Optimal Stopping with Randomly Arriving Stopping Opportunities*

Josha A. Dekker
Dept. of Quantitative Economics
University of Amsterdam
and Tinbergen Institute
J.A.Dekker@uva.nl

Roger J. A. Laeven†
Dept. of Quantitative Economics
University of Amsterdam, EURANDOM
and CentER
R.J.A.Laeven@uva.nl

John G. M. Schoenmakers
Stochastic Algorithms & Nonparametric Statistics
Weierstrass Institute Berlin
schoenma@wias-berlin.de

Michel H. Vellekoop
Dept. of Quantitative Economics
University of Amsterdam
M.H.Vellekoop@uva.nl

August 29, 2026

Abstract

We develop simulation-based methods to solve general optimal stopping problems with opportunities to stop that arrive randomly. Such problems occur naturally in a wide variety of applications with market frictions, such as limited liquidity, incomplete information and transaction costs. Our approach employs random time scales to map the original problem to an integer-time stopping problem with (possibly) infinite horizon. We introduce an infinite-horizon policy iteration algorithm to generate lower-biased estimates of the value function and establish a martingale dual representation to generate upper-biased estimates. Furthermore, we extend a family of backward dynamic programming methods that includes least-squares Monte Carlo. We illustrate the performance of our methods, their general applicability and the managerial implications of randomly arriving opportunities to stop in three examples: a stopped jump-diffusion that can be analyzed in closed form, a random-exercise version of a multi-dimensional max-call contract, and a real options problem in an illiquid market.

Keywords: Optimal stopping on random times; Infinite-horizon optimal stopping problems; Duality; Policy iteration; Regression.

AMS 2020 Subject Classification: *Primary*: 60G40; 65C05; *Secondary*: 93E24; 91G60.

*We are very grateful to Albert Menkveld, Peter Spreij, and Florian Wagener for helpful suggestions on an early draft of this paper. We are also grateful to conference and seminar participants at the SIAM FM23 Conference in Philadelphia, the Stochastic Analysis and Stochastics of Financial Markets Seminar at Humboldt University Berlin, the 26th International IME Congress in Edinburgh, the Amsterdam–Leuven–London joint PhD workshop 2024 in Leuven, the 15th AFMath Conference in Brussels, the ICCF24 Conference in Amsterdam, DynStoch 2024 in Kiel, the 3rd Dutch Mathematical Finance Afternoon in Amsterdam, Stochastics and Computational Finance 2025 in Lisbon, the 2026 World Congress of the Bachelier Finance Society in Bologna, the University of Amsterdam, the Tinbergen Institute and the Weierstrass Institute for their comments and suggestions. This research was funded in part by the Netherlands Organization for Scientific Research (NWO) under grant NWO-Vici 2020–2027 (Laeven) and by the DFG Excellence Cluster Math+ Berlin, project AA4-2 (Schoenmakers).

†Corresponding author. University of Amsterdam, PO Box 15867, 1001 NJ Amsterdam, The Netherlands.

1 Introduction

Since the early work of Wald ([55]) and Snell ([54]), a vast literature has developed the theory of optimal stopping. Its areas of application encompass a wide variety of fields, including statistics (sequential testing), operations research (secretary and search problems), and financial (Bermudan and American option pricing) and insurance mathematics (risk process optimization); see, e.g., [48] for a textbook treatment.

A standard assumption in this literature is that opportunities to stop arise deterministically. For example, for financial options that allow for early exercise timing, i.e., exercise before or at a given maturity date, stopping is typically restricted to either pre-specified time instances (Bermudan options) or to any time within a pre-specified time interval (American options). More recently, motivated by market frictions, problems in which the stopping opportunities form a random set have been studied. Early examples assume the stopping opportunities to be generated by a standard Poisson process, independent of the reward process ([3, 18]).

This paper develops simulation-based numerical methods to solve general optimal stopping problems with randomly arriving opportunities to stop. Pivotal to our approach is to proceed on a random time scale, from one randomly arriving opportunity to the next.[1] This enables us to convert the original optimal stopping problem into a discrete-time optimal stopping problem, with respect to integer-valued stopping times with filtration generated by the consecutively arriving random times, on an infinite horizon. To encompass the wide variety of problems in which randomly arriving stopping opportunities naturally occur, our setting allows for general dynamics of the stopping opportunities and general dynamics underlying the reward process, including rich forms of stochastic dependence between the stopping opportunities and the reward process. Moreover, we consider applications in which the stopping problem is multi-dimensional (easily up to dimension 100) and in which state transition densities are not available. In such settings, existing numerical solution methods typically become intractable.

We provide both lower and upper bounds to the value function of the stopping problem. To obtain a lower-biased estimate of the value function, we develop a policy improvement procedure using forward induction. This non-trivially extends [33], which was developed in a finite-horizon setting. More specifically, since the theory in [33] is entirely based on backward induction from a finite horizon, we extend their main statements to an infinite-horizon setting and give new, essentially different proofs. To obtain an upper bound to the value function, we first establish a suitable martingale dual representation, and next employ it in a corresponding Monte Carlo simulation procedure that we design. We also extend a family of backward dynamic programming methods, including least-squares Monte Carlo (LSMC, [41]). If the dimension of the underlying process is low, the resulting random times LSMC method is usually accurate and fast. It suffers, however, from the curse of dimensionality, in contrast to policy iteration. Therefore, for high-

[1] That is, instead of deploying a — possibly infinitesimally fine — deterministic time grid, our approach consists in considering a random time scale. On the random time scale, the order of the arrivals of stopping opportunities is encoded in the set of natural numbers rather than designated by the actual arrival times on the positive real line.

dimensional underlying processes, backward dynamic programming serves primarily as an input for our policy improvement approach.

We illustrate the good performance of our methods in three applications: a geometric Brownian motion with jumps stopped at random times, which can be analyzed in explicit closed form; a random-exercise version of a multi-dimensional max-call option pricing problem; and a real options problem in a scarce, interconnected environment. In particular, our methods are shown to deliver (very) small duality gaps, and to be applicable under a wide range of problem specifications allowing for general multivariate dynamics generating the stopping opportunities and reward. Furthermore, we show that, compared to optimal stopping problems with deterministic stopping opportunities, problems with randomly arriving stopping opportunities feature two additional sources of risk: *timing risk*, referring to the change in value due to the randomness in the inter-arrival times between stopping opportunities; and *opportunity risk*, referring to the change in value due to the possibility that no further opportunities to stop are seen. From a managerial point of view, ignoring either source of risk will lead to incorrect pricing and decision making. Whereas *opportunity risk* typically warrants a more prudent stopping decision and reduces the value function, we show that *timing risk* may be beneficial or detrimental, depending on the specific application.

Potential applications of our methods are numerous. When *information* is *scarce*, it is natural to consider optimal stopping problems where stopping is restricted to the random times at which an information signal is received; see e.g., [18, 39, 44, 43] for applications to American option valuation. In such a setting, both cross-sectional and temporal dependence between the process generating the signal and the process generating the reward occurs naturally: a signal from the release of information may affect underlying prices and hence the reward, which in turn may generate further signals. Furthermore, in the presence of *limited liquidity*, or *transaction costs*, stopping opportunities may be linked to the arrival of (attractive) trading opportunities. Indeed, [51] studies an optimal consumption and investment problem where trades are executed at a Poisson rate instead of continuously, representing, for example, the arrival of buyers. Due to the existence of asset-liquidity spirals (see, e.g., [29]), dependence between underlying asset prices and liquidity risk, hence the arrival of trading opportunities, arises naturally in this setting.

In *real option* pricing, constraints of various forms can lead to randomly arriving stopping opportunities. Examples are the arrival of scarce goods or investment opportunities ([35, 37]), or unexpected technological progress ([3]). In these problems, it is again natural to allow for dependence between the constraints affecting the stopping opportunities and the reward process; we illustrate this in Section 6, in the context of selling a valuable object, employing multivariate point processes allowing for feedback including self- and cross-excitation ([27, 2]) to generate the stopping opportunities and reward. In the *bank run* literature, [28, 40] suppose that short-term credit matures randomly, at which point a creditor needs to make a stopping decision: either roll-over debt or run. This gives rise to a debt run model based on the classical staggered pricing model of [11] and the credit risk model by [38]. Other financial and environmental stopping problems also lend themselves naturally to randomly arriving opportunities. Examples in *energy economics*

include the exploration of natural resources ([47, 16]), weather derivatives ([10, 9]) or emissions reduction projects ([21]).

Optimal stopping problems with constraints may be represented in terms of a Hamilton-Jacobi-Bellman (HJB) equation or a free boundary formulation; see e.g., [42], for a detailed account. For general Markov processes, an explicit solution of the continuous-time HJB equation is generally not available; the solution may then be obtained, theoretically and in some cases analytically, through a discrete-time HJB associated to the continuous-time problem. Stopping problems with randomly arriving opportunities also naturally give rise to a sequence of partial differential equations (PDEs), which may then be solved iteratively. Indeed, if we consider a sequence of problems in which stopping is only allowed at the first k opportunities (in 'problem k'), then this problem may be formulated as a PDE where the boundary condition is determined by the solution to 'problem $k-1$'. This value function iteration approach is taken for infinite-horizon stopping problems in, e.g., [35, 37]. Their 'POST' algorithm exhibits geometric convergence of the iterated value functions to the true value function, and [35] indicate how it may be adapted for finite-horizon problems by including time explicitly as a state-variable. Value function iteration has also been used for control problems with random control opportunities, both in an infinite-horizon ([45, 46]) and a finite-horizon setting ([22, 23]). The generality of our setting — in terms of the processes generating the stopping opportunities, the reward processes, the dependences among them, or the dimensionality of our problem — poses challenges to obtaining numerical solutions using HJB-based methods or value function iteration.

For specific dynamics of the reward and stopping opportunities, explicit solutions may be obtained for stationary (originally) infinite-horizon stopping problems. Indeed, for rewards driven by geometric Brownian motion and independent Poissonian arrivals of stopping opportunities (hence, inter-arrival times that are exponential), explicit solutions may be found through variational inequalities and smooth-fit principles ([18, 25]), resolvent methods ([39]) or Green representations ([3, 40]). Furthermore, recent developments in the fluctuation theory of Lévy processes may be used to solve stopping problems where rewards follow a Lévy process, using scale functions and other properties of Lévy processes; see, e.g., [44, 43]. For a more extensive overview of recent developments in the Poissonian case, see [52]. When arrivals occur according to an inhomogenous Poisson process with state-dependent rate, stochastic coupling and a change of time scale may be used to derive theoretical solutions ([30]), while if the rate may be controlled, one may obtain a differential equation for the value function that may then be solved ([31]).

For finite-horizon problems, the explicit time-dependence makes it harder to obtain solutions, theoretically and numerically. For a problem with state-dependent Cox arrivals, [40] use Green representations to obtain a free-boundary problem for the value function; see also [42, 35] discussed above, and [19] who treat a related problem where stopping is only allowed while the state process is in a subset of the state space.

We summarize our main contributions. Methodologically, we show how general finite-horizon stopping problems with randomly arriving opportunities to stop can be transformed into integer-

time, infinite-horizon problems, by mapping deterministic into random time scales, and establish equivalence between their solutions. For the induced infinite-horizon problem, we develop and formally substantiate a policy improvement procedure using forward induction; this significantly extends [33] that is entirely based on backward induction from a finite horizon. Furthermore, we establish and design a martingale dual representation and a corresponding Monte Carlo procedure. Besides, we also provide a suitable truncation of the horizon that allows a class of backward dynamic programming methods to be extended to our problem, including the least-squares Monte Carlo method in [41] and the threshold methods in [4]. Our simulation-based methods are demonstrated to yield good performance in a wide variety of problems and are generally applicable, also in settings in which existing methods become impractical.[2] From a managerial perspective, we show, across a variety of model specifications, how the randomness of the stopping opportunities impacts both the optimal value and the strategies of the stopping problem, by analyzing their *timing* and *opportunity* risk; these can be (very) substantial.

This paper is organized as follows. In Section 2, we formally introduce the problem and establish equivalence, Bellman and duality results. In Section 3, we develop a policy iteration procedure that uses forward induction. Section 4 describes the algorithms that we propose. In Section 5, we present extensive numerical illustrations to assess the accuracy and efficiency of our algorithms, and in Section 6, we analyze a real options example to illustrate the relevance of randomly arriving opportunities for optimal stopping. Conclusions are in Section 7. Proofs that are not included in the main text, additional technical details and additional numerical results are contained in the Appendix.

2 Optimal Stopping on Random Times

We consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ with filtration $\mathbb{F} := (\mathcal{F}_t)_{0\leq t\leq\infty}$ satisfying "the usual conditions" and suppose that we are given a strictly increasing, random sequence of $\mathbb{F}$-stopping times satisfying

$$\tau_0 = 0\,, \quad \tau_i < \tau_{i+1} < \infty \quad \text{for} \quad i \in \mathbb{N}_0\,, \quad \text{and} \quad \lim_{i\to\infty} \tau_i = \infty\,, \quad \text{a.s.} \tag{2.1}$$

Here, $\mathbb{N}$ denotes the (strictly positive) natural numbers and $\mathbb{N}_0$ the natural numbers including zero. On each event ω, we denote the first $K \in \mathbb{N} \cup \{\infty\}$ realized stopping times in (2.1) by $\mathcal{T}_K(\omega) := \{\tau_i(\omega) : i \in \mathbb{N}_0,\ i \leq K\}$.

[2] Our simulation-based approach does not require solving HJBs or free boundary problems, sequences of (integro-)PDEs, or deriving explicit expressions for state transition densities; this may be hard (or even impossible) in cases where the reward process is inhomogeneous, the arrival process for opportunities to stop is non-trivial or not independent of the reward process or when the dimensionality of the problem becomes large.

2.1 Problem definition

We consider an $\mathbb{F}$-adapted reward process $(Z_t)_{0\leq t<\infty}$ that is càdlàg on $[0,T]$ and satisfies $Z_t = 0$ a.s. for $t > T$, where $T > 0$ is the finite time horizon of our problem. Our objective is to optimally stop Z on one of the elements of $\mathcal{T}_K(\omega)$, i.e., we consider the optimal stopping problem

$$Y_0^{(K)} := \sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t}\in\mathcal{T}_K} \mathbb{E}\left[Z_{\mathfrak{t}}\right] . \tag{2.2}$$

For ease of exposition, we will assume throughout that Z is non-negative. We impose the following (standing) integrability condition:

$$\mathbb{E}\left[\sup_{i\in\mathbb{N}_0} Z_{\tau_i}^{\alpha}\right] \leq B^{\alpha}\,, \quad \text{for some } B > 0 \text{ and } \alpha > 1\,. \tag{2.3}$$

For $K < \infty$, it suffices to consider the weaker bound $\mathbb{E}\left[\sup_{i\leq K} Z_{\tau_i}^{\alpha}\right] \leq B^{\alpha}$.

Depending on the specific application, the reward process (Z_t) may represent, e.g., the payoff of exercising a financial contract, the market value of a piece of real estate or an art object, the value of running when short-term credit matures, or the payoff of exploiting a natural resources project. Furthermore, the first K stopping times $\mathcal{T}_K$ may represent, for example, the arrival times of a financial exercise or trading opportunity, the release dates of valuable information, the arrival times of scarce goods or natural resources investment opportunities, or the expiry times of short-term credit.

2.2 An equivalent integer-time problem with infinite horizon

To solve the optimal stopping problem (2.2) on a random sequence of stopping times, we consider the induced discrete-time reward process, $(U_i)_{i\in\mathbb{N}_0}$, and study the corresponding infinite horizon integer-time stopping problem, where

$$U_i := Z_{\tau_i}\,, \quad i \in \mathbb{N}_0\,.$$

More specifically, due to the order of the stopping times (2.1), $\mathcal{G}_i := \mathcal{F}_{\tau_i} \subset \mathcal{F}$ defines a discrete filtration $\mathbb{G} := (\mathcal{G}_i)_{i\in\mathbb{N}_0}$, where the stop σ-algebras

$$\mathcal{F}_{\tau_i} := \{A \in \mathcal{F}_\infty : A \cap \{\tau_i \leq t\} \in \mathcal{F}_t\,,\ t \geq 0\} \tag{2.4}$$

satisfy

$$\mathcal{F}_{\tau_0} \subset \mathcal{F}_{\tau_1} \subset \mathcal{F}_{\tau_2} \subset \ldots .$$

As the discrete reward process $(U_i)_{i\in\mathbb{N}_0}$ is $\mathbb{G}$-adapted (see, e.g., [49, Thm. 6] or [20]), we may formally relate our original stopping problem (2.2) to a stopping problem over the set of discrete

$\mathbb{G}$-stopping times, $\mathcal{N}$, given by

$$\mathcal{N} := \{\mathfrak{n} : \Omega \to \mathbb{N}_0 \text{ such that } \{\mathfrak{n} = i\} \in \mathcal{G}_i \text{ for all } i \in \mathbb{N}_0\} .$$

Proposition 1 *Recall* (2.2). *We have that*

$$Y_0^{(K)} = \sup_{\mathfrak{n}\in\mathcal{N},\ \mathfrak{n}\leq K} \mathbb{E}\left[U_{\mathfrak{n}}\right] . \tag{2.5}$$

Proof. See Appendix A.1. ■

In view of (2.5), we may thus analyze (2.2) by studying $\sup_{\mathfrak{n}\in\mathcal{N},\ \mathfrak{n}\leq K} \mathbb{E}\left[U_{\mathfrak{n}}\right]$ instead. Furthermore, for the process U, the following bound follows from (2.3):

$$\mathbb{E}\left[\sup_{i\in\mathbb{N}_0} U_i^{\alpha}\right] \leq B^{\alpha} , \tag{2.6}$$

for some $B > 0$ and $\alpha > 1$.

Remark 2 *The bound* (2.3)*, hence* (2.6)*, may be determined by simulation. An explicit, analytic bound may be obtained using a slightly stronger condition in conjunction with the Burkholder-Davis-Gundy inequality; see Remark 20 in Appendix A.1.*

Denote by N_T the number of random times until we cross T, i.e.,

$$N_T := \min\{m \in \mathbb{N} : \tau_m > T\} . \tag{2.7}$$

It is clear that, for $i \geq N_T$, $U_i = Z_{\tau_i} = 0$ by our assumptions on Z, and in fact

$$\mathbb{P}(\exists \kappa < \infty : \forall i > \kappa\,,\ U_i = 0) = \mathbb{P}(N_T < \infty) = 1 . \tag{2.8}$$

In general, the sequence (2.1) may contain one or more deterministic exercise possibilities. For example, as a relevant case in applications, the decision-maker may be given the right to stop either at a random time before T or at the terminal time T, which means that we need to ensure that (2.1) satisfies a.s.,

$$\tau_{i(\omega)}(\omega) = T \text{ for a unique } i\,(\omega) \in \mathbb{N} . \tag{2.9}$$

One may establish (2.9) by intertwining the deterministic time T with a particular sequence of increasing stopping times, see Lemma 18 and Corollary 19 in Appendix A.1; these results also establish that (2.3) will continue to hold for the new sequence of stopping times if it held for the old sequence of stopping times and $\mathbb{E}[Z_T^{\alpha}]$ is finite.

2.3 Bellman principle and optimal stopping time

To formulate a Bellman principle for Problem (2.5), we begin by generalizing it to a stopping problem as seen from a generic point in discrete time $i \in \mathbb{N}_0$, which is given in terms of the

Snell envelope $Y_i^{(K)}$ and continuation value process $C_i^{(K)}$ for the reward process $(U_i)_{i\in\mathbb{N}_0}$, with $K \in \mathbb{N}\cup\{\infty\}$:

$$Y_i^{(K)} := \operatorname*{ess\,sup}_{\mathfrak{n}\in\mathcal{N},\ i\le\mathfrak{n}\le K} \mathbb{E}_{\mathcal{G}_i}\left[U_{\mathfrak{n}}\right], \quad i\in\mathbb{N}_0, \tag{2.10}$$

$$C_i^{(K)} := \operatorname*{ess\,sup}_{\mathfrak{n}\in\mathcal{N},\ i+1\le\mathfrak{n}\le K} \mathbb{E}_{\mathcal{G}_i}\left[U_{\mathfrak{n}}\right], \quad i\in\mathbb{N}_0, \quad 0\le i<K, \tag{2.11}$$

$$C_K^{(K)} := 0, \quad K<\infty.$$

It satisfies the following properties:

Proposition 3 *Under assumption* (2.3)*, we have for all* $K\in\mathbb{N}\cup\{\infty\}$ *and for all* $i\in\mathbb{N}_0$ *with* $0\le i\le K$ *that:*

(i) The continuation value process (2.11) *is connected with* (2.10) *via the Bellman principle*

$$Y_i^{(K)} = \max\left(U_i, C_i^{(K)}\right).$$

(ii) The process $Y^{(K)}$ *is a* $\mathcal{G}$*-super-martingale, i.e.,*

$$Y_i^{(K)} \ge \mathbb{E}_{\mathcal{G}_i}\left[Y_{i+1}^{(K)}\right], \tag{2.12}$$

and it is the smallest super-martingale that dominates U*.*

(iii) The random time $\mathfrak{n}_i^{K,*}$*, defined by*

$$\mathfrak{n}_i^{K,*} := \inf\left\{i\le j\le K:\ Y_j^{(K)} = U_j\right\}, \tag{2.13}$$

is an optimal stopping time for $\mathcal{G}$ *with* $\mathfrak{n}_i^{K,*}<\infty$ *a.s., hence,*

$$Y_i^{(K)} = \mathbb{E}_{\mathcal{G}_i}\left[U_{\mathfrak{n}_i^{K,*}}\right]. \tag{2.14}$$

In particular, for all $K<\infty$*,*

$$Y_i^{(K)} \le \lim_{K'\to\infty} Y_i^{(K')} = Y_i^{(\infty)} = \operatorname*{ess\,sup}_{i\le\mathfrak{n}<\infty} \mathbb{E}_{\mathcal{G}_i}[U_{\mathfrak{n}}] = \mathbb{E}_{\mathcal{G}_i}[U_{\mathfrak{n}_i^{\infty,*}}],$$

with $\mathfrak{n}_i^{\infty,*}$ *given by* (2.13) *for* $K=\infty$*.*

Proof. See Appendix A.1. ■

Corollary 4 *Under assumption* (2.3)*, there exists an optimal stopping time for Problem* (2.2)*, which may be obtained by taking* $\mathfrak{t}^* := \mathfrak{t}_0^{K,*} := \tau_{\mathfrak{n}^*} := \tau_{\mathfrak{n}_0^{K,*}}$*.*

Proof. See Appendix A.1 ■

2.4 Duality

In this subsection, we present a dual representation of the discrete-time optimal stopping problem in our present setting, which essentially goes back to [50] and [26] (see also the early [15]). We first provide a result for the case $K = \infty$ in Proposition 1, which makes use of a uniform integrability condition that arises naturally for stopping problems on random times. The case $K < \infty$ follows analogously, and moreover does not require this extra uniform integrability argument; see Corollary 7.

The next proposition tells us that (2.5) may be reformulated as a penalized optimization problem with full foresight, in the terminology of perfect information relaxation.

Proposition 5 *Let $\mathcal{M}_0^{\mathsf{UI}}$ be the collection of uniformly integrable $\mathbb{G}$-martingales $(M_i)_{i\in\mathbb{N}_0}$ with $M_0 = 0$. Assume* (2.3) *and*

$$\mathbb{E}\left[N_T \sup_{i\in\mathbb{N}_0} U_i\right] < \infty\,. \tag{2.15}$$

We then have the following duality results.
(i) Weak duality: It holds that

$$Y_0 \equiv Y_0^{(\infty)} = \inf_{M\in\mathcal{M}_0^{\mathsf{UI}}} \mathbb{E}\left[\sup_{i\geq 0}(U_i - M_i)\right]\,. \tag{2.16}$$

(ii) Strong duality: Let us consider, for $Y_i \equiv Y_i^{(\infty)}$ given in (2.10)*, the $\mathcal{G}$-martingale, a.k.a. the Doob martingale of the Snell envelope,*

$$M_i^\circ := \sum_{j=1}^{i} Y_j - \mathbb{E}_{\mathcal{G}_{j-1}}[Y_j]\,,\quad i\geq 1\,,\quad M_0^\circ = 0\,. \tag{2.17}$$

Then it holds that $M^\circ \in \mathcal{M}_0^{\mathsf{UI}}$ and

$$Y_0 = \sup_{i\geq 0}(U_i - M_i^\circ)\,,\quad \text{a.s.}$$

Proof. See Appendix A.2. ■

Remark 6 *Proposition 5-(ii) has a beneficial consequence in numerical applications: If the Doob martingale M° is approximated 'close enough' by some martingale M, then the variance of the random variable $\sup_{i\geq 0}(U_i - M_i)$ in the dual representation* (2.16) *will be 'close' to zero. The quantification of this phenomenon is omitted, as it is similar to the finite horizon case (see, for example, [6, 36]).*

Corollary 7 *If, for some $\alpha > 1$,* (2.3) *holds, a sufficient condition for* (2.15) *is that the survival distribution of N_T is sufficiently regular in the following sense: $\sum_{j=1}^{\infty} j^{\frac{1}{\alpha-1}}\mathbb{P}(N_T > j) < \infty$. This holds in particular if one of the following holds:*

(i) $\mathbb{P}(N_T > j) \lesssim j^{-\gamma}$ *for some* $\gamma > \alpha/(\alpha-1)$ *and* $j \geq j_0$.

(ii) There is some $L \in \mathbb{N}$ *such that* $\mathbb{P}(\tau_L > T) = 1$*; condition* (2.15) *is then trivially fulfilled. However, we are then dealing with Problem* (2.2) *for* $K = L$ *in fact.*

Furthermore, if the following condition:

$$0 \leq \sup_{i \in \mathbb{N}_0} U_i \leq B < \infty \text{ a.s.}$$

is satisfied, condition (2.15) *holds if* $\mathbb{E}[N_T] < \infty$*. (In fact, one then has* (A.14) *for* $\alpha = \infty$*.)*

Similarly, if $K < \infty$*, it suffices for Proposition 5 that* (2.3) *holds with* $\alpha = 1$*, as a finite family of integrable martingales is uniformly integrable.*

Proof. Follows from Corollary 22 and Lemma 23 in Appendix A.2. ■

3 Optimal Stopping via Policy Improvement

By passing to discrete time in Section 2.2, our problem becomes an integer-time, infinite-horizon problem. In this section, we introduce a method based on forward induction to solve (2.5) that extends the iterative policy improvement procedure in [33]; see also [7]. As the results in [33] assume a finite horizon with a deterministic, fixed and finite number of opportunities to stop, and all proofs exploit the backward recursive nature with initialization at the last opportunity to stop, we extend their Theorem 3.1 and Propositions 4.1, 4.3 and 4.4 to the present infinite-horizon setting with essentially different proofs.

Let us consider a generic stopping family σ, being a family of stopping times $\sigma := (\sigma_i)_{i \in \mathbb{N}_0}$, with respect to the discrete filtration $\mathbb{G}$ that satisfies the consistency conditions

$$i \leq \sigma_i < \infty \text{ and } \sigma_i > i \Longrightarrow \sigma_i = \sigma_{i+1}\,, \quad i \in \mathbb{N}_0\,. \tag{3.1}$$

For the stopping family σ, we define the process

$$Y_i^{\sigma} := \mathbb{E}_{\mathcal{G}_i}[U_{\sigma_i}]\,. \tag{3.2}$$

Our aim is to construct, from the consistent stopping policy σ in (3.1), an improved policy. Let us fix a window parameter $\kappa \in \mathbb{N} \cup \{\infty\}$ and introduce a new stopping family $\widehat{\sigma} := (\widehat{\sigma}_i)_{i \in \mathbb{N}_0}$ by

$$\widehat{\sigma}_i := \inf\left\{ j \geq i :\ U_j \geq \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j}[U_{\sigma_k}] \right\}. \tag{3.3}$$

Clearly, the family $(\widehat{\sigma}_i)$ also satisfies (3.1). Moreover, $(\widehat{\sigma}_i)$ improves over (σ_i) in the following sense:

Proposition 8 *Let* (σ_i) *satisfy* (3.1) *and* $(\widehat{\sigma}_i)$ *be given by* (3.3)*. Then it holds that*

$$Y_i^{\sigma} \leq \widetilde{Y}_i^{\sigma} \leq Y_i^{\widehat{\sigma}} \leq Y_i\,,$$

with

$$\widetilde{Y}_j^{\sigma} := \max_{j \le k \le j+\kappa} \mathbb{E}_{\mathcal{G}_j}\left[U_{\sigma_k}\right] .$$

Proof. See Appendix A.3 for details. The proof first applies a result from [33] to a suitably truncated reward, and then considers and analyzes a limiting procedure. ■

The role of κ is to provide the window over which we search for possible improvements, where Part (ii) of Corollary 9 indicates that a larger window allows for more improvement:

Corollary 9 *(i) It holds that $Y_i^{\widehat{\sigma}} \ge U_i$ for all $i \in \mathbb{N}_0$.*
(ii) Let $\widehat{\sigma}_i$ and $\widehat{\sigma}'_i$ be improvements of σ_i due to Proposition 8 that correspond to κ' and κ, respectively. Then, if $\kappa' > \kappa$, one has $Y_i^{\widehat{\sigma}'} \ge Y_i^{\widehat{\sigma}}$.

Proof. See Appendix A.3. ■

To solve (2.5), we may construct a sequence of stopping families starting from some initial stopping family $\sigma^{(0)}$ that satisfies (3.1). A myopic choice is $\sigma_i^{(0)} = i$, $i \in \mathbb{N}_0$, for example. The construction is done as follows: if $\sigma^{(m)}$, $m \in \mathbb{N}_0$, has already been constructed, then set

$$\sigma_i^{(m+1)} := \inf\left\{ j \ge i : \; U_j \ge \max_{j \le k \le j+\kappa} \mathbb{E}_{\mathcal{G}_j}\left[U_{\sigma_k^{(m)}}\right] \right\} . \tag{3.4}$$

By Proposition 8, each stopping family improves over the previous one, and Theorem 10 below establishes that the sequence converges to an optimal stopping family.

Theorem 10 *Let (σ_i^*) denote the family of (first) optimal stopping times. That is,*

$$\sigma_i^* := \inf\left\{ j \ge i : \; U_j \ge \mathbb{E}_{\mathcal{G}_j}\left[Y_{j+1}\right] \right\} .$$

For the sequence $\left(\sigma^{(m)}\right)_{m \in \mathbb{N}_0}$ defined by iterating (3.4), we have

$$\sigma_i^{(m)} \le \sigma_i^{(m+1)} \le \sigma_i^* < \infty , \quad m \in \mathbb{N}_0 , \tag{3.5}$$

and moreover it holds that

$$\sigma_i^* = \uparrow \lim_{m \to \infty} \sigma_i^{(m)} \text{ and } Y_i = \uparrow \lim_{m \to \infty} Y_i^{(m)}, \tag{3.6}$$

with $Y_i^{(m)} \equiv Y_i^{\sigma^{(m)}}$. The up-arrows indicate that the respective sequences are non-decreasing.

Proof. See Appendix A.3 for details. The proof covers the infinite-horizon case and therefore differs essentially from the corresponding one in [33]. The latter proof was based on backward induction with trivial initialization at the finite horizon. ■

4 Simulation-Based Primal and Dual Methods

In this section, we introduce three simulation-based methods to solve, theoretically and pseudo-algorithmically, problems (2.5) (primal) and (2.16) (dual) in a Markovian environment specified by

the following assumption.

Assumption 11 *There exists a càdlàg process X in $\mathbb{R}^d$ and a measurable map $(t,x) \mapsto Z(t,x)$ such that $Z_t = Z(t, X_t)$. Furthermore, there exists an auxiliary càdlàg process Θ in $\mathbb{R}^q$ such that the process (X, Θ) in $\mathbb{R}^d \times \mathbb{R}^q$ generates the (augmented) filtration $\mathbb{F}$, and the $\mathbb{G}$-adapted discrete-time process*

$$(\tau_i, X_{\tau_i}, \Theta_{\tau_i})_{i \in \mathbb{N}_0}\,, \quad (\tau_0, X_{\tau_0}, \Theta_{\tau_0}) = (0, X_0, \Theta_0)\,,$$

is a discrete-time Markov chain in the state space $\mathbb{R}_{\geq 0} \times \mathbb{R}^d \times \mathbb{R}^q$.

Assumption 11 may be seen as a kind of a strong Markov property, and its interpretation is that the joint process (X, Θ) determines the sequence of stopping times (τ_i) and that the reward process is given as $Z_{\tau_i} = Z(\tau_i, X_{\tau_i})$. In Appendix A.5, we establish that this assumption suffices to express conditional expectations as Borel functions of the Markov chain and provide some specific examples of possible structures that (X, Θ) may have.

In Section 4.1, we propose two backward dynamic programming approaches to problem (2.5). In particular, we propose a least-squares Monte Carlo method in the spirit of [41] and a threshold optimization method in the spirit of [4], for the case $K < \infty$. For the case $K = \infty$, we first establish a finite-horizon approximation to (2.5), for an arbitrary, prespecified level of accuracy, and then proceed with the finite-horizon problem. Next, in Section 4.2 we develop a simulation algorithm based on the policy improvement method presented in Section 3 for problem (2.5). Finally, for problem (2.16), we extend the dual method of Andersen-Broadie [5] to the infinite-horizon setting in Section 4.3, by exploiting that each path has finite length due to (2.8).

4.1 Backward dynamic programming, finite-horizon approximation

For numerically solving problem (2.5) under Assumption 11 in the case $K < \infty$, we may take advantage of backward dynamic programming methods, such as the least-squares Monte Carlo algorithm in [41], or some stochastic optimization approach, such as the threshold-based approach in [4]. In this context, one carries out a backward recursion over the stopping opportunities, trivially initialized at the last opportunity K. If $K = \infty$, there is no such natural terminal condition, however. We may then approximate the infinite-horizon problem by a finite-horizon problem, and next proceed with the latter one. The following lemma shows that this can be done with arbitrary accuracy.

Lemma 12 *Assume* (2.3). *For any fixed tolerance level $\varepsilon > 0$, there exists a finite K such that*

$$Y_0^{(K)} \leq Y_0^{(\infty)} \leq Y_0^{(K)} + \varepsilon\,.$$

Proof. See Appendix A.4. ■

Thus, for any pre-specified tolerance level ε, $Y_0^{(K)}$ yields a lower bound that is "ε-close" to $Y_0^{(\infty)}$. A suitable candidate for such a K, given a fixed tolerance level $\varepsilon > 0$, is the smallest K'

such that

$$\mathbb{E}\left[1_{\{\tau_{K'}\leq T\}}\sup_{i\in\mathbb{N}_0,\ i>K'}Z_{\tau_i}\right]<\varepsilon\,, \tag{4.1}$$

which may be estimated through simulation.

4.1.1 Random times least-squares Monte Carlo, pseudo-algorithm

Let us suppose we are in a Markovian environment $(X_t,\Theta_t)_{t\geq 0}$, under Assumption 11, with a reward of the form $Z_t := Z(t,X_t)$. Then, $Y_0^{(K)}$, with $K<\infty$, may be numerically approximated by regression, based on a given set of measurable basis functions

$$\psi_l : \mathbb{R}_{\geq 0}\times\mathbb{R}^d\times\mathbb{R}^q\to\mathbb{R},\quad l\in\mathbb{N}.$$

We propose the following pseudo-algorithm for $n=1,\dots,N$:

(i) Simulate $\left(X^{(n)}_{\tau_1^{(n)}},\Theta^{(n)}_{\tau_1^{(n)}}\right),\dots,\left(X^{(n)}_{\tau_K^{(n)}},\Theta^{(n)}_{\tau_K^{(n)}}\right)$.

(ii) Initialize $\widehat{Y}_K^{K,n}=U_K^{(n)}=Z(\tau_K^{(n)},X^{(n)}_{\tau_K^{(n)}})$, $\widehat{\mathfrak{n}}_K^{(n)}=K$.

(iii) If, for $0<k\leq K$, the values $\widehat{Y}_k^{K,n}$ and $\widehat{\mathfrak{n}}_k^{(n)}$ have been constructed, then solve the regression problem

$$(\widehat{\alpha}_{k-1,l})_{l=1,\dots,L} := \arg\min_{\alpha\in\mathbb{R}^L}\sum_{n=1}^{N}\left(\widehat{Y}_k^{K,n}-\sum_{l=1}^{L}\alpha_l\psi_l\left(\tau_{k-1}^{(n)},X^{(n)}_{\tau_{k-1}^{(n)}},\Theta^{(n)}_{\tau_{k-1}^{(n)}}\right)\right)^2,$$

and define

$$\widehat{C}_{k-1}(t,x,\theta) := \sum_{l=1}^{L}\widehat{\alpha}_{k-1,l}\psi_l\left(t,x,\theta\right).$$

Next, set

$$\widehat{\mathfrak{n}}_{k-1}^{(n)} = \begin{cases} k-1 & \text{if}\quad Z(\tau_{k-1}^{(n)},X^{(n)}_{\tau_{k-1}^{(n)}})>\widehat{C}_{k-1}(\tau_{k-1}^{(n)},X^{(n)}_{\tau_{k-1}^{(n)}},\Theta^{(n)}_{\tau_{k-1}^{(n)}}) \\ \widehat{\mathfrak{n}}_k^{(n)} & \text{else.}\end{cases}$$

Having constructed the functions $\widehat{C}_0(t,x,\theta),\dots,\widehat{C}_{K-1}(t,x,\theta)$ (with $\widehat{C}_K\equiv 0$), one thus has constructed a stopping policy

$$\sigma_k := \inf\left\{k\leq j\leq K : Z(\tau_j,X_{\tau_j})\geq\widehat{C}_j(\tau_j,X_{\tau_j},\Theta_{\tau_j})\right\},\quad 0\leq k\leq K, \tag{4.2}$$

and we may generate, by a new simulation, i.i.d. copies

$$\left(\widetilde{\tau}_1^{(n)},\widetilde{X}^{(n)}_{\widetilde{\tau}_1^{(n)}},\widetilde{\Theta}^{(n)}_{\widetilde{\tau}_1^{(n)}}\right),\dots,\left(\widetilde{\tau}_K^{(n)},\widetilde{X}^{(n)}_{\widetilde{\tau}_K^{(n)}},\widetilde{\Theta}^{(n)}_{\widetilde{\tau}_K^{(n)}}\right),\quad\text{for}\quad n=1,\dots,\widetilde{N},$$

of the Markov chain $(\tau_1, X_{\tau_1}, \Theta_{\tau_1}), \ldots, (\tau_K, X_{\tau_K}, \Theta_{\tau_K})$, in order to obtain a lower-biased estimate $\widetilde{Y}_0^{(K)}$ for $Y_0^{(K)}$ by setting

$$\widetilde{\mathfrak{n}}^{(n)} := \inf\left\{0 \le j \le K : Z(\widetilde{\tau}_j^{(n)}, \widetilde{X}^{(n)}_{\widetilde{\tau}_j^{(n)}}) \ge \widehat{C}_j(\widetilde{\tau}_j^{(n)}, \widetilde{X}^{(n)}_{\widetilde{\tau}_j^{(n)}}, \widetilde{\Theta}^{(n)}_{\widetilde{\tau}_j^{(n)}})\right\},$$

and defining

$$\widetilde{Y}_0^{(K)} := \frac{1}{\widetilde{N}} \sum_{n=1}^{\widetilde{N}} Z(\widetilde{\tau}^{(n)}_{\widetilde{\mathfrak{n}}^{(n)}}, \widetilde{X}^{(n)}_{\widetilde{\tau}^{(n)}_{\widetilde{\mathfrak{n}}^{(n)}}}).$$

Theorems 29 and 30 in Appendix A.6 establish convergence to the true value of the optimal stopping problem and the optimal stopping rule, respectively. The policy (4.2) can and will be used in Section 4.3 as input policy for the well-known Andersen-Broadie algorithm [5] (see also [24]) in order to obtain a dual estimate of an upper bound on $Y_0^{(K)}$ (in mean).

Remark 13 *Over the past decades, various refinements have been proposed for the implementation of the standard, deterministic times LSMC algorithm. We have included in our random times LSMC implementation the (counterparts of the) most obvious and usual refinements: (i) the cash-flow itself is included in the regression; and (ii) the regression is only carried out at in-the-money trajectories.*

4.1.2 Threshold optimization, pseudo-algorithm

The threshold optimization method proposed by Andersen [4] can be seen as a special case of stochastic optimization of the stopping region for problem (2.5) with $K < \infty$. Formally and adapted to our setting, it aims at constructing an optimal "threshold policy" (σ_k) due to thresholds $h_0, \ldots, h_K \ge 0$ of the form

$$\sigma_k = \inf\left\{k \le j \le K : Z_{\tau_j} \ge h_j\right\}, \quad k = 0, \ldots, K, \tag{4.3}$$

which is initialized by $h_K = 0$ and $\sigma_K = K$. Then, given $h_{k+1}, \ldots, h_K$ and σ_{k+1}, one defines, recursively,

$$\begin{aligned} h_k &:= \arg\max_{h \ge 0} \mathbb{E}\left[Z_{\tau_k} 1_{\{Z_{\tau_k} \ge h\}} + Z_{\tau_{\sigma_{k+1}}} 1_{\{Z_{\tau_k} < h\}}\right], \quad \text{and} \\ \sigma_k &:= k 1_{\{Z_{\tau_k} \ge h_k\}} + \sigma_{k+1} 1_{\{Z_{\tau_k} < h_k\}}. \end{aligned}$$

The Monte Carlo analogue of this method due to sample trajectories $n = 1, \ldots, N$ is obtained as follows. Initialize $h_{K,N} = 0$ and $\sigma_{K,n} = K$, $n = 1, \ldots, N$. Then, given $h_{k+1,N}, \ldots, h_{K,N}$ and $\sigma_{k+1,n} \in \left\{k+1, \ldots, \bar{K}\right\}$, $n = 1, \ldots, N$, set

$$\begin{aligned} h_{k,N} &:= \arg\max_{h \ge 0} \frac{1}{N} \sum_{n=1}^{N} \left(Z^{(n)}_{\tau_k} 1_{\{Z^{(n)}_{\tau_k} \ge h\}} + Z^{(n)}_{\tau_{\sigma_{k+1,n}}} 1_{\{Z^{(n)}_{\tau_k} < h\}}\right), \quad \text{and} \\ \sigma_{k,n} &:= k 1_{\{Z^{(n)}_{\tau_k} \ge h_{k,N}\}} + \sigma_{k+1,n} 1_{\{Z^{(n)}_{\tau_k} < h_{k,N}\}}, \quad \text{for} \quad n = 1, \ldots, N. \end{aligned}$$

It should be noted that in general the optimal threshold policy may be far away from a real optimal policy, but serves in many practical examples as a "not too bad" approximation of it.

4.2 Policy improvement

Proposition 8 and Theorem 10 suggest a natural algorithm: starting from some initial policy σ that satisfies (3.1), apply (3.3) with some window $\kappa \geq 1$, where the conditional expectations are replaced by Monte Carlo estimates. The initial policy may be obtained, for example, by the least-squares approach of Section 4.1.1 or as a threshold policy constructed in Section 4.1.2. Depending on the underlying structure, other constructions are also possible. As the most trivial initial policy, one could take the myopic policy $\sigma_i = i$. In the next subsection, we present a corresponding pseudo-algorithm.

4.2.1 Pseudo-algorithm

Under Assumption 11, the conditional expectations in (3.3) take the form

$$\mathbb{E}_{\mathcal{G}_j}\left[U_{\sigma_k}\right] = \mathbb{E}_{\left(\tau_j, X_{\tau_j}, \Theta_{\tau_j}\right)}\left[Z(\tau_{\sigma_k}, X_{\tau_{\sigma_k}})\right], \quad j \leq k \leq j+\kappa,$$

and can be estimated in the following way. Construct for $n = 1, \ldots, N$ the "outer" trajectories $(X_t^{(n)}, \Theta_t^{(n)})_{t\geq 0}$, and for trajectory number n the sequence

$$\left(\tau_j, X_{\tau_j}, \Theta_{\tau_j}\right)^{(n)} \equiv \left(\tau_j^{(n)}, X^{(n)}_{\tau_j^{(n)}}, \Theta^{(n)}_{\tau_j^{(n)}}\right), \qquad j \geq 0.$$

Then simulate on the outer trajectory with number n, for each $m = 1, \ldots, M$ and $j \geq 0$, the sub-trajectory

$$\left(X_s^{n,m,j}, \Theta_s^{n,m,j}\right)_{s\geq \tau_j^{(n)}} \quad \text{with} \quad \left(X^{n,m,j}_{\tau_j^{(n)}}, \Theta^{n,m,j}_{\tau_j^{(n)}}\right) = \left(X^{(n)}_{\tau_j^{(n)}}, \Theta^{(n)}_{\tau_j^{(n)}}\right),$$

along with a strictly increasing sequence of $\mathbb{R}_{\geq 0}$-valued stopping times $\left(\tau_k^{n,m,j}\right)_{k\geq j}$ with $\tau_j^{n,m,j} = \tau_j^{(n)}$. Next, construct on this sequence the integer-valued stopping times (or indices) $\sigma_k^{n,m,j}$, $j \leq k \leq j+\kappa$, according to the input stopping family (3.1) and set

$$E_k^{n,j} \equiv \frac{1}{M}\sum_{m=1}^{M} Z(\tau_{\sigma_k^{n,m,j}}, X^{n,m,j}_{\tau_{\sigma_k^{n,m,j}}}) \approx \mathbb{E}_{\left(\tau_j, X_{\tau_j}, \Theta_{\tau_j}\right)^{(n)}}\left[Z(\tau_{\sigma_k}, X_{\tau_{\sigma_k}})\right], \quad j \leq k \leq j+\kappa. \tag{4.4}$$

The improved policy (3.3) is now path-wise estimated by

$$\widehat{\sigma}_0^{(n)} := \inf\left\{j \geq 0 : Z(\tau_j^{(n)}, X^{n,j}) \geq \max_{j\leq k\leq j+\kappa} E_k^{n,j}\right\},$$

and we finally obtain the estimate

$$Y_0^\sigma \leq Y_0^{\widehat{\sigma}} = \mathbb{E}\left[U_{\widehat{\sigma}_0}\right] \approx \frac{1}{N}\sum_{n=1}^{N} Z(\widehat{\sigma}_0^{(n)}, X^{(n)}_{\widehat{\sigma}_0^{(n)}}),$$

which is naturally a lower-biased estimate in the sense that $Y_0^{\widehat{\sigma}} \leq Y_0$.

Remark 14 *In the above pseudo-algorithm, the number of sub-simulations is taken to be equal for any $j \geq 0$, for simplicity. Of course, the algorithm may be trivially refined in this respect, to further improve its efficiency.*

Remark 15 *In the implementation of the policy improvement algorithm, we have used the estimate due to the initial policy employing a form of variance reduction as proposed in Remark 6.2 in [33].*

4.3 Dual estimation

To obtain an upper-biased estimate of Y_0 from the dual representation (2.16), we proceed by adapting the algorithm from [5] to an infinite-horizon setting. For a given stopping family $\sigma = (\sigma_i)_{i\in\mathbb{N}_0}$ satisfying (3.1), we write the Doob martingale corresponding to (3.2) as

$$M_i^\sigma := \sum_{j=1}^{i} Y_j^\sigma - \mathbb{E}_{\mathcal{G}_{j-1}}\left[Y_j^\sigma\right], \quad i \geq 1, \quad M_0^\sigma = 0, \tag{4.5}$$

which satisfies the following recursion:

$$M_j^\sigma - M_{j-1}^\sigma = Y_j^\sigma - \mathbb{E}_{\mathcal{G}_{j-1}}\left[Y_j^\sigma\right] = Y_j^\sigma - Y_{j-1}^\sigma - 1_{\{\sigma_{j-1}=j-1\}}\left(\mathbb{E}_{\mathcal{G}_{j-1}}\left[Y_j^\sigma\right] - Y_{j-1}^\sigma\right), \tag{4.6}$$

by the consistency of the family σ. We thus obtain M_i^σ by summing (4.6) from $j = 1$ to $j = i$, and then define

$$\begin{aligned}\delta_i^\sigma &:= U_i - M_i^\sigma - Y_0^\sigma = U_i - Y_i^\sigma + \sum_{j=1}^{i} 1_{\{\sigma_{j-1}=j-1\}}\left(\mathbb{E}_{\mathcal{G}_{j-1}}\left[Y_j^\sigma\right] - Y_{j-1}^\sigma\right) \\ &= 1_{\{\sigma_i>i\}}\left(U_i - \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_{i+1}}\right]\right) + \sum_{j=0}^{i-1} 1_{\{\sigma_j=j\}}\left(\mathbb{E}_{\mathcal{G}_j}\left[U_{\sigma_{j+1}}\right] - U_j\right).\end{aligned} \tag{4.7}$$

Now, a dual upper bound due to the policy σ is obtained by Proposition 5(i), as

$$\begin{aligned}Y_0^{\sigma,up} &:= \mathbb{E}\left[\sup_{i\geq 0}(U_i - M_i^\sigma)\right] = Y_0^\sigma + \mathbb{E}\left[\sup_{i\geq 0}\delta_i^\sigma\right] \\ &=: Y_0^\sigma + \Delta_0^\sigma,\end{aligned} \tag{4.8}$$

provided that M^σ is uniformly integrable. Note that

$$\delta^\sigma_{i+1} - \delta^\sigma_i = 1_{\{\sigma_{i+1}>i+1\}} \left(U_{i+1} - \mathbb{E}_{\mathcal{G}_{i+1}}\left[U_{\sigma_{i+2}}\right]\right) + \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_{i+1}}\right] - U_i, \quad i \geq 0, \tag{4.9}$$

which allows for a path-wise upward recursive computation of the δ^σ_i in the pseudo-algorithm presented below.

4.3.1 Pseudo-algorithm

Due to Assumption 11, we may estimate (4.8) by using nested Monte Carlo in the spirit of Section 4.2.1. In view of the monotonicity of the (conditional) expectation, it is easily verified that, upon replacing the expectation and conditional expectations in (4.8), and hence in (4.7), by nested Monte Carlo simulation due to outer trajectories and sub-trajectories, respectively, one obtains a "biased high" estimator. That is, one obtains an estimator with expectation larger than $Y_0^{\sigma,up}$. In pseudo-algorithmic terms, we proceed in the following way.

- For $n = 1$ to N:
  - Simulate trajectory number n with rewards $(U_i^{(n)})_{i=0,\ldots,N_T^{(n)}}$, where $U_i^{(n)} = Z(\tau_i^{(n)}, X^{(n)}_{\tau_i^{(n)}})$ and N_T is path-wise defined in (2.7);
  - Following the construction of (4.4) in the pseudo-algorithm of Section 4.2.1, estimate, for all i, $0 \leq i \leq N_T$, the conditional expectations in (4.7) by

    $$E^{n,i}_{i+1} \approx \mathbb{E}_{\mathcal{G}_i}[U^{(n)}_{\sigma^{(n)}_{i+1}}],$$

    due to M sub-simulations;
  - Construct initially $\widehat{\delta}^{\sigma,n}_0 := 1_{\left\{\sigma^{(n)}_0>0\right\}} \left(U_0 - E^{n,0}_1\right) \approx \delta^{\sigma,n}_0$. Next, after having constructed $\widehat{\delta}^{\sigma,n}_i$ for $0 \leq i < N_T^{(n)}$, construct

    $$\widehat{\delta}^{\sigma,n}_{i+1} := \widehat{\delta}^{\sigma,n}_i + 1_{\left\{\sigma^{(n)}_{i+1}>i+1\right\}} \left(U^{(n)}_{i+1} - E^{n,i+1}_{i+2}\right) + E^{n,i}_{i+1} - U^{(n)}_i \approx \delta^{\sigma,n}_{i+1}$$

    in view of (4.9);
- Estimate the duality gap in (4.8) by

$$\widehat{\Delta}^\sigma_0 := \frac{1}{N} \sum_{n=1}^{N} \max_{0\leq i\leq N_T^{(n)}} \widehat{\delta}^{\sigma,n}_i \,. \tag{4.10}$$

We now obtain as an upper-biased estimate for the dual upper bound,

$$\widehat{Y}_0^{\sigma,up} = \widehat{Y}_0^\sigma + \widehat{\Delta}^\sigma_0 \,, \tag{4.11}$$

where the lower-biased estimate $\widehat{Y}_0^\sigma$ of Y_0^σ may be obtained by an unbiased non-nested Monte Carlo estimation due to a sufficiently large number of trajectories. We finally note that, concerning the number of sub-simulations involved, a remark similar to Remark 14 also applies here.

5 Numerical Illustrations

To illustrate the performance and general applicability of the numerical methods introduced in Section 4, we begin by considering an example that admits an analytical solution in Section 5.1. Next, we extend in Section 5.2 the max-call example from [5] to a random exercise opportunities setting, allowing for various dependence structures between the process generating the stopping opportunities and the process generating the reward.

5.1 Stopped geometric Brownian motion with jumps

To analyze the performance of our algorithms for a problem admitting an explicit, analytical solution, which can serve as a benchmark, we consider a geometric Brownian motion with jumps:

$$X_t = X_0 e^{\left(\mu - \frac{1}{2}\sigma^2\right)t + \sigma W_t + N_t \ln(1+j)}\,, \tag{5.1}$$

where $X_0 > 0,\ \mu \in \mathbb{R},\ \sigma > 0$ and $j > -1$ are known constants. Here, W is a standard Brownian motion and N is a homogeneous Poisson process with intensity $\lambda > 0$, assumed to be mutually independent. We consider the problem of stopping the reward

$$Z_t = X_t^\eta 1_{\{t \le T\}}, \tag{5.2}$$

for a given power preference parameter $\eta > 0$ and time horizon $T > 0$, on $\{\tau_i, \ldots, \tau_{i+K}\}$ with $i \in \mathbb{N}_0$, the set of τ_i and the first $K \ge 1$ jump times of N that occur strictly after i. We wish to determine, for $i \in \mathbb{N}_0$,

$$Y_i^{(K)} := \operatorname*{ess\,sup}_{i \le n \le i+K} \mathbb{E}_{\mathcal{G}_i}\left[Z_{\tau_n}\right], \qquad C_i^{(K)} := \operatorname*{ess\,sup}_{i+1 \le n \le i+K} \mathbb{E}_{\mathcal{G}_i}\left[Z_{\tau_n}\right],$$

with $\tau_0 := 0$. These definitions are similar to the ones in (2.10)–(2.11) but for ease of exposition we specify in this subsection K remaining stopping opportunities instead of K opportunities in total. Letting

$$\xi := \mu\eta + \frac{1}{2}\sigma^2\eta(\eta-1) - \lambda\,, \qquad \psi := \lambda(1+j)^\eta\,, \tag{5.3}$$

a straightforward application of Itô's formula shows that if $\xi + \psi \le 0$, $(Z_{\tau_k})_{k\in\mathbb{N}}$ is a $\mathbb{G}$-super-martingale. Since it equals its own Snell envelope in this case, it is then always optimal to stop at the first opportunity. Note that we do not need to impose that Z is an $\mathbb{F}$-super-martingale at all times to conclude that we should stop as soon as possible, only that $(Z_{\tau_k})_{k\in\mathbb{N}}$ is a $\mathbb{G}$-super-

martingale. If that is the case, the following is immediate:

Lemma 16 *If $\xi + \psi \leq 0$, then $\xi < 0$ and, for any $K \in \mathbb{N} \cup \{\infty\}$, it is optimal to stop immediately, so $Y_i^{(K)} = Z_{\tau_i}$. Furthermore, we have that the continuation value $C_i^{(K)}$ equals the expected value of Z when stopped at the first opportunity, so*

$$C_i^{(K)} = \mathbb{E}_{\mathcal{G}_i}\left[Z_{\tau_{i+1}}\right] = \int_0^{T-\tau_i} \psi e^{\xi s} X_t^\eta \, \mathrm{d}s = X_t^\eta \frac{\psi}{\xi}\left(e^{\xi(T-\tau_i)} - 1\right).$$

By induction on K and a limiting argument, we may also find a solution for the case when $\xi > 0$:

Theorem 17 *Assume $\xi > 0$. For finite $K \geq 1$ and $K = \infty$, it is optimal to stop at time $\tau_i \leq T$ if and only if*

$$\tau_i \;\geq\; t^* := T - \frac{\ln(1 + \frac{\xi}{\psi})}{\xi}.$$

The corresponding optimal values are $Y_i^{(K)} = X_{\tau_i}^\eta \vee C_i^{(K)}$ with:

$$\begin{aligned}
C_i^{(1)} &= X_{\tau_i}^\eta \frac{\psi}{\xi}\left(e^{\xi(T-\tau_i)} - 1\right), \\
C_i^{(K)} &= 1_{\{\tau_i \geq t^*\}} C_i^{(1)} + 1_{\{\tau_i < t^*\}} X_{\tau_i}^\eta e^{\xi(t^* - \tau_i)} z_K(t^* - \tau_i), \quad 1 < K < \infty, \\
C_i^{(\infty)} &= 1_{\{\tau_i \geq t^*\}} C_i^{(1)} + 1_{\{\tau_i < t^*\}} X_{\tau_i}^\eta e^{(\xi+\psi)(t^* - \tau_i)},
\end{aligned}$$

where

$$z_K(s) = \frac{\Gamma_K(\psi s)}{(K-1)!} e^{\psi s} + \frac{\gamma_K(-\xi s)}{(K-1)!}\left(-\frac{\psi}{\xi}\right)^K e^{-\xi s},$$

with Γ_K and γ_K the upper and lower Gamma functions of order K, respectively.

Proof. See Appendix A.7. ■

In Figure 1, we visualize the scaled solutions of Theorem 17 for different characteristics of the arrival process inducing the exercise opportunities and different parameters of the geometric Brownian motion with jumps. We notice that for the parameter values used here, the effect of allowing for at most 1, 3 or an unlimited number of randomly arriving remaining exercise opportunities is much smaller than the effect of changing the size of the downward jumps or the expected number of exercise opportunities before maturity. The optimal stopping times in terms of the remaining time to maturity have been indicated by circles in Figure 1. They correspond to the point in time at which the scaled value function equals one. As expected, stopping tends to take place later (hence, $T - t^*$ is smaller) when the expected number of remaining exercise opportunities is larger and/or when future values of the geometric Brownian motion with jumps have a higher expected value.

The explicit form of the value function makes this optimal stopping problem a suitable benchmark to test our random times primal-dual approach. In Table 1, we display the corresponding estimates provided by the algorithms from Sections 4.1.1 and 4.3.1. The estimates from the random

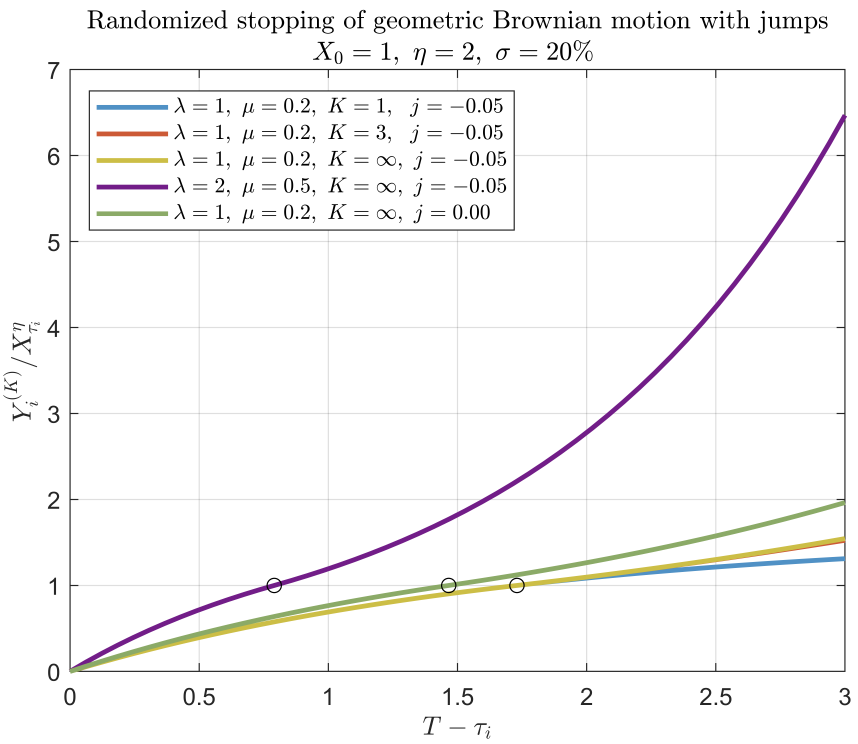

Figure 1: Randomized stopping of geometric Brownian motions with jumps.

Notes. The scaled value function $Y_i^{(K)}/X_{\tau_i}^{\eta}$ is displayed for different remaining times to maturity $T-\tau_i$. The different lines correspond to the cases shown in Table 1; see there for a more detailed description of the model and assumptions. The black circles indicate the remaining time to maturities $T - t^*$ where it becomes optimal to stop, i.e., the times where $Y_i^{(K)} = X_{\tau_i}^{\eta}$. We notice that the second and third cases (the red and yellow lines) are almost indistinguishable.

times LSMC and dual algorithms are close to the exact value functions from Theorem 17 across various specifications, and duality gaps are of the order 0.1%. As the least-squares method requires a finite horizon, we have used the finite-horizon approximation due to Lemma 12 for all problems with $K = \infty$ and for all methods, to ensure a fair comparison. Specifically, we have determined a suitable K' by simulating (4.1) with 100,000 simulated paths and $\varepsilon = 0.001$. All simulated paths have been generated in antithetic pairs.

5.2 Max-call option

In view of the good performance of our random times LSMC primal-dual algorithm in the benchmark example, we now analyze our algorithms in a more complex example for which analytic reference values are no longer available. To illustrate the flexibility of our set-up, we consider several extensions of the max-call example in [5]. The objective is to optimally exercise a max-call contract on d assets with maturity date T, for which the payoff if exercised at time t is given by

$$Z_t = e^{-rt}\left(\max_{j=1,\dots,d} X_t^j - \mathcal{K}\right)^+ 1_{\{t\le T\}}\,. \tag{5.4}$$

j	K	μ	λ	T	True	LSMC	(s.e.)	Dual	(s.e.)
-0.05	1	0.2	1	3	1.31125	1.31103	(0.00066)	1.31103	(0.00000)
-0.05	3	0.2	1	3	1.52503	1.52484	(0.00096)	1.52507	(0.00006)
-0.05	∞	0.2	1	3	1.54483	1.54443	(0.00099)	1.54498	(0.00012)
-0.05	∞	0.2	1	1	0.69104	0.69101	(0.00042)	0.69101	(0.00000)
-0.05	∞	0.5	2	3	6.46847	6.47332	(0.00482)	6.47560	(0.00041)
0	∞	0.2	1	3	1.96392	1.96316	(0.00135)	1.96399	(0.00024)

Table 1: Randomized stopping of geometric Brownian motion with jumps.
Notes. This table displays estimated values of $Y_0^{(K)}$ under homogeneous Poisson opportunities with rate λ. X_t follows (5.1) with $X_0 = 1$, $\sigma = 20\%$ and various specifications for the remaining model parameters, whereas Z_t follows (5.2) with $\eta = 2$. Column True contains the true values (from Theorem 17). Column LSMC contains the lower-biased estimates based on the random times LSMC algorithm of Section 4.1 with 200,000 and 2,000,000 independent paths for the backward and forward passes, respectively, along with its standard error. The continuation values are estimated by using the basis functions $(t, x) \mapsto t^i x^j$, $i + j \leq 3$. Column Dual contains the upper-biased estimates based on the dual algorithm in Section 4.3 using 1,500 paths and Doob martingales (4.5) based on the estimated LSMC policy, where conditional expectations have been estimated using 10,000 sub-simulations, along with its standard error. All simulated paths have been generated in antithetic pairs.

Here, the X_t^j are jump-diffusions, interpreted as underlying asset prices driving the value of potential investment projects, with dynamics governed by the following stochastic differential equations:

$$\mathrm{d}X_t^j = (r - \delta^j - \mu_J^j \lambda_J^j) X_t^j \,\mathrm{d}t + \sigma^j X_t^j \,\mathrm{d}W_t^j + \mu_J^j X_{t-}^j \,\mathrm{d}N_t^j \,, \quad j = 1, \ldots, d \,,$$

where the W_t^j are independent standard Brownian motions, the N_t^j are independent homogeneous Poisson processes with arrival rates $\lambda_J^j \geq 0$, with the W_t^j and the N_t^j assumed to be mutually independent, and $r,\ \delta^j,\ \mu_J^j \in (-1, \infty)$ and $\sigma^j > 0$ are deterministic constants.

5.2.1 Reward-independent Poissonian arrivals of stopping opportunities

We begin by considering the stopping problem (2.2) for the reward (5.4) with stopping opportunities generated according to a reward-independent homogeneous Poisson process with rate λ:

$$\tau_{i+1} = \tau_i + \eta_i \,, \quad \eta_i \sim \mathrm{Exp}(\lambda) \text{ i.i.d}, \quad \tau_0 = 0 \,. \tag{5.5}$$

Hence, this example can be seen as (a multi-dimensional, jump-diffusion extension of) the finite-horizon counterpart of the perpetual call considered e.g., in [18].

Table 2 displays value function estimates obtained by using the random times LSMC method of Section 4.1, the dual method of Section 4.3 and the policy iteration (PI) method of Section 4.2, respectively, for the case of $d = 5$ assets with $T = 3,\ r = 5\%,\ \delta^j = 10\%$, $\sigma^j = 20\%$ and $\lambda_J^j = \mu_J^j = 0$. Note that the duality gaps, i.e., the differences between random times LSMC and dual estimates, are typically (very) small (less than 0.4% of the option value), and are generally slightly increasing with the arrival rate λ of stopping opportunities. The fact that random times LSMC already performs very well, leaving little room for further improvement, motivates our use of a trivial initial policy for the policy iteration method in this table: if we were to consider policy iteration over random times

LSMC initial policies, standard errors would dominate the improvements of the policy iteration method. The trivial policy consists in stopping as soon as $Z_{\tau_i} > 0$ or $\tau_i \geq T$.

We consider the case with jumps ($\lambda_J^j = 1,\ \mu_J^j = 0.06$) in Table 3. For comparison, we also provide Tables 10–11 and 12–13 for the cases with $d = 1$ and $d = 2$ assets in Appendix B. The duality gaps and standard errors are (even) smaller there, which is to be expected when the dimension of the problem is smaller. The results in Tables 2 and 11 may be viewed as the random times counterpart of Table 2 in [5], which considers a Bermudan max-call option on the same underlying assets. More specifically, let us consider the case $d = 2, 5$ in Tables 2 and 11 and compare our results to those in Table 2 of [5]. They consider 9 equidistant exercise opportunities on the interval $[0, 3]$, whereas in our setting there are in expectation $\lambda T = \{3, 9, 15\}$ times at which the contract may be exercised. It is remarkable that, even for the case $\lambda = 5$, i.e., when there are 15 exercise opportunities in expectation, our estimated values are substantially lower than those in [5]: the difference is about 0.4\$ ($d = 2$) and 1.0\$ ($d = 5$), which amounts to 3% to 4% of the asset price. For $\lambda = 1, 3$, this difference is even (much) larger. This entails that there is a substantial premium on "market frictions".

To analyze these effects in more detail, we compute the probability of not exercising, $\mathbb{P}(\tau^* > T)$, the expected time to maturity at exercise, $\mathbb{E}[T - \tau^* | \tau^* \leq T]$, and the *opportunity risk* and *timing risk* mentioned in the Introduction, and formally defined as

$$R_O := \sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t} \in \mathcal{T}_K} \mathbb{E}\left[Z_{\mathfrak{t}}\right] - \sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t} \in \mathcal{T}_K \cup \{T\}} \mathbb{E}\left[Z_{\mathfrak{t}}\right],$$

$$R_T := \sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t} \in \mathcal{T}_K \cup \{T\}} \mathbb{E}\left[Z_{\mathfrak{t}}\right] - \sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t} \in t_1, \ldots, t_{\tilde{K}}} \mathbb{E}\left[Z_{\mathfrak{t}}\right], \quad \tilde{K} = \left\lceil \mathbb{E}\left[\sum_{\tau_i} 1_{\tau_i \leq T} \right] \right\rceil,$$

where the problems $\sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t} \in \mathcal{T}_K \cup \{T\}} \mathbb{E}\left[Z_{\mathfrak{t}}\right]$ and $\sup_{\mathbb{F}\text{-stopping time } \mathfrak{t},\ \mathfrak{t} \in t_1, \ldots, t_{\tilde{K}}} \mathbb{E}\left[Z_{\mathfrak{t}}\right]$ may be computed by the (deterministic-times) counterparts of the algorithms used. We pick $t_i = \frac{i}{\tilde{K}} T$, such that the last term in R_T corresponds to a Bermudan stopping problem with equidistant stopping opportunities.

Table 4 illustrates how the timing risk and the opportunity risk are sizeable for the max-call example, and vary with arrival rates and initial conditions. We note that for this problem, both opportunity and timing risk increase with λ, whereas the expected time to maturity at exercise and the probability of stopping after T decrease with λ.

We note also that our methods are not limited to dimension $d = 5$ and show in Table 5 that good results may be obtained for dimensions easily up to $d = 100$. As expected, when the dimension increases, the duality gaps between random times LSMC and Dual become slightly larger, now leaving some room for improvement by invoking the policy iteration method. Specifically, we find that a single step of policy iteration over the random times LSMC based policy yields results that are (very) close to the Dual estimates.

λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Trivial	(s.e.)	Trivial-PI	(s.e.)
1	90	11.67443	(0.01032)	11.68861	(0.00307)	9.16316	(0.00838)	11.65284	(0.03108)
	100	19.53414	(0.01284)	19.55874	(0.00417)	14.74861	(0.01012)	19.50832	(0.04418)
	110	28.64464	(0.01503)	28.69280	(0.00669)	24.45769	(0.01185)	28.59180	(0.04894)
3	90	15.01530	(0.01087)	15.06641	(0.01824)	7.28141	(0.00574)	14.73748	(0.04027)
	100	24.15326	(0.01308)	24.22701	(0.02013)	10.88004	(0.00670)	23.76937	(0.05106)
	110	34.44991	(0.01486)	34.55357	(0.02302)	21.12589	(0.00757)	33.96516	(0.05650)
5	90	15.76541	(0.01098)	15.81797	(0.00450)	5.87923	(0.00451)	15.23885	(0.03972)
	100	25.09820	(0.01317)	25.19821	(0.00827)	8.74741	(0.00528)	24.30068	(0.04866)
	110	35.56730	(0.01488)	35.70128	(0.00905)	19.02219	(0.00595)	34.61747	(0.05376)

Table 2: Value of a max-call option on random times ($d = 5$).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5) with $d = 5$ independent geometric Brownian motions (without jumps, $\mu_J^j = \lambda_J^j = 0$) and with varying arrival rates λ and initial asset prices X_0. Column LSMC contains the lower-biased estimates based on the random times LSMC algorithm of Section 4.1 and column Dual contains the upper-biased estimates (4.11) based on the dual algorithm in Section 4.3 and the estimated LSMC policy; standard errors for the dual estimates are the standard errors of the duality gap (4.10). Both algorithms use the same settings as in Table 1, except for the basis functions: as in [5], we use basis functions in the ordered asset prices $X_t^{(1)} \geq \ldots \geq X_t^{(d)}$. The continuation values in the LSMC algorithm are estimated by using the basis functions $(t, x^{(1)}, x^{(2)}) \mapsto t^i \left(x^{(1)}\right)^j \left(x^{(2)}\right)^k$, $i + j + k \leq 5$. Column Trivial corresponds to stopping as soon as $Z_{\tau_i} > 0$ or $\tau_i \geq T$ and column Trivial-PI contains the results of the policy iteration method (3.4) with 100,000 outer paths, 500 inner paths and a window length $\kappa = K'$. All simulated paths have been generated in antithetic pairs.

λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Trivial	(s.e.)	Trivial-PI	(s.e.)
1	90	12.76951	(0.01119)	12.78354	(0.00372)	9.97336	(0.00903)	12.71436	(0.03361)
	100	20.79173	(0.01380)	20.81643	(0.00473)	15.64744	(0.01082)	20.76846	(0.04722)
	110	29.99001	(0.01605)	30.03046	(0.00536)	25.40365	(0.01261)	29.96612	(0.05241)
3	90	16.37745	(0.01179)	16.41311	(0.00401)	7.85547	(0.00616)	16.11181	(0.04371)
	100	25.70427	(0.01408)	25.76604	(0.00716)	11.50560	(0.00716)	25.26734	(0.05468)
	110	36.12987	(0.01597)	36.21073	(0.00697)	21.76685	(0.00808)	35.60821	(0.06066)
5	90	17.18631	(0.01192)	17.36115	(0.09173)	6.33664	(0.00484)	16.69437	(0.04318)
	100	26.71960	(0.01418)	26.97080	(0.10666)	9.24555	(0.00564)	25.97468	(0.05221)
	110	37.32898	(0.01603)	37.63602	(0.12062)	19.52885	(0.00636)	36.37071	(0.05770)

Table 3: Value of a max-call option on random times with jumps ($d = 5$).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5) with $d = 5$ independent geometric Brownian motions with jumps, $\mu_J^j = 0.06$ and $\lambda_J^j = 1$ and with varying arrival rates λ and initial asset prices X_0. The respective algorithms are carried out as described in the caption of Table 2.

5.2.2 Reward-dependent Poissonian arrivals of stopping opportunities

We now introduce a form of dependence between the stopping opportunities τ_i and the process X_t driving the reward. Specifically, we consider the payoff (5.4) with *potential* stopping opportunities generated according to a Poisson process with rate λ. The potential opportunities to stop constitute only *genuine* opportunities to stop when the lowest asset price among d asset prices, $X_t^{(d)}$, lies above

λ	X_0	$\mathbb{E}[T-\tau^*\|\tau^*\le T]$	(s.e.)	$\mathbb{P}(\tau^*>T)$	(s.e.)	R_O	R_T
1	90	1.18156	(0.00062)	0.43962	(0.00017)	-4.03581	-0.28296
	100	1.33435	(0.00056)	0.27743	(0.00014)	-5.25015	-0.46380
	110	1.45362	(0.00053)	0.18089	(0.00010)	-6.35926	-0.63461
3	90	0.99291	(0.00053)	0.34714	(0.00016)	-1.34970	-0.25629
	100	1.11190	(0.00049)	0.18858	(0.00011)	-1.59880	-0.36390
	110	1.20112	(0.00046)	0.09639	(0.00006)	-1.81204	-0.46686
5	90	0.93253	(0.00053)	0.33563	(0.00016)	-0.80653	-0.19892
	100	1.03873	(0.00049)	0.17846	(0.00010)	-0.95009	-0.27320
	110	1.12509	(0.00047)	0.08612	(0.00006)	-1.06166	-0.34326

Table 4: Various properties and risks of a max-call option on random times ($d = 5$).

Notes. This table displays the expected time to maturity at exercise, $\mathbb{E}[T-\tau^*|\tau^*\le T]$, the probability of not exercising, $\mathbb{P}(\tau^*>T)$, and the opportunity risk, R_O, and timing risk, R_T, for the stopping problem analyzed in Table 2. $\mathbb{E}[T-\tau^*|\tau^*\le T]$ and $\mathbb{P}(\tau^*>T)$ are based on Monte Carlo estimates using the 2,000,000 paths of the forward pass of the LSMC algorithm. R_O and R_T are computed as differences between LSMC estimates of both components in their respective definitions. All simulated paths have been generated in antithetic pairs.

λ	d	LSMC	(s.e.)	LSMC-PI	(s.e.)	Dual	(s.e.)
1	10	28.01203	(0.01470)	28.05150	(0.02168)	28.06987	(0.00719)
	20	36.87827	(0.01666)	36.88446	(0.02498)	36.94684	(0.00932)
	50	48.54796	(0.01966)	48.59752	(0.02934)	48.61604	(0.00929)
	100	57.29451	(0.02220)	57.33810	(0.03228)	57.35563	(0.00947)
3	10	34.81370	(0.01437)	34.88107	(0.02419)	34.88875	(0.00653)
	20	46.15149	(0.01539)	46.24209	(0.02698)	46.24795	(0.00892)
	50	61.23947	(0.01692)	61.28060	(0.03034)	61.34735	(0.00969)
	100	72.62265	(0.01847)	72.61737	(0.03395)	72.75048	(0.01125)
5	10	36.34189	(0.01446)	36.46260	(0.02644)	36.50452	(0.01622)
	20	48.40340	(0.01527)	48.53107	(0.02914)	48.57309	(0.01144)
	50	64.49434	(0.01636)	64.58681	(0.03271)	64.69898	(0.01402)
	100	76.67618	(0.01754)	76.76847	(0.03541)	76.93326	(0.02325)

Table 5: Value of a max-call option on random times (varying d).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5) with d independent geometric Brownian motions (without jumps, $\mu_J^j = \lambda_J^j = 0$), with initial asset prices $X_0 = 100.0$, and with varying arrival rates λ. Column LSMC contains the lower-biased estimates based on the random times LSMC algorithm of Section 4.1 and column Dual contains the upper-biased estimates (4.11) based on the dual algorithm in Section 4.3 and the estimated LSMC policy; standard errors for the dual estimates are the standard errors of the duality gap (4.10). Both algorithms use the same settings as in Table 1, except for the basis functions: as in [5], we use basis functions in the ordered asset prices $X_t^{(1)} \ge \ldots \ge X_t^{(d)}$. The continuation values in the LSMC algorithm are estimated by using the basis functions $(t, x^{(1)}, x^{(2)}) \mapsto t^i \left(x^{(1)}\right)^j \left(x^{(2)}\right)^k$, $i+j+k \le 5$. Column LSMC-PI contains the results of the policy iteration method (3.4), starting from the LSMC estimates, with 100,000 outer paths, 500 inner paths and a window length $\kappa = K'$. All simulated paths have been generated in antithetic pairs.

a given threshold value $X_{min} > 0$ at that time, i.e.:

$$\begin{aligned} \nu_{j+1} &= \nu_j + \eta_j\,, \quad \eta_j \sim \mathrm{Exp}(\lambda) \text{ i.i.d.}, \quad \nu_0 = 0\,, \\ \tau_{i+1} &= \inf\{\nu_j : \nu_j > \tau_i,\ X_{\nu_j}^{(d)} \ge X_{min}\}\,, \quad \tau_0 = 0\,. \end{aligned} \tag{5.6}$$

d	λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Threshold	(s.e.)	Threshold-PI	(s.e.)
5	5	90	10.70117	(0.00872)	11.16820	(0.03041)	10.23629	(0.00819)	11.05121	(0.02260)
		100	21.02792	(0.01158)	21.99106	(0.04818)	19.90658	(0.01098)	21.45137	(0.03356)
		110	32.59407	(0.01380)	33.76811	(0.05761)	30.77168	(0.01311)	32.93245	(0.04233)
	10	90	11.39417	(0.00874)	12.10415	(0.03554)	10.84847	(0.00819)	11.75694	(0.02403)
		100	21.80430	(0.01154)	23.16668	(0.05569)	20.66690	(0.01095)	22.19446	(0.03490)
		110	33.43591	(0.01373)	35.07910	(0.06711)	31.74492	(0.01311)	33.69887	(0.04304)
10	5	90	12.56657	(0.00848)	13.27405	(0.04561)	12.12012	(0.00823)	13.11287	(0.02374)
		100	26.00339	(0.01161)	27.40261	(0.06667)	24.75869	(0.01108)	26.93033	(0.03735)
		110	40.34155	(0.01417)	42.38168	(0.08850)	38.29985	(0.01360)	41.35958	(0.04856)
	10	90	13.58689	(0.00847)	14.64311	(0.05502)	13.03969	(0.00813)	14.26980	(0.02578)
		100	27.20783	(0.01150)	29.39729	(0.08570)	25.94181	(0.01110)	28.13476	(0.03892)
		110	41.60750	(0.01398)	44.69199	(0.10200)	39.67943	(0.01358)	42.49327	(0.04899)

Table 6: Value of a max-call option on reward-dependent random times (varying d).
Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and reward-dependent Poissonian arrivals of stopping opportunities defined in (5.6) with d independent geometric Brownian motions (without jumps, $\mu_J^j = \lambda_J^j = 0$) and with varying arrival rates λ and initial asset prices X_0, where $X_{min} = 60$. All columns are determined as in Table 2, except for the policy iteration columns: Threshold and PI-threshold are based on the backward dynamic programming method with a threshold policy (4.3) rather than a trivial initial policy as in Table 2.

To illustrate the performance of the policy iteration method when the initial policy is already somewhat informed, we now use the backward dynamic programming procedure of Section 4.1 with the threshold-based strategy in (4.3) to determine the initial policy. Table 6 displays the results.

Table 6 shows that a single step of policy iteration over the threshold policy based on (4.3) may already lead to estimates that are close to those obtained by the random times LSMC method (4.2) with many basis functions, and appear to be always larger than the random times LSMC lower bound. If the dimension cardinality of the underlying process is relatively low, the random times LSMC method will usually be computationally faster, since the simulation due to the iterated policy requires one degree of nesting. However, it is known that the LSMC method suffers from the curse of dimensionality, in contrast to the policy iteration method.[3]

This phenomenon is in line with findings from the existing literature on deterministic stopping opportunities, and illustrates how policy iteration may help mitigate the curse of dimensionality. For instance, in [8] it was demonstrated that for an exotic cancelable coupon swap with 20 underlying assets, the classical Longstaff-Schwartz LSMC algorithm with 300 basis functions did not give satisfactory results (duality gaps of around 5% in relative terms). This stands in stark contrast with a one step improvement of an input policy, obtained via the classical Longstaff-Schwartz algorithm with a low number of basis functions, which yields duality gaps of under 1.5% in relative terms, and moreover is up to 8 times faster.

It is likely that for (very) high-dimensional underlying processes, a combination of random

[3]In this respect, we note that the relatively large duality gaps for LSMC may partially be attributed to the (somewhat coarse) choice of basis functions: although the arrivals of the opportunities depend on $X_t^{(d)}$, no basis functions in this term are used for $d \geq 3$. As the effect of this is less pronounced for small λ and the PI method uses fewer paths due to its increased computational complexity, the associated larger standard deviations partially nullify the improvement for smaller λ, whereas the gains from policy iteration become more sizeable as λ, and correspondingly the duality gaps for LSMC, increases.

times LSMC with policy improvement may be (much) more efficient than running the random times LSMC algorithm alone with many basis functions. We consider this an interesting topic for a detailed further study, but beyond the scope of the present paper.

6 Real options application

In this section, we analyze a real options pricing problem, to further illustrate the managerial implications of randomly arriving opportunities to stop. We suppose that an agent owns an illiquid asset, e.g., a piece of real estate or an art object, where illiquid means that a potential buyer is not always directly available and may arrive randomly. The agent attaches a monetary time-invariant value of $\mathcal{K}$ to the asset, but she may try to sell the asset at the prevailing time-varying market price X_t, if a buyer appears at that time instance, i.e., at $t = \tau_i$. Of course, she will only choose to do so if the market price is sufficiently large, and in any case larger than $\mathcal{K}$.[4]

The overall value of the asset for the agent thus consists of two parts: the monetary value $\mathcal{K}$ attached to it and the option value of selling the asset at a price X_{τ_i} beyond $\mathcal{K}$. We suppose that X_t evolves according to a geometric Brownian motion with log-normal jump sizes J_i, with mean 1 and that occur at the arrival times of buyers, i.e., at $t = \tau_i$, hence:

$$\mathrm{d}X_t = (r-\delta)X_t\,\mathrm{d}t + \sigma X_t\,\mathrm{d}W_t + X_{t-}\sum_{t=\tau_i}(J_i - 1)\,, \quad J_i \sim \mathrm{LogNormal}(1, \sigma_J) \text{ i.i.d.}, \tag{6.1}$$

with $X_0 >$, r, δ, $\sigma > 0$, $\sigma_J > 0$ known constants. These dynamics underscore the advantage of simulation-based methods over Markov generator-based methods: the infinitesimal Markov generator for (6.1) would be non-local, which means that its discretized matrix is non-sparse; in such cases, finite-difference methods typically become numerically costly in higher dimensions.[5]

We will study the option value, which means that the reward process of interest in this section is given by $Z_t = e^{-rt}(X_t - \mathcal{K})^+$, corresponding to the discounted payoff of a call option contract. We suppose that the agent is only considering to sell the asset over a three years time window, i.e., $T = 3$; if the asset is not sold by that time, the agent will retain the asset. In the previous section, we have seen that in a low-dimensional setting random times LSMC performs comparably to policy iteration. As this example is meant to illustrate the implications of randomly arriving stopping opportunities, and in particular the role of feedback, we focus only on lower- and upper-biased estimates based on the random times LSMC and dual algorithms, respectively.

In this application, it is natural to consider various dynamics for the arrivals of prospective buyers and thus for the stopping opportunities in (2.1). Buyers may arrive independently of everything else, or with a rate that depends on previous arrivals or the current market price of the asset. We will illustrate how the probability of not exercising, $\mathbb{P}(\tau^* > T)$, the expected time to maturity at exercise, $\mathbb{E}[T - \tau^* | \tau \leq T]$, and the opportunity and timing risks, R_O and R_T, vary

[4] This application follows the spirit of [17] but with randomly arriving stopping opportunities.

[5] See [14] for an in-depth analysis.

with the arrival dynamics of stopping opportunities and the dependence structure between arrivals of stopping opportunities and asset dynamics.

6.1 Reward-independent arrival of buyers

We begin by considering two cases in which buyers arrive either at fixed moments, say at the end of each year, i.e.,

$$\tau_{i+1}^{\mathrm{Ber}} = \tau_i^{\mathrm{Ber}} + \frac{T}{K}\,, \quad \tau_0 = 0\,, \quad i = 0, \ldots, K-1, \tag{6.2}$$

with $K = T = 3$, or according to a Poisson process that is independent of the reward Z_t, i.e.,

$$\tau_{i+1}^{\mathrm{Poi}} = \tau_i^{\mathrm{Poi}} + \eta_i\,, \quad \tau_0 = 0\,, \quad \eta_i \sim \mathrm{Exp}(\lambda) \text{ i.i.d.}, \quad i = 0, 1, \ldots, \tag{6.3}$$

with $\lambda = 1$. Note that when opportunities follow (6.3), there is both timing risk and opportunity risk involved in the stopping decision at time τ_i^{Poi}: the length of the time increment $\tau_{i+1}^{\mathrm{Poi}} - \tau_i^{\mathrm{Poi}}$ is *ex ante* unknown and it may happen that $\tau_{i+1}^{\mathrm{Poi}}$ occurs after T.

The presence of timing risk and opportunity risk does not entail that decision problems where opportunities follow (6.2) are in general always to be preferred to decision problems where (6.3) applies: as a simple example, suppose $Z_t = 1 + \sin(\frac{2K\pi}{T}t + a)\,,\ a \in \{0, \pi\}$, and consider $\tau_i^{\mathrm{Ber}} < T$. If $a = 0$, $Z_{\tau_i^{\mathrm{Ber}}} \equiv 0$, whereas if $a = \pi$, $Z_{\tau_i^{\mathrm{Ber}}} \equiv 2$, which are the minimal and maximal rewards possible, respectively. So, if $a = 0$, the agent clearly prefers Poissonian opportunities despite the presence of opportunity risk, as these attain a positive reward with positive probability. However, if $a = \pi$, the agent prefers Bermudan opportunities, which achieve higher rewards than Poissonian opportunities with probability one as both the timing risk and opportunity risk make Poissonian arrivals less favorable.

Both forms of risk are also present, and sizeable, under reward-independent Poissonian arrivals in the context of the reward in the present section, as illustrated in the first two rows of Table 7, leading on average to stopping earlier.

6.2 Arrival rates with feedback mechanisms

Next, we analyze the case that the arrival of buyers is affected by previous arrivals or by the current market price of the asset. This may be modeled by a càdlàg point process N_t for the number of buyers with (left-continuous, predictable) stochastic jump intensity λ_t, where, as $\Delta \downarrow 0$,

$$\begin{aligned} &\mathbb{P}(N_{t+\Delta} = N_t | \mathcal{F}_t) = 1 - \lambda_t \Delta + o(\Delta)\,, \quad \mathbb{P}(N_{t+\Delta} - N_t = 1 | \mathcal{F}_t) = \lambda_t \Delta + o(\Delta)\,, \\ &\mathbb{P}(N_{t+\Delta} - N_t \geq 2 | \mathcal{F}_t) = o(\Delta)\,, \quad \tau_i^{\lambda} = \inf\{t : N_t = i\}\,. \end{aligned}$$

To let the arrival rate depend upon previous arrivals, we suppose that (N_t, λ_t) follows a Hawkes process with exponential kernel ([27]), which introduces a self-exciting feature:

$$\mathrm{d}\lambda_t^{\mathrm{Haw}} = \alpha(\psi - \lambda_t^{\mathrm{Haw}})\,\mathrm{d}t + \beta\,\mathrm{d}N_t\,. \tag{6.4}$$

Arrivals	LSMC	(s.e.)	Dual	(s.e.)	$\mathbb{E}$[TtM]	$\mathbb{P}(\tau^* > T)$	R_O	R_T
Bermudan	7.54039	(0.00891)	7.54071	(0.00014)	0.42816	0.00000	0.00000	0.00000
Poisson	6.04603	(0.00792)	6.04657	(0.00022)	1.58655	0.66616	-1.19929	-0.29507
Hawkes(0.3)	5.68566	(0.00773)	5.68705	(0.00052)	1.60190	0.68344	-1.48159	-0.37314
Hawkes(0.9)	4.08941	(0.00656)	4.09257	(0.00144)	1.74170	0.76196	-2.79049	-0.66049
Feedback(-1)	2.07946	(0.00383)	2.07947	(0.00001)	1.62801	0.76158	-4.18977	-1.27115
Feedback(1)	6.99131	(0.00815)	6.99558	(0.00126)	1.55637	0.64199	-0.60036	0.05128

Table 7: Option value of the asset with buyers arriving at random times.

Notes. This table displays estimated values of $Y_0^{(K)}$ and its properties in the optimal stopping problem for reward $Z_t = e^{-rt}(X_t - 100)^+$ with price process (6.1) (without jumps, $\sigma_J = 0\%$) and specifications (6.2)–(6.5) for (2.1). Column LSMC contains the lower-biased estimates based on the random times LSMC algorithm of Section 4.1 and column Dual contains the upper-biased estimates based on the dual algorithm in Section 4.3 and the estimated LSMC policy as input for (4.5); standard errors for the dual estimates are the standard errors of the duality gap (4.10). Both algorithms use the same settings as in Table 1, except for the choice of basis functions: denote by d_{eff} the effective state dimension (number of assets + time variable + any state variables needed for the arrival process), then we use $\{(y_1, \ldots, y_{d_{\text{eff}}}) \mapsto \prod_{i=1}^{d_{\text{eff}}} y^{j_i} : \sum_{i=1}^{d_{\text{eff}}} j_i \leq 5\}$ as basis functions. $\mathbb{E}$[TtM] $:= \mathbb{E}[T - \tau^* | \tau \leq T]$ is the expected time to maturity at exercise, $\mathbb{P}(\tau^* > T)$ the probability of not exercising, R_O measures the opportunity risk and R_T the timing risk. The measures $\mathbb{E}$[TtM], $\mathbb{P}(\tau^* > T)$, R_O and R_T are computed as in Table 4.

Arrivals	LSMC	(s.e.)	Dual	(s.e.)	$\mathbb{E}$[TtM]	$\mathbb{P}(\tau^* > T)$	R_O	R_T
Bermudan	8.79715	(0.01045)	8.79767	(0.00023)	0.40733	0.00000	0.00000	0.00000
Poisson	7.53311	(0.00937)	7.53705	(0.00127)	1.62419	0.65759	-1.19040	-0.07364
Hawkes(0.3)	7.12238	(0.00914)	7.13228	(0.00537)	1.64638	0.67436	-1.46015	-0.21462
Hawkes(0.9)	5.37431	(0.00817)	5.38934	(0.00515)	1.77864	0.75885	-2.78202	-0.64082

Table 8: Option value of the asset with buyers arriving at random times and where arrivals cause price jumps.

Notes. This table displays estimated values of $Y_0^{(K)}$ and its properties in the optimal stopping problem for reward $Z_t = e^{-rt}(X_t - 100)^+$ with price process (6.1) (with jumps, $\sigma_J = 10\%$) and specifications (6.2)–(6.4) for (2.1). Details for the respective algorithms are provided in the caption of Table 7. $\mathbb{E}$[TtM] $:= \mathbb{E}[T - \tau^* | \tau \leq T]$ is the expected time to maturity at exercise, $\mathbb{P}(\tau^* > T)$ the probability of not exercising, R_O measures the opportunity risk and R_T the timing risk.

That is, upon the occurrence of arrivals, the arrival intensity ramps up, by $\beta > 0$, making future arrivals temporarily more likely, after which there is mean reversion at rate $\alpha > \beta$ to a level $\psi > 0$. For the sake of comparison, we consider two parameter specifications for λ_t such that $\mathbb{E}[\lambda_t] = 1\,,\ \forall t$, which means that the expected arrival rate at any point in time is the same as under the Poissonian arrivals (6.3). In turn this means that $\mathbb{E}[N_T] = 3$, the (expected) number of buyers under the Bermudan arrivals (6.2). In particular, we consider either $\alpha = 1,\ \beta = 3/10,\ \psi = 7/10,\ \lambda_0 = 1$ (denoted Hawkes(0.3)) or $\alpha = 1,\ \beta = 9/10,\ \psi = 1/10,\ \lambda_0 = 1$ (denoted Hawkes(0.9)).

In Tables 7 and 8, we observe that the option value under Hawkes arrivals is lower than that under Poissonian arrivals, which is due to both a larger (absolute value of) opportunity and timing risk, both of which are moreover increasing with the degree of clustering. Indeed, due to the self-exciting nature of Hawkes arrivals and the condition that $\mathbb{E}[\lambda_t] = 1$, the variation of inter-arrival times increases with the degree of clustering, i.e., with β; hence, the timing risk is larger. Furthermore, under the condition $\mathbb{E}[\lambda_t] = 1$, one may show that $\mathbb{P}(\tau_1^{\text{Poi}} > T) < \mathbb{P}(\tau_1^{\text{Haw}} > T)$, which

explains why the opportunity risk is larger under Hawkes arrivals. From a managerial point of view, when only $\mathbb{E}[\lambda_t]$ or $\mathbb{E}[N_T]$ is available, it is thus important not to confuse self-exciting arrivals with Poissonian or Bermudan arrivals: Hawkes arrivals reduce the option value and require stopping earlier on average. Upon comparing Tables 7 and 8, we see that the option value increases overall when arrivals of stopping opportunities lead to price jumps. Indeed, the expected call option payoff benefits from upward potential due to the price jumps.

We may also let the arrival rate of stopping opportunities depend upon current market prices by setting

$$\lambda_t^\phi := \lambda(t, X_t, \phi), \quad \text{where} \quad \lambda(t,x,\phi) := (1-\phi) + 2\phi F_{X_t}(x)\,, \quad \phi \in [-1,1]\,, \tag{6.5}$$

assuming $\sigma_J \equiv 0$ (i.e., no price jumps), with F_{X_t} denoting the cumulative distribution of X_t. This entails that λ_t^ϕ and X_t move in the same direction if $\phi > 0$ and in opposite directions if $\phi < 0$. Note that $\mathbb{E}[\lambda_t^\phi] = 1$. We denote the extremal cases $\phi = -1$ and $\phi = 1$ by Feedback(-1) and Feedback(1), respectively.

The associated results are displayed in the last two rows of Table 7. Note that Poissonian arrivals correspond to Feedback(0). Upon comparing the values we see that both the opportunity risk and the timing risk become more pronounced (i.e., more negative) as the feedback becomes more negative. Under positive feedback, when the option becomes attractive (i.e., high market price of the asset), the corresponding increased arrival rate of stopping opportunities is beneficial. These insights underscore the importance of the ability to account for the relevant dependence structure and feedback mechanisms.

6.3 Cross-exciting arrivals

We finally analyze a more complex form of dependence between arrivals and prices. Specifically, we assume price jumps and arrivals of stopping opportunities are jointly driven by a two-dimensional Hawkes process with exponential kernel ([27, 1, 2]), which allows for cross-excitation between the number of stopping opportunities, N_t^1, and the number of price jumps, N_t^2:

$$\begin{aligned}
&\mathbb{P}(N_{t+\Delta}^i = N_t^i | \mathcal{F}_t) = 1 - \lambda_t^i \Delta + o(\Delta)\,, \quad \mathbb{P}(N_{t+\Delta}^i - N_t^i = 1 | \mathcal{F}_t) = \lambda_t^i \Delta + o(\Delta)\,,\\
&\mathbb{P}(N_{t+\Delta}^i - N_t^i \geq 2 | \mathcal{F}_t) = o(\Delta)\,, \quad \tau_i = \inf\{t : N_t^1 = i\}\,,\\
&\mathrm{d}\lambda_t^i = \alpha(\psi - \lambda_t^i)\,\mathrm{d}t + \beta\,\mathrm{d}N_t^i + \rho\,\mathrm{d}N_t^{3-i}\,, \quad i = 1,2\,, \qquad (6.6)\\
&\mathrm{d}X_t = (r-\delta)X_t\,\mathrm{d}t + \sigma X_t\,\mathrm{d}W_t + X_{t-} J_{N_t^2}\,\mathrm{d}N_t^2\,, \quad J_i \sim \text{LogNormal}(1, \sigma_J) \text{ i.i.d.} \qquad (6.7)
\end{aligned}$$

We set $r = 5\%,\ \delta = 10\%,\ \sigma = 20\%$ as before and fix $\alpha = 1$ and $\psi = 0.1$. We will vary the cross-excitation parameter $\rho > 0$ while keeping $\mathbb{E}[\lambda_t^i] = 1$, $i = 1, 2$, which dictates that the self-excitation parameter $\beta = 0.9 - \rho$ ([2]). In Table 9, we denote different parameters (β, ρ) as Hawkes(β, ρ). Upon comparing rows 2–4, it is clearly visible that an increase in the degree of cross-excitation, i.e., larger ρ, leads to larger values for the stopping problem, in line with (but expanding the scope

Arrivals	LSMC	(s.e.)	Dual	(s.e.)	$\mathbb{E}$[TtM]	$\mathbb{P}(\tau^* > T)$	R_O	R_T
Bermudan	8.79715	(0.01045)	8.79767	(0.00023)	0.40733	0.00000	0.00000	0.00000
Hawkes(0.9,0)	4.67110	(0.00773)	4.68102	(0.00423)	1.71095	0.76522	-3.40758	-0.71847
Hawkes(0.6,0.3)	5.58933	(0.00853)	5.62453	(0.01209)	1.63582	0.73352	-2.61380	-0.59402
Hawkes(0.3,0.6)	5.85595	(0.00864)	5.88006	(0.00655)	1.61496	0.72337	-2.39815	-0.54305

Table 9: Option value of the asset with buyers arriving at random times where price and arrivals are driven by a Hawkes process.

Notes. This table displays estimated values of $Y_0^{(K)}$ and its properties in the optimal stopping problem for reward $Z_t = e^{-rt}(X_t - 100)^+$ with price process and arrivals given by (6.6)–(6.7), where we set $\sigma_J = 0.1$. Details for the respective algorithms are provided in the caption of Table 7. $\mathbb{E}$[TtM] $:= \mathbb{E}[T - \tau^* | \tau \leq T]$ is the expected time to maturity at exercise, $\mathbb{P}(\tau^* > T)$ the probability of not exercising, R_O measures the opportunity risk and R_T the timing risk.

of) the positive feedback case in the previous setting.

An interesting observation about the role of dependence can be made by comparing rows Hawkes(0.9) in Table 8 and Hawkes$(0.9, 0)$ in Table 9: in either case, the relevant jump times are the event times of a 1-dimensional exponential Hawkes process with the same parameters. In Table 8, both the price jumps and stopping opportunities are generated by the same Hawkes process, whereas in Table 9, both are generated by a different, independent (i.e., non-cross-exciting), Hawkes process; under this independence, the value of the stopping problem is substantially lower. This reinforces our earlier finding that the dependence between the reward and the arrival of stopping opportunities plays an important role in optimal stopping with randomly arriving stopping opportunities, and warrants the development of methods that can handle them.

7 Conclusion

In this paper, we have studied general optimal stopping problems with randomly arriving opportunities to stop. Such problems occur naturally when markets display frictions (e.g., limited liquidity, incomplete information, transaction costs) in a wide variety of settings (e.g., financial contracts, real options, bank runs, natural resources). To encompass these, we allow for general dynamics of the stopping opportunities and the reward process, including general forms of stochastic dependence among them. We have recast the original problem into an equivalent integer-time, infinite-horizon problem. In this context, we have developed, theoretically and pseudo-algorithmically, a policy improvement approach, a dual martingale method, and a regression procedure for optimal stopping problems with infinite horizon.

We have illustrated the good performance of our algorithms in a variety of applications, efficiently yielding (very) small duality gaps, with cardinality of the state space that can easily go up to 100. We have shown that, due in part to the presence of timing risk and opportunity risk, both the optimal value of the stopping problem and the optimal timing of the stopping decisions are substantially affected when stopping opportunities become risky. In the context of a real options problem, this is illustrated by analyzing several models of how prospective buyers may present themselves

for a valuable asset: at deterministic times, following Poissonian arrivals, self-exciting (Hawkes) arrivals and price-dependent arrivals. Furthermore, we have shown that in our flexible framework, problems with feedback, in which the arrival of buyers may lead to sudden price changes, can also be treated.

In order to deal with *multiple* stopping with randomly arriving opportunities, one may follow a similar path. In particular, the corresponding approaches for multiple stopping with finite horizon in the literature, see for example [7] for policy iteration, [53] for the dual martingale representation, and [12] for regression methods, may be extended to infinite-horizon problems in the spirit of this paper.

The methods in this paper were tailored to the specific properties of *stopping* problems; in future work, we intend to extend our simulation-based approach to general *control* problems with random control opportunities.

References

[1] Aït-Sahalia, Y., Laeven, R. J. A., and Pelizzon, L. (2014). Mutual excitation in Eurozone sovereign CDS. *Journal of Econometrics* 183, 151–167.

[2] Aït-Sahalia, Y., Cacho-Diaz, J. A., and Laeven, R. J. A. (2015). Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics* 117, 585–606.

[3] Alvarez, L. H. R., and Stenbacka, R. (2001). Adoption of uncertain multi-stage technology projects: A real options approach. *Journal of Mathematical Economics* 35, 71–97.

[4] Andersen, L. (2000). A simple approach to the pricing of Bermudan swaptions in the multifactor LIBOR market model. *The Journal of Computational Finance* 3, 5–32.

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

[6] Belomestny, D., Bender, C., and Schoenmakers, J. G. M. (2009). True upper bounds for Bermudan products via non-nested Monte Carlo. *Mathematical Finance*19, 53–71.

[7] Bender, C., and Schoenmakers, J. G. M. (2006). An iterative method for multiple stopping: Convergence and stability. *Advances in Applied Probability* 38, 729–749.

[8] Bender, C., Kolodko, A., and Schoenmakers, J. G. M. (2008). Enhanced policy iteration for American options via scenario selection. *Quantitative Finance* 8, 135–146.

[9] Benth, F. E., Di Persio, L., and Lavagnini, S. (2018). Stochastic modeling of wind derivatives in energy markets. *Risks* 6, 56.

[10] Benth, F. E., Lange, N., and Myklebust, T. Å. (2015). Pricing and hedging quanto options in energy markets. *Journal of Energy Markets* 8, 1–35.

[11] Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics* 12, 383–398.

[12] Carmona, R., and Touzi, N. (2008). Optimal multiple stopping and valuation of swing options. *Mathematical Finance* 18, 239–268.

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

[14] Cont, R., and Voltchkova, E. (2005). A finite difference scheme for option pricing in jump diffusion and exponential Lévy models. *SIAM Journal on Numerical Analysis* 43, 1596–1626.

[15] Davis, M., and Karatzas, I. (1994). A deterministic approach to optimal stopping. In: Kelly, F. P. (ed.) *Probability, statistics and optimisation. A tribute to Peter Whittle.* Wiley, Chichester. Wiley Series in Probability and Mathematical Statistics. 455–466.

[16] Dias, M. A. G. (2004). Valuation of exploration and production assets: An overview of real options models. *Journal of Petroleum Science and Engineering* 44, 93–114.

[17] Dixit, A. K., and Pindyck, R.S. (1994). *Investment under Uncertainty*, Princeton University Press, Princeton.

[18] Dupuis, P., and Wang, H. (2002). Optimal stopping with random intervention times. *Advances in Applied Probability* 34, 141–157.

[19] Egloff, D., and Leippold, M. (2009). The valuation of American options with stochastic stopping time constraints. *Applied Mathematical Finance* 16, 287–305.

[20] Fischer, T. (2013). On simple representations of stopping times and stopping time sigma-algebras. *Statistics and Probability Letters* 83, 345–349.

[21] Fouilloux, J., Moraux, F., and Viviani, J.-L. (2015). Investing in finite-life carbon emissions reduction program under risk and idiosyncratic uncertainty. *Energy Policy* 82, 310–320.

[22] Fujimoto, K., Nagai, H., and Runggaldier, W. J. (2013). Expected power-utility maximization under incomplete information and with Cox-process observations. *Applied Mathematics & Optimization* 67, 33–72.

[23] Fujimoto, K., Nagai, H., and Runggaldier, W. J. (2014). Expected log-utility maximization under incomplete information and with Cox-process observations. *Asia-Pacific Financial Markets* 21, 35–66.

[24] Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering.* Springer-Verlag, New York.

[25] Guo, X., and Liu, J. (2005). Stopping at the maximum of geometric Brownian motion when signals are received. *Journal of Applied Probability* 42, 826–38.

[26] Haugh, M. B., and Kogan, L. (2004). Pricing American options: A duality approach. *Operations Research* 52, 258–270.

[27] Hawkes, A. G. (11971). Spectra of some self-exciting and mutually exciting point processes. *Biometrika* 58, 83–90.

[28] He, Z., and Xiong, W. (2012). Dynamic debt runs. *The Review of Financial Studies* 25, 1799–1843.

[29] He, Z., and Milbradt, K. (2014). Endogenous liquidity and defaultable bonds. *Econometrica* 82, 1443–1508.

[30] Hobson, D. (2021). The shape of the value function under Poisson optimal stopping. *Stochastic Processes and their Applications* 133, 229–246.

[31] Hobson, D., and Zeng, M. (2022). Constrained optimal stopping, liquidity and effort. *Stochastic Processes and their Applications* 150, 819–843.

[32] Karatzas, I., and Shreve, S. E. (1998). *Brownian Motion and Stochastic Calculus*. Springer, Berlin.

[33] Kolodko, A., and Schoenmakers, J. G. M. (2006). Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics* 10, 1177–1209.

[34] Krätschmer, V., Ladkau, M., Laeven, R. J. A., Schoenmakers, J. G. M., and Stadje, M. (2018). Optimal stopping under uncertainty in drift and jump intensity. *Mathematics of Operations Research* 43, 27–49.

[35] Lange, R. J., Ralph, D., and Store, K. (2020). Real-option valuation in multiple dimensions using Poisson optional stopping times. *Journal of Financial and Quantitative Analysis* 55, 653–677.

[36] Laeven, R. J. A., Schoenmakers, J. G. M., Schweizer, N. F. F., and Stadje, M. (2025). Robust multiple stopping — A duality approach. *Mathematics of Operations Research* 50, 1250–1276.

[37] Lange, R. J., and Teulings, C. N. (2021). The option value of vacant land: Don't build when demand for housing is booming, SSRN Electronic Journal, `https://ssrn.com/abstract=3846109`.

[38] Leland, H. E. (1998). Agency costs, risk management, and capital structure. *The Journal of Finance* 53, 1213–1243.

[39] Lempa, J. (2012). Optimal stopping with information constraint. *Applied Mathematics and Optimization* 66, 147–173.

[40] Liang, G., Lütkebohmert, E., and Wei, W. (2015). Funding liquidity, debt tenor structure, and creditor's belief: An exogenous dynamic debt run model. *Mathematics and Financial Economics* 9, 271–302.

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

[42] Menaldi, J. L., and Robin, M. (2016). On some optimal stopping problems with constraint. *SIAM Journal on Control and Optimization* 54, 2650–2671.

[43] Palmowski, Z., Pérez, J. L., and Yamazaki, K. (2021). Double continuation regions for American options under Poisson exercise opportunities. *Mathematical Finance* 31, 722–771.

[44] Pérez, J. L., and Yamazaki, K. (2018). American options under periodic exercise opportunities. *Statistics and Probability Letters* 135, 92–101.

[45] Pham, H., and Tankov, P. (2008). A model of optimal consumption under liquidity risk with random trading times. *Mathematical Finance* 18, 613–627.

[46] Pham, H., and Tankov, P. (2009). A coupled system of integrodifferential equations arising in liquidity risk model. *Applied Mathematics and Optimization* 59, 147–173.

[47] Paddock, J. L., Siegel, D. R., and Smith, J. L. (1988). Option valuation of claims on real assets: The case of offshore petroleum leases. *The Quarterly Journal of Economics* 103, 479–508.

[48] Peskir, G., and Shiryaev, A. N. (2006). *Optimal Stopping and Free-Boundary Problems*. Birkhäuser Verlag, Basel.

[49] Protter, P. (1990). *Stochastic Integration and Differential Equations: A New Approach*. Springer Verlag, Berlin.

[50] Rogers, L. C. G. (2002). Monte Carlo valuation of American options. *Mathematical Finance* 12, 271–286.

[51] Rogers, L. C. G., and Zane, O. (2002). A simple model of liquidity effects. In: Sandmann, K., and Schönbucher, P. (eds.). *Advances in Finance and Stochastics. Essays in Honour of Dieter Sondermann*. Springer, Berlin. 161–176.

[52] Saarinen, H. (2023). *On Poisson Constrained Control of Linear Diffusions*. Doctoral Thesis, University of Turku, Finland.

[53] Schoenmakers, J. G. M. (2012). A pure martingale dual for multiple stopping. *Finance and Stochastics* 16, 319–334.

[54] Snell, L. J. (1952). Applications of martingale system theorems. *Transactions of the American Mathematical Society* 73, 293–312.

[55] Wald, A. (1950). *Statistical Decision Functions*. Wiley, New York.

A Proofs and Auxiliary Technical Results

A.1 Properties of Problem (2.2)

Proof of Proposition 1. (i) For fixed $\mathfrak{n} \in \mathcal{N}$ with $\mathfrak{n} \leq K$ a.s., $\mathfrak{t} := \tau_{\mathfrak{n}}$ is an $\mathbb{F}$-stopping time with $\mathfrak{t} \in \mathcal{T}_K$, so $\sup_{\mathfrak{n}\in\mathcal{N},\ \mathfrak{n}\leq K} \mathbb{E}[U_{\mathfrak{n}}] \leq Y_0^{(K)}$. Indeed, one trivially has that $\tau_{\mathfrak{n}(\omega)}(\omega) \in \mathcal{T}_K(\omega)$, and moreover,

$$\{\mathfrak{t} \leq t\} = \{\tau_{\mathfrak{n}} \leq t\} = \bigcup_{i=0}^{K} \{\tau_{\mathfrak{i}} \leq t\} \cap \{\mathfrak{n} = i\} \in \mathcal{F}_t\,,$$

since, for any $i \in \mathbb{N}_0$, $\{\tau_i \leq t\} \in \mathcal{F}_t$ and $\{\mathfrak{n} = i\} \in \mathcal{F}_{\tau_{\mathfrak{i}}}$, and so $\{\tau_{\mathfrak{i}} \leq t\} \cap \{\mathfrak{n} = i\} \in \mathcal{F}_t$, due to (2.4).

(ii) Now let $\mathfrak{t}$ be any $\mathbb{F}$-stopping time with a.s. $\mathfrak{t}(\omega) \in \mathcal{T}_K(\omega)$. Define $\mathfrak{n} : \Omega \to \mathbb{N}_0$ by

$$\mathfrak{n}(\omega) = \sum_{i=0}^{K} i 1_{\{\mathfrak{t}(\omega)=\tau_i(\omega)\}}\,.$$

Then, for any j, since both $\mathfrak{t}$ and τ_j are $\mathbb{F}$-stopping times, it holds that

$$\{\mathfrak{t}(\omega) = \tau_j(\omega)\} \in \mathcal{F}_{\mathfrak{t}} \cap \mathcal{F}_{\tau_j}\,;$$

see e.g., [32, Lemma 2.16]. Hence,

$$\{\mathfrak{n}(\omega) = j\} = \{\mathfrak{t}(\omega) = \tau_j(\omega)\} \in \mathcal{F}_{\tau_j} = \mathcal{G}_j\,,$$

i.e., $\mathfrak{n}$ is a discrete $\mathbb{G}$-stopping time with $\mathfrak{n} \leq K$ a.s., and moreover

$$\mathbb{E}[Z_{\mathfrak{t}}] = \mathbb{E}\left[\sum_{i=1}^{K} U_{\mathfrak{i}} 1_{\{\mathfrak{t}=\tau_{\mathfrak{i}}\}}\right] = \mathbb{E}\left[\sum_{i=1}^{K} U_i 1_{\{\mathfrak{n}=i\}}\right] = \mathbb{E}[U_{\mathfrak{n}}]\,.$$

Thus, $Y_0^{(K)} \leq \sup_{\mathfrak{n}\in\mathcal{N},\ \mathfrak{n}\leq K} \mathbb{E}[U_{\mathfrak{n}}]$. ■

Lemma 18 *Consider a sequence (τ_i') that satisfies* (2.1)*. Next, construct (τ_i) by setting $\tau_0 = 0$ and, for $i \geq 1$,*

$$\tau_i = \tau_i' 1_{\{\tau_i' < T\}} + T 1_{\{\tau_{i-1}' < T \wedge \tau_i' \geq T\}} + \tau_i' 1_{\{\tau_{i-1}' \geq T\}}\,. \tag{A.8}$$

Then, (τ_i) defined by (A.8) *satisfies both* (2.1) *and* (2.9)*.*

Proof. We first show that (A.8) is a stopping time. For $s < T$, we have that

$$\{\tau_i \leq s\} = \left\{\tau_i' \leq s\right\} \in \mathcal{F}_s\,, \tag{A.9}$$

since $\tau_i \in \{T, \tau_i'\}$. Furthermore, we have that

$$\begin{aligned}
\{\tau_i = T\} &= \{\tau_{i-1}' < T\} \cap \{\tau_i' < T\}^c \\
&= \left(\bigcup_n \left\{\tau_{i-1}' \le T - \frac{1}{n}\right\}\right) \cap \left(\bigcup_n \left\{\tau_i' \le T - \frac{1}{n}\right\}\right)^c \\
&\in \mathcal{F}_T\,, \quad \text{hence} \\
\{\tau_i \le T\} &= \{\tau_i = T\} \cup \{\tau_i < T\} \\
&= \{\tau_i = T\} \cup \bigcup_n \left\{\tau_i \le T - \frac{1}{n}\right\} \in \mathcal{F}_T\,,
\end{aligned}$$

using (A.9). For $s > T$, it holds that

$$\begin{aligned}
\{\tau_i \le s\} &= \{\tau_i \le T\} \cup \{T < \tau_i \le s\} \\
&= \{\tau_i \le T\} \cup \left(\{\tau_i' \le T\}^c \cap \{\tau_i' \le s\}\right) \in \mathcal{F}_s\,.
\end{aligned}$$

We next show that $\tau_{i+1} > \tau_i$ for all $i \in \mathbb{N}_0$. Suppose that $\tau_{i+1} \le \tau_i$. Since $\tau_i \in \{T, \tau_i'\}$, $\tau_{i+1} \in \{T, \tau_{i+1}'\}$, and by assumption $\tau_{i+1}' > \tau_i'$, we must have that (a) $(\tau_i = T) \wedge (\tau_{i+1} \le T)$ or that (b) $(\tau_{i+1} = T) \wedge (\tau_i \ge T)$. On the one hand, (a) implies $\tau_i' \ge T$ by (A.8), and this in turn implies $\tau_{i+1} = \tau_{i+1}'$ by replacing i by $i+1$ in (A.8). Hence, we get $\tau_{i+1} = \tau_{i+1}' > T$, a contradiction. On the other hand, (b) implies $\tau_i' < T$ by (A.8), and this in turn implies $\tau_i = \tau_i'$ by (A.8). Hence, we get $\tau_i < T$, a contradiction. Finally, clearly, $\lim_{i\to\infty} \tau_i = \infty$, and so (2.1) and (2.9) follow. ■

Corollary 19 *If in Lemma 18 (τ_i') satisfies* (2.3) *and $\mathbb{E}[Z_T^\alpha]$ is finite, then (τ_i) also satisfies* (2.3)*, albeit possibly with a larger value of B.*

Proof. Note that

$$\sup_{i\in\mathbb{N}_0} Z_{\tau_i}^\alpha \le \sup_{i\in\mathbb{N}_0} Z_{\tau_i'}^\alpha + Z_T^\alpha\,,$$

and then take expectations on both sides. ■

Remark 20 *In practice, the bound* (2.3)*, hence* (2.6)*, may be determined by simulation. For an analytic alternative, let us observe that condition* (2.3) *is implied by the somewhat stronger condition*

$$\mathbb{E}\left[\sup_{0\le t\le T} Z_t^\alpha\right] \le B^\alpha\,, \quad \text{for some } B > 0 \text{ and } \alpha > 1\,.$$

In the case in which the reward can be written in the form $Z_t = g(t, M_t)$ for some d-dimensional martingale M, and the function g satisfies

$$|g(t, M_t)| \le c_1 \|M_t\|_\infty^q + c_2\,, \quad 0 \le t \le T\,, \quad q \ge 1\,,$$

where $\|\cdot\|_\infty$ *denotes the max-norm in* $\mathbb{R}^d$*, it follows that*

$$\begin{aligned}
\mathbb{E}\left[\sup_{0\le t\le T} Z_t^\alpha\right] &\le \mathbb{E}\left[\sup_{0\le t\le T}\left(c_1\|M_t\|_\infty^q + c_2\right)^\alpha\right] \le \mathbb{E}\left[\left(c_1\sup_{0\le t\le T}\|M_t\|_\infty^q + c_2\right)^\alpha\right] \\
&\le 2^{\alpha-1}c_1^\alpha\mathbb{E}\left[\sup_{0\le t\le T}\|M_t\|_\infty^{\alpha q}\right] + 2^{\alpha-1}c_2^\alpha \\
&\text{(with } M=(M^i)_{i=1,\dots,d}) \le 2^{\alpha-1}c_1^\alpha d^{\alpha q-1}\sum_{i=1}^d \mathbb{E}\left[\sup_{0\le t\le T}\left|M_t^i\right|^{\alpha q}\right] + 2^{\alpha-1}c_2^\alpha \\
&\le 2^{\alpha-1}c_1^\alpha d^{\alpha q-1}C_{\alpha q}^{BDG}\sum_{i=1}^d \mathbb{E}\left[\langle M_T^i, M_T^i\rangle^{\alpha q/2}\right] + 2^{\alpha-1}c_2^\alpha\,,
\end{aligned}$$

using the Burkholder-Davis-Gundy inequality. The constant $C_{\alpha q}^{BDG}$ *is universal and explicitly known for continuous martingales and jump martingales, respectively. Furthermore, if* $\alpha q = 2$ *for example,* $\mathbb{E}\left[\langle M_T^i, M_T^i\rangle\right]$ *may usually be estimated explicitly for a given specific model.*

Lemma 21 *Let*

$$\mathfrak{m}_i := \min\{j \ge i : \tau_j > T\}\,, \quad \text{for} \quad i \in \mathbb{N}_0.$$

Then,

(i) for all $i \in \mathbb{N}_0$*, we have that* $U_j = 0$ *a.s. for all* $j \ge \mathfrak{m}_i$*,*

(ii) $i \le \mathfrak{m}_i < \infty$ *a.s., and*

(iii) $\mathfrak{m}_i$ *is a* $\mathbb{G}$*-stopping time.*

Proof. (i) and (ii) follow from the assumption that $Z_t = 0$ a.s. for all $t > T$ and assumption (2.3).

(iii): Consider, for all $n \in \mathbb{N}_0$,

$$\{\mathfrak{m}_i = n\} = \{\tau_n > T\} \cap \bigcap_{i\le m<n}\{\tau_m \le T\}\,,$$

and so, for all $s \ge 0$,

$$\begin{aligned}
\{\mathfrak{m}_i = n\} \cap \{\tau_n \le s\} &= \{\tau_n > T\} \cap \bigcap_{i\le m<n}\{\tau_m \le T\} \cap \{\tau_n \le s\} \\
&= \begin{cases} \varnothing \in \mathcal{F}_s, & s < T \\ \{\tau_n > T\} \cap \bigcap_{i\le m<n}\{\tau_m \le T\} \cap \{\tau_n \le s\} \in \mathcal{F}_s, & s \ge T \end{cases}.
\end{aligned}$$

That is, $\{\mathfrak{m}_i = n\} \in \mathcal{F}_{\tau_n} = \mathcal{G}_n$. ■

Proof of Proposition 3. We note that (2.6), which follows from assumption (2.3), implies that condition [48, Ch. 1, (1.1.3))] is satisfied. Following [48, Ch. 1], we then have (i) and (ii), and that

(2.13) is an optimal stopping time that satisfies (2.14). We next prove for the infinite horizon case $K = \infty$ that $\mathfrak{n}_i^{\infty,*} < \infty$ a.s. Consider the event

$$A_i := \left\{\mathfrak{n}_i^{\infty,*} = \infty\right\} = \left\{\forall_{j\geq i} : U_j < \mathbb{E}_{\mathcal{G}_j}\left[\operatorname*{ess\,sup}_{j+1\leq \mathfrak{n}<\infty} \mathbb{E}_{\mathcal{G}_{j+1}}\left[U_{\mathfrak{n}}\right]\right]\right\}$$
$$\subset \left\{\forall_{j\geq i} : U_j < \mathbb{E}_{\mathcal{G}_j}\left[\sup_{j+1\leq r<\infty} U_r\right]\right\} =: \widetilde{A}_i,$$

for some $i \in \mathbb{N}_0$. Then, with $\mathfrak{m}_i$ defined in Lemma 21, we have on the set $\widetilde{A}_i$,

$$0 = U_{\mathfrak{m}_i} = \sum_{l=i}^{\infty} 1_{\{\mathfrak{m}_i=l\}} U_l < \sum_{l=i}^{\infty} 1_{\{\mathfrak{m}_i=l\}} \mathbb{E}_{\mathcal{G}_l}\left[\sup_{l+1\leq r<\infty} U_r\right]$$
$$= \sum_{l=i}^{\infty} \mathbb{E}_{\mathcal{G}_l}\left[1_{\{\mathfrak{m}_i=l\}} \sup_{l+1\leq r<\infty} U_r\right] = 0 \quad \text{a.s.},$$

using that $\mathfrak{m}_i$ is a stopping time due to Lemma 21. This implies that $\mathbb{P}(A_i) \leq \mathbb{P}\left(\widetilde{A}_i\right) = 0$. Hence, it follows that $\mathfrak{n}_i^{\infty,*} < \infty$ a.s. ■

Proof of Corollary 4. Indeed, from item (i) in the proof of Proposition 1, we see that $\mathfrak{t}^*$ is an $\mathbb{F}$-stopping time that satisfies $\mathfrak{t}^* \in \mathcal{T}_K$ and that moreover satisfies

$$Y_0^{(K)} = \sup_{\mathfrak{n}\in\mathcal{N},\ \mathfrak{n}\leq K} \mathbb{E}\left[U_{\mathfrak{n}}\right] = \mathbb{E}\left[U_{\mathfrak{n}^*}\right] = \mathbb{E}\left[Z_{\mathfrak{t}^*}\right].$$

■

A.2 Duality

Proof of Proposition 5. By Proposition 1 and Doob's sampling theorem we have, for any $M \in \mathcal{M}_0^{\mathsf{UI}}$,

$$Y_0 = \sup_{\mathfrak{n}\in\mathcal{N}} \mathbb{E}\left[U_{\mathfrak{n}}\right] = \sup_{\mathfrak{n}\in\mathcal{N}} \mathbb{E}\left[U_{\mathfrak{n}} - M_{\mathfrak{n}}\right] \leq \mathbb{E}\left[\sup_{i\geq 0}\left(U_i - M_i\right)\right]. \tag{A.10}$$

Note that for any $i \geq 0$ (empty sums being zero),

$$U_i - M_i^{\circ} = U_i - \sum_{j=1}^{i} Y_j + \sum_{j=1}^{i} \mathbb{E}_{\mathcal{G}_{j-1}}\left[Y_j\right] \tag{A.11}$$
$$= Y_0 + U_i - Y_i - \sum_{j=0}^{i-1} Y_j + \sum_{j=0}^{i-1} \mathbb{E}_{\mathcal{G}_j}\left[Y_{j+1}\right]$$
$$\leq Y_0,$$

due to (2.12). Hence, we have

$$\sup_{i\geq 0}\left(U_i - M_i^\circ\right) \leq Y_0, \tag{A.12}$$

by (A.11). We now only need to show that $M^\circ \in \mathcal{M}_0^{\mathsf{UI}}$. Indeed, if $M^\circ \in \mathcal{M}_0^{\mathsf{UI}}$, we have

$$Y_0 \leq \mathbb{E}\left[\sup_{i\geq 0}\left(U_i - M_i^\circ\right)\right],$$

by inequality (A.10), and then by (A.12) we obtain statement (ii) due to the sandwich property. Finally, (ii) combined with (A.10) yields statement (i). Let us show that $M^\circ \in \mathcal{M}_0^{\mathsf{UI}}$. Note that, by (2.17),

$$|M_i^\circ| \leq \sum_{j=1}^{\infty}\left|Y_j - \mathbb{E}_{\mathcal{G}_{j-1}}[Y_j]\right| \quad \text{a.s. for all } i \geq 0\,. \tag{A.13}$$

Since $Z_t \geq 0$ for $t \geq 0$, and $Z_t = 0$ for $t > T$, we have that $Y_j = Y_j 1_{\{\tau_{j-1}\leq T\}} \geq 0$ a.s. Hence, by the monotone convergence theorem,

$$\begin{aligned}
\mathbb{E}\left[\sum_{j=1}^{\infty}\left|Y_j - \mathbb{E}_{\mathcal{G}_{j-1}}[Y_j]\right|\right] &= \mathbb{E}\left[\sum_{j=1}^{\infty}\left|Y_j - \mathbb{E}_{\mathcal{G}_{j-1}}[Y_j]\right| 1_{\{\tau_{j-1}\leq T\}}\right] \\
&\leq \mathbb{E}\left[\sum_{j=1}^{\infty} Y_j 1_{\{\tau_{j-1}\leq T\}} + \mathbb{E}_{\mathcal{G}_{j-1}}\left[Y_j 1_{\{\tau_{j-1}\leq T\}}\right]\right] \\
&= 2\mathbb{E}\left[\sum_{j=1}^{\infty} Y_j 1_{\{\tau_{j-1}\leq T\}}\right] \leq 2\mathbb{E}\left[\sup_{i\in\mathbb{N}_0} U_i \sum_{j=1}^{\infty} 1_{\{\tau_{j-1}\leq T\}}\right] \\
&= 2\mathbb{E}\left[N_T \sup_{i\in\mathbb{N}_0} U_i\right] < \infty,
\end{aligned}$$

due to (2.15) and using that $\{\tau_{j-1} \leq T\} = \{\mathfrak{m}_j \leq j-1\}^c$ is $\mathcal{G}_{j-1}$-measurable since $\mathfrak{m}_j$ is a stopping time by Lemma 21. So the family $(M_i^\circ)_{i\geq 0}$ is uniformly integrable since it has an integrable dominator (A.13). ■

Corollary 22 *Condition* (2.6) *implies that* $\sup_{i\in\mathbb{N}_0} U_i \in L^\alpha$. *By Hölder's inequality,*

$$\mathbb{E}\left[N_T \sup_{i\in\mathbb{N}_0} U_i\right] \leq B\mathbb{E}\left[N_T^{\alpha'}\right]^{1/\alpha'},$$

where $\alpha' := \alpha/(\alpha-1) > 1$ *is the Hölder conjugate of* α. *So a sufficient condition for* (2.15) *is that* $N_T \in L^{\alpha'}$ *i.e.,*

$$\mathbb{E}\left[N_T^{\frac{\alpha}{\alpha-1}}\right] < \infty\,. \tag{A.14}$$

The following lemma provides a convenient way to check (A.14) by checking the tail probabilities of N_T.

Lemma 23 *Let N be an $\mathbb{N}$-valued random variable and let $\alpha > 1$. Then,*

$$N \in L^{\frac{\alpha}{\alpha-1}} < \infty \text{ if and only if } \sum_{j=1}^{\infty} j^{\frac{1}{\alpha-1}} \mathbb{P}(N > j) < \infty \,.$$

Proof. We may write (since $N \geq 1$ a.s.),

$$\begin{aligned}
\mathbb{E}\left[N^{\frac{\alpha}{\alpha-1}}\right] &= \sum_{j=1}^{\infty} j^{\frac{\alpha}{\alpha-1}} \mathbb{P}(N = j) \\
&= \sum_{j=1}^{\infty} j^{\frac{\alpha}{\alpha-1}} \left(\mathbb{P}(N > j-1) - \mathbb{P}(N > j)\right) \\
&= 1 + \sum_{j=1}^{\infty} \left((j+1)^{\frac{\alpha}{\alpha-1}} - j^{\frac{\alpha}{\alpha-1}}\right) \mathbb{P}(N > j) \,.
\end{aligned}$$

Let us now write, with $\alpha' = \alpha/(\alpha-1)$,

$$(j+1)^{\frac{\alpha}{\alpha-1}} - j^{\frac{\alpha}{\alpha-1}} = j^{\alpha'} \left(\left(1 + \frac{1}{j}\right)^{\alpha'} - 1\right).$$

Then, due to elementary inequalities, we have that

$$\frac{\alpha'}{j} \leq \left(1 + \frac{1}{j}\right)^{\alpha'} - 1 \leq e^{\frac{\alpha'}{j}} - 1 \leq \frac{2\alpha'}{j} \quad \text{for } j > \alpha'.$$

Hence, $(j+1)^{\frac{\alpha}{\alpha-1}} - j^{\frac{\alpha}{\alpha-1}} \asymp j^{\alpha'-1} = j^{\frac{1}{\alpha-1}}$, which shows the equivalence. ∎

A.3 Policy iteration

Proof of Proposition 8. Consider for $N \in \mathbb{N}$ the truncated reward $U_i^N := U_i 1_{\{i<N\}}$ and the stopped policy $\sigma_i^N := \sigma_i \wedge N$. It is easy to see that, for $0 \leq i \leq N$,

$$i \leq \sigma_i^N \leq N \quad \text{and} \quad \sigma_i^N > i \quad \Longrightarrow \quad \sigma_i^N = \sigma_{i+1}^N \,,$$

i.e., σ_i^N is consistent in the sense of (3.1). We next consider the iteration step

$$\widehat{\sigma}_i^N := \inf\left\{j \geq i : U_j^N \geq \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j}\left[U_{\sigma_k^N}^N\right]\right\}$$

and define

$$
\begin{aligned}
Y_i^{\sigma^N} &:= \mathbb{E}_{\mathcal{G}_i}\left[U^N_{\sigma_i^N}\right], \\
\widetilde{Y}_i^{\sigma^N} &:= \max_{i\le k\le i+\kappa} \mathbb{E}_{\mathcal{G}_i}\left[U^N_{\sigma_k^N}\right], \\
Y_i^{\widehat{\sigma}^N} &:= \mathbb{E}_{\mathcal{G}_i}\left[U^N_{\widehat{\sigma}_i^N}\right].
\end{aligned}
$$

Note that $i \le \widehat{\sigma}_i^N \le N$. Due to [33, Theorem 3.1], it holds that

$$
Y_i^{\sigma^N} \le \widetilde{Y}_i^{\sigma^N} \le Y_i^{\widehat{\sigma}^N} \le Y_i^{(N)}, \quad \text{for} \quad \le i \le N, \tag{A.15}
$$

where $\left(Y_i^{(N)}\right)$ is the Snell envelope for the truncated reward $\left(U_i^N\right)$. We now fix i and show that for $N \to \infty$, a.s.,

(i) $Y_i^{\sigma^N} \to Y_i^{\sigma}$,
(ii) $\widetilde{Y}_i^{\sigma^N} \to \widetilde{Y}_i^{\sigma}$,
(iii) $Y_i^{\widehat{\sigma}^N} \to Y_i^{\widehat{\sigma}}$,

from which the statement follows by (A.15) and the obvious fact that $Y_i^{(N)} \le Y_i$.

Ad (i)+(ii): For any fixed k one has $\sigma_k^N \to \sigma_k < \infty$ a.s. for $N \to \infty$, hence

$$
U^N_{\sigma_k^N} = U_{\sigma_k^N} 1_{\{\sigma_k^N < N\}} = U_{\sigma_k} 1_{\{\sigma_k < N\}} \to U_{\sigma_k}, \quad \text{a.s.}
$$

It then follows by conditional dominated convergence that

$$
\begin{aligned}
Y_i^{\sigma^N} &= \mathbb{E}_{\mathcal{G}_i}\left[U^N_{\sigma_i^N}\right] \to \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_k}\right] = Y_i^{\sigma}, \quad \text{and} \\
\widetilde{Y}_i^{\sigma^N} &= \max_{i\le k\le i+\kappa} \mathbb{E}_{\mathcal{G}_i}\left[U^N_{\sigma_k^N}\right] \to \max_{i\le k\le i+\kappa} \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_k}\right] = \widetilde{Y}_i^{\sigma}.
\end{aligned}
$$

Ad (iii): Note that by the definition of U_i^N,

$$
\widehat{\sigma}_i^N := \inf\left\{ j \ge i : U_j 1_{\{j\le N\}} \ge \max_{j\le k\le j+\kappa} \mathbb{E}_{\mathcal{G}_j}\left[U_{\sigma_k} 1_{\{\sigma_k < N\}}\right] \right\}.
$$

Let us suppose that we are on the set $\{\omega : \widehat{\sigma}_i(\omega) = l\}$. For $N > l$, one then has

$$
U_l(\omega)\, 1_{\{l<N\}} = U_l(\omega) \ge \max_{l\le k\le l+\kappa} \mathbb{E}_{\mathcal{G}_l}\left[U_{\sigma_k}\right](\omega) \ge \max_{l\le k\le l+\kappa} \mathbb{E}_{\mathcal{G}_l}\left[U_{\sigma_k} 1_{\{\sigma_k<N\}}\right](\omega)
$$

(recall that by assumption $U_\cdot \ge 0$). Hence, for $N \ge l$, one has $\widehat{\sigma}_i^N(\omega) \le l$. Since l was arbitrary, we conclude that

$$
\widehat{\sigma}_i^N \le \widehat{\sigma}_i \text{ a.s.} \quad \text{for } N \ge \mathcal{N}(\omega) \quad \text{say}.
$$

Now assume that $\widehat{\sigma}_i^N \nrightarrow \widehat{\sigma}_i$ with positive probability. Then with positive probability, there must exist numbers r and l with $0 \le r < l$ and a sequence $(N_m)_{m\in\mathbb{N}}$ with $N_m \uparrow \infty$ for $m \uparrow \infty$, such that

$\widehat{\sigma}_i = l$ and $\widehat{\sigma}_i^{N_m} = r < l$ for $m \in \mathbb{N}$. This implies

$$U_r 1_{\{r \leq N_m\}} \geq \max_{r \leq k \leq r+\kappa} \mathbb{E}_{\mathcal{G}_r} \left[U_{\sigma_k} 1_{\{\sigma_k < N_m\}} \right], \quad \forall m \in \mathbb{N},$$

on a set of positive probability. By conditional dominated convergence it then follows that

$$U_r \geq \max_{r \leq k \leq r+\kappa} \mathbb{E}_{\mathcal{G}_r} \left[U_{\sigma_k} \right]$$

on this set, which contradicts the definition of $\widehat{\sigma}_i$. That is, $\widehat{\sigma}_i^N \to \widehat{\sigma}_i$ a.s., and then (iii) follows by conditional dominated convergence. ■

Proof of Corollary 9. Ad (i): By Proposition 8 we may write

$$\begin{aligned} 1_{\left\{Y_i^{\widehat{\sigma}} < U_i\right\}} Y_i^{\widehat{\sigma}} &= 1_{\left\{Y_i^{\widehat{\sigma}} < U_i\right\}} 1_{\left\{\widetilde{Y}_i^{\sigma} \leq U_i\right\}} \mathbb{E}_{\mathcal{G}_i} \left[U_{\widehat{\sigma}_i} \right] \\ \text{(definition of } \widehat{\sigma}_i) &= 1_{\left\{Y_i^{\widehat{\sigma}} < U_i\right\}} 1_{\left\{\widetilde{Y}_i^{\sigma} \leq U_i\right\}} U_i = 1_{\left\{Y_i^{\widehat{\sigma}} < U_i\right\}} U_i, \quad \text{a.s.} \end{aligned}$$

That is, on the set $\left\{Y_i^{\widehat{\sigma}} < U_i\right\}$ one has $Y_i^{\widehat{\sigma}} = U_i$, a.s., which implies $\mathbb{P}(Y_i^{\widehat{\sigma}} < U_i) = 0$.

Ad (ii): First note that $\widehat{\sigma}_i{}' \geq \widehat{\sigma}_i$ by construction. Hence, we may write

$$\begin{aligned} Y_i^{\widehat{\sigma}'} = \mathbb{E}_{\mathcal{G}_i} \left[U_{\widehat{\sigma}_i'} \right] &= \sum_{m=i}^{\infty} \mathbb{E}_{\mathcal{G}_i} \left[1_{\{\widehat{\sigma}_i = m\}} U_{\widehat{\sigma}_i'} \right] \\ \text{(using (3.1))} &= \sum_{m=i}^{\infty} \mathbb{E}_{\mathcal{G}_i} \left[1_{\{\widehat{\sigma}_i = m\}} U_{\widehat{\sigma}_m'} \right] = \sum_{m=i}^{\infty} \mathbb{E}_{\mathcal{G}_i} \left[1_{\{\widehat{\sigma}_i = m\}} Y_m^{\widehat{\sigma}'} \right] \\ &\geq \sum_{m=i}^{\infty} \mathbb{E}_{\mathcal{G}_i} \left[1_{\{\widehat{\sigma}_i = m\}} U_m \right] = \mathbb{E}_{\mathcal{G}_i} \left[U_{\widehat{\sigma}_i} \right] = Y_i^{\widehat{\sigma}}, \end{aligned}$$

where the already proved statement (i) is used. ■

Proof of Theorem 10. The fact that $\sigma_i^* < \infty$ follows from Proposition 3. Suppose that $\sigma_i^{(m)} > \sigma_i^*$ for some $m \in \mathbb{N}_0$. Then we have

$$\begin{aligned} U_{\sigma_i^*} &< \max_{\sigma_i^* \leq k \leq \sigma_i^* + \kappa} \mathbb{E}_{\mathcal{G}_{\sigma_i^*}} \left[U_{\sigma_k^{(m)}} \right] = \sum_{j=i}^{\infty} 1_{\{\sigma_i^* = j\}} \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m)}} \right] \\ \text{(by Proposition 8)} &\leq \sum_{j=i}^{\infty} 1_{\{\sigma_i^* = j\}} Y_j = Y_{\sigma_i^*}, \end{aligned}$$

which contradicts the optimality of σ_i^*. Suppose next that $\sigma_i^{(m+1)} < \sigma_i^{(m)}$ for some $m \in \mathbb{N}_0$. Then,

since $\sigma_i^{(0)} = i$, we must have that $m \geq 1$, and due to the definition of $\sigma_i^{(m)}$ that

$$\begin{aligned} U_{\sigma_i^{(m+1)}} &< \max_{\sigma_i^{(m+1)} \leq k \leq \sigma_i^{(m+1)} + \kappa} \mathbb{E}_{\mathcal{G}_{\sigma_i^{(m+1)}}} \left[U_{\sigma_k^{(m-1)}} \right] \\ &= \sum_{j=i}^{\infty} 1_{\left\{ \sigma_i^{(m+1)} = j \right\}} \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m-1)}} \right] . \end{aligned}$$

On the other hand, due to the definition of $\sigma_i^{(m+1)}$ we must have that

$$\begin{aligned} U_{\sigma_i^{(m+1)}} &\geq \max_{\sigma_i^{(m+1)} \leq k \leq \sigma_i^{(m+1)} + \kappa} \mathbb{E}_{\mathcal{G}_{\sigma_i^{(m+1)}}} \left[U_{\sigma_k^{(m)}} \right] \\ &= \sum_{j=i}^{\infty} 1_{\left\{ \sigma_i^{(m+1)} = j \right\}} \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m)}} \right] , \end{aligned}$$

which in turn yields

$$\sum_{j=i}^{\infty} 1_{\left\{ \sigma_i^{(m+1)} = j \right\}} \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m)}} \right] < \sum_{j=i}^{\infty} 1_{\left\{ \sigma_i^{(m+1)} = j \right\}} \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m-1)}} \right] .$$

However, since for all $j \geq i$,

$$\max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m)}} \right] \geq \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_j^{(m)}} \right] = Y_j^{(m)} \geq \max_{j \leq k \leq j+\kappa} \mathbb{E}_{\mathcal{G}_j} \left[U_{\sigma_k^{(m-1)}} \right] \quad \text{a.s.} ,$$

by Proposition 8, we get a contradiction. Thus, (3.5) is proved.

Due to (3.5) and Proposition 8 we may define

$$\sigma_i^{\infty} = \uparrow \lim_{m \to \infty} \sigma_i^{(m)} \quad \text{and} \quad Y_i^{\infty} = \uparrow \lim_{m \to \infty} Y_i^{(m)} \quad \text{a.s.}$$

Note that $\sigma_i^{\infty} \leq \sigma_i^* < \infty$, and that, since $\sigma_i^{(m)}$ is integer valued, one must have that $\sigma_i^{(m)}(\omega) = \sigma_i^{(\infty)}(\omega)$ for $m > N(i, \omega)$ say. Hence, $U_{\sigma_i^{(m)}} \to U_{\sigma_i^{\infty}}$ a.s., and then by dominated convergence we get

$$Y_i^{\infty} = \uparrow \lim_{m \to \infty} Y_i^{(m)} = \uparrow \lim_{m \to \infty} \mathbb{E}_{\mathcal{G}_i} \left[U_{\sigma_i^{(m)}} \right] = \mathbb{E}_{\mathcal{G}_i} \left[U_{\sigma_i^{\infty}} \right] .$$

Obviously, we have that $Y_i \geq Y_i^{\infty} \geq U_i$ a.s. Since (Y_i) is the smallest supermartingale that dominates the reward (U_i) by Proposition 3, we can prove that $\sigma_i^{\infty} = \sigma_i^*$ and $Y_i^{\infty} = Y_i$, and hence (3.6), by showing that (Y_i^{∞}) is a supermartingale. It is easy to see that (σ_i^{∞}) is a consistent $\mathcal{G}$-stopping family, i.e., (σ_i^{∞}) satisfies (3.1). So for any $i \geq 0$ one has

$$\begin{aligned} \mathbb{E}_{\mathcal{G}_i} \left[Y_{i+1}^{\infty} \right] &= \mathbb{E}_{\mathcal{G}_i} \left[U_{\sigma_{i+1}^{\infty}} \right] = Y_i^{\infty} + \mathbb{E}_{\mathcal{G}_i} \left[\left(U_{\sigma_{i+1}^{\infty}} - U_{\sigma_i^{\infty}} \right) \right] \\ &= Y_i^{\infty} + \mathbb{E}_{\mathcal{G}_i} \left[1_{\{\sigma_i^{\infty} = i\}} \left(U_{\sigma_{i+1}^{\infty}} - U_{\sigma_i^{\infty}} \right) \right] \\ (\text{by (3.1)}) &= Y_i^{\infty} + 1_{\{\sigma_i^{\infty} = i\}} \mathbb{E}_{\mathcal{G}_i} \left[U_{\sigma_{i+1}^{\infty}} \right] - 1_{\{\sigma_i^{\infty} = i\}} U_i . \end{aligned} \tag{A.16}$$

Note that by (3.5) and the definition of σ_i^∞ one has $\{\sigma_i^\infty = i\} = \bigcap_{m=0}^{\infty} \left\{\sigma_i^{(m)} = i\right\}$. Hence, by the definition of $\sigma_i^{(m')}$, for all $m' \geq 1$,

$$\begin{aligned} 1_{\left\{\sigma_i^\infty = i\right\}} U_i &\geq 1_{\left\{\sigma_i^\infty = i\right\}} \max_{i \leq k \leq i+\kappa} \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_k^{(m'-1)}}\right] \\ &\geq 1_{\left\{\sigma_i^\infty = i\right\}} \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_i^{(m'-1)}}\right] . \end{aligned}$$

Since $\uparrow \lim_{m' \uparrow \infty} \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_i^{(m'-1)}}\right] = \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_i^\infty}\right]$ a.s., it then follows that

$$1_{\left\{\sigma_i^\infty = i\right\}} U_i \geq 1_{\left\{\sigma_i^\infty = i\right\}} \mathbb{E}_{\mathcal{G}_i}\left[U_{\sigma_i^\infty}\right] ,$$

hence by (A.16) we get $\mathbb{E}_{\mathcal{G}_i}\left[Y_{i+1}^\infty\right] \leq Y_i^\infty$. That is, (Y_i^∞) is a super-martingale. ■

A.4 Finite-horizon approximation

Proof of Lemma 12. The first inequality is trivial. To prove the second inequality, we write:

$$\begin{aligned} Y_0^{(\infty)} &= \sup_{\mathfrak{n} \in \mathcal{N}} \mathbb{E}\left[U_{\mathfrak{n}} 1_{\{\mathfrak{n} \leq K\}} + U_{\mathfrak{n}} 1_{\{\mathfrak{n} > K\}}\right] \\ &\leq \sup_{\mathfrak{n} \in \mathcal{N}} \mathbb{E}\left[U_{\mathfrak{n} \wedge K} 1_{\{\mathfrak{n} \leq K\}}\right] + \sup_{\mathfrak{n} \in \mathcal{N}} \mathbb{E}\left[U_{\mathfrak{n} \vee K} 1_{\{\mathfrak{n} > K\}}\right] \\ &\leq \sup_{\mathfrak{n} \in \mathcal{N}, 0 \leq \mathfrak{n} \leq K} \mathbb{E}\left[U_{\mathfrak{n}}\right] + \sup_{\mathfrak{n} \in \mathcal{N}, \mathfrak{n} > K} \mathbb{E}\left[U_{\mathfrak{n}}\right] \leq Y_0^{(K)} + \sup_{\mathfrak{n} \in \mathcal{N}, \mathfrak{n} > K} \mathbb{E}\left[U_{\mathfrak{n}}\right] , \end{aligned}$$

by Proposition 1. Now, for any $\mathfrak{n}$ with $\mathfrak{n} > K$, we have

$$\begin{aligned} \mathbb{E}\left[U_{\mathfrak{n}}\right] &= \mathbb{E}\left[Z_{\tau_{\mathfrak{n}}} 1_{\{\tau_{\mathfrak{n}} \leq T\}}\right] + \mathbb{E}\left[Z_{\tau_{\mathfrak{n}}} 1_{\{\tau_{\mathfrak{n}} > T\}}\right] \\ &\leq \mathbb{E}\left[Z_{\tau_{\mathfrak{n}}} 1_{\{\tau_K \leq T\}}\right] + 0 \\ &\leq \mathbb{E}\left[1_{\{\tau_K \leq T\}} \sup_{i \in \mathbb{N}_0,\ i > K} Z_{\tau_i}\right] \\ &\leq B \mathbb{P}\left(\tau_K \leq T\right)^{1 - 1/\alpha} , \end{aligned}$$

by a Hölder estimate using (2.3). Because, by (2.1), $\tau_k \to \infty$ a.s. as $k \to \infty$, dominated convergence implies that $\mathbb{P}(\tau_K \leq T) \to 0$ as $K \to \infty$, $K \in \mathbb{N}$. Hence, to bound $\sup_{\mathfrak{n} \in \mathcal{N}, \mathfrak{n} > K} \mathbb{E}\left[U_{\mathfrak{n}}\right]$ by ε, it suffices to pick a value of K such that $\mathbb{P}(\tau_K \leq T) < (\varepsilon / B)^{\alpha/(\alpha-1)}$, with B and α specified in (2.3). ■

A.5 Markovian structure

The following results are readily obtained.

Lemma 24 *If Assumption 11 is fulfilled, then for any non-negative Borel measurable* $g : (t, x, \theta) \in$

$\mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^q \to g(t,x,\theta)$, *we have that*

$$\begin{aligned}\mathbb{E}_{\mathcal{F}_{\tau_k}}\left[g(\tau_{k+1}, X_{\tau_{k+1}}, \Theta_{\tau_{k+1}})\right] &= \mathbb{E}_{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})}\left[g(\tau_{k+1}, X_{\tau_{k+1}}, \Theta_{\tau_{k+1}})\right] \\ &= c(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\,,\end{aligned}$$

for some non-negative Borel function $c(\cdot,\cdot,\cdot)$.

Lemma 25 *Suppose that Assumption 11 applies,* X *and* Θ *are independent, and* $\tau_{k+1} - \tau_k$ *is independent of* Θ_{τ_k}. *Then we have that*

$$\mathbb{E}_{\mathcal{F}_{\tau_k}}\left[g(\tau_{k+1}, X_{\tau_{k+1}})\right] = c(\tau_k, X_{\tau_k})\,,$$

for some non-negative Borel function $c(\cdot,\cdot)$.

Proof. By Assumption 11,

$$\begin{aligned}\mathbb{E}_{\mathcal{F}_{\tau_k}}\left[g(\tau_{k+1}, X_{\tau_{k+1}})\right] &= \mathbb{E}_{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})}\left[g(\tau_k + \tau_{k+1} - \tau_k, X_{\tau_k + \tau_{k+1} - \tau_k})\right] \\ &= \mathbb{E}_{(\tau_k, X_{\tau_k})}\left[g(\tau_{k+1}, X_{\tau_{k+1}})\right] =: c(\tau_k, X_{\tau_k})\,.\end{aligned}$$

■

Example 26 *Take* X *independent of* Θ, *where* Θ *is a continuous-time Markov chain with state space* $\mathbb{N}_0$, *generator*

$$\begin{bmatrix} -\lambda_1 & \lambda_1 & 0 & \dots \\ 0 & -\lambda_2 & \lambda_2 & \dots \\ 0 & 0 & -\lambda_3 & \dots \\ \dots & \dots & \dots & \dots \end{bmatrix}, \quad \lambda_i > 0\,, \quad i = 1, 2, \dots\,,$$

and $\Theta_0 = 0$. *One thus has that* $\tau_k = \inf\{s \geq 0 : \Theta_s = k\}$.

Example 27 *As a more general example, consider* (X, Θ) *to be given by a jump-diffusion, more specifically, a strong Markov process of the form*

$$\begin{aligned}\mathrm{d}X_t &= \sigma(t, X_t) X_t \mathrm{d}W\,, \quad X_0 = x_0\,, \\ \mathrm{d}\Theta_t &= \int_{\mathbb{R}^q} z\mathcal{N}\left(X_t, \mathrm{d}t, \mathrm{d}z\right)\,, \quad \Theta_0 = 0\,,\end{aligned}$$

where $\mathcal{N}(x, \mathrm{d}t, \mathrm{d}z, \omega)$ *is some Poisson random measure on* $\mathbb{N}_0$, *independent of* W, *with*

$$\mathbb{P}\left(\mathcal{N}\left(x, (s,t], B\right) = k\right) := \exp(-v(x,B)\,(t-s))\frac{v(x,B)^k\,(t-s)^k}{k!}, \quad k \in \mathbb{N}_0\,,$$

where for simplicity $v(x,B)$ is of finite activity in the sense that

$$v(x,\{0\}) = 0 \quad \textit{and} \quad \int v(x,\mathrm{d}z) < \infty\,, \quad x \in \mathbb{R}^d\,.$$

The stopping times are determined by $\tau_0 = 0$ and $\tau_k = \inf\{s > 0 : \Theta_s = k\}$, $k > 0$.

(i): Simple case with $d = q = 1$ and $v(x,B) = \lambda\delta_1(B) := \lambda 1_B(1)$:

$$\mathbb{P}\left(\mathcal{N}\left(x,(0,t],B\right) = k\right) = \exp\left(-\lambda t\delta_1(B)\right)\frac{\lambda^k t^k \delta_1(B)^k}{k!}\,,$$

with $0^0 := 1$, and hence we get

$$\Theta_t = \int_{(0,t]}\int_{\mathbb{R}} z\mathcal{N}\left(\mathrm{d}s,\mathrm{d}z\right) = \mathcal{N}\left((0,t],\{1\}\right)\,, \quad \textit{where}$$
$$\mathbb{P}\left(\mathcal{N}\left((0,t],\{1\}\right) = k\right) = \exp\left(-\lambda t\right)\frac{\lambda^k t^k}{k!}\,,$$

i.e., $\Theta_t = \sharp$ jumps in the interval $(0,t]$ observed by the agent. With $\tau_0 = 0$ and $\tau_k = \inf\{s > 0 : \Theta_s = k\}$, $k > 0$, one has that $\tau_k - \tau_{k-1}$, $k \geq 1$, is $\mathrm{Exp}(\lambda)$ distributed and independent of τ_{k-1}. In particular, the conclusion of Lemma 24 applies.

(ii): Suppose $d = q = 1$ and that the Poisson random measure is only active during times that the process X is visiting a certain interval $I^{\pm} := [c_-, c_+]$ with $0 < c_- < c_+$. That is, we consider

$$v(x,B) := \lambda 1_{[c_-,c_+]}(x)\delta_1(B)\,, \quad \textit{hence}$$
$$\mathbb{P}\left(\mathcal{N}\left(x,(0,t],B\right) = k\right) = \exp\left(-\lambda t 1_{[c_-,c_+]}(x)\delta_1(B)\right)\frac{\lambda^k t^k 1_{[c_-,c_+]}(x)^k \delta_1(B)^k}{k!}$$

with $0^0 := 1$, and

$$\begin{aligned}\Theta_t &= \int_{(0,t]} 1_{[c_-,c_+]}(X_{s-})\int_{\mathbb{R}} z\mathcal{N}\left(X_{s-},\mathrm{d}s,\mathrm{d}z\right)\\ &= \int_{(0,t]} 1_{[c_-,c_+]}(X_s)\mathcal{N}_0\left(\mathrm{d}s,\{1\}\right)\,,\end{aligned}$$

with $\mathcal{N}_0$ as in (i). Note that X is continuous. So $\Theta_t = \sharp$ jumps in the interval $(0,t]$ while X is in the interval $[c_-,c_+]$, which is observable by the agent. Let $\tau_0 = 0$, and $\tau_k = \inf\{s > 0 : \Theta_s = k\}$, $k > 0$. Then, obviously, $\Theta_{\tau_k} = k$, $k \geq 0$, and

$$\begin{aligned}\tau_{k+1} - \tau_k &\in \sigma\left\{\left(X_s, \mathcal{N}_0\left((\tau_k, s],\{1\}\right)\right)1_{\{s \geq \tau_k\}}\right\}\\ &\subset \sigma\left\{\left(X_s,\Theta_s\right)1_{\{s\geq\tau_k\}}\right\}\,.\end{aligned}$$

Hence, in particular Assumption 11 applies, i.e., by the Markov property,

$$\begin{aligned}
&\mathbb{E}_{\mathcal{F}_{\tau_k}}\left[g(\tau_{k+1}, X_{\tau_{k+1}}, \Theta_{\tau_{k+1}})\right] \\
&= \mathbb{E}_{\mathcal{F}_{\tau_k}}\left[g(\tau_k + \tau_{k+1} - \tau_k, X_{\tau_k+\tau_{k+1}-\tau_k}, \Theta_{\tau_k+\tau_{k+1}-\tau_k})\right] \\
&= \mathbb{E}_{\mathcal{F}_{\tau_k}}\left[g(\tau_k + \tau_{k+1} - \tau_k, X_{\tau_k+\tau_{k+1}-\tau_k}, \Theta_{\tau_k+\tau_{k+1}-\tau_k})\right] \\
&= \mathbb{E}_{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})}\left[g(\tau_k + \tau_{k+1} - \tau_k, X_{\tau_k+\tau_{k+1}-\tau_k}, \Theta_{\tau_k+\tau_{k+1}-\tau_k})\right] \\
&= \mathbb{E}_{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})}\left[g(\tau_{k+1}, X_{\tau_{k+1}}, \Theta_{\tau_{k+1}})\right] \\
&=: c(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\,.
\end{aligned}$$

So, this is a simple example in which the continuation functions are dependent on the evaluation of both X and Θ.

A.6 Convergence of random times least-squares Monte Carlo

In the setting of Assumption 11, consider measurable functions

$$\psi_l : \mathbb{R}_{\geq 0} \times \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}\,, \quad l \in \mathbb{N}\,. \tag{A.17}$$

To establish convergence results, we require the following assumptions on the set of basis functions that are used in our regression:

Assumption 28 *For the collection of measurable functions in* (A.17)*, let P_k^L denote the L^2-projection of $\sigma\{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\}$-measurable random variables onto the subspace*

$$\mathcal{S}_L = \overline{\mathsf{span}\left\{\psi_l\left(\tau_k, X_{\tau_k}, \Theta_{\tau_k}\right),\ l = 1, \ldots, L\right\}}\,.$$

We assume the following. For each $k = 1, \ldots, K-1$:

1. *The sequence $\left(\psi_l\left(\tau_k, X_{\tau_k}, \Theta_{\tau_k}\right)\right)_{l\in\mathbb{N}}$ is total in $L^2\left(\Omega, \sigma\left\{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\right\}, \mathbb{P}\right)$, i.e., $\mathcal{S}_\infty = L^2\left(\Omega, \sigma\left\{(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\right\}, \mathbb{P}\right)$.*
2. *For $L \in \mathbb{N}$, if $\sum_{l=1}^L \lambda_l \psi_l\left(\tau_k, X_{\tau_k}, \Theta_{\tau_k}\right) = 0$ a.s., then $\lambda_l = 0$ for $l = 1, \ldots L$.*
3. *We have*

$$\mathbb{P}\left(\sum_{l=1}^L \alpha_{k,l}\psi_l\left(\tau_k, X_{\tau_k}, \Theta_{\tau_k}\right) = Z\left(\tau_k, X_{\tau_k}\right)\right) = 0\,,$$

where

$$\sum_{l=1}^L \alpha_{k,l}\psi_l\left(\tau_k, X_{\tau_k}, \Theta_{\tau_k}\right) = P_k^L\left(Z(\tau_{\mathfrak{n}_k^L}, X_{\tau_{\mathfrak{n}_k^L}})\right)\,,$$

and $\left(\mathfrak{n}_k^L\right)_{k=0,\ldots,K}$ is a sequence of discrete stopping times with values in $\{1, \ldots, K\}$ solving the

recursion $\mathfrak{n}_K^L = K$,

$$\mathfrak{n}_k^L = \begin{cases} k, & Z\left(\tau_k, X_{\tau_k}\right) \geq P_k^L\left(Z(\tau_{\mathfrak{n}_k^L}, X_{\tau_{\mathfrak{n}_k^L}})\right) \\ \mathfrak{n}_{k+1}^L, & Z_k < P_k^L\left(Z(\tau_{\mathfrak{n}_k^L}, X_{\tau_{\mathfrak{n}_k^L}})\right) \end{cases}, \quad 1 \leq k < K\,.$$

Under Assumptions 11 and 28, we have the following convergence result:

Theorem 29 *([13, Theorem 3.2]) Let* $\mathfrak{n}_k^*$ *be an optimal stopping index after the (physical) times* $\tau_1, \ldots, \tau_{k-1}$ *have passed. It then holds a.s. that*

$$\lim_{L\to\infty}\lim_{N\to\infty}\frac{1}{N}\sum_{n=1}^{N} Z(\widehat{\tau}^{(n)}_{\widehat{\mathfrak{n}}_k^{L,n}}, X^{(n)}_{\widehat{\tau}^{(n)}_{\widehat{\mathfrak{n}}_k^{L,n}}}) = \mathbb{E}\left[Z(\tau_{\mathfrak{n}_k^*}, X_{\tau_{\mathfrak{n}_k^*}})\right]\,.$$

[13] further implies that one may obtain a sequence $\widehat{C}_k^{(r)}(t,x,\theta)$, $r \in \mathbb{N}$, of approximate continuation functions from a sequence of LS-training procedures due to L_r basis functions and N_r training trajectories, such that

$$\widehat{C}_k^{(r)}(\tau_k, X_{\tau_k}, \Theta_{\tau_k}) \xrightarrow{P} C_k^*(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\,, \quad k = 1, \ldots, K\,, \tag{A.18}$$

for $r \to \infty$,, where C_k^* denotes the optimal continuation function at time k, and the process $(\tau_k, X_{\tau_k}, \Theta_{\tau_k})_{k=1,\ldots,K}$ is independent of the training procedures. Under the same assumptions, one next has the following convergence result on stopping times.

Theorem 30 *For each* $r \in \mathbb{N}$ *there exists an optimal stopping time* $(\mathfrak{n}^{*,r})$ *such that*

$$\mathbb{P}(\mathfrak{n}^{(r)} \neq \mathfrak{n}^{*,r}) \to 0 \text{ for } r \to \infty\,,$$

where

$$\mathfrak{n}^{(r)} := \min\left\{k \geq 0 : Z(\tau_k, X_{\tau_k}) \geq \widehat{C}_k^{(r)}(\tau_k, X_{\tau_k}, \Theta_{\tau_k})\right\}\,, \quad r \in \mathbb{N}\,.$$

If in particular $\mathfrak{n}^*$ *is unique, one has that* $\mathfrak{n}^{(r)} \xrightarrow{P} \mathfrak{n}^*$.

Proof. The result follows, in principle, from the convergence (A.18) of the continuation values. For a detailed proof in the context of robust optimal stopping, we refer to [34, Section 6]. ■

A.7 Solution for stopped geometric Brownian motion with jumps

Proof of Theorem 17. We begin by solving the problem for $K < \infty$ and will denote the value function at a time $t \in [0,T)$ by $v_K(t,x,d)$, where $x \in \mathbb{R}$ is the value of the state variable at time t, d is a binary variable that indicates whether we may stop ($d = 1$) or not ($d = 0$) at this given moment and K denotes the number of stopping opportunities remaining. This implies that if $\mathcal{T}_{K,t}$

denotes the set of the first K stopping times that occur strictly after time t, we can write

$$v_K(t,x,0) = \sup_{\tau\in\mathcal{T}_{K,t}} \mathbb{E}[Z_\tau|X_t = x], \qquad v_K(t,x,1) = \sup_{\tau\in(\mathcal{T}_{K,t}\cup\{t\})} \mathbb{E}[Z_\tau|X_t = x].$$

The values that are characterized in the theorem, $Y_i^{(K)}$ and $C_i^{(K)}$, are given by $v_K(\tau_i, X_{\tau_i}, 1)$ and $v_K(\tau_i, X_{\tau_i}, 0)$, respectively. By definition, $v_0 \equiv 0$. If we are allowed to stop at time t, i.e., $d = 1$, we either receive x^η or continue with the problem with $K-1$, so for $K \geq 1$,

$$v_K(t,x,1) = x^\eta 1_{\{t\leq T\}} \vee v_{K-1}(t,x,0)\,. \tag{A.19}$$

If we are not allowed to stop at time t, the value function at that time equals the discounted version of the value at the following possible stopping time, if this stopping times occurs before maturity:

$$v_K(t,x,0) = \int_0^{T-t} \lambda e^{-\lambda u}\mathbb{E}\left[v_K\left(t+u, x(1+j)e^{(\mu-\frac{1}{2}\sigma^2)u+\xi\sigma\sqrt{u}}, 1\right)\right]\,\mathrm{d}u\,, \tag{A.20}$$

with $\xi \sim \mathcal{N}(0,1)$.

We now make an *ansatz* to solve (A.19) and (A.20):

$$v_K(t,x,0) = x^\eta e^{\xi(T-t)} f_K(T-t)\,, \tag{A.21}$$

with ξ as defined in (5.3). We substitute (A.19) and (A.21) in (A.20) and see that it suffices to find a sequence of functions $(f_K)_{K\in\mathbb{N}_0}$ that satisfies

$$f_{K+1}(s) = \psi\int_0^s \left(e^{-\xi u} \vee f_K(u)\right)\,\mathrm{d}u\,, \tag{A.22}$$

where $\psi = \lambda(1+j)^\eta$ and $f_0 \equiv 0$. This immediately gives that, for all $s \geq 0$, we have $f_1(s) = \frac{\psi}{\xi}(1-e^{-\xi s})$. Let s^* be the solution to $f_1(s) = e^{-\xi s}$ so $s^* = \ln(1+\xi/\psi)/\xi$ and define $\zeta := f_1(s^*) = e^{-\xi s^*} = \psi/(\psi+\xi)$. We use a further ansatz

$$f_{K+1}(s) = f_1(s)1_{\{s\leq s^*\}} + \zeta\, z_{K+1}(s-s^*)1_{\{s>s^*\}}$$

and conclude that this is a solution to (A.22) for $K > 0$ if

$$\partial_s z_{K+1}(s) = \partial_s f_{K+1}(s+s^*) = \psi f_K(s+s^*) = \psi z_K(s), \qquad z_{K+1}(0) = 1,$$

while for $K = 0$ we have

$$z_1(s) = \frac{1}{\zeta}f_1(s+s^*) = \frac{\psi}{\zeta\xi}(1-e^{-\xi s}e^{-\xi s^*}) = \frac{\psi}{\zeta\xi}(1-\zeta e^{-\xi s}) = 1+\frac{\psi}{\xi}(1-e^{-\xi s}).$$

A power series expansion gives

$$z_{K+1}(s) = \sum_{m=0}^{K} \frac{(\psi s)^m}{m!} + (\frac{-\psi}{\xi})^{K+1} \sum_{m=K+1}^{\infty} \frac{(-\xi s)^m}{m!} \;=\; \frac{\Gamma_{K+1}(\psi s)}{K!} e^{\psi s} + \frac{\gamma_{K+1}(-\xi s)}{K!} (-\frac{\psi}{\xi})^{K+1} e^{-\xi s},$$

with Γ_{K+1} and γ_{K+1} the upper and lower Gamma functions of order $K+1$, respectively.

Combining this with our ansatz (A.21), we arrive at an expression for the value function in terms of $t^* := T - s^*$:

$$v_K(t, x, 0) = x^\eta e^{\xi(T-t)} \left(1_{\{t \geq t^*\}} f_1(T-t) + 1_{\{t < t^*\}} \zeta z_K(t^* - t) \right) . \tag{A.23}$$

For finite K, equations (A.19) and (A.21) show that it is optimal to stop if $1 \geq e^{\xi(T-t)} f_{K-1}(T-t)$, which we can rewrite as $e^{\xi s^*} f_{K-1}(s^*) \geq e^{\xi(T-t)} f_{K-1}(T-t)$ or $t \geq t^* = T - s^*$. For finite K, the expressions in the theorem for $C_i^{(K)}$ then follow by taking $t = \tau_i$ and $x = X_{\tau_i}$ in (A.21) and rewriting $\zeta e^{\xi(T-t)} = e^{\xi(t^*-t)}$.

Taking a limit $K \to \infty$ in (A.23), we find that

$$\lim_{K \to \infty} v_K(t, x, 0) = x^\eta \left(1_{\{t \geq t^*\}} e^{\xi(T-t)} f_1(T-t) + 1_{\{t < t^*\}} \zeta e^{\xi(T-t)} e^{\psi(t^*-t)} \right),$$

and $\lim_{K\to\infty} v_K(\tau_i, X_{\tau_i}, 0) = C_i^{(\infty)}$ follows from letting $\varepsilon \to 0$ in Lemma 12. It is clear that it is optimal to stop if and only if $t \geq t^*$ for this case as well. This concludes the proof. ■

B Additional Tables

λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Trivial	(s.e.)	Trivial-PI	(s.e.)
1	90	3.07696	(0.00578)	3.07707	(0.00006)	2.75785	(0.00516)	3.06875	(0.01232)
	100	6.04603	(0.00792)	6.04657	(0.00022)	5.32202	(0.00695)	6.06463	(0.01883)
	110	10.63174	(0.00986)	10.63399	(0.00065)	10.02295	(0.00903)	10.61661	(0.02016)
3	90	3.97787	(0.00619)	3.98025	(0.00074)	2.77611	(0.00414)	3.93131	(0.01836)
	100	7.46066	(0.00811)	7.46832	(0.00162)	4.96859	(0.00514)	7.38342	(0.02606)
	110	12.57331	(0.00951)	12.57662	(0.00069)	10.43446	(0.00659)	12.42274	(0.02737)
5	90	4.17705	(0.00621)	4.18727	(0.00236)	2.47309	(0.00343)	4.09642	(0.01923)
	100	7.73938	(0.00807)	7.75021	(0.00290)	4.34751	(0.00416)	7.54921	(0.02649)
	110	12.93490	(0.00932)	12.95111	(0.00460)	10.19641	(0.00546)	12.58117	(0.02676)

Table 10: Value of a max-call option on random times ($d = 1$).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5), with $d = 1$ independent geometric Brownian motion (without jumps, $\mu_J^j = \lambda_J^j = 0$) and with varying arrival rates λ and initial asset prices X_0. Column LSMC contains the lower-biased estimates based on the random times LSMC algorithm of Section 4.1 and column Dual contains the upper-biased estimates (4.11) based on the dual algorithm in Section 4.3 and the estimated LSMC policy; standard errors for the dual estimates are the standard errors of the duality gap (4.10). Both algorithms use the same settings as in Table 1, except for the basis functions: the continuation values in the LSMC algorithm are estimated by using the basis functions $(t, x) \mapsto t^i x^j$, $i + j \leq 5$. Column Trivial corresponds to stopping as soon as $Z_{\tau_i} > 0$ or $\tau_i \geq T$ and column Trivial-PI contains the results of the policy iteration method (3.4) with 100,000 outer paths, 500 inner paths and a window length $\kappa = K'$. All simulated paths have been generated in antithetic pairs.

λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Trivial	(s.e.)	Trivial-PI	(s.e.)
1	90	5.68421	(0.00765)	5.68804	(0.00116)	4.91301	(0.00661)	5.66093	(0.01842)
	100	10.55799	(0.01011)	10.55963	(0.00053)	8.68818	(0.00839)	10.55060	(0.02845)
	110	17.13788	(0.01224)	17.14160	(0.00096)	15.49008	(0.01041)	17.12895	(0.03247)
3	90	7.33534	(0.00818)	7.34869	(0.00221)	4.58246	(0.00498)	7.22310	(0.02632)
	100	12.99966	(0.01048)	13.02089	(0.00395)	7.24893	(0.00580)	12.77270	(0.03689)
	110	20.33221	(0.01235)	20.35280	(0.00243)	14.78042	(0.00707)	19.95245	(0.04102)
5	90	7.69256	(0.00826)	7.76684	(0.03866)	3.95536	(0.00403)	7.51463	(0.02730)
	100	13.46328	(0.01047)	13.51186	(0.00739)	6.04586	(0.00460)	13.04278	(0.03635)
	110	20.89358	(0.01222)	21.02209	(0.06126)	13.91635	(0.00568)	20.23306	(0.03970)

Table 11: Value of a max-call option on random times ($d = 2$).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5), with $d = 2$ independent geometric Brownian motions (without jumps, $\mu_J^j = \lambda_J^j = 0$) and with varying arrival rates λ and initial asset prices X_0. Column LSMC contains the lower-biased estimates based on the random times LSMC algorithm of Section 4.1 and column Dual contains the upper-biased estimates (4.11) based on the dual algorithm in Section 4.3 and the estimated LSMC policy; standard errors for the dual estimates are the standard errors of the duality gap (4.10). Both algorithms use the same settings as in Table 1, except for the basis functions: as in [5], we use basis functions in the ordered asset prices $X_t^{(1)} \geq X_t^{(2)}$. The continuation values in the LSMC algorithm are estimated by using the basis functions $(t, x^{(1)}, x^{(2)}) \mapsto t^i \left(x^{(1)}\right)^j \left(x^{(2)}\right)^k$, $i + j + k \leq 5$. Column Trivial corresponds to stopping as soon as $Z_{\tau_i} > 0$ or $\tau_i \geq T$ and column Trivial-PI contains the results of the policy iteration method (3.4) with 100,000 outer paths, 500 inner paths and a window length $\kappa = K'$. All simulated paths have been generated in antithetic pairs.

λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Trivial	(s.e.)	Trivial-PI	(s.e.)
1	90	3.37559	(0.00626)	3.37832	(0.00080)	3.03057	(0.00559)	3.37001	(0.01349)
	100	6.42173	(0.00846)	6.42563	(0.00128)	5.64728	(0.00739)	6.40339	(0.01983)
	110	10.98917	(0.01043)	10.99165	(0.00061)	10.30297	(0.00947)	11.00701	(0.02226)
3	90	4.36383	(0.00667)	4.36687	(0.00081)	3.03411	(0.00445)	4.32993	(0.01991)
	100	7.91895	(0.00863)	7.92280	(0.00141)	5.25481	(0.00547)	7.78849	(0.02775)
	110	13.01151	(0.01018)	13.01976	(0.00245)	10.62043	(0.00690)	12.82926	(0.02978)
5	90	4.57511	(0.00678)	4.58881	(0.00237)	2.70413	(0.00370)	4.47377	(0.02088)
	100	8.20940	(0.00865)	8.22161	(0.00203)	4.60046	(0.00443)	7.97062	(0.02823)
	110	13.37356	(0.01002)	13.39197	(0.00342)	10.32018	(0.00570)	13.01693	(0.02926)

Table 12: Value of a max-call option on random times with jumps ($d = 1$).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5) with $d = 1$ independent geometric Brownian motion, with jumps, $\mu_J^j = 0.06$ and $\lambda_J^j = 1$ and with varying arrival rates λ and initial asset prices X_0. The respective algorithms are carried out as described in the caption of Table 10.

λ	X_0	LSMC	(s.e.)	Dual	(s.e.)	Trivial	(s.e.)	Trivial-PI	(s.e.)
1	90	6.23507	(0.00828)	6.24121	(0.00197)	5.36950	(0.00713)	6.21446	(0.02009)
	100	11.21449	(0.01083)	11.22446	(0.00226)	9.20470	(0.00893)	11.21218	(0.03066)
	110	17.81153	(0.01306)	17.81885	(0.00166)	15.97752	(0.01098)	17.78208	(0.03489)
3	90	8.02699	(0.00881)	8.04256	(0.00331)	4.98881	(0.00536)	7.92356	(0.02882)
	100	13.80539	(0.01116)	13.82755	(0.00459)	7.66957	(0.00618)	13.58178	(0.03947)
	110	21.18165	(0.01310)	21.21136	(0.00554)	15.10441	(0.00747)	20.82862	(0.04389)
5	90	8.40994	(0.00884)	8.44444	(0.00501)	4.30081	(0.00434)	8.23108	(0.02953)
	100	14.30547	(0.01111)	14.35577	(0.00655)	6.39494	(0.00490)	13.95940	(0.03905)
	110	21.77564	(0.01296)	21.83676	(0.00848)	14.15769	(0.00600)	21.09714	(0.04272)

Table 13: Value of a max-call option on random times with jumps ($d = 2$).

Notes. This table displays estimated values of $Y_0^{(\infty)}$ in the optimal stopping problem for payoff (5.4) with strike price $\mathcal{K} = 100$ and independent Poissonian arrivals of stopping opportunities defined in (5.5) with $d = 2$ independent geometric Brownian motions, with jumps, $\mu_J^j = 0.06$ and $\lambda_J^j = 1$ and with varying arrival rates λ and initial asset prices X_0. The respective algorithms are carried out as described in the caption of Table 11.