weighted expectations of option payoffs was developed by Fournié et al. (1999). The algorithm described in Fournié et al. (1999) is based on Malliavin calculus. Although both algorithms demonstrate that a derivative can be expressed in terms of a weighted expectation of the payoff, the algorithms are not identical. The algorithm based on Malliavin calculus may produce more than one weight function for each sensitivity and can lead to easier analytic computations when the coefficients of the diffusion process are not constant. On the other hand, the likelihood-ratio algorithm depends only on the regularity of the transition density and can be used in cases when the underlying stochastic process includes jumps. The weights provided by Fournié et al. (1999) for the case of geometric Brownian motion with constant coefficients are identical to the ones calculated using the likelihood-ratio algorithm. In this paper, we focus on the representation of the sensitivities of option prices as weighted expectations of the option payoffs and are agnostic regarding the algorithm used to derive the weights.

5. These options are also called Bermudan, or, in the limit where the time between possible exercises tends to zero, American.

6. We refer the reader to Davis and Karatzas (1994), Rogers (2002), Haugh and Kogan (2004), and Andersen and Broadie (2004) for the proof.

7. We note that the algorithm for the estimation of a confidence interval for the option price described in Haugh and Kogan (2004) appears, according to the numerical results in Haugh and Kogan (2004), to be faster than the one used in Andersen and Broadie (2004). One of the differences between the algorithms is that the Haugh and Kogan (2004) algorithm is based on neural networks and low-discrepancy sequences. The use of such methods results in a considerable speed-up, but, unfortunately, makes it difficult to estimate the complexity of the algorithm. In this paper, we use the Andersen and Broadie (2004) algorithm because it allows us to optimize the choice of simulation parameters for a given computer budget.

8. The exclusion of exercise at time t_0 is not a restriction because it is trivial to calculate the sensitivity of an option price for an option that is exercised immediately.

9. The choice of N_L is restricted to be less than $O(\sqrt{\mathcal{B}})$, but does not otherwise influence the width of the confidence interval for the option price sensitivities, to first order in the computer budget.

10. The estimation of the asymptotic width of the confidence interval for option prices follows arguments similar to the ones used to estimate the asymptotic width of the optimal confidence interval for the option price sensitivities.

11. The values A^i play the role of the payoff $h(S(T))$ in Equation (6).

12. Fournié et al. (1999) provide the appropriate weights for the case of an Asian option.

13. We thank the associate editor for suggesting this alternative.

14. Unlike the LRD algorithm, Algorithm 4 does not provide rigorous confidence intervals for the sensitivities of option prices. To obtain such estimates, one would need to show a relation of the form

$$\mathbb{E}_0[(e^{-\int_0^{t_1} r(t) dt} \theta(S(t_1)))^2]^{1/2} \leq U_0 - L_0, \quad (18)$$

where U_0, L_0 are the upper and lower bounds for the option price at time t_0 . Then Algorithm 4, combined with the estimates of U_0, L_0 , would produce confidence intervals for the sensitivities. However, we point out that verifying (18) can be as time consuming as the LRD algorithm itself: as discussed in §3.2, for the same computer budget, the width of the confidence interval achieved for $U_0 - L_0$ is asymptotically of the same order as the width of the confidence interval for the sensitivities computed by the LRD algorithm. This means that the width of confidence intervals computed using Algorithm 4 and an estimate of $U_0 - L_0$ would decrease at the same speed as the width of the confidence interval computed using the LRD algorithm.

15. The estimates of the lower and upper bounds and the standard errors reported in Table 1 take into account the symmetry in the payoff of the options, and in the dynamics of the asset prices—for example, only one value for Δ and only two values for Γ are reported.

16. When calculating an expected value using Monte Carlo simulation in parallel, one does not need to communicate all the option values calculated along each path. Instead, only the number of paths used by each processor and the average value and variance need to be communicated.

Acknowledgments

The authors thank Mark Broadie, Rama Cont, Michael Gallmayer, and Martin Haugh for helpful comments and discussions. They also thank seminar participants at the University of Texas at Austin, University of Michigan, Princeton University, Universidad Autonoma Metropolitana, Instituto Tecnológico Autónomo de México, University of Pittsburgh, Duke University, the American Mathematical Society meeting at Orlando, Florida, the 13th Annual Derivative Securities Conference, the Banff International Research Station Workshop on Semimartingale Theory and Practice in Finance, and the Third World Congress of the Bachelier Society for suggestions. In addition, the authors thank the Texas Advanced Computing Center for providing computing resources.

References

- Andersen, L., M. Broadie. 2004. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Sci.* **50** 1222–1234.
- Broadie, M., P. Glasserman. 1996. Estimating security price derivatives using simulation. *Management Sci.* **42** 269–285.
- Carriere, J. F. 1996. Valuation of the early-exercise price for derivative securities using simulation and splines. *Insurance: Math. Econom.* **19** 19–30.
- Davis, M., I. Karatzas. 1994. A deterministic approach to optimal stopping, with applications. P. Kelly, ed. *Probability, Statistics, and Optimization: A Tribute to Peter Whittle*. John Wiley & Sons, New York and Chichester, 455–466.
- Fournié, E., J.-M. Lasry, J. Lebuchoux, P.-L. Lions, N. Touzi. 1999. Applications of Malliavin calculus to Monte Carlo methods in finance. *Finance Stochastics* **3** 391–412.
- Glasserman, P. 2003. *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- Glynn, P. W. 1989. Optimization of stochastic systems via simulation. *Proc. 1989 Winter Simulation Conf.*, Society for Computer Simulation, San Diego, 90–105.
- Haugh, M. B., L. Kogan. 2004. Pricing American options: A duality approach. *Oper. Res.* **52** 258–270.
- Longstaff, F., E. Schwartz. 2001. Valuing American options by simulation: A simple least squares approach. *Rev. Financial Stud.* **14** 113–147.
- Rogers, L. C. G. 2002. Monte Carlo valuation of American options. *Math. Finance* **12** 271–286.
- Tsitsiklis, J. N., B. Van Roy. 1999. Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms and an application to pricing high-dimensional financial derivatives. *IEEE Trans. Automatic Control* **44** 1840–1851.