JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 55, No. 2, Mar. 2020, pp. 653–677
COPYRIGHT 2019, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109019000048

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

Real-Option Valuation in Multiple Dimensions Using Poisson Optional Stopping Times

Rutger-Jan Lange, Daniel Ralph, and Kristian Støre*

Abstract

We provide a new framework for valuing multidimensional real options where opportunities to exercise the option are generated by an exogenous Poisson process, which can be viewed as a liquidity constraint on decision times. This approach, which we call the Poisson optional stopping times (POST) method, finds the value function as a monotone sequence of lower bounds. In a case study, we demonstrate that the frequently used quasi-analytic method yields a suboptimal policy and an inaccurate value function. The proposed method is demonstrably correct, straightforward to implement, reliable in computation, and broadly applicable in analyzing multidimensional option-valuation problems.

I. Introduction

There is a large and growing body of literature focused on the valuation of options with multiple state variables (e.g., Trigeorgis (1991), (1993), Rogers (2002), Stentoft (2004), Andersen and Broadie (2004), Bally and Printems (2005), and Cortazar, Gravet, and Urzua (2008)). This article focuses on the valuation of real options of the American type with no expiration date (i.e., perpetual options), but also encompasses options with finite maturity. The resulting option-valuation problems are notoriously difficult to solve, because the optimal policy and resulting value function are intricately linked and analytic solutions can be found only in special cases. Researchers may resort to approximate analytic solutions, such as for the American put option in Bunch and Johnson (2000) and the European spread option in Bjerksund and Stensland (2014). Alternatively, they may employ numerical methods, such as classic tree-/lattice approaches as in Cox, Ross, and Rubinstein (1979), finite-difference schemes as in Schwartz (1977), or simulation approaches as in Longstaff and Schwartz (2001).

There is a trade-off between these two approaches: While analytic approximations are usually faster and easier to inspect, numerical solution methods tend

*Lange (corresponding author), lange@ese.eur.nl, Erasmus University Rotterdam Econometric Institute; Ralph, d.ralph@jbs.cam.ac.uk, University of Cambridge Judge Business School; and Støre, kristian.store@nord.no, Nord University Business School. We thank two anonymous referees, Jan van Casteren, Dick van Dijk, Øystein Gjerde, David Mauer (a referee), and Coen Teulings for helpful comments. We also thank Hendrik Bessembinder (the editor) for his helpful guidance.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

to be more flexible and possibly more accurate. This flexibility allows researchers to explore different models, thereby reducing the incentive to (over)simplify or approximate the original model for the sake of maintaining analytic tractability, which may come at the expense of model error. This view is also taken in the recently developed work in Lange, Ralph, and van Casteren (2018), which seeks to solve the classic real-options case, where decision makers may exercise the option at any point in time. We contribute an algorithm for computing explicit solutions in a straightforward and reliable manner. In this way we respond to Strulovici and Szydlowski's (2015) call for "a better understanding of the properties of optimal policies and value functions with a multidimensional state space" by "constructing explicit solutions."

Specifically, we explore the trade-off between accuracy and computational simplicity and provide a new numerical method with a clear financial interpretation (and attractive algorithmic properties) for valuing real options with multiple state variables. We permit the decision maker to exercise an option at random moments in time that arrive at a fixed Poisson arrival rate $\lambda < \infty$. In between any two Poisson arrival times, the option cannot be exercised. Points in time when the option can be exercised are known as "optional stopping times" (see, e.g., Svenstrup (2005), Biagini and Björk (2007), Stroock and Varadhan ((2007), p. 24), and Cont and Fournié (2013)), and are, in our case, generated by a Poisson process. Hence, we name this approach the Poisson optional stopping times (POST) method. In financial terms, the imposed restriction can be viewed as a liquidity constraint, which prevents the decision maker from exercising the option at (what would otherwise be) the optimal time.

To our knowledge, the case $\lambda < \infty$ has been used only rarely; by Rogers and Zane (2002) to model liquidity constraints, by Dupuis and Wang (2002) and Lempa (2012) to model discretely monitored American call options, and by Lange and Teulings (2018) to compute the option value of vacant land when opportunities to build arrive at the Poisson rate λ. Our approach differs from that of Carr (1998), who randomizes the option's expiry date using a Poisson process, but does not restrict exercise times prior to maturity. In our setup, the classic real-options case is obtained in the limit $\lambda \to \infty$. Our generalization to $\lambda < \infty$ has not only theoretical but also practical appeal, since the problem corresponding to λ large but finite is more tractable both mathematically and numerically than the problem corresponding to $\lambda = \infty$.

Our method finds the solution as a monotone (nondecreasing) sequence of lower bounds; moreover, this property persists through the discretization. These desirable properties are related to those observed in the literature on the "penalty method," which has been used for some time in the financial engineering literature (e.g., Zvan, Forsyth, and Vetzal (1998), Forsyth and Vetzal (2002), d'Halluin, Forsyth, and Labahn (2004), Zhang, Yang, and Teo (2008), and Zhang, Wang, Yang, and Teo (2009)). It has apparently gone unnoticed, however, that a penalty term emerges naturally when the decision maker is permitted to stop only at a set of Poisson arrival times, as in this article.

To shed light on the aforementioned trade-off, we compare the proposed method against the frequently used quasi-analytic (QA) method for valuing real options as introduced by Adkins and Paxson (2011b) and used in various

subsequent articles of theirs (Adkins and Paxson (2011a), (2013a), (2013b), (2014), (2017a), (2017b), (2017c)) and others (Rohlfs and Madlener (2011), Heydari, Ovenden, and Siddiqui (2012), Dockendorf and Paxson (2013), Armada, Pereira, and Rodrigues (2013), Farzan, Mahani, Gharieh, and Jafari (2015), Stutzman, Weiland, Preckel, and Wetzstein (2017), and Støre, Fleten, Hagspiel, and Nunes (2018), etc.). The QA method is light on computation but, as we will see, may result in a large approximation error. As we show, the correct optimal policy and value function differ substantially from those derived in Adkins and Paxson (2011b).

The remainder of this article is organized as follows: Section II introduces stopping problems where the decision maker is permitted to stop only at a set of Poisson arrival times. Section III generalizes these ideas to compound options and, as a case study, presents the renewal option of Adkins and Paxson (2011b). Section IV contains a critical review of Adkins and Paxson's (2011b) QA method. Section V compares the solutions obtained by our method versus that obtained by the QA method, while Section VI discusses the implications of our findings.

II. Valuation Using Poisson Optional Stopping Times

We consider a project that produces a stochastic cash flow $f(X,Y)\,\mathrm{d}t$, which represents an ongoing profit stream that may be either positive or negative. We assume the cash flow depends on two stochastic state variables (X,Y) taking values in the first quadrant, denoted as $\mathbb{R}^2_{\geq 0}$, but our method allows straightforward generalizations to three or more state variables. The owner of the project may exercise the option to "stop" receiving the cash flow in exchange for receiving the stopping gain $g(X,Y)$, a one-off payment that represents the value of the project after some irreversible decision. This setup is the same as that in Lange et al. (2018), on whose framework we build. The question is how to find the optimal exercise policy and the corresponding (optimal) value of the project. Such "real" options have received much interest in the literature; classic references include Brennan and Schwartz (1985), McDonald and Siegel (1986), Majd and Pindyck (1987), Ekern (1988), Pindyck (1990), and Abel and Eberly (1996).

In our problem setup, state variables are considered to be stochastic in general, although deterministic processes may be included as well. For example, the deterministic variable "time" could be included as a "stochastic" process by setting its volatility equal to 0, and its drift equal to 1. The inclusion of time as a state variable allows us to deal with finite-horizon problems, in which case the functions f and g may be set to unattractive values beyond the option's expiration date. When implementing this in practice, it is convenient to truncate the state space at the option's maturity date and impose a "terminal" boundary condition.

Our approach differs from the literature in assuming that opportunities to exercise the option are generated by an exogenous Poisson process with arrival rate $\lambda < \infty$. The decision maker is forced to continue prior to the first arrival time and in between any two arrival times. Because we consider optional stopping times generated by a Poisson process, we call this the Poisson optional stopping times (POST) method. The (optimal) value of the project is denoted by $V_\lambda(X,Y)$; this includes the value of the option to stop. The dependence on the Poisson arrival rate λ

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

is indicated by the subscript. The optimal policy partitions the state space $\mathbb{R}^2_{\geq 0}$ into a region where the decision maker optimally continues, called the "continuation" region, and a region where she stops, if and when given the opportunity. Optimality demands that any point (X,Y) is part of the continuation region if and only if $V_\lambda(X,Y) > g(X,Y)$; otherwise, the decision maker would be better off stopping if possible. Similarly, any point (X,Y) is part of the stopping region if and only if $V_\lambda(X,Y) \leq g(X,Y)$, such that stopping is (weakly) preferable to continuing.

Invoking the standard assumption[1] of contingent-claims analysis, the following equation must be satisfied on the entire state space:

$$(1) \qquad rV_\lambda(X,Y) = \mathrm{L}V_\lambda(X,Y) + f(X,Y) + \lambda[g(X,Y) - V_\lambda(X,Y)]^+,$$

for all $(X,Y) \in \mathbb{R}^2_{\geq 0}$, where $r > 0$ is the risk-free rate, L is the infinitesimal generator of the process $\{(X_t, Y_t)\}$ under the risk-neutral measure, and $[\cdot]^+ = \max(\cdot, 0)$. Intuitively, equation (1) says that the (risk-neutral) return on the project, $rV_\lambda(X,Y)$, equals the expected change in $V_\lambda(X,Y)$ caused by the change in the state variables (X,Y), as measured by $\mathrm{L}V_\lambda(X,Y)$, plus the cash flow $f(X,Y)$ generated by the project, plus the additional gain $g(X,Y) - V_\lambda(X,Y)$ if and only if i) this quantity is positive, which occurs if and only if the process is in the stopping region, and ii) an opportunity to stop is generated, which happens at the Poisson arrival rate λ. More formally, equation (1) can be derived as in Appendix A using Bellman's dynamic programming approach, in which case L is understood as the infinitesimal generator under the physical measure, while r is the risk-adjusted discount rate (see, e.g., Dixit and Pindyck (1994)).

As an aside, we note that the term $\lambda[g(X,Y) - V_\lambda(X,Y)]^+$ in equation (1) can also be viewed as a "penalty term," because it penalizes the value function $V_\lambda(X,Y)$ for being exceeded by the stopping gain $g(X,Y)$. This penalty term is "active" only in the stopping region, where the decision maker wants to stop and the constraint of being permitted to stop only at Poisson arrival times is binding. In fact, such penalty terms have been used for some time in the financial engineering literature as an ad hoc method for ensuring numerical stability (e.g., Bensoussan and Lions (1982), Forsyth and Vetzal (2002), d'Halluin et al. (2004), Wang, Yang, and Teo (2006), Birge and Linetsky (2007), Glowinski (2008), and Zhang et al. (2008)). By pointing out that penalty methods solve "liquidity-constrained" stopping problems, we hope to enhance their appeal in the finance community.

For diffusion processes, the infinitesimal generator L is a partial differential operator, such that equation (1) is a partial differential equation (PDE). If $\{(X_t, Y_t)\}$ is a 2-dimensional geometric Brownian motion, as in our case study in Section III, then L is given by

$$(2) \qquad \mathrm{L} = \theta_X X \frac{\mathrm{d}}{\mathrm{d}X} + \theta_Y Y \frac{\mathrm{d}}{\mathrm{d}Y} + \frac{\sigma_X^2 X^2}{2} \frac{\mathrm{d}^2}{\mathrm{d}X^2} + \frac{\sigma_Y^2 Y^2}{2} \frac{\mathrm{d}^2}{\mathrm{d}Y^2} + \rho\sigma_X\sigma_Y XY \frac{\mathrm{d}^2}{\mathrm{d}X\mathrm{d}Y},$$

[1] That is, markets are complete, so that we can construct a portfolio which perfectly replicates the project to be valued.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

where θ_X and θ_Y are risk-neutral drifts, σ_X^2 and σ_Y^2 represent the variances of the innovations in each direction, and $|\rho| < 1$ captures the correlation between the innovations. Using the terminology of PDEs, generator (2) is classified as "elliptical" if $\sigma_X, \sigma_Y > 0$. While elliptic PDEs form a particularly well-studied class of problems, our theory development below applies equally to "parabolic" PDEs. For example, a parabolic generator L is obtained if X_t represents time, such that $X_t = t$, while Y_t represents a geometric Brownian motion with drift, in which case equation (2) is changed to

$$\text{L} = \frac{\mathrm{d}}{\mathrm{d}X} + \theta_Y Y \frac{\mathrm{d}}{\mathrm{d}Y} + \frac{\sigma_Y^2 Y^2}{2} \frac{\mathrm{d}^2}{\mathrm{d}Y^2}. \tag{3}$$

Here, the process $\{X_t\}$ drifts deterministically to the right at unit speed. Solving elliptic or parabolic PDEs with boundary conditions can be challenging, particularly when the location of the boundary is, a priori, unknown (see, e.g., Peskir and Shiryeav (2006)). Such "free-boundary" problems arise in option-valuation problems, because the boundary separating the continuation and stopping regions must be found as part of the solution.

Regardless of the type of PDE, equation (1) can be rewritten as

$$(r + \lambda - \text{L}) V_\lambda(X, Y) = f(X, Y) + \lambda \max\{g(X, Y), V_\lambda(X, Y)\}, \tag{4}$$

for all $(X, Y) \in \mathbb{R}^2_{\geq 0}$, which motivates the following POST algorithm for computing $V_\lambda(X, Y)$ as the limit of a sequence of functions $\{V_\lambda^{(j)}(X, Y)\}$ indexed by the superscript j and defined recursively:

$$\begin{aligned} (r + \lambda - \text{L}) V_\lambda^{(1)}(X, Y) &= f(X, Y) + \lambda g(X, Y), \\ (r + \lambda - \text{L}) V_\lambda^{(j)}(X, Y) &= f(X, Y) + \lambda \max\{g(X, Y), V_\lambda^{(j-1)}(X, Y)\}, \quad j \geq 2. \end{aligned} \tag{5}$$

Mathematically, equation (4) defines the function $V_\lambda(X, Y)$ as a "fixed point" that can be found by the POST algorithm (5). However, algorithm (5) can also be understood intuitively as follows. The (suboptimal) value $V_\lambda^{(1)}(X, Y)$ corresponds to a situation in which the decision maker is forced to stop at the first Poisson arrival time. The (suboptimal) value $V_\lambda^{(2)}(X, Y)$ corresponds to the situation in which she may (or may not) stop at the first Poisson arrival time and, if she has not stopped already, will be forced to stop at the second Poisson arrival time. The (suboptimal) value $V_\lambda^{(j)}(X, Y)$ takes into account that the decision maker may postpone the moment of exercise $j - 1$ times. Since each iteration of the POST algorithm (5) adds a single (Poisson) optional stopping time, the value function $V_\lambda^{(j)}(X, Y)$ is nondecreasing in the superscript j. If the value function $V_\lambda(X, Y)$ is bounded above, then the monotonicity of algorithm (5) also delivers convergence. This result is formalized in the following theorem:

Theorem 1 (Monotone Convergence of POST Algorithm). The sequence $\{V_\lambda^{(j)}(X, Y)\}$ defined by algorithm (5) converges in a monotone (nondecreasing) fashion to $V_\lambda(X, Y)$ as $j \to \infty$ if one technical condition, which ensures $V_\lambda(X, Y) < \infty$, is satisfied.

In fact, the monotonicity of algorithm (5) also gives the following result:

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

Theorem 2 (Geometric Convergence of POST Algorithm). Under one technical condition ensuring $V_\lambda(X,Y) < \infty$, the distance from the known quantity $V_\lambda^{(j)}(X,Y)$ to the unknown quantity $V_\lambda(X,Y)$ satisfies a computable bound as follows:

$$0 \le V_\lambda(X,Y) - V_\lambda^{(j)}(X,Y) \quad \le \quad \left(\frac{\lambda}{r+\lambda}\right)^{j-1} \frac{c(X,Y)}{r} < \infty, \tag{6}$$

where the constant $c(X,Y)$ depends only on the problem data and the location (X,Y).

Theorems 1 and 2 together demonstrate the reliability of the POST algorithm. The proof of Theorems 1 and 2 can be found in Appendix B. Theorem 2 implies that $V_\lambda^{(j)}(X,Y)$ converges to $V_\lambda(X,Y)$ at the geometric rate $\lambda/(r+\lambda)$, such that any desired precision level can be guaranteed by choosing some sufficiently large number of iterations j. The convergence of algorithm (5) may be accelerated if $f(X,Y)=0$, in which case $g(X,Y)$ can be replaced by $\max\{g(X,Y),0\}$, because the decision maker can ensure a minimum payoff equal to 0 by continuing indefinitely.

While the POST algorithm operates in function space, for practical purposes we must work in some finite-dimensional (vector) space. Hence we take a bounded subset of $\mathbb{R}^2_{\ge 0}$ as our computational domain, which we discretize using M grid points. This discretization turns the functions $f(X,Y)$, $g(X,Y)$ and $V_\lambda^{(j)}(X,Y)$ into $M \times 1$ vectors $\mathbf{f}$, $\mathbf{g}$ and $\mathbf{V}_\lambda^{(j)}$, respectively, while the generator L becomes the $M \times M$ matrix $\mathbf{L}$ (vectors and matrices are bold). The discretized version of POST then reads

$$\begin{aligned} [(r+\lambda)\mathbf{I} - \mathbf{L}]\mathbf{V}_\lambda^{(1)} &= \mathbf{f} + \lambda \mathbf{g}, \\ [(r+\lambda)\mathbf{I} - \mathbf{L}]\mathbf{V}_\lambda^{(j)} &= \mathbf{f} + \lambda \max\{\mathbf{g}, \mathbf{V}_\lambda^{(j-1)}\}, \quad j \ge 2, \end{aligned} \tag{7}$$

where $\mathbf{I}$ denotes an identity matrix of appropriate size, and the max operator is applied elementwise. This algorithm is, to the best of our knowledge, new in both the penalty literature and the optimal stopping literature. As demonstrated in Appendix C, the monotone (nondecreasing) and geometric convergence properties of algorithm (7) are guaranteed if two technical conditions regarding the matrix $\mathbf{L}$ are satisfied, irrespective of the resolution of the grid used for the discretization. This makes algorithm (7) essentially foolproof; we need only choose a grid with a sufficiently high resolution to ensure the discretized solution closely approximates the true solution. Further, each iteration of algorithm (7) can be performed efficiently using standard linear solvers if the matrix $(r+\lambda)\mathbf{I} - \mathbf{L}$ is sparse, which is typically the case for diffusion processes. The finite-difference stencil used to discretize the infinitesimal generator in equation (2) is given in Appendix D.

For jump-diffusion processes, the infinitesimal generator L is a partial *integro*-differential operator (due the integral nature of the jump process), such that equation (1) is a partial integro-differential equation (PIDE). In this case, the integral part of L may be discretized using a straightforward Riemann approximation. The resulting matrix $\mathbf{L}$ automatically satisfies the requirements that guarantee the convergence of algorithm (7). However, the matrix $\mathbf{L}$ will be dense rather

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

than sparse, making the inversion (or factorization) of the matrix $[(r+\lambda)\mathbf{I}-\mathbf{L}]$ in algorithm (7) more time-consuming. Part of this additional computational burden may be avoided by an "operator-splitting" method (see, e.g., Feng and Linetsky (2008)) that suggests decomposing $\mathbf{L}$ in a sparse and dense part as $\mathbf{L}=\mathbf{L}_{\text{sparse}}+\mathbf{L}_{\text{dense}}$, in which case algorithm (7) may be amended to read

$$[(r+\lambda)\mathbf{I}-\mathbf{L}_{\text{sparse}}]\mathbf{V}_\lambda^{(j)} = \mathbf{L}_{\text{dense}}\mathbf{V}_\lambda^{(j-1)}+\mathbf{f}+\lambda\max\{\mathbf{g},\mathbf{V}_\lambda^{(j-1)}\}, \quad j\geq 2. \tag{8}$$

In equation (8) we have placed the computable quantity $\mathbf{L}_{\text{dense}}\mathbf{V}_\lambda^{(j-1)}$ on the right-hand side to compensate for omitting the unknown quantity $-\mathbf{L}_{\text{dense}}\mathbf{V}_\lambda^{(j)}$ that would otherwise have appeared on the left-hand side. The operator-splitting method gives rise to what is known as an implicit–explicit (IMEX) scheme, because algorithm (8) treats the matrices $\mathbf{L}_{\text{sparse}}$ and $\mathbf{L}_{\text{dense}}$ implicitly and explicitly, respectively, for the purpose of computing $\mathbf{V}_\lambda^{(j)}$.

III. Case Study: The AP Renewal Option

In some applications, the decision maker is faced not with a single but rather multiple sequential stopping problems (e.g., Carmona and Dayanik (2008), Carmona and Touzi (2008)), known as compound options. Here we consider the "renewal" option introduced by Adkins and Paxson (AP) (2011b). AP assume the owner of the project receives the cash flow $f(X,Y)=Y-X$, where $\{Y_t\}$ and $\{X_t\}$ are (potentially correlated) geometric Brownian motions that represent revenues and costs generated by the project, respectively.[2] The infinitesimal generator corresponding to this process is given in our equation (2). Revenues are expected to fall ($\theta_Y<0$), while costs are expected to increase ($\theta_X>0$), such that the cash flow generated by the project, that is $Y-X$, declines on average. The owner of the project may pay the amount K to restart the bivariate stochastic process $\{(X_t,Y_t)\}$ at some favorable "initial" location (X_I,Y_I), where $Y_I>X_I$. Exercising this option is termed *renewal* of the project. The decision maker may renew the project as often as desired.

We differ from AP in assuming that opportunities to renew the project arrive at the Poisson arrival rate $\lambda<\infty$, where λ is large but finite. We denote the (optimal) value of the project by $V_\lambda(X,Y)$; this includes the value of the perpetual compound option to renew the project. In our setting, when the decision maker "stops", she actually renews the project and hence receives the project value at the renewal location, that is $V_\lambda(X_I,Y_I)$, minus the renewal cost K. Thus the problem of finding $V_\lambda(X,Y)$ amounts to solving a stopping problem with a stopping gain $g(X,Y)=V_\lambda(X_I,Y_I)-K$, which is *itself* unknown. It follows that the value of the project satisfies PDE (1) with $g(X,Y)=V_\lambda(X_I,Y_I)-K$, that is

$$rV_\lambda(X,Y) = \mathrm{L}V_\lambda(X,Y)+f(X,Y)+\lambda[V_\lambda(X_I,Y_I)-K-V_\lambda(X,Y)]^+. \tag{9}$$

To solve this PDE, we set up the following iterative scheme by analogy with algorithm (7):

$$[(r+\lambda)\mathbf{I}-\mathbf{L}]\mathbf{V}_\lambda^{(j)} = \mathbf{f}+\lambda\max\{[V_\lambda^{(j-1)}(X_I,Y_I)-K]\mathbf{1},\mathbf{V}_\lambda^{(j-1)}\}, \quad j\geq 2, \tag{10}$$

[2] AP denote costs and revenues by C_t and P_t, respectively, but for consistency with Section II we use X_t and Y_t.

where $\mathbf{1}$ is an $M \times 1$ vector containing ones, the max operator is applied elementwise, and $V_\lambda^{(j-1)}(X_I, Y_I)$ denotes the (scalar) element in the vector $\mathbf{V}_\lambda^{(j-1)}$ that represents the value at the renewal location (X_I, Y_I). Each iteration of algorithm (10) updates both the (discretized) value function $\mathbf{V}_\lambda^{(j)}$ as well as the (discretized) stopping gain $[V_\lambda^{(j)}(X_I, Y_I) - K]\mathbf{1}$. If the initialization $\mathbf{V}_\lambda^{(1)}$ can be chosen such that algorithm (10) is nondecreasing for the first iteration, then it is nondecreasing for all subsequent iterations.

In some cases a solution may be desirable for a situation in which the decision maker can renew the project at most $N < \infty$ times. In this case we use $V_{\lambda,n}(X, Y)$ for $n = 0, 1, 2, \ldots, N$ to denote the value of the project when a maximum of n renewals remain. The project value with a maximum of n renewals remaining, that is $V_{\lambda,n}(X, Y)$, satisfies PDE (1) with stopping gain $g(X, Y) = V_{\lambda,n-1}(X_I, Y_I) - K$, that is

$$
\begin{aligned}
(11) \qquad r V_{\lambda,n}(X,Y) &= \mathrm{L} V_{\lambda,n}(X,Y) + f(X,Y) \\
&\quad + \lambda [V_{\lambda,n-1}(X_I, Y_I) - K - V_{\lambda,n}(X,Y)]^+.
\end{aligned}
$$

In this case, algorithm (7) for fixed n may be amended to read

$$
\begin{aligned}
(12) \qquad [(r+\lambda)\mathbf{I} - \mathbf{L}]\mathbf{V}_{\lambda,n}^{(1)} &= \mathbf{f} + \mathbf{1}\lambda [V_{\lambda,n-1}(X_I, Y_I) - K] \\
[(r+\lambda)\mathbf{I} - \mathbf{L}]\mathbf{V}_{\lambda,n}^{(j)} &= \mathbf{f} + \lambda \max\{\mathbf{1}[V_{\lambda,n-1}(X_I, Y_I) - K], \mathbf{V}_{\lambda,n}^{(j-1)}\}, \quad j \geq 2,
\end{aligned}
$$

where $V_{\lambda,n-1}(X_I, Y_I) - K$ is assumed known for $n = 1$, and, for $n \geq 1$, can be based on the output of the algorithm corresponding to $n - 1$. In using algorithm (12), we iterate over the number of possible renewals $n = 1, 2, \ldots, N$ and, for each $n \leq N$, over fixed-point iterations j. We approximate the case $N = \infty$ by taking N to be large but finite.

For its initialization, algorithm (12) relies on the project with 0 renewals remaining, that is $V_{\lambda,0}(X, Y)$, being defined. For the case considered by AP, the net present value (NPV) of continuing the project indefinitely (when no renewal options remain) equals

$$
(13) \qquad V_{\lambda,0}(X,Y) = \mathrm{E}_{(X,Y)} \int_0^\infty (Y_t - X_t) \exp(-rt)\, \mathrm{d}t = \frac{Y}{r - \theta_Y} - \frac{X}{r - \theta_X},
$$

where $\mathrm{E}_{(X,Y)}$ denotes an expectation conditional on the current state (X, Y). Following AP, here we assume that $\theta_X, \theta_Y < r$; otherwise, the integral in equation (13) would be unbounded. Of course, the stopping problem with $\theta_Y \geq r > \theta_X$ would be trivial, since the decision maker would continue forever and reap infinite benefits (revenues increase faster than costs). However, the problem with $\theta_X \geq r > \theta_Y$ is nontrivial as costs increase faster than revenues. In this case, the decision maker will renew the project frequently to avoid "infinite" costs. Although the NPV of continuing indefinitely, that is $V_{\lambda,0}(X, Y)$, no longer exists (it is negative infinity), the value of the project may still be finite, even positive. In this case, algorithm (10) must be used instead of algorithm (12); that is, the value of the project with an infinite possible number of renewals must be found directly (i.e., without looping over the number of renewals n).

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

IV. Critical Review of the QA Method

The quasi-analytic (QA) method for valuing real options was introduced by Adkins and Paxson (2011b) and used by numerous subsequent articles. This section contains a critical review of the QA method for treating the renewal option discussed in Section III. AP denote the (optimal) value of the project by $V(X,Y)$. It is assumed the decision maker can continuously exercise the option to renew; hence, there is no subscript λ. No-arbitrage arguments require $V(X,Y)$ to satisfy the following PDE:

$$rV(X,Y) = \mathrm{L}V(X,Y) + Y - X, \tag{14}$$

which holds for any location (X,Y) in the continuation region, and where the infinitesimal generator L is given in our equation (2). We refer to this equation as the "no-arbitrage PDE." In the stopping region, where renewal is imminent, the classic value-matching condition requires $V(X,Y) = V(X_I,Y_I) - K = \text{constant}$. Importantly for our argument below, this constant depends only on the problem setup; that is, it does not vary with the coordinates (X,Y). To solve the no-arbitrage PDE, AP assume $\theta_X, \theta_Y < r$ and propose the following candidate solution:

$$V(X,Y) = \frac{Y}{r-\theta_Y} - \frac{X}{r-\theta_X} + AX^{\eta}Y^{\beta}, \tag{15}$$

which is assumed to hold for any (X,Y) in the continuation region. At this point it is assumed that the parameters A, β and η are constant, although AP later violate this assumption. Candidate solution (15) splits the project value $V(X,Y)$ into an intrinsic part, $Y/(r-\theta_Y) - X/(r-\theta_X)$, which represents the NPV of continuing the project indefinitely, and an option value, $AX^{\eta}Y^{\beta}$, that captures the value of being able to renew the project. If the parameters A, β, and η are assumed to be constant, substituting candidate solution (15) into the no-arbitrage PDE (14) gives rise to the following characteristic root equation for β and η:

$$\tfrac{1}{2}\sigma_X^2\beta(\beta-1) + \tfrac{1}{2}\sigma_X^2\eta(\eta-1) + \rho\sigma_X\sigma_Y\beta\eta + \theta_Y\beta + \theta_X\eta - r = 0. \tag{16}$$

We follow AP in using hatted coordinates $(\widehat{X},\widehat{Y})$ to denote points on the stopping boundary, which separates the stopping region from the continuation region. To determine the unknown parameters A, β, and η, AP propose using the standard value-matching and smooth-pasting conditions for all coordinates $(\widehat{X},\widehat{Y})$ on the stopping boundary as follows:

$$\frac{\widehat{Y}}{r-\theta_Y} - \frac{\widehat{X}}{r-\theta_X} + A\widehat{X}^{\eta}\widehat{Y}^{\beta} = \frac{Y_I}{r-\theta_Y} - \frac{X_I}{r-\theta_X} + AX_I^{\eta}Y_I^{\beta} - K = \text{constant}, \tag{17}$$

$$\frac{\mathrm{d}V(\widehat{X},\widehat{Y})}{\mathrm{d}\widehat{X}} = \eta A\widehat{X}^{\eta-1}\widehat{Y}^{\beta} - \frac{1}{r-\theta_X} = 0, \tag{18}$$

$$\frac{\mathrm{d}V(\widehat{X},\widehat{Y})}{\mathrm{d}\widehat{Y}} = \beta A\widehat{X}^{\eta}\widehat{Y}^{\beta-1} + \frac{1}{r-\theta_Y} = 0. \tag{19}$$

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

AP suppose that exactly one critical value of $\widehat{Y}$ corresponds to each value of X, such that $\widehat{Y} = \widehat{Y}(X)$, and proceed to numerically solve the four equations (16)–(19) for each possible value of X, to obtain four quantities $\widehat{Y}(X)$, A, β, and η. Unfortunately, it is impossible to solve this equation set for different values of X if constant values of A, β, and η are imposed. Hence, candidate solution (15) cannot technically be correct. AP resort to solving equations (16)–(19) by means of approximation, allowing the purported constants A, β, and η to vary with X. This implies the constants in the QA method are only quasi constant (i.e., it would be more precise to write $A(X)$, $\beta(X)$, and $\eta(X)$).

While AP acknowledge that their constants vary with X (indeed, their Figure 3 shows considerable variation), they stop short of acknowledging four important drawbacks. First, the value $V(X,Y)$ in (15) remains undefined for any (X,Y) inside the continuation region, because the constants A, β, and η are determined only along the stopping boundary. Hence, the QA method does not provide a value in the entire state space (e.g., $V(X_1,Y_1)$ is left undefined). Second, the presence of the varying constants means that the candidate solution no longer satisfies the no-arbitrage PDE, because the derivation of equation (16) is invalidated. This concern was also raised in a recent working paper by Compernolle, Huisman, Kort, Lavrutich, Nunes, and Thijssen (2018). As a general rule, substantial variation in the purported constants along the stopping boundary ought to set off alarm bells about the potential inaccuracy of the QA method. Third, the value along the stopping boundary, $V(\widehat{X},\widehat{Y}) = V(X_1,Y_1) - K$ (which should be a single number), is no longer uniquely identified, as the right-hand side of equation (17) now varies with X (because A, β, and η vary with X). Indeed, for the base-case parameters used by AP, rather than being constant, the right-hand side of equation (17) varies in the range from 260 to 1,032 as X varies along the positive real line.[3] Large as this range is, it does not actually contain the true value of the project at the stopping boundary, which we find to be in excess of 1,200 (see Section V). Fourth, the impact of variable constants on the derived policy in the AP exposition is unclear, but turns out to be considerable (again, see Section V).

V. Comparison between QA and Proposed Solution Method

This section compares the stopping boundary and value function obtained using the QA method against those obtained by our proposed solution method. We consider the AP renewal option with their base-case parameters as in our footnote 3. We compute the project value $V_{\lambda,n}(X,Y)$ using algorithm (12) based on a Poisson intensity $\lambda = 512$ (implying that there are, on average, 512 renewal opportunities per year) for $n = 0,1,\ldots,N$ with $N = 30$ (implying a maximum of 30 renewals). Both numbers proved to be sufficiently large for the project value to be stable as a function of λ and n.[4]

[3] In their base case, AP take $(X_1,Y_1) = (80,20)$, $\theta_Y = -0.01$, $\theta_X = 0.04$, $\sigma_Y = \sigma_X = 0.30$, $\rho = 0$, $K = 100$, and $r = 0.07$ on an annual time scale (see their Table 1, p. 795).

[4] Appendix E contains further details on the grid and boundary conditions. Further, while our solution $V_{\lambda,n}(X,Y)$ depends on the chosen values for λ and n, Appendix F demonstrates that the stopping boundary and project value $V_{\lambda,n}(X,Y)$ are robust as a function of λ and n when both are sufficiently large.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

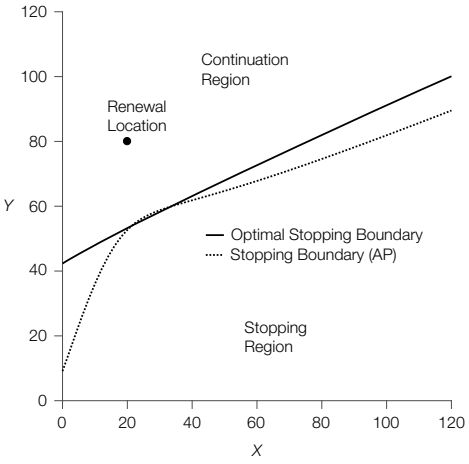
Graph A of Figure 1 shows the optimal stopping boundary (solid line), which turns out to resemble a straight line. In contrast, the boundary produced by the QA method (dotted line) is curved and mostly lies below the optimal boundary, except in a small section for X between 21 and 34, where it marginally lies above the optimal boundary. This implies the decision maker should generally renew the project earlier than suggested by the QA method. The dotted line decreases sharply for low costs X; this is not visible in Figure 3 of AP, because their horizontal axis starts at $X = 20$. AP remark that "the boundary could be represented

FIGURE 1

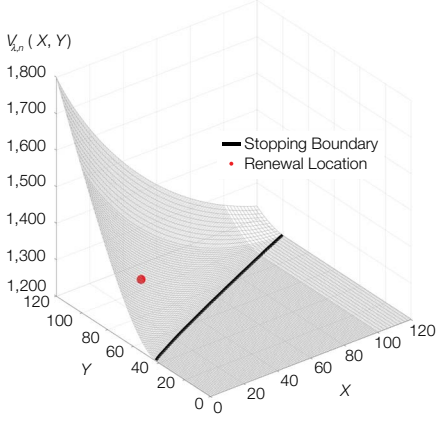
Optimal Policy and Value Function with Base-Case Parameters

Graph A of Figure 1 displays the stopping boundary (solid line) using the proposed solution method with $\lambda = 512$ and $n = 30$ versus that produced by the QA method (dotted line). As in Section III, the variables X and Y denote the cost and revenue generated by the project, respectively. The stopping boundary (solid curve) is shown in both Graphs A and B. Graph B shows the project value $V_{\lambda,n}(X, Y)$ for $\lambda = 512$ and $n = 30$. The value at the renewal location is indicated by the dot at $(X_l; Y_l; V_\lambda(X_l, Y_l)) = (20; 80; 1{,}302)$. A comparison with AP is impossible because the QA method does not produce a value function.

Graph A. Stopping Boundaries



Graph B. Value Function



https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

by its least squares line, resulting in a simple, approximate renewal rule" (p. 796). As the optimal stopping boundary is indeed nearly linear, this is not incorrect as such. However, as Graph A of Figure 1 shows, a straight line based on the QA method would probably be suboptimal; not only for small but also for large costs, where the slope appears to be too small.

Graph B of Figure 1 contains a surface plot of the project value $V_{\lambda,n}(X,Y)$ as computed on our grid (only part of the grid is shown). The dot located at $(X_1; Y_1; V_{\lambda,n}(X_1,Y_1)) = (20; 80; 1{,}302)$ shows that the value of the project at the renewal location is just over 1,300. Since the QA method cannot be used to compute project values in the continuation region, a direct comparison of this number with AP is impossible. The QA method does produce project values on the stopping boundary via the value matching condition (17), which gives rise to values within the range from 260 to 1,032 for different points on the stopping boundary. As can be seen in Graph B, however, we find the value on the stopping boundary to be 1,202 (exactly $K = 100$ lower than the value at the renewal location), which falls outside of the AP range.

We deviate from AP in considering, in addition to the base-case parameters as in our footnote 3, the case where the risk-free rate is lowered from $r = 0.07$ to $r = 0.04$ at an annual time scale. This has the consequence that the intrinsic value of the project as in equation (13) (the NPV of continuing forever without renewing the project) no longer exists, since $\theta_X = 0.04$ (i.e., the costs grow at a rate that equals the discount rate). The QA method's candidate solution (15), which includes the term $-X/(r-\theta_X)$, no longer applies; indeed, the intrinsic value of continuing indefinitely without renewing is negative infinity. However, the option value of being able to renew the project, which allows the decision maker to avoid this undesirable situation, is positive infinity. The project value, which is the sum of the intrinsic value and the option value, is bounded, and even positive, as we will see. Moreover, algorithm (10) still applies. The resulting optimal stopping boundary and value function are shown in Graphs A and B of Figure 2, respectively. We find that the project value (with $r = 0.04$) is greater than that in the base case (where $r = 0.07$). As the dot in Graph B of Figure 2 shows, the value at the renewal location is 2,888, which is more than double the corresponding value in the base case, due to the lower discounting of future cash flows. While the intrinsic value of the project decreases (to negative infinity), the decision maker can achieve a positive project value by renewing more frequently than in the base case. This can be seen from the fact that the stopping boundary in Graph A of Figure 2 lies higher than that in Graph A of Figure 1. A comparison with the QA method is not possible, as candidate solution (15) is undefined.

VI. Conclusion

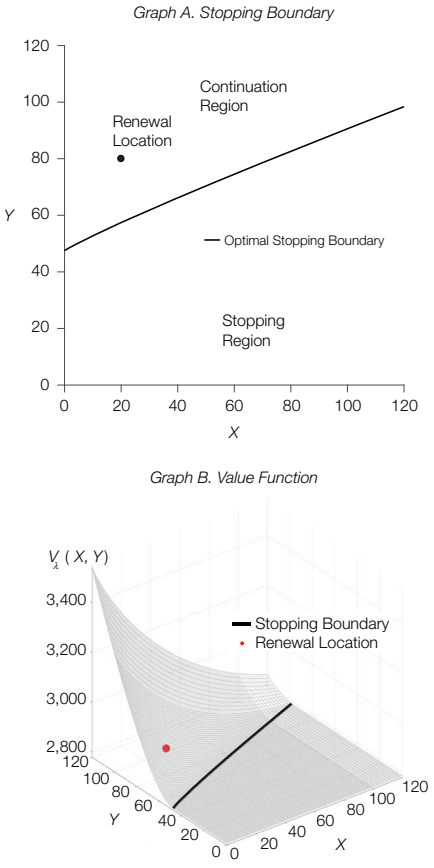
We have presented a new numerical method, inspired by Lange et al. (2018), which uses Poisson optional stopping times (POST) to value liquidity-constrained real options with multiple state variables. In our problem setup, the decision maker is permitted to exercise the option only at a set of Poisson arrival times that arrive at rate $\lambda < \infty$. For some real options, such a liquidity constraint may be realistic (e.g., off-shore operations that require particular weather conditions).

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

FIGURE 2

Optimal Policy and Value Function with $r = 0.04$

Graph A of Figure 2 shows the stopping boundary (solid curve) obtained by algorithm (10) using $\lambda = 512$, $n = 30$ and the base-case parameters (see our footnote 3), except $r = 0.04$. As in Section III, the variables X and Y denote the cost and revenue generated by the project, respectively. A comparison with the QA method is not possible as candidate solution (15) is not defined. The stopping boundary (solid curve) is shown in both Graphs A and B. Graph B shows the project value $V_\lambda(X, Y)$. The value at the renewal location is indicated by the dot at $(X_I; Y_I; V_\lambda(X_I, Y_I)) = (20; 80; 2{,}888)$.



Alternatively, the case $\lambda < \infty$ may be viewed as a mathematical tool that makes the problem corresponding to $\lambda = \infty$ more amenable.

The POST algorithm finds the project value as the limit of a monotone (nondecreasing) sequence of lower bounds. Moreover, this property persists through the discretization. In essence, each iteration of POST adds a single (Poisson) optional stopping time, at which point the decision maker may (but need not) exercise the option, thereby yielding the monotonicity of the algorithm. While related to "penalty methods" in the financial engineering literature, the proposed algorithm is, to the best of our knowledge, new.

For diffusion processes, each POST iteration can be performed efficiently using a sparse linear solver and convergence is guaranteed (at a geometric rate) irrespective of the resolution of the grid used for the discretization. The algorithm

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

remains valid for jump-diffusion processes, although the sparsity of the linear system is lost, such that solving the problem becomes more computationally intensive (although part of this additional computational burden may be alleviated using operator-splitting methods).

As a case study, we considered the renewal option of Adkins and Paxson (2011b), which they solve using a quasi-analytic (QA) method. We showed that the QA method is internally inconsistent and therefore only an approximate method for determining optimal stopping boundaries. This fact has been explicitly recognized in two recent articles interrogating the accuracy of the QA method. Dockendorf and Paxson (2013) compare it with an alternative method and find the QA approximation to be acceptable, for one specific case, although they do not consider its accuracy more generally. Compernolle et al. (2018), moreover, demonstrate that the QA method produces an incorrect exercise boundary for a stopping problem with a cash flow $f(X,Y)=0$ and stopping gain $g(X,Y)=X+Y-K$ (using our notation from Section II). As we have shown, the QA method computes a policy, albeit inaccurately, but no value function. For these reasons we would urge caution in using this method. Our proposed solution method enabled us to derive a more accurate stopping boundary and compute project values at points in the state space away from the stopping boundary.

In short, the combination of interpretability, flexibility, and reliability make the POST algorithm particularly attractive from a practical viewpoint. We expect this method to have broad applicability in the analysis of real options with multiple state variables.

Appendix A. Derivation of Equation (1)

For any (X,Y) in the continuation region, the decision maker receives the infinitesimal cash flow $f(X,Y)\mathrm{d}t$ as well as the expected value of the value of the project at time $t+\mathrm{d}t$, the present value of which equals $(1-r\,\mathrm{d}t)\mathrm{E}V_\lambda(X+\mathrm{d}X,Y+\mathrm{d}Y)$, where r is the risk-adjusted discount rate. By Bellman's optimality principle, the sum of both terms must be equal to the value of the project itself, and thus we require

$$\begin{aligned} V_\lambda(X,Y) &= f(X,Y)\mathrm{d}t+(1-r\,\mathrm{d}t)\mathrm{E}V_\lambda(X+\mathrm{d}X,Y+\mathrm{d}Y), \\ &= f(X,Y)\mathrm{d}t+(1-r\,\mathrm{d}t)[V_\lambda(X,Y)+\mathrm{L}V_\lambda(X,Y)\mathrm{d}t]+O(\mathrm{d}t^2), \\ &= V_\lambda(X,Y)+[(\mathrm{L}-r)V_\lambda(X,Y)+f(X,Y)]\mathrm{d}t+O(\mathrm{d}t^2), \end{aligned}$$

where, in line two, we have used

$$\mathrm{E}V_\lambda(X+\mathrm{d}X,Y+\mathrm{d}Y) = V_\lambda(X,Y)+\mathrm{L}V_\lambda(X,Y)\mathrm{d}t+O(\mathrm{d}t^2),$$

which holds by Itô's lemma, and we ignore terms of order $O(\mathrm{d}t^2)$. It follows that

$$(\mathrm{L}-r)V_\lambda(X,Y)+f(X,Y) = 0$$

for all (X,Y) in the continuation region.

Similarly, for any (X,Y) in the stopping region, the decision maker receives $g(X,Y)$ if an option to stop presents itself, which happens with probability $\lambda\,\mathrm{d}t$, and the expected present value of continuing otherwise. By Bellman's optimality principle, the weighted

average of both possibilities must be equal to the value of the project, and thus

$$\begin{aligned} V_\lambda(X,Y) & \\ &= g(X,Y)\lambda\,\mathrm{d}t + (1-\lambda\,\mathrm{d}t)[f(X,Y)\,\mathrm{d}t + (1-r\,\mathrm{d}t)\mathrm{E}V_\lambda(X+\mathrm{d}X, Y+\mathrm{d}Y)], \\ &= g(X,Y)\lambda\,\mathrm{d}t + (1-\lambda\,\mathrm{d}t)[f(X,Y)\,\mathrm{d}t + (1-r\,\mathrm{d}t)(V_\lambda(X,Y) + \mathrm{L}V_\lambda(X,Y)\,\mathrm{d}t)] + O(\mathrm{d}t^2), \\ &= V_\lambda(X,Y) + [(\mathrm{L}-r)V_\lambda(X,Y) + f(X,Y) + \lambda(g(X,Y) - V_\lambda(X,Y))]\,\mathrm{d}t + O(\mathrm{d}t^2). \end{aligned}$$

It follows that

$$(\mathrm{L}-r)V_\lambda(X,Y) + f(X,Y) + \lambda(g(X,Y) - V_\lambda) \quad = \quad 0$$

for all (X,Y) in the stopping region.

By optimality, (X,Y) is part of the stopping region if and only if $g(X,Y) - V_\lambda(X,Y) \geq 0$. Combining both arguments above, it follows that on the entire state space $\mathbb{R}^2_{\geq 0}$ it must hold that

$$(\mathrm{L}-r)V_\lambda(X,Y) + f(X,Y) + \lambda[g(X,Y) - V_\lambda(X,Y)]^+ \quad = \quad 0, \quad \forall (X,Y) \in \mathbb{R}^2_{\geq 0},$$

which is equivalent to equation (1) in the main text. It can be shown that this equation is not only necessary but, in fact, also sufficient (see Lange et al. (2018)). Any algorithm that solves equation (1) thus solves the original problem.

Appendix B. Proof of Theorems 1 and 2

This section shows that algorithm (5) converges both monotonically and geometrically as long as the constant $c(X,Y)$, which appears in the computable bound (6), is itself bounded, where

$$\text{(B-1)} \qquad c(X_0,Y_0) \quad := \quad \sup_{t\geq 0} \mathrm{E}_0 \sup_{s\geq 0} \mathrm{e}^{-rs}\{\mathrm{E}_t[(\mathrm{L}-r)g(X_{t+s},Y_{t+s}) + f(X_{t+s},Y_{t+s})]\}^+.$$

Here $\mathrm{E}_t[\cdot] := \mathrm{E}[\cdot|\mathcal{F}_t]$, such that the operator E_0 takes an expectation conditional on the current state (X_0,Y_0). Condition (B-1) is hardly restrictive,[5] because the decision maker would continue forever without stopping and reap infinite benefits if it were not satisfied. To explain intuitively why the constant $c(X,Y)$ should be bounded, we consider the infinitesimal gain of a small delay in stopping. By postponing the decision to stop for a short time $\mathrm{d}t$, the decision maker gains $[f(X,Y) + \mathrm{L}g(X,Y)]\,\mathrm{d}t$ relative to stopping right away; that is, the received cash flow and the expected change in the stopping gain. If, instead, the decision maker had stopped right away, then the additional gain after time $\mathrm{d}t$ would have been $r\,g(X,Y)\mathrm{d}t$; that is, the risk-free rate r times the stopping gain. Hence the infinitesimal gain of postponing for a short time $\mathrm{d}t$ relative to stopping right away equals $f(X,Y) + (\mathrm{L}-r)g(X,Y)$. For the project value at location (X_t,Y_t) to be bounded, the positive part of $\mathrm{E}_t[f(X_{t+s},Y_{t+s}) + (\mathrm{L}-r)g(X_{t+s},Y_{t+s})]$ cannot grow faster than e^{rs}; otherwise, the decision maker would continue forever and reap infinite benefits. However, even if the project value at location (X_t,Y_t) is unbounded, $c(X_0,Y_0)$ may remain finite if it is improbable that location (X_t,Y_t) is reached when starting from location (X_0,Y_0). That is, there may exist locations (X_t,Y_t) for which $V_\lambda(X_t,Y_t) = \infty$, as long as the set containing such locations has measure 0 under the expectation operator E_0. This explains the appearance of two expectation operators in (B-1).

[5] Note that $\mathrm{L}g(X,Y)$ itself need not be bounded above, which is relevant when L contains second derivatives and $g(X,Y)$ contains kinks, in which case $\mathrm{L}g(X_t,Y_t)$ can be interpreted in a (Schwarz) distributional sense under the expectation (i.e., integral) operator E_t.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

For the renewal option considered in Section III, where $f(X,Y)=Y-X$, for condition (B-1) to hold it is sufficient that $\theta_X > \theta_Y$, such that costs increase faster than revenues.

Having explained the intuition behind condition (B-1), we proceed to the formal proofs of Theorems 1 and 2.

Proof of Theorem 1. To prove the convergence of the algorithm (5), we first note that for any $j \geq 2$, it can be written as follows:

$$\begin{aligned} \text{(B-2)} \quad V_\lambda^{(j)}(X_0,Y_0) &= (r+\lambda-\mathrm{L})^{-1}[f(X_0,Y_0)+\lambda\max\{g(X_0,Y_0),V_\lambda^{(j-1)}(X_0,Y_0)\}], \\ &= \int_0^\infty \mathrm{e}^{-(r+\lambda)t}\mathrm{E}_0[f(X_t,Y_t)+\lambda\max\{g(X_t,Y_t),V_\lambda^{(j-1)}(X_t,Y_t)\}]\,\mathrm{d}t, \end{aligned}$$

where the first line is purely formal, while the second holds by the resolvent formalism (see, e.g., Revuz and Yor ((1999), pp. 89–90, 290–291) or Rogers and Williams ((2000), pp. 234–238)). The probability density $\lambda\exp(-\lambda t)\,\mathrm{d}t$ is recognized as the (marginal) probability density of the first Poisson arrival time. From equation (B-2), it is clear that $V_\lambda^{(j)}(X,Y)$ is equal to the value of receiving the cash flow $f(X_t,Y_t)$ until the first Poisson arrival time, while the decision maker chooses the maximum of $g(X_t,Y_t)$ and $V_\lambda^{(j-1)}(X_t,Y_t)$ if and when the first Poisson arrival time occurs. By the (time-homogeneous) Markov property, equation (B-2) implies

$$\begin{aligned} \text{(B-3)} \quad & V_\lambda^{(j)}(X_t,Y_t) \\ &= \int_0^\infty \mathrm{e}^{-(r+\lambda)s}\mathrm{E}_t[f(X_{t+s},Y_{t+s})+\lambda\max\{g(X_{t+s},Y_{t+s}),V_\lambda^{(j-1)}(X_{t+s},Y_{t+s})\}]\,\mathrm{d}s, \end{aligned}$$

for all $j \geq 2$ and all (X_t,Y_t), as we will use below.

To prove that the iterations in algorithm (B-2) produce a nondecreasing sequence, we use the fact that $\max\{g(X_t,Y_t),V_\lambda^{(1)}(X_t,Y_t)\}$ cannot be exceeded by $g(X_t,Y_t)$, which implies

$$\begin{aligned} V_\lambda^{(2)}(X_0,Y_0) &:= \int_0^\infty \mathrm{e}^{-(r+\lambda)t}\mathrm{E}_0[f(X_t,Y_t)+\lambda\max\{g(X_t,Y_t),V_\lambda(X_t,Y_t)\}]\,\mathrm{d}t, \\ &\geq \int_0^\infty \mathrm{e}^{-(r+\lambda)t}\mathrm{E}_0[f(X_t,Y_t)+\lambda g(X_t,Y_t)]\,\mathrm{d}t =: V_\lambda^{(1)}(X_0,Y_0), \end{aligned}$$

hence $V_\lambda^{(2)}(X_0,Y_0) \geq V_\lambda^{(1)}(X_0,Y_0)$ and thus the algorithm is nondecreasing at the first step. By induction it follows that the algorithm is nondecreasing at all subsequent steps.

To prove that the algorithm converges at a geometric rate $\lambda/(r+\lambda)$, we define the difference between two iterations of (B-2) as follows:

$$\begin{aligned} \text{(B-4)} \quad \Delta_\lambda^{(j)}(X_0,Y_0) &:= V_\lambda^{(j+1)}(X_0,Y_0)-V_\lambda^{(j)}(X_0,Y_0), \\ &= \int_0^\infty \lambda\,\mathrm{e}^{-(r+\lambda)t}\mathrm{E}_0\left[\max\left\{g(X_t,Y_t),V_\lambda^{(j)}(X_t,Y_t)\right\}\right. \\ &\qquad \left.-\max\left\{g(X_t,Y_t),V_\lambda^{(j-1)}(X_t,Y_t)\right\}\right]\mathrm{d}t, \\ &\qquad \text{for } j\geq 2, \text{ which follows from equation (B-2)} \\ &\leq \int_0^\infty \lambda\,\mathrm{e}^{-(r+\lambda)t}\mathrm{E}_0\left[V_\lambda^{(j)}(X_t,Y_t)-V_\lambda^{(j-1)}(X_t,Y_t)\right]\mathrm{d}t, \\ &= \int_0^\infty \lambda\,\mathrm{e}^{-(r+\lambda)t}\mathrm{E}_0\Delta_\lambda^{(j-1)}(X_t,Y_t)\mathrm{d}t. \end{aligned}$$

The inequality follows from $\max\{a,c\}-\max\{b,c\} \leq a-b$ for $a \geq b$ and $a,b,c \in \mathbb{R}$. The last line holds by the definition of $\Delta_\lambda^{(j-1)}(X_t,Y_t)$. Using the (time-homogeneous) Markov

property, the last line of equation (B-4) holds not only at location (X_0, Y_0), but rather for all locations (X_t, Y_t) as follows:

$$\text{(B-5)} \qquad \Delta_\lambda^{(j)}(X_t, Y_t) \quad \leq \quad \int_0^\infty \lambda\, \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \Delta_\lambda^{(j-1)}(X_{t+u}, Y_{t+u})\, \mathrm{d}u, \quad j \geq 2, t \geq 0,$$

where the expectation on the right-hand side is conditional on the information set at time t, and the integral is taken over all times $u \geq 0$. In turn, as we will show below, inequality (B-5) implies

$$\text{(B-6)} \qquad \Delta_\lambda^{(j)}(X_0, Y_0) \quad \leq \quad \left(\frac{\lambda}{r+\lambda}\right)^{j-1} \sup_{t \geq 0} \mathrm{E}_0 \Delta_\lambda^{(1)}(X_t, Y_t), \quad j \geq 2.$$

Equation (B-6) says that the maximum possible increase in the approximate solution shrinks at the geometric rate $\lambda/(r+\lambda)$ as the number of iterations of the algorithm increases. Of course, this holds if and only if $\sup_{t \geq 0} \mathrm{E}_0 \Delta_\lambda^{(1)}(X_t, Y_t)$ is bounded, as will be shown below. To prove that inequality (B-5) implies inequality (B-6), we take $j = 3$ in equation (B-5) and write

$$\begin{aligned}
\text{(B-7)} \quad \Delta_\lambda^{(3)}(X_0, Y_0) \quad &\leq \quad \int_0^\infty \lambda\, \mathrm{e}^{-(r+\lambda)t} \mathrm{E}_0 \Delta_\lambda^{(2)}(X_t, Y_t)\, \mathrm{d}t, \\
&\leq \quad \int_0^\infty \lambda\, \mathrm{e}^{-(r+\lambda)t} \mathrm{E}_0 \left[\int_0^\infty \lambda\, \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \Delta_\lambda^{(1)}(X_{t+u}, Y_{t+u})\, \mathrm{d}u \right] \mathrm{d}t, \\
&\qquad \text{which follows by applying inequality (B-5) for } j = 2 \\
&\leq \quad \int_0^\infty \int_0^\infty \lambda^2\, \mathrm{e}^{-(r+\lambda)(t+u)} \mathrm{E}_0 \Delta_\lambda^{(1)}(X_{t+u}, Y_{t+u})\, \mathrm{d}u\, \mathrm{d}t \\
&\qquad \text{by the tower property and a rearrangement} \\
&\leq \quad \int_0^\infty \int_0^\infty \lambda^2\, \mathrm{e}^{-(r+\lambda)(t+u)} \sup_{t,u \geq 0} \mathrm{E}_0 \Delta_\lambda^{(1)}(X_{t+u}, Y_{t+u})\, \mathrm{d}u\, \mathrm{d}t, \\
&\qquad \text{by inserting two suprema} \\
&\leq \quad \left(\frac{\lambda}{r+\lambda}\right)^2 \sup_{t \geq 0} \mathrm{E}_0 \Delta_\lambda^{(1)}(X_t, Y_t), \\
&\qquad \text{performing the integrals and merging both suprema into one.}
\end{aligned}$$

By repeatedly applying this argument we obtain inequality (B-6).

To establish that algorithm (5) converges, what remains to be done is to confirm that $\sup_{t \geq 0} \mathrm{E}_0 \Delta_\lambda^{(1)}(X_t, Y_t)$ is bounded. To this end, we write the stopping gain $g(X_t, Y_t)$ as

$$\text{(B-8)} \qquad g(X_t, Y_t) \quad = \quad \int_0^\infty \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t[(r + \lambda - \mathrm{L}) g(X_{t+u}, Y_{t+u})] \mathrm{d}u,$$

which follows by the Markov property in combination with the resolvent formalism as long as function g is sufficiently well behaved at spatial infinity.[6] Using equation (B-8), we show that $\mathrm{E}_0 \Delta_\lambda^{(1)}(X_t, Y_t)$ is indeed bounded uniformly for $t \geq 0$ as follows:

[6] It is sufficient if $\mathrm{E}_0[g(X_t, Y_t)] = o(\mathrm{e}^{-rt})$ as $t \to \infty$. Further, the quantity $\mathrm{L}g(X_t, Y_t)$ does not necessarily need to exist as a (bounded) function; it needs only to exist, possibly as a distribution function, under the expectation (i.e., integral) operator E_t for almost all $t \geq 0$ in a Lebesgue sense.

https://doi.org/10.1017/S0021109019000048 Published online by Cambridge University Press

(B-9) $$\sup_{t\geq 0} \mathrm{E}_0\, \Delta_\lambda^{(1)}(X_t, Y_t)$$

$$= \sup_{t\geq 0} \mathrm{E}_0 \left[V_\lambda^{(2)}(X_t, Y_t) - V_\lambda^{(1)}(X_t, Y_t) \right],$$

$$= \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda\, \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \left[\max \left\{ g(X_{t+u}, Y_{t+u}), V_\lambda^{(1)}(X_{t+u}, Y_{t+u}) \right\} - g(X_{t+u}, Y_{t+u}) \right] \mathrm{d}u,$$

which follows from (B-3)

$$= \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \left\{ V_\lambda^{(1)}(X_{t+u}, Y_{t+u}) - g(X_{t+u}, Y_{t+u}) \right\}^+ \mathrm{d}u,$$

$$= \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \left\{ \int_0^\infty \mathrm{e}^{-(r+\lambda)s} \mathrm{E}_{t+u} \left[\lambda g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s}) \right] \mathrm{d}s - g(X_{t+u}, Y_{t+u}) \right\}^+ \mathrm{d}u,$$

$$= \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \left\{ \int_0^\infty \mathrm{e}^{-(r+\lambda)s} \mathrm{E}_{t+u} [\lambda g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s})] \mathrm{d}s - \int_0^\infty \mathrm{e}^{-(r+\lambda)s} \mathrm{E}_{t+u} [(r+\lambda-\mathrm{L}) g(X_{t+u+s}, Y_{t+u+s})] \mathrm{d}s \right\}^+ \mathrm{d}u,$$

which follows from (B-8)

$$= \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \left\{ \int_0^\infty \mathrm{e}^{-(r+\lambda)s} \mathrm{E}_{t+u} \left[(\mathrm{L}-r) g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s}) \right] \mathrm{d}s \right\}^+ \mathrm{d}u,$$

which follows by combining terms

$$\leq \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \int_0^\infty \mathrm{e}^{-(r+\lambda)s} \left\{ \mathrm{E}_{t+u} \left[(\mathrm{L}-r) g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s}) \right] \right\}^+ \mathrm{d}s \mathrm{d}u,$$

which follows from $\left\{ \int [\cdot] \right\}^+ \leq \int [\{\cdot\}^+]$

$$\leq \sup_{t\geq 0} \mathrm{E}_0 \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_t \int_0^\infty \mathrm{e}^{-\lambda s} \sup_{s\geq 0} \mathrm{e}^{-rs} \left\{ \mathrm{E}_{t+u} \left[(\mathrm{L}-r) g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s}) \right] \right\}^+ \mathrm{d}s \mathrm{d}u,$$

by inserting a sup over $s \geq 0$ inside the integral over s

$$= \frac{1}{\lambda} \sup_{t\geq 0} \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \mathrm{E}_0 \sup_{s\geq 0} \mathrm{e}^{-rs} \left\{ \mathrm{E}_{t+u} \left[(\mathrm{L}-r) g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s}) \right] \right\}^+ \mathrm{d}u,$$

by the tower property ($\mathrm{E}_0 \mathrm{E}_t = \mathrm{E}_0$), and integral over s gives λ^{-1}

$$\leq \frac{1}{\lambda} \int_0^\infty \lambda \mathrm{e}^{-(r+\lambda)u} \sup_{t,u\geq 0} \mathrm{E}_0 \sup_{s\geq 0} \mathrm{e}^{-rs} \left\{ \mathrm{E}_{t+u} \left[(\mathrm{L}-r) g(X_{t+u+s}, Y_{t+u+s}) + f(X_{t+u+s}, Y_{t+u+s}) \right] \right\}^+ \mathrm{d}u,$$

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

https://doi.org/10.1017/S0021109019000048 Published online by Cambridge University Press

by moving the $\sup_{t\geq 0}$ inside the integral, and inserting an additional $\sup_{u\geq 0}$

$$= \frac{1}{\lambda}\int_0^\infty \lambda e^{-(r+\lambda)u} \sup_{t\geq 0} E_0 \sup_{s\geq 0} e^{-rs}\left\{E_t\left[(L-r)g(X_{t+s},Y_{t+s}) + f(X_{t+s},Y_{t+s})\right]\right\}^+ du,$$

by replacing $t+u$ by t, and merging suprema over $t,u\geq 0$

$$= \frac{1}{\lambda}\int_0^\infty \lambda e^{-(r+\lambda)u} c(X_0,Y_0)\,du,$$

by definition of $c(X_0,Y_0)$ in (B-1)

$$= \frac{1}{r+\lambda} c(X_0,Y_0).$$

In sum, we have shown that the right-hand side of (B-6) is bounded. In turn, this means that the maximum possible increase in the approximate solution shrinks at the geometric rate $\lambda/(r+\lambda)$ as the number of POST iterations increases; thus algorithm (5) is convergent. This concludes the proof of Theorem 1.

Proof of Theorem 2. Continuing with the same notation as in the proof of Theorem 1, we note that $V_\lambda(X,Y) = V_\lambda^{(j)}(X,Y) + \sum_{k=j}^\infty \Delta_\lambda^{(k)}(X,Y)$, which implies

$$\begin{aligned}
0 &\leq V_\lambda(X_0,Y_0) - V_\lambda^{(j)}(X_0,Y_0)\\
&= \sum_{k=j}^\infty \Delta_\lambda^{(k)}(X_0,Y_0),\\
&\leq \sum_{k=j}^\infty \left(\frac{\lambda}{r+\lambda}\right)^{k-1} \sup_{t\geq 0} E_0\Delta_\lambda^{(1)}(X_t,Y_t), \text{ from equation (B-6)}\\
&\leq \sum_{k=j}^\infty \left(\frac{\lambda}{r+\lambda}\right)^{k-1} \frac{1}{r+\lambda} c(X_0,Y_0), \text{ from the last line of equation (B-11)}\\
&= \frac{1}{\lambda}\sum_{k=j}^\infty \left(\frac{\lambda}{r+\lambda}\right)^{k} c(X_0,Y_0)\\
&= \frac{1}{\lambda}\left(\frac{\lambda}{r+\lambda}\right)^{j} \sum_{k=j}^\infty \left(\frac{\lambda}{r+\lambda}\right)^{k-j} c(X_0,Y_0)\\
&= \frac{1}{\lambda}\left(\frac{\lambda}{r+\lambda}\right)^{j} \frac{1}{1-\frac{\lambda}{r+\lambda}} c(X_0,Y_0)\\
&= \frac{1}{\lambda}\left(\frac{\lambda}{r+\lambda}\right)^{j} \frac{r+\lambda}{r} c(X_0,Y_0),\\
&= \left(\frac{\lambda}{r+\lambda}\right)^{j-1} \frac{c(X_0,Y_0)}{r}.
\end{aligned}$$

Hence the computable bound (6) holds and any desired precision level can be achieved by choosing the number of iterations j sufficiently large.

Appendix C. Convergence of Algorithm (7)

Convergence of algorithm (7) as $j \to \infty$ is guaranteed at the geometric rate $\lambda/(r+\lambda)$ if the matrix $\mathbf{L}$ satisfies the following conditions:

1. The matrix $\mathbf{L}$ is weakly diagonally dominant. That is, in each row the absolute value of the diagonal element is greater than or equal to the sum of the absolute values of the off-diagonal elements.
2. The matrix $\mathbf{L}$ contains nonpositive diagonal entries and nonnegative off-diagonal entries.

If these conditions hold, then for any $r, \lambda > 0$ the inverse $((r+\lambda)\mathbf{I}-\mathbf{L})^{-1}$ exists and contains only nonnegative entries (Plemmons (1977)), while its operator norm (with respect to the infinity norm under $\mathbb{R}^M$, where M is the number of grid points) is bounded above by $1/(r+\lambda)$ (Ahlberg and Nilson (1963)). It can easily be shown that the nonnegativity yields the monotonicity of algorithm (7), while the norm yields the contraction.

Appendix D. Discretization

In the base case considered by AP (see our footnote 3), $\rho = 0$. In this case, the 5-point stencil below generates a matrix $\mathbf{L}$ that satisfies both conditions in Appendix C:

$$\begin{array}{|c|c|c|}
\hline
\cdot & \dfrac{Y^2\sigma_Y^2}{2\mathrm{d}Y^2} + \dfrac{Y\theta_Y^+}{\mathrm{d}Y} & \cdot \\
\hline
\dfrac{X^2\sigma_X^2}{2\mathrm{d}X^2} + \dfrac{X^2\theta_X^-}{\mathrm{d}X} & -\dfrac{X^2\sigma_X^2}{\mathrm{d}X^2} - \dfrac{Y^2\sigma_Y^2}{\mathrm{d}Y^2} - \dfrac{X|\theta_X|}{\mathrm{d}X} - \dfrac{Y|\theta_Y|}{\mathrm{d}Y} & \dfrac{X^2\sigma_X^2}{2\mathrm{d}X^2} + \dfrac{X\theta_X^+}{\mathrm{d}X}, \\
\hline
\cdot & \dfrac{Y^2\sigma_Y^2}{2\mathrm{d}Y^2} + \dfrac{Y\theta_Y^-}{\mathrm{d}Y} & \cdot \\
\hline
\end{array}$$

where $(\cdot)^+ = \max(0, \cdot)$ and $(\cdot)^- = \max(0, -\cdot)$ denote the positive and negative parts of a real number, while $\mathrm{d}X$ and $\mathrm{d}Y$ denote the horizontal and vertical spacing of X and Y. For presentation purposes we use a homogeneous grid; in computation we do not. Condition 1 in Appendix C is satisfied since the value in the center, which ends up on the diagonal of $\mathbf{L}$, is not exceeded in absolute value by sum of the other values in the stencil. Condition 2 in Appendix C is satisfied because the center value, which is placed on the diagonal of $\mathbf{L}$, is nonpositive, while the values corresponding to its four neighbors, which are placed on off-diagonal elements of $\mathbf{L}$, are nonnegative.

In the approximation of second derivatives, the 5-point stencil uses a central difference scheme, which uses grid points to both sides of the center point. In the approximation of first derivatives, we only use the neighbor on one side; which one is chosen depends on the direction of the drift. In the horizontal direction, the nearest neighbor on the right gets a positive value if the drift is towards the right. Conversely, the nearest neighbor on the left gets a positive value if the drift is towards the left. In both cases, negative values are guaranteed to end up at the center of the stencil, such that condition 2 in Appendix C is guaranteed.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

While $\rho = 0$ in the base case of AP, a nonzero correlation can be dealt with by using the following stencil

$$\begin{array}{c|c|c} -\dfrac{\rho\sigma_X\sigma_Y XY}{4\,\mathrm{d}X\,\mathrm{d}Y} & \cdot & \dfrac{\rho\sigma_X\sigma_Y XY}{4\,\mathrm{d}X\,\mathrm{d}Y} \\ \hline \cdot & \cdot & \cdot \\ \hline \dfrac{\rho\sigma_X\sigma_Y XY}{4\,\mathrm{d}X\,\mathrm{d}Y} & \cdot & -\dfrac{\rho\sigma_X\sigma_Y XY}{4\,\mathrm{d}X\,\mathrm{d}Y} \end{array},$$

see, for example, Abramovitz and Stegun ((1972), p. 884, eq. 25.3.26). While this stencil does not satisfy the diagonal dominance and sign requirements in Appendix C, convergence is found to be fast regardless. In any case, correlated processes may be avoided altogether by framing the problem in terms of two uncorrelated processes; something that for correlated geometric Brownian motions is always possible.

Appendix E. Grid and Boundary Conditions

Our inhomogeneous rectilinear grid contains $M = 201$ grid points and ranges from 10^{-6} to 3,000 in both directions, where the grid points are placed at

$$[10^{-6}, (1:1:100), (102:2:200), (210:10:400), (425:25:500), (550:50:1{,}000), (1{,}500:100:3{,}000)].$$

This grid proved to be sufficiently large and fine-grained to ensure a discretized solution of acceptable precision and resolution.

With respect to the boundary conditions, when we reach the edge of our grid, some points in our stencil may not be "available." The simplest method for dealing with nonexistent grid points is to set the corresponding value to 0, which leads to Dirichlet boundary conditions. Alternatively, we may set the corresponding value equal to that of the center value, leading to Neumann boundary conditions. Finally, we may use a linear extrapolation to obtain the value corresponding to the nonexistent grid point. For our results, we impose Dirichlet boundary conditions on the bottom, left-, and right-hand sides of our grid, while we perform a linear extrapolation at the top, where the project value is particularly high (and increasing). We experimented with Dirichlet and Neumann boundary conditions on all four sides and found the differences to be sufficiently small.

Appendix F. Robustness

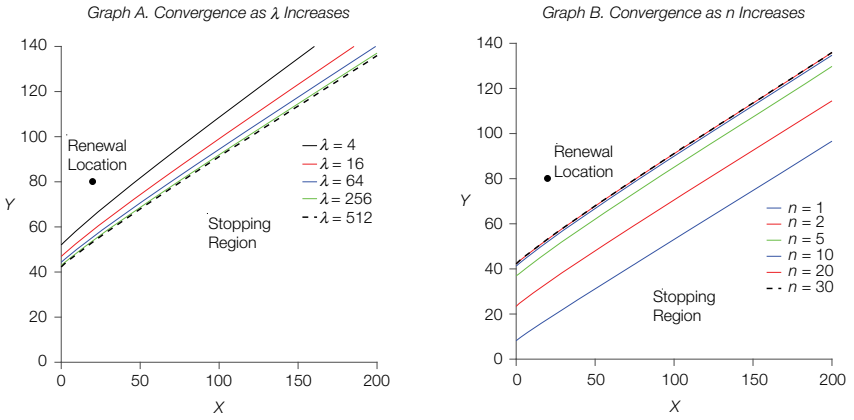
Graph A of Figure F1 shows that, as λ increases, the stopping boundary moves downwards, so that the decision maker waits longer before renewing the project. That is, if opportunities to renew arrive more frequently, the decision maker can afford to wait longer to renew. In contrast, Graph B shows that, as n increases, the stopping boundary moves upwards, so that the decision maker renews the project earlier. That is, if a larger number of renewals is allowed, the decision maker is more inclined to renew. The black dotted line, which corresponds to $\lambda = 512$ and $n = 30$, is the same in both figures. It is clear that these values of λ and n are sufficiently large for the stopping boundary to be stable.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

FIGURE F1
Convergence of the Stopping Boundary

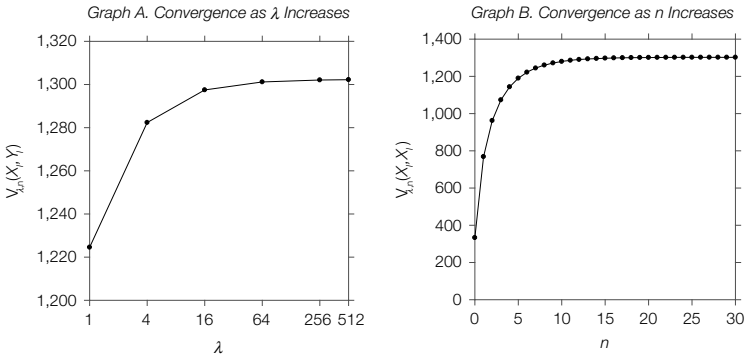
Graph A of Figure F1 shows the stopping boundary of the AP renewal option discussed in Section III for $n=30$ and different values of λ. Graph B shows the stopping boundary for $\lambda=512$ and different values of n. The black dotted line, which represents the optimal stopping boundary for $\lambda=512$ and $n=30$, is shown in both Graphs A and B.



Graph A of Figure F2 shows that the AP project value at the renewal location, that is $V_{\lambda,n=30}(X_I, Y_I)$, is relatively insensitive as a function of the Poisson intensity λ for $\lambda \geq 64$. Similarly, Graph B shows that $V_{\lambda=512,n}(X_I, Y_I)$ is relatively insensitive as a function of n for $n \geq 15$. The project value $V_{\lambda=512,n=30}(X_I, Y_I) = 1,302$ is shown in both figures and referenced in the main text.

FIGURE F2
Convergence of the Value Function

Graph A of Figure F2 shows the project value $V_{\lambda,n}(X_I, Y_I)$ with $n=30$ as a function of λ. Graph B shows the project value $V_{\lambda,n}(X_I, Y_I)$ with $\lambda=512$ as a function of n.



References

Abel, A. B., and J. C. Eberly. "Optimal Investment with Costly Reversibility." *Review of Economic Studies*, 63 (1996), 581–593.

Abramovitz, M., and I. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs and Mathematical Tables,* Vol. 55, 10th ed. Washington, DC: National Bureau of Standards (1972).

Adkins, R., and D. Paxson. "Reciprocal Energy-Switching Options." *Journal of Energy Markets*, 4, 1 (2011a), 91–120.

Adkins, R., and D. Paxson. "Renewing Assets with Uncertain Revenues and Operating Costs." *Journal of Financial and Quantitative Analysis*, 46 (2011b), 785–813.

Adkins, R., and D. Paxson. "Deterministic Models for Premature and Postponed Replacement." *Omega*, 41 (2013a), 1008–1019.

Adkins, R., and D. Paxson. "The Effect of Tax Depreciation on the Stochastic Replacement Policy." *European Journal of Operational Research*, 229 (2013b), 155–164.

Adkins, R., and D. Paxson. "Stochastic Equipment Capital Budgeting with Technological Progress." *European Financial Management*, 20 (2014), 1031–1049.

Adkins, R., and D. Paxson. "The Effects of an Uncertain Abandonment Value on the Investment Decision." *European Journal of Finance*, 23 (2017a), 1083–1106.

Adkins, R., and D. Paxson. "Replacement Decisions with Multiple Stochastic Values and Depreciation." *European Journal of Operational Research*, 257 (2017b), 174–184.

Adkins, R., and D. Paxson. "Sequential Investments with Stage-Specific Risks and Drifts." *European Journal of Finance*, 23 (2017c), 1150–1175.

Ahlberg, J., and E. Nilson. "Convergence Properties of the Spline Fit." *Journal of the Society for Industrial and Applied Mathematics*, 11 (1963), 95–104.

Andersen, L., and M. Broadie. "Primal-Dual Simulation Algorithm for Pricing Multidimensional American Options." *Management Science*, 50 (2004), 1222–1234.

Armada, M. J. R.; P. J. Pereira; and A. Rodrigues. "Optimal Investment with Two-Factor Uncertainty." *Mathematics and Financial Economics*, 7 (2013), 509–530.

Bally, V., and J. Printems. "A Quantization Tree Method for Pricing and Hedging Multidimensional American Options." *Mathematical Finance*, 15 (2005), 119–168.

Bensoussan, A., and J.-L. Lions. *Applications of Variational Inequalities in Stochastic Control*, Vol. 12, Amsterdam, Netherlands: Elsevier (1982).

Biagini, F., and T. Björk. "On the Timing Option in a Futures Contract." *Mathematical Finance*, 17 (2007), 267–283.

Birge, J. R., and V. Linetsky. *Handbooks in Operations Research and Management Science: Financial Engineering*, Vol. 15, Amsterdam, Netherlands: Elsevier (2007).

Bjerksund, P., and G. Stensland. "Closed Form Spread Option Valuation." *Quantitative Finance*, 14 (2014), 1785–1794.

Brennan, M. J., and E. S. Schwartz. "Evaluating Natural Resource Investments." *Journal of Business*, 58 (1985), 135–157.

Bunch, D. S., and H. Johnson. "The American Put Option and Its Critical Stock Price." *Journal of Finance*, 55 (2000), 2333–2356.

Carmona, R., and S. Dayanik. "Optimal Multiple Stopping of Linear Diffusions." *Mathematics of Operations Research*, 33 (2008), 446–460.

Carmona, R., and N. Touzi. "Optimal Multiple Stopping and Valuation of Swing Options." *Mathematical Finance*, 18 (2008), 239–268.

Carr, P. "Randomization and the American Put." *The Review of Financial Studies*, 11 (1998), 597–626.

Compernolle, T.; K. Huisman; P. Kort; M. Lavrutich; C. Nunes; and J. Thijssen. "Investment Decisions with Two-Factor Uncertainty." CentER Discussion Paper Series No. 2018-003 (2018).

Cont, R., and D.-A. Fournié. "Functional Itô Calculus and Stochastic Integral Representation of Martingales." *Annals of Probability*, 41 (2013), 109–133.

Cortazar, G.; M. Gravet; and J. Urzua. "The Valuation of Multidimensional American Real Options Using the LSM Simulation Method." *Computers & Operations Research*, 35 (2008), 113–129.

Cox, J. C.; S. A. Ross; and M. Rubinstein. "Option Pricing: A Simplified Approach." *Journal of Financial Economics*, 7 (1979), 229–263.

d'Halluin, Y.; P. A. Forsyth; and G. Labahn. "A Penalty Method for American Options with Jump Diffusion Processes." *Numerische Mathematik*, 97 (2004), 321–352.

Dixit, A., and R. Pindyck. *Investment under Uncertainty*. Princeton, NJ: Princeton University Press (1994).

Dockendorf, J., and D. Paxson. "Continuous Rainbow Options on Commodity Outputs: What Is the Real Value of Switching Facilities?" *European Journal of Finance*, 19 (2013), 645–673.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

Dupuis, P., and H. Wang. "Optimal Stopping with Random Intervention Times." *Advances in Applied Probability*, 34 (2002), 141–157.
Ekern, S. "An Option Pricing Approach to Evaluating Petroleum Projects." *Energy Economics*, 10 (1988), 91–99.
Farzan, F.; K. Mahani; K. Gharieh; and M. A. Jafari. "Microgrid Investment under Uncertainty: A Real Option Approach Using Closed Form Contingent Analysis." *Annals of Operations Research*, 235 (2015), 259–276.
Feng, L., and V. Linetsky. "Pricing Options in Jump-Diffusion Models: An Extrapolation Approach." *Operations Research*, 56 (2008), 304–325.
Forsyth, P. A., and K. R. Vetzal. "Quadratic Convergence for Valuing American Options Using a Penalty Method." *SIAM Journal on Scientific Computing*, 23 (2002), 2095–2122.
Glowinski, R. *Lectures on Numerical Methods for Non-Linear Variational Problems*. Berlin, Germany: Springer Science & Business Media (2008).
Heydari, S.; N. Ovenden; and A. Siddiqui. "Real Options Analysis of Investment in Carbon Capture and Sequestration Technology." *Computational Management Science*, 9 (2012), 109–138.
Lange, R.-J.; D. Ralph; and J. van Casteren. "A Vanilla Framework for Solving Multidimensional Options Using Poisson Optional Stopping Times." Working Paper, Erasmus University Rotterdam (2018).
Lange, R.-J., and C. Teulings. "The Option Value of Vacant Land and the Optimal Timing of City Extensions." CEPR Discussion Paper No. DP12847 (2018).
Lempa, J. "Optimal Stopping with Information Constraint." *Applied Mathematics & Optimization*, 66 (2012), 147–173.
Longstaff, F. A., and E. S. Schwartz. "Valuing American Options by Simulation: A Simple Least-Squares Approach." *Review of Financial Studies*, 14 (2001), 113–147.
Majd, S., and R. S. Pindyck. "Time to Build, Option Value, and Investment Decisions." *Journal of Financial Economics*, 18 (1987), 7–27.
McDonald, R., and D. Siegel. "The Value of Waiting to Invest." *Quarterly Journal of Economics*, 101 (1986), 707–727.
Peskir, G., and A. Shiryeav. *Optimal Stopping and Free-Boundary Problems*. Basel, Switzerland: Birkhäuser (2006).
Pindyck, R. S. "Irreversibility, Uncertainty, and Investment." Technical Report, National Bureau of Economic Research (1990).
Plemmons, R. J. "*M*-Matrix Characterizations. I—Nonsingular *M*-Matrices." *Linear Algebra and Its Applications*, 18 (1977), 175–188.
Revuz, D., and M. Yor. *Continuous Martingales and Brownian Motion*, Vol. 293, Berlin, Germany: Springer (1999).
Rogers, L. C. "Monte Carlo Valuation of American Options." *Mathematical Finance*, 12 (2002), 271–286.
Rogers, L. C., and D. Williams. *Diffusions, Markov Processes, and Martingales*, Vol. 1, Cambridge, UK: Cambridge University Press (2000).
Rogers, L. C., and O. Zane. "A Simple Model of Liquidity Effects." In *Advances in Finance and Stochastics*, K. Sandmann and P. J. Schönbucher, eds. Berlin, Germany: Springer (2002), 161–176.
Rohlfs, W., and R. Madlener. "Valuation of CCS-Ready Coal-Fired Power Plants: A Multi-Dimensional Real Options Approach." *Energy Systems*, 2 (2011), 243–261.
Schwartz, E. S. "The Valuation of Warrants: Implementing a New Approach." *Journal of Financial Economics*, 4 (1977), 79–93.
Stentoft, L. "Convergence of the Least Squares Monte Carlo Approach to American Option Valuation." *Management Science*, 50 (2004), 1193–1203.
Støre, K.; S.-E. Fleten; V. Hagspiel; and C. Nunes. "Switching from Oil to Gas Production in a Depleting Field." *European Journal of Operational Research*, 271 (2018), 710–719.
Stroock, D. W., and S. R. S. Varadhan. *Multidimensional Diffusion Processes*. Berlin, Germany: Springer (2007).
Strulovici, B., and M. Szydlowski. "On the Smoothness of Value Functions and the Existence of Optimal Strategies in Diffusion Models." *Journal of Economic Theory*, 159 (2015), 1016–1055.
Stutzman, S.; B. Weiland; P. Preckel; and M. Wetzstein. "Optimal Replacement Policies for an Uncertain Rejuvenated Asset." *International Journal of Production Economics*, 185 (2017), 21–33.
Svenstrup, M. "On the Suboptimality of Single-Factor Exercise Strategies for Bermudan Swaptions." *Journal of Financial Economics*, 78 (2005), 651–684.
Trigeorgis, L. "A Log-Transformed Binomial Numerical Analysis Method for Valuing Complex Multi-Option Investments." *Journal of Financial and Quantitative Analysis*, 26 (1991), 309–326.
Trigeorgis, L. "The Nature of Option Interactions and the Valuation of Investments with Multiple Real Options." *Journal of Financial and Quantitative Analysis*, 28 (1993), 1–20.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press

Wang, S.; X. Yang; and K. Teo. "Power Penalty Method for a Linear Complementarity Problem Arising from American Option Valuation." *Journal of Optimization Theory and Applications*, 129 (2006), 227–254.

Zhang, K.; S. Wang; X. Yang; and K. L. Teo. "A Power Penalty Approach to Numerical Solutions of Two-Asset American Options." *Numerical Mathematics: Theory, Methods and Applications*, 2 (2009), 202–223.

Zhang, K.; X. Yang; and K. L. Teo. "Convergence Analysis of a Monotonic Penalty Method for American Option Pricing." *Journal of Mathematical Analysis and Applications*, 348 (2008), 915–926.

Zvan, R.; P. A. Forsyth; and K. R. Vetzal. "Penalty Methods for American Options with Stochastic Volatility." *Journal of Computational and Applied Mathematics*, 91 (1998), 199–218.

https://doi.org/10.1017/S0022109019000048 Published online by Cambridge University Press