

Robust Multiple Stopping — A Duality Approach

Roger J. A. Laeven*

University of Amsterdam

Mail: University of Amsterdam

Dept. of Quantitative Economics

PO Box 15867

1001 NJ Amsterdam, The Netherlands

Email: R.J.A.Laeven@uva.nl

John G. M. Schoenmakers

Weierstrass Institute Berlin

Mail: Weierstrass Institute Berlin

Stochastic Algorithms and Nonparametric Statistics

Mohrenstrasse 39

D-10117 Berlin, Germany

Email: schoenma@wias-berlin.de

Nikolaus Schweizer

Tilburg University

Mail: Tilburg University

Dept. of Econometrics and Operations Research

PO Box 90153

5000 LE Tilburg, The Netherlands

Email: N.F.F.Schweizer@tilburguniversity.edu

Mitja Stadje

Ulm University

Mail: Ulm University

Faculty of Mathematics and Economics

Helmholtzstrasse 20

D-89069 Ulm, Germany

Email: Mitja.Stadje@uni-ulm.de

This Version: November 13, 2023

Abstract

We develop a method to solve, theoretically and numerically, general optimal stopping problems. Our general setting allows for multiple exercise rights, i.e., optimal multiple stopping, for a robust evaluation that accounts for model uncertainty with a dominated family of priors, and for general reward processes driven by multi-dimensional jump-diffusions. Our approach relies on first establishing robust martingale dual representation results for the multiple stopping problem that satisfy appealing almost sure pathwise optimality properties. Next, we exploit these theoretical results to develop upper and lower bounds that, as we formally show, not only converge to the true solution asymptotically, but also constitute genuine pre-limiting upper and lower bounds. We illustrate the applicability of our approach in a few examples and analyze the impact of model uncertainty on optimal multiple stopping strategies.

Keywords: Optimal stopping; Multiple stopping; Robustness; Model uncertainty; Ambiguity; Pathwise duality; g -expectations; BSDEs; Regression.

AMS 2010 Subject Classification: *Primary:* 49L20; 60G40; 62L15; *Secondary:* 91B06; 91B16.

OR/MS Classification: Dynamic programming/optimal control: Models; Dynamic programming/optimal control: Applications; Decision analysis: Risk.

*Corresponding author.

1 Introduction

In this paper we analyze general optimal stopping problems of the following form:

$$Y_t^{*,L} := \sup_{\substack{t \leq \tau_1 < \dots < \tau_L \\ (\tau_1, \dots, \tau_L) \in \mathcal{T}_t(L)}} \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{t=1}^L H_{\tau_t} \middle| \mathcal{F}_t \right], \quad 0 \leq t \leq T, \quad (1.1)$$

where $\mathcal{T}_t(L)$ is a family of stopping time vectors, L is a number of exercise rights, $T < \infty$ is a fixed time horizon, \mathcal{Q} is a family of probabilistic models, and H is a general \mathcal{F}_t -adapted reward process. The operator \sup is to be understood as ess sup if it applies to an uncountable family of random variables. The optimal stopping problem (1.1) features generality along three dimensions: (i) it allows for optimal multiple stopping (when $L > 1$), (ii) it allows for a robust evaluation that explicitly takes probabilistic model uncertainty (i.e., ambiguity) into account (when \mathcal{Q} is not a singleton), and (iii) it allows for general reward processes that will be driven by multi-dimensional jump-diffusion processes. The process $Y_t^{*,L}$ is referred to as the upper Snell envelope of H due to L exercise rights after the seminal work of Snell [81]. Problems of this type, or special cases thereof, occur naturally in a wide variety of applications in probability, operations research, economics and finance.

Our aim is to develop upper and lower bounds on $Y_t^{*,L}$ that satisfy several desirable properties. We achieve this by first establishing suitable martingale dual representations for problem (1.1) that can be viewed as significant generalizations of the classical additive dual representations for *standard* (i.e., $\mathcal{Q} = \{\mathbb{Q}\}$ as opposed to *robust*) optimal stopping problems, developed independently by Rogers [70] and Haugh and Kogan [48] (see also the early Davis and Karatzas [34]) and their extension to standard multiple stopping problems in Schoenmakers [77]. A multiplicative dual representation for standard optimal stopping problems was proposed by Jamshidian [50]. Our dual representations take the form of an infimum over (robust) martingales, with no appearance of stopping times.

An appealing feature—both theoretically and for numerical stability—of the dual representations we establish is their (almost sure) pathwise optimality. Already when $L = 1$ these results are new and of independent interest for robust optimal single stopping. They are developed here in the general setting of robust optimal multiple stopping (1.1). Throughout, we assume a fixed reference measure \mathbb{P} to which the probabilistic models in \mathcal{Q} are absolutely continuous. Hence, our dual representation results hold \mathbb{P} -almost surely and not *quasi surely* (see e.g., Soner, Touzi and Zhang [82]). While bearing in mind this important distinction, we adopt the common practice (e.g., Rogers [71]) of using the term “pathwise” to indicate the pathwise maxima over the exercise dates in the dual representation results we establish. The almost sure nature of the dual representations suggests that finding a ‘good’ martingale that is ‘close’ to a ‘surely optimal’ martingale will yield tight and nearly constant upper bounds. The target can be the unique (robust) Doob martingale, to be constructed from an approximation to the upper Snell envelope or, more generally, a martingale for which the dual representation’s infimum is attained and the almost sure property is satisfied. While this phenomenon of tightness and constancy is known in the case of standard, non-robust single stopping problems (i.e., when $\mathcal{Q} = \{\mathbb{Q}\}$ and $L = 1$, see Rogers [72] and Schoenmakers, Zhang and Huang [78]), we will analyze it in our general setting of robust stopping. We will show in particular that a low (vanishing in probability) robust variance implies a tight (converging in L^1) approximation. The mathematical details of these results are delicate.

These new theoretical results justify and enable us to next develop a numerically implementable method to obtain upper and lower bounds to $Y_t^{*,L}$ with desirable properties. Our lower bound, derived from the proposed exercise strategy, will, as we formally show, not only converge to the optimal solution asymptotically but also be ‘biased low’ at the pre-limiting level in a Brownian-Poisson filtration. This is not the case for the initially proposed upper bound: it converges to the true solution but is not in general ‘biased high’. We therefore also develop a second upper bound that as we prove both converges to the true solution asymptotically and is biased high in a Brownian-Poisson filtration. It is based on a Lipschitzian L^2 -approximation, with a Lipschitz constant that we explicitly derive, and a suitable (reversed) application of Jensen’s inequality. We will refer to this second upper bound as our *genuine* upper bound. The computational complexity of both upper bounds is only linear in the number of exercise rights, and our method does not require nested simulation.

We provide extensive numerical examples, including single and multiple stopping problems, univariate and multivariate stochastic drivers, increasing and decreasing reward functions, and pure diffusion and jump-diffusion models, to illustrate the applicability and generality of our approach. They demonstrate that our approach yields tight upper and lower bounds that, due to almost sure properties, moreover have low standard errors. They also analyze the impact of multiple vs. single stopping rights and reveal that employing a robust evaluation that takes ambiguity into account is highly relevant for optimal stopping, especially in the presence of multiple exercise rights.

When embedded in a Brownian-Poisson filtration, the problems we analyze are naturally represented as stopping problems with respect to g -expectations (Peng [65, 66]), leading to backward stochastic differential equations (BSDEs). Hence, as a contribution to the BSDE literature of independent interest, we explicitly construct novel genuine upper and lower bounds to BSDE solutions with positively homogeneous convex drivers in a Brownian-Poisson filtration. Bender, Schweizer and Zhuo [16], when analyzing solutions to discrete-time (reflected) BSDEs rather than the continuous-time BSDEs we consider, develop upper and lower bounds applying techniques different from the techniques we employ. Bender, Gärtner and Schweizer [17] construct Monte Carlo upper and lower bounds for a class of discrete-time stochastic dynamic programs which includes discretizations of multiple stopping problems. Our genuine upper and lower bounds apply directly to the original continuous-time problem. Our genuine lower bound takes advantage of an almost sure property of a ‘second kind’ that we formally establish in order to reduce its variance—‘second kind’ to distinguish it from the additive dual representation’s almost sure property. This almost sure property entails that the difference between the BSDEs terminal condition and the associated (robust) martingale is constant almost surely. Our genuine upper bound for the continuous-time problem is based on forward simulation of an approximate BSDE solution. The construction is somewhat related to the a posteriori criterion for error evaluation introduced in Bender and Steiner [14] in a Brownian filtration, and developed here to obtain explicit genuine upper bounds for BSDEs in a Brownian-Poisson filtration.

The development of numerically implementable methods to obtain approximations to problems of the type (1.1) but with \mathcal{Q} a singleton (no ambiguity), $L = 1$ (single stopping), and with H multi-dimensional but satisfying strong conditions, started with the regression-based Monte Carlo methods of Carriere [28] and Longstaff and Schwartz [59]; see also Tsitsiklis and Van Roy [85] and Clément, Lambertson and Protter [32]. These methods yield lower bounds to $Y_t^{*,1}$ by approximating the optimal stopping time using regression and are commonly referred to as

“primal” approaches. An important example of a non-regression based primal approach is the stochastic mesh method of Broadie and Glasserman [22] (see, for further details, Glasserman [43] and also Belomestny, Kaledin and Schoenmakers [13]). “Dual” algorithms that exploit additive dual representations to numerically compute upper bounds were first proposed by Andersen and Broadie [1] in the standard single stopping problem and were further developed by e.g., Belomestny, Bender and Schoenmakers [10] to allow for non-nested simulation. While primal methods rely in a sense on constructing an appropriate stopping time, dual methods rely on constructing an appropriate martingale. Brown, Smith and Sun [23] in an innovative paper enlarge the information on which an exercise decision may depend in dual optimization, yielding tight upper bounds.

Model uncertainty, and the distinction between risk and ambiguity, has received much attention in recent years. Under the Bayesian paradigm, as adopted in Savage’s [76] subjective expected utility model, this distinction is, in a sense, nullified, through subjective probabilities resulting from a subjective prior probability over probabilistic models that quantifies model uncertainty. A popular approach beyond the Bayesian paradigm is provided by the multiple priors model of Gilboa and Schmeidler [42], which is a decision-theoretic formalization of the classical Waldian maxmin decision rule (Wald [86]; see also Huber [49]) and experimentally motivated by the Ellsberg [38] paradox. These models are intimately related to coherent, convex and entropy convex measures of risk in financial risk measurement (Föllmer and Schied [39, 40], Frittelli and Rosazza Gianin [41], Ruszczyński and Shapiro [73, 74], and Laeven and Stadje [55]). They explicitly recognize that probabilistic models may be misspecified and are often referred to as *robust* approaches (Hansen and Sargent [47]). The literature on robust single stopping theory is rapidly growing; it includes Riedel [68], Krätschmer and Schoenmakers [53], Bayraktar, Karatzas and Yao [5], Bayraktar and Yao [6], Cheng and Riedel [29], Øksendal, Sulem and Zhang [64], Belomestny and Krätschmer [11, 12], Bayraktar and Yao [7, 8, 9], Ekren, Touzi, and Zhang [37], Matoussi, Piozin, and Possamaï [61], Matoussi, Possamaï, and Zhou [62], and Nutz and Zhang [63]. However, numerically implementable methods to solve general optimal stopping problems of the form (1.1) have not been well-developed as yet. Krätschmer *et al.* [54] propose a numerically implementable method for single stopping problems under uncertainty in drift and jump intensity. Their approach is dual but not pathwise, i.e., it does not rely on a dual representation with the appealing (almost sure) pathwise optimality property, and cannot handle multiple stopping problems.

The multiple stopping problem can be viewed as L nested single stopping problems, where the decision-maker first chooses between stopping at time τ_1 on the one hand, thus collecting the reward and entering into a new contract with $L-1$ exercise rights, and retaining L exercise rights on the other hand, and so on. Multiple exercise rights occur naturally in many applications across various fields. For example, in environmental economics, a swing option gives the investor the right to change his purchased energy quantity a number of times per time period; in finance, a flexible interest rate cap gives the investor the right to exercise at each interest rate reset date a number of times over the life of the contract; and in insurance, a partial surrender option provides a payoff to the policyholder each time he partially surrenders his life insurance contract; see e.g., Carmona and Dayanik [24] and Carmona and Touzi [25] and the references therein. Kobylanski, Quenez and Rouy-Mironescu [52] analyze the standard multiple stopping problem (without ambiguity) allowing the payoff to be a general functional of an ordered sequence of stopping times. Bender, Schoenmakers and Zhang [15] develop a dual approach to generalized multiple stopping problems with respect to standard conditional expectations that

is intimately related to the information relaxation approach of Brown, Smith and Sun [23]. A primal-dual algorithm for standard multiple stopping with respect to standard conditional expectations in the context of flexible interest rate caps has been proposed in Balder, Mahayni and Schoenmakers [3].

As an important application, our approach may be used for robust no-arbitrage pricing (Hansen and Jagannathan [46], Cochrane and Saá-Requejo [33]) of American-style derivatives with possibly multiple exercise rights, via superhedging. This entails a significant advancement of the standard approach, where in a usually incomplete market the corresponding stopping problem is solved with respect to an arbitrarily chosen (local equivalent martingale) measure. Our results can also be applied to indifference valuation (seller's perspective; Carmona [26], Laeven and Stadje [56]) of general optimally stopped reward processes under the multiple priors model. Another application is that of robust risk measurement (Ben-Tal and Nemirovski [18], Bertsimas and Brown [19], Föllmer and Schied [40]) to determine e.g., the risk capital required to cover optimally stopped reward processes.

From an abstract perspective, single and multiple stopping problems bear similarities with reinforcement learning problems and Markov decision processes (in a Markovian setting). There exists by now a large literature on reinforcement learning and robust optimization in machine learning and operations research. Markov decision processes in which the parameters are uncertain are studied in Xu and Mannor [88], Wiesemann, Kuhn and Rustem [87], Grand-Clément and Kroer [44] and Liu *et al.* [58]; see also the references therein. The duality approach developed in this paper may therefore also prove to be fruitful in this related area of research.

The remainder of this paper is organized as follows. In Section 2 we recall some basic notions, establish some general properties, introduce the robust optimal multiple stopping problem, and provide some examples. In Section 3, we present our pathwise dual representations and establish our results on surely optimal (robust) martingales. In Section 4, we outline a general primal-dual algorithm and prove its convergence. Section 5 presents explicit upper and lower bounds in a Brownian-Poisson filtration. Section 6 provides extensive numerical results. All proofs and several auxiliary results are in the Online Appendix.

2 Robust Optimal Multiple Stopping

2.1 Basic Notions and General Properties

We start by considering a general stochastic setup. We let $(\Omega, (\mathcal{F}_t)_{t \in \{0, \dots, T\}}, \mathbb{P})$ be a filtered probability space and let \mathfrak{X} be a linear subspace of $L^0(\Omega, \mathcal{F}, \mathbb{P})$ with $\mathcal{F} := \mathcal{F}_T$. We further assume that \mathfrak{X} has a lattice structure, i.e., \mathfrak{X} is closed under the operations \wedge (min) and \vee (max), and that \mathfrak{X} contains all indicator functions 1_A , $A \in \mathcal{F}$. We emphasize that we work throughout under a given reference probability measure \mathbb{P} . (In)equalities between random variables are understood in the \mathbb{P} -almost sure sense, often without explicit mention. Furthermore, convergences in L^1 and L^2 are understood as $L^1(\mathbb{P})$ - and $L^2(\mathbb{P})$ -convergences.

To represent preferences, we consider a family of mappings $\rho := (\rho_t)_{t=0, \dots, T}$,

$$\rho_t : \mathfrak{X} \rightarrow \mathfrak{X} \cap L^0(\Omega, \mathcal{F}_t, \mathbb{P}).$$

It is referred to as a monotone, regular, recursive, conditional translation invariant *dynamic monetary utility functional*, henceforth *DMU* for short, if it satisfies the following conditions:

- (C1) $\rho_t(X) \leq \rho_t(Y)$ for all $X, Y \in \mathfrak{X}$ with $X \leq Y$ and $t \in \{0, \dots, T\}$ (*monotonicity*).
- (C2) $\rho_t(1_A X) = 1_A \rho_t(X)$ for all $X \in \mathfrak{X}$, $A \in \mathcal{F}_t$ and $t \in \{0, \dots, T\}$ (*regularity*).
- (C3) $\rho_t = \rho_t \circ \rho_{t+1}$ for all $t \in \{0, \dots, T-1\}$ (*recursiveness*).
- (C4) $\rho_t(X + Y) = \rho_t(X) + Y$ for all $X, Y \in \mathfrak{X}$ with $Y \in \mathcal{F}_t$ and $t \in \{0, \dots, T\}$ (*conditional translation invariance*).

As additional properties we consider:

- (P1) $\rho_t(X + Y) \leq \rho_t(X) + \rho_t(Y)$ for all $X, Y \in \mathfrak{X}$ and $t \in \{0, \dots, T\}$ (*subadditivity*).
- (P2) $[X \leq 0 \text{ and } \rho_t(X) \geq 0] \implies X = 0$, for all $X \in \mathfrak{X}$ and $t \in \{0, \dots, T\}$ (*sensitivity*).
- (P3) $\rho_t(\lambda X) = \lambda \rho_t(X)$ for all $X \in \mathfrak{X}$, $\lambda \geq 0$ and $t \in \{0, \dots, T\}$ (*positive homogeneity*).

Conditions (C1)–(C4) will always be assumed. In the sequel, we will mention explicitly which of the properties (P1)–(P3) is required. Properties (P1) and (P2) also entail the implication $[X \geq 0 \text{ and } \rho_t(X) \leq 0] \implies X = 0$; see Lemma A.1 in Appendix A.1. DMUs that satisfy (P1)–(P3), in addition to (C1)–(C4), take the form of robust, or worst case, expectations and have been widely used in applied probability, operations research, economics and finance; see the references in the Introduction and Section 2.3 below.

In this paper, we will frequently use the following implications of (C2) and (C4):

- (C5) $\rho_t(0) = 0$ for all $t \in \{0, \dots, T\}$ (*normalization*).
- (C6) $\rho_t(X) = X$ for all $X \in \mathfrak{X}$ with $X \in \mathcal{F}_t$ and $t \in \{0, \dots, T\}$ (\mathcal{F}_t -*invariance*).

Let \mathcal{H} be the set of adapted processes $(U_t)_{t \in \{0, \dots, T\}}$ such that $U_t \in \mathfrak{X} \cap L^0(\Omega, \mathcal{F}_t, \mathbb{P})$. A process $M = (M_t)_{t \in \{0, \dots, T\}} \in \mathcal{H}$ is said to be a ρ -martingale if

$$M_t = \rho_t(M_{t+1}), \quad 0 \leq t < T. \quad (2.1)$$

We present two auxiliary lemmas. The first lemma provides a generalization of Doob's optional sampling theorem towards our setup:

Lemma 2.1 (*Doob*) *Suppose ρ satisfies (C1)–(C4). Then, for any ρ -martingale M and any stopping time τ_i , $i \leq \tau_i \leq T$, it holds that $\rho_i(M_{\tau_i}) = M_i$, $0 \leq i \leq T$.*

Due to the next lemma, the properties of recursiveness (C3) and conditional translation invariance (C4) carry over to stopping times, as we will exploit later:

Lemma 2.2 *Let ρ satisfy (C1)–(C4), and let $t \in \{0, \dots, T\}$ be fixed. Consider, for any stopping time τ , $t \leq \tau \leq T$, the functional*

$$\rho_\tau(X) := \sum_{j=t}^T 1_{\{\tau=j\}} \rho_j(X).$$

Then, ρ_τ acts from $\mathcal{F}_T \rightarrow \mathcal{F}_\tau \supset \mathcal{F}_t$, and

- (i) ρ_τ satisfies $\rho_t = \rho_t \circ \rho_\tau$;

(ii) $\rho_\tau(X + Y) = X + \rho_\tau(Y)$, for $X \in \mathcal{F}_\tau$, $Y \in \mathcal{F}_T$.

Remark 2.3 We note that common g -expectations (with a positively homogeneous Lipschitz-continuous driver g ; see Jiang [51]) satisfy our axioms (C1)–(C4) and (P1)–(P3). Moreover, in a continuous-time setting, if our axioms hold for all $t \in [0, T]$ (and an additional domination or weak compactness condition is valid), then ρ is necessarily a g -expectation. However, in principle ρ only needs to be defined in discrete time. In discrete time, the class of risk measures satisfying (C1)–(C4) and (P1)–(P3) is much larger than the class of g -expectations. Specifically, every conditional coherent one-step risk measure, say $\bar{\rho}_t : L^2(\mathcal{F}_{t+1}) \rightarrow L^2(\mathcal{F}_t)$, gives rise to a dynamic risk measure, $(\rho_t)_{t=0,1,\dots,T} : L^2(\mathcal{F}_T) \rightarrow L^2(\mathcal{F}_t)$, satisfying (C1)–(C4) and (P1)–(P3), by defining recursively $\rho_T(X) = X$ and $\rho_t(X) := \bar{\rho}_t(\rho_{t+1}(X))$; see Cheridito, Delbaen and Kupper [30].

2.2 The Stopping Problem

Consider a fixed adapted reward, or (discounted) cash-flow, process $H = (H_t)_{t \in \{0, \dots, T\}} \in \mathcal{H}$ and a DMU decision-maker with L exercise rights that have to be exercised at different exercise dates. For each fixed t and L , $0 \leq t \leq T$, let $\mathcal{T}_t(L)$ be the family of stopping vectors (τ_1, \dots, τ_L) such that $\tau_1 \geq t$ and $\tau_l \geq \tau_{l-1} + 1$ for all l , $1 < l \leq L$. The decision-maker faces the following robust optimal multiple stopping problem:

$$Y_t^{*,L} := \operatorname{ess\,sup}_{t \leq \tau_1 < \tau_2 < \dots < \tau_L} \rho_t \left(\sum_{l=1}^L H_{\tau_l} \right), \quad t \in \{0, \dots, T\}, \quad (2.2)$$

for a DMU functional ρ that satisfies (C1)–(C4). We note that problem (2.2) is even slightly more general than problem (1.1), which arises when additionally (P1)–(P3) are satisfied. Henceforth, we write \sup (and \inf) instead of $\operatorname{ess\,sup}$ (and $\operatorname{ess\,inf}$) for convenience, understanding that they apply to an uncountable family of random variables. For a clean formulation of the multiple stopping problem (2.2), we extend the cash-flow process by setting $H_j \equiv 0$ and $\mathcal{F}_j \equiv \mathcal{F}_T$, for $j = T + 1, T + 2, \dots$. That is, the subset of rights l , $l = 2, \dots, L$, not exercised by time T becomes valueless. Hence, for any ρ -martingale M , $M_j = M_T$, $j > T$.

When $L \equiv 1$, the single stopping problem

$$Y_t^* \equiv Y_t^{*,1} = \sup_{\tau \in \mathcal{T}_t} \rho_t(H_\tau), \quad t \in \{0, \dots, T\}, \quad (2.3)$$

occurs as a special case, where the family of stopping times $\mathcal{T}_t \equiv \mathcal{T}_t(1)$ takes values in the set $\{t, \dots, T\}$.

The multiple stopping problem can be viewed as L nested single stopping problems with only a single exercise right. Indeed, setting $Y^{*,0} \equiv 0$, $Y^{*,1} \equiv Y^*$ is the upper Snell envelope of H due to a single exercise right. Then, for multiple exercise rights $L \geq 1$, $Y^{*,L}$ can be viewed as the upper Snell envelope of the process

$$H_t + \rho_t \left(Y_{t+1}^{*,L-1} \right), \quad t \in \{0, \dots, T-1\},$$

due to only a single exercise right.

Let us denote the set of ρ -martingales M with $M_0 = 0$ by \mathcal{M}_0^ρ . There exists a unique ρ -martingale $M^{*\rho} \in \mathcal{M}_0^\rho$ and a non-decreasing predictable $A^{*\rho} \in \mathcal{H}$ such that

$$Y_t^* = Y_0^* + M_t^{*\rho} - A_t^{*\rho}, \quad t \in \{0, \dots, T\}, \quad (2.4)$$

which represents the ρ -Doob decomposition of $Y^* = (Y_t^*)_{t \in \{0, \dots, T\}}$. It is easy to verify that, for $t \in \{0, \dots, T-1\}$,

$$M_{t+1}^{*\rho} - M_t^{*\rho} = Y_{t+1}^* - \rho_t(Y_{t+1}^*), \quad \text{and} \quad A_{t+1}^{*\rho} - A_t^{*\rho} = Y_t^* - \rho_t(Y_{t+1}^*). \quad (2.5)$$

Henceforth, the ρ -martingale $M^{*\rho}$ will often be referred to as the ρ -Doob martingale and we often suppress its superscript ρ to simplify notation.

In Appendix A.2, we establish some auxiliary results for problem (2.3) that will be exploited in the proofs of the results that follow.

2.3 Examples

We provide the following examples in which specific versions of the robust optimal multiple stopping problem of the general form (2.2) occur naturally. The probabilistic models in the sets \mathcal{Q} considered below are absolutely continuous with respect to the reference probability measure \mathbb{P} .

- (A.) No-arbitrage pricing: Let \mathcal{Q} be the set of local equivalent martingale measures. (Only if markets are complete \mathcal{Q} is a singleton, i.e., $\mathcal{Q} = \{\mathbb{Q}\}$.) Then, the *superhedging* price π^L of a contract with $L \geq 1$ exercise rights and associated payoff $\sum_{l=1}^L H_{\tau_l}$ is given by

$$\pi^L = \sup_{\tau_1 < \tau_2 < \dots < \tau_L} \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{l=1}^L H_{\tau_l} \right].$$

Many different approaches to no-arbitrage pricing have been proposed in the literature; see, e.g., the good-deal bounds of Cochrane and Saá-Requejo [33], Hansen and Jagannathan [46] and Björk and Slinko [20], or the acceptable opportunities of Carr, Geman and Madan [27]. All these approaches yield prices of the form

$$\tilde{\pi}^L = \sup_{\tau_1 < \tau_2 < \dots < \tau_L} \sup_{\mathbb{Q} \in \mathcal{Q}_{\text{restricted}}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{l=1}^L H_{\tau_l} \right],$$

where $\mathcal{Q}_{\text{restricted}} \subset \mathcal{Q}$.

Prototypical situations leading to single and multiple stopping problems in economics and finance are the pricing and exercising of American-style, Bermudan-style, and swing options. American options give the holder the right to exercise the option (once) on any preferred trading day before expiration. Different from American options, Bermudan options prescribe a set of trading days on which the option can be exercised (once). Swing options, more generally, give the holder the right to exercise the option multiple times, at a pre-specified set of exercise dates. With $L \geq 1$ exercise rights, exercised at $\tau_1 < \tau_2 < \dots < \tau_L$, the payoff equals $\sum_{l=1}^L H_{\tau_l}$ for a cash-flow process $H \in \mathcal{H}$. Swing options are particularly popular in energy markets to manage the risk of fluctuations in oil, gas, or electricity prices.

- (B.) Indifference valuation—the seller’s perspective: Suppose that the seller of a contract has a max-min utility functional of the form

$$U(H) = \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}}[H],$$

for a family of probabilistic models (i.e., priors) \mathcal{Q} and adopts a utility indifference valuation approach (Carmona [26], Laeven and Stadje [56]). Then, the value V^L of a contract with $L \geq 1$ exercise rights and associated payoff $\sum_{l=1}^L H_{\tau_l}$ is determined from the indifference relation

$$U(0) = \inf_{\tau_1 < \tau_2 < \dots < \tau_L} U \left(- \sum_{l=1}^L H_{\tau_l} + V^L \right) = \inf_{\tau_1 < \tau_2 < \dots < \tau_L} \inf_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}} \left[- \sum_{l=1}^L H_{\tau_l} + V^L \right].$$

Hence,

$$V^L = \sup_{\tau_1 < \tau_2 < \dots < \tau_L} \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{l=1}^L H_{\tau_l} \right].$$

(C.) Robust risk measurement: Suppose that ρ is a robust, or worst case, expectation, that is,

$$\rho(H) = \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}}[H], \quad (2.6)$$

for a family of probabilistic models \mathcal{Q} . In financial risk measurement, (2.6) is referred to as a coherent risk measure and \mathcal{Q} as a set of generalized scenarios (Artzner *et al.* [2] and Föllmer and Schied [40]); see also Ben-Tal and Nemirovski [18] for the intimately connected robust optimization paradigm. It determines the minimal amount of risk capital required to be added to the financial position H to make it ‘safe’ from the viewpoint of the regulatory authority. Applications of coherent risk measures and generalized scenarios to decision and optimization include Lesnevski, Nelson and Staum [57], Bertsimas and Brown [19], Choi, Ruszczyński and Zhao [31], Philpott, de Matos and Finardi [67] and Tekaya, Shapiro, Soares and da Costa [84]. Assume now that H_{τ_l} is a payout obligation (i.e., liability) at time τ_l , where τ_l is a stopping time, due to e.g., a flexible interest rate cap in interest rate markets or a partial surrender option in life insurance to be paid to a policyholder who decides to partially surrender his insurance contract. Then, the required amount of risk capital due to $L \geq 1$ stopping rights is given by

$$\sup_{\tau_1 < \tau_2 < \dots < \tau_L} \rho \left(\sum_{l=1}^L H_{\tau_l} \right) = \sup_{\tau_1 < \tau_2 < \dots < \tau_L} \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{l=1}^L H_{\tau_l} \right].$$

3 Pathwise Duality

3.1 Pathwise Dual Representation

The following theorem establishes our (almost sure) pathwise additive dual representation for general multiple stopping problems of the form (2.2). We emphasize that the pathwise dual representation holds \mathbb{P} -almost surely, with \mathbb{P} the reference measure.

Theorem 3.1 *Suppose ρ satisfies (C1)–(C4) and is subadditive (P1). Then, for any adapted process $H = (H_t)_{t \in \{0, \dots, T\}} \in \mathcal{H}$ and each fixed $t \in \{0, \dots, T\}$,*

(i) *we have the dual representation*

$$Y_t^{*,L} = \inf_{M^{(1)}, \dots, M^{(L)} \in \mathcal{M}_0^{\rho}} \rho_t \left(\max_{t \leq j_1 < j_2 < \dots < j_L} \sum_{k=1}^L \left(H_{j_k} + M_{j_{k-1}}^{(k)} - M_{j_k}^{(k)} \right) \right); \quad (3.1)$$

(ii) the dual representation's infimum is attained:

$$Y_t^{*,L} = \rho_t \left(\max_{t \leq j_1 < j_2 < \dots < j_L} \sum_{k=1}^L \left(H_{j_k} + M_{j_{k-1}}^{*,L-k+1} - M_{j_k}^{*,L-k+1} \right) \right); \quad (3.2)$$

(iii) if in addition ρ is sensitive (P2), we have the pathwise dual representation

$$Y_t^{*,L} = \max_{t \leq j_1 < j_2 < \dots < j_L} \sum_{k=1}^L \left(H_{j_k} + M_{j_{k-1}}^{*,L-k+1} - M_{j_k}^{*,L-k+1} \right), \quad \text{almost surely}; \quad (3.3)$$

where the ρ -martingales $M^{*,L-k+1}$ satisfy

$$M_{r+1}^{*,L-k+1} - M_r^{*,L-k+1} = Y_{r+1}^{*,L-k+1} - \rho_r \left(Y_{r+1}^{*,L-k+1} \right), \quad (3.4)$$

and $Y^{*,L-k+1}$ is the upper Snell envelope due to $L-k+1$ exercise rights, satisfying the Bellman principle,

$$Y_r^{*,L-k+1} = \max \left[H_r + \rho_r \left(Y_{r+1}^{*,L-k} \right), \rho_r \left(Y_{r+1}^{*,L-k+1} \right) \right]. \quad (3.5)$$

Already when $L \equiv 1$, Theorem 3.1 is new and of significant independent interest. In this case it simplifies to:

Corollary 3.2 *Suppose ρ satisfies (C1)–(C4) and (P1). Then, for any adapted process $H \in \mathcal{H}$ and each fixed $t \in \{0, \dots, T\}$, we have the dual representation*

$$Y_t^* = \inf_{M \in \mathcal{M}_0^\rho} \rho_t \left(\max_{t \leq j \leq T} (H_j + M_t - M_j) \right) \quad (3.6)$$

$$= \rho_t \left(\max_{t \leq j \leq T} (H_j + M_t^* - M_j^*) \right), \quad (3.7)$$

where M^* is the ρ -Doob martingale in Eqn. (2.4). If in addition ρ is sensitive (P2), we have the almost sure property:

$$Y_t^* = \max_{t \leq j \leq T} (H_j + M_t^* - M_j^*), \quad \text{almost surely}. \quad (3.8)$$

Remark 3.3 *We note that the single stopping problem also admits an alternative but non-pathwise additive dual representation for functionals ρ satisfying (C1)–(C4); see Proposition A.2 in the Appendix. In Theorem 3.1 and Corollary 3.2, the subadditivity property (P1) of ρ is required, and exploited through application of Lemma B.1 in the proof of Theorem 3.1. A dual representation theorem in the spirit of Theorem 3.1 and Corollary 3.2 without assuming (P1) seems not possible to us.*

3.2 Surely Optimal ρ -Martingales

The ρ -Doob martingale in Eqn. (2.4) plays a special role in the set of ρ -martingales \mathcal{M}_0^ρ as its appearance in Corollary 3.2, Eqn. (3.7) (and indirect appearance in Theorem 3.1, (ii)) confirms. In our numerically implementable method developed and applied in Sections 4–6 we rely on the ρ -Doob martingale. From a theoretical perspective, however, and as a general

justification of our pathwise dual, martingale-based approach, we develop in this section several results on so-called *surely optimal ρ -martingales*. To achieve this, we generalize the concept of standard surely optimal martingales (see Schoenmakers, Zhang and Huang [78] in the context of standard conditional expectations and optimal single stopping problems) to subadditive DMU functionals. The results in this section show formally that if a general ρ -martingale—not necessarily the ρ -Doob martingale—induces ‘small’ (robust) variance, then the associated bounds obtained from the dual representation can be expected to be ‘tight’ and nearly constant.

Our results on surely optimal ρ -martingales can also serve as a diagnostic device to assess the quality of the estimated ρ -Doob martingale, derived from an (input) approximation to the upper Snell envelope. If the (robust) variance the estimate induces fails to be small, then it must be far from the ρ -Doob martingale. If, on the other hand, this variance is small, then the estimate will be close to an optimal ρ -martingale (attaining the dual representation’s infimum), even though not necessarily close to the ρ -Doob martingale.

For ease of exposition, we focus attention first on optimal single stopping problems. The next theorem generalizes the analogous key measurability result for standard conditional expectations and optimal single stopping problems to DMU functionals satisfying (C1)–(C4) and (P1).

Theorem 3.4 *Let Y_i^* be the upper Snell envelope of the cash-flow process H with respect to a subadditive DMU functional ρ satisfying (C1)–(C4) and (P1) as in Corollary 3.2 and let M be a ρ -martingale. Then, for any $i \in \{0, \dots, T\}$,*

$$\max_{i \leq j \leq T} (H_j - M_j + M_i) \in \mathcal{F}_i \Rightarrow \max_{i \leq j \leq T} (H_j - M_j + M_i) = Y_i^*.$$

The following lemma will later allow for a generalization of the results in this section to multiple stopping.

Lemma 3.5 *Let Y^* , H , M and ρ be as in Theorem 3.4. Then, for any fixed $0 \leq i < T$,*

$$\begin{aligned} \theta_{i+} &:= \max_{i < j \leq T} (H_j - M_j + M_i) \in \mathcal{F}_i \Rightarrow \\ (i) \quad \theta_{i+} &= \rho_i(Y_{i+1}^*) \quad \text{and} \quad (ii) \quad M_{i+1} - M_i = Y_{i+1}^* - \rho_i(Y_{i+1}^*), \end{aligned}$$

hence $M_{i+1} - M_i$ is a ρ -Doob martingale increment. Note the strict first inequality under the max operator. Thus, in particular, if $\theta_{i+} \in \mathcal{F}_i$ for every $0 \leq i < T$, then M is the ρ -Doob martingale.

Let us define the conditional ρ -variance as follows:

$$\text{Var}_{\rho_i}(X) := \rho_i\left((X - \rho_i(X))^2\right). \tag{3.9}$$

It admits a conditional Chebyshev inequality, exploited in the proof of Theorem 3.8 below, as follows:

Proposition 3.6 *Assume (C1)–(C4). If ρ is positively homogeneous (P3), then*

$$\rho_i\left(\mathbf{1}_{\{|X - \rho_i(X)| \geq \epsilon\}}\right) \leq \frac{\text{Var}_{\rho_i}(X)}{\epsilon^2}. \tag{3.10}$$

Next, we state the following lemma:

Lemma 3.7 *Assume (C1)–(C4). Let ρ be subadditive (P1) and sensitive (P2). Then,*

$$\text{Var}_{\rho_i}(X) = 0 \iff X \in \mathcal{F}_i.$$

By virtue of Lemma 3.7, Theorem 3.4 implies that if a ρ -martingale M is such that, for some $i \leq j \leq T$, the conditional ρ -variance

$$\text{Var}_{\rho_i}(\theta_i(M)) := \text{Var}_{\rho_i} \left(\max_{i \leq j \leq T} (H_j - M_j + M_i) \right)$$

is zero a.s., then $\theta_i(M) = Y_i^*$ a.s. In that case, we say that the ρ -martingale M is *surely optimal* at i . Note that, in particular, the ρ -Doob martingale in (2.4) is surely optimal.

We then present a stability result for ρ -martingales M that are, in loose terms, ‘close’ to be surely optimal, in the sense that the conditional ρ -variance $\text{Var}_{\rho_i}(\theta_i(M))$ is ‘small’. In particular, for a sequence of ρ -martingales $(M^{(n)})_{n \geq 1}$ that induces vanishing conditional ρ -variance, we establish weak conditions guaranteeing that the corresponding upper bounds converge to the upper Snell envelope (in L^1), even though the sequence of ρ -martingales $(M^{(n)})_{n \geq 1}$ itself does not necessarily converge.

Theorem 3.8 *Assume (C1)–(C4). Let ρ be subadditive (P1) and positively homogeneous (P3). Suppose that*

$$\text{Var}_{\rho_i}(\theta_i^{(n)}) \xrightarrow{P} 0, \quad \text{with} \quad \theta_i^{(n)} = \max_{i \leq j \leq T} (H_j - M_j^{(n)} + M_i^{(n)}).$$

If, in addition, for every i and every $\epsilon > 0$ there exists $K_\epsilon > 0$ such that

$$\sup_{n \geq 1} \mathbb{E} \left[\rho_i \left(\left| M_i^{(n)} \right| 1_{\{|M_i^{(n)}| > K_\epsilon\}} \right) \right] < \epsilon, \quad (3.11)$$

then

$$\rho_i(\theta_i^{(n)}) \xrightarrow{L^1} Y_i^*.$$

Note that, if $\rho_i \equiv \mathbb{E}_{\mathcal{F}_i}$, i.e., the standard conditional expectation, (3.11) boils down to a standard uniform integrability condition. More generally, we have the following:

Proposition 3.9 *Assume (C1)–(C4). Let ρ be subadditive (P1) and positively homogeneous (P3). If, for some $\eta > 0$,*

$$\sup_{n \geq 1} \mathbb{E} \left[\rho_i \left(\left| M_i^{(n)} \right|^{1+\eta} \right) \right] < \infty,$$

then $(M_i^{(n)})_{n \geq 1}$ satisfies (3.11).

Under an additional Lipschitz continuity condition, Theorem 3.8 may readily be applied as follows. Specifically, let us assume that, for some number p ,

$$\mathbb{E} [|\rho_i(Z)|^p] \leq C_p \mathbb{E} [|Z|^p], \quad (3.12)$$

with $C_p > 0$. In particular, assuming that (3.12) holds for $p = 1$, one has

$$\mathbb{E} [\text{Var}_{\rho_i}(X)] := \mathbb{E} \left[\rho_i \left((X - \rho_i(X))^2 \right) \right] \leq C_1 \mathbb{E} \left[(X - \rho_i(X))^2 \right] = C_1 \mathbb{E} \left[\widetilde{\text{Var}}_{\rho_i}(X) \right], \quad (3.13)$$

where $\widetilde{\text{Var}}_{\rho_i}(X) := \mathbb{E}_{\mathcal{F}_i} \left[(X - \rho_i(X))^2 \right]$. Now suppose that we achieve in an algorithm that $\widetilde{\text{Var}}_{\rho_i}(\theta_i^{(n)}) \xrightarrow{L^1} 0$, where the sequence $(M_i^{(n)})_{n \geq 1}$ is uniformly integrable in the standard sense, i.e.,

$$\sup_{n \geq 1} \mathbb{E} \left[\left| M_i^{(n)} \right| 1_{\{|M_i^{(n)}| > K_\epsilon\}} \right] < \epsilon.$$

Then, obviously (by Lipschitzianity), $(M_i^{(n)})_{n \geq 1}$ satisfies the notion of ρ_i -uniform integrability in (3.11), and $\text{Var}_{\rho_i}(\theta_i^{(n)}) \xrightarrow{L^1} 0$, due to (3.13). The latter implies $\text{Var}_{\rho_i}(\theta_i^{(n)}) \xrightarrow{P} 0$, and then Theorem 3.8 implies that

$$\rho_i(\theta_i^{(n)}) \xrightarrow{L^1} Y_i^*.$$

The next theorem generalizes Theorem 3.4 and Lemma 3.5 to multiple stopping:

Theorem 3.10 *Assume (C1)–(C4) and (P1). Let us define for a set of ρ -martingales $M^{(k)}$, $k = 1, \dots, L$,*

$$\Theta_{i+}^q := \max_{i < j_1 < j_2 < \dots < j_q} \sum_{k=1}^q \left(H_{j_k} + M_{j_{k-1}}^{(q-k+1)} - M_{j_k}^{(q-k+1)} \right) \quad \text{for } q = 1, \dots, L,$$

with $j_0 := i$. Note the strict first inequality under the max operator. Then it holds that

$$\begin{aligned} \Theta_{i+}^q \in \mathcal{F}_i \quad \text{for } q = 1, \dots, L, \quad 0 \leq i < T & \implies \\ \begin{cases} (i) & \Theta_{i+}^q = \rho_i(Y_{i+1}^{*,q}) \\ (ii) & M_{i+1}^{(q)} - M_i^{(q)} = Y_{i+1}^{*,q} - \rho_i(Y_{i+1}^{*,q}) \end{cases} \quad \text{for } q = 1, \dots, L, \quad 0 \leq i < T. \end{aligned}$$

Remark 3.11 *Without doubt it is also possible to derive a version of Theorem 3.8 for the multiple stopping setting. However, as our algorithm in Section 4 below aims at approximative construction of ρ -Doob martingale increments associated with the upper Snell envelopes $Y^{*,l}$ of the generalized cash-flows*

$$U_j^{*,l} := H_j + \rho_j(Y_{j+1}^{*,l-1}), \quad l = 1, \dots, L, \quad (3.14)$$

respectively, rather than at approximative construction of merely surely optimal ρ -martingales, we refrain from such an analysis.

4 A General Primal-Dual Pseudo Algorithm

In this section, we develop a primal-dual pseudo algorithm for robust multiple stopping (henceforth called *algorithm* for short). Our treatment in this section applies to DMUs satisfying (C1)–(C4), (P1) and weak continuity conditions, and to general reward processes in a Markovian environment; in particular, our treatment in this section is not restricted to g -expectations. The following lemma will serve as a cornerstone in our construction.

Lemma 4.1 *Let ρ satisfy (C1)–(C4), (P1) and be Lipschitz continuous in the sense of (3.12) for $p = 2$. Furthermore, let $\mathcal{C}^N \in \mathcal{F}_j$, $\mathcal{Y} \in \mathcal{F}_{j+1}$, and let $\mathbf{m}^N \in \mathcal{F}_{j+1}$ be a ρ_j -martingale increment, that is, $\rho_j(\mathbf{m}^N) = 0$, for $j = 0, \dots, T-1$, $N \in \mathbb{N}$, such that*

$$\mathbb{E} \left[(\mathcal{Y} - \mathbf{m}^N - \mathcal{C}^N)^2 \right] \rightarrow 0, \quad \text{for } N \rightarrow \infty. \quad (4.1)$$

Then,

$$\mathcal{C}^N \xrightarrow{L^2} \rho_j(\mathcal{Y}) \quad \text{and} \quad \mathbf{m}^N \xrightarrow{L^2} \mathcal{Y} - \rho_j(\mathcal{Y}). \quad (4.2)$$

In Lemma 4.1 as well as in the following corollary, N can be thought of as the number of Monte Carlo simulations, in line with our algorithm below. The corollary shows that, if we consider the ρ_j -martingale and \mathcal{F}_j -measurable parts in a converging estimator of a given \mathcal{F}_{j+1} -measurable target random variable (think of an approximation to the upper Snell envelope), then these converge to the ρ_j -martingale and \mathcal{F}_j -measurable parts of the target, respectively.

Corollary 4.2 *Let Y_{j+1}^l be an already constructed approximation to a random variable $Y_{j+1}^{*,l}$. Furthermore, let ρ , the ρ -martingale increment $\mathbf{m}_{j+1}^{l,N} \in \mathcal{F}_{j+1}$, and $\mathcal{C}_j^{l,N} \in \mathcal{F}_j$ be as in Lemma 4.1, such that*

$$\mathbb{E} \left[\left(Y_{j+1}^l - \mathbf{m}_{j+1}^{l,N} - \mathcal{C}_j^{l,N} \right)^2 \right] \rightarrow 0, \quad \text{for } N \rightarrow \infty.$$

Then,

$$\mathcal{C}_j^{l,N} \xrightarrow{L^2} \rho_j \left(Y_{j+1}^l \right) \quad \text{and} \quad Y_{j+1}^l - \mathbf{m}_{j+1}^{l,N} \xrightarrow{L^2} \rho_j \left(Y_{j+1}^l \right).$$

Guided by Lemma 4.1 and Corollary 4.2, we now develop a primal-dual algorithm in the context of a Markovian underlying process X with state space \mathbb{R}^d , possibly in continuous time, that is monitored at the exercise dates as X_j , $j = 0, \dots, T$. We assume that ρ_j is Markovian in the sense that for any j and any random variable in \mathcal{X} of the form $f_j(X_j, \dots, X_T)$ for some measurable f_j , we have that $\rho_j(f_j(X_j, \dots, X_T)) = g_j(X_j)$ for some measurable $g_j : \mathbb{R} \rightarrow \mathbb{R}$. The assumption that ρ_j is Markovian is satisfied, for example, if ρ_j is conditional law invariant (i.e., for each j and $H \in \mathcal{X}$, $\rho_j(H)$ depends only on the (conditional) distribution of H under $\mathbb{P}|_{\mathcal{F}_j}$) or if ρ_j is the solution of a BSDE as in Sections 5–6. As usual, we assume that \mathcal{F}_j is the σ -field generated (directly or, as in the next section, indirectly) by the process X up to exercise date j . Furthermore, we assume that the cash-flows are of the form

$$H_j = f_j(X_j), \quad \text{for } j = 0, \dots, T,$$

where $f_j : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, $j = 0, \dots, T$, are given nonnegative payoff functions such that $H \in \mathcal{H}$. Note that, due to the Bellman principle (3.5), $Y^{*,l}$ can be seen as the upper Snell envelope corresponding to the generalized cash-flow

$$U_j^{*,l} := H_j + \rho_j \left(Y_{j+1}^{*,l-1} \right) =: H_j + c_j^{*,l-1}(X_j), \quad l = 1, \dots, L, \quad j = 0, \dots, T-1, \quad (4.3)$$

where the so-called “continuation functions” $c_j^{*,l-1}$ exist by Markovianity. In particular, $c_j^{*,l-1}(x)$ denotes the value at time j in state x due to $l-1$ remaining exercise rights. We also assume to have a set of Monte Carlo simulated training trajectories

$$X^{(n)} \equiv X^n, \quad n = 1, \dots, N.$$

A main ingredient of our algorithm is the following three-step-procedure, based on an outer loop over exercise rights and a backward loop over exercise dates.

1. Initialize $\bar{M}^0 = \bar{Y}^0 = \bar{c}^0 = 0$, for $l = 0$.
2. Suppose that, for a particular l with $0 \leq l < L$:
 - (i) we have constructed a set of (approximate) continuation functions $\bar{c}_j^{l'} : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, $1 \leq j \leq T$, $0 \leq l' \leq l$, hence an (approximate) continuation value process (for up to l exercise rights) of the form

$$\bar{C}_j^{l'} = \bar{c}_j^{l'}(X_j), \quad 0 \leq l' \leq l;$$

- (ii) we have constructed a (true) ρ -martingale \bar{M}_j^l ; and
- (iii) we have constructed, on each trajectory n , a path

$$\bar{Y}_j^{l,n} := \max \left[\bar{U}_j^{l,n}, \bar{c}_j^l(X_j^n) \right] \text{ if } j \in \{0, 1, \dots, T\}, \quad \text{where} \quad (4.4)$$

$$\bar{U}_j^{l,n} := \begin{cases} f_j(X_j^n) + \bar{c}_j^{l-1}(X_j^n) & \text{if } l > 0 \\ 0 & \text{if } l = 0 \end{cases}, \quad 0 \leq j \leq T, \quad (4.5)$$

as an approximation to $Y^{*,l,n}$.

3. Now construct, using these trajectories, a subsequent (true) ρ -martingale \bar{M}^{l+1} , a subsequent set of continuation functions \bar{c}_j^{l+1} , $j = 0, \dots, T$, and subsequent trajectories $\bar{Y}^{l+1,n}$, $n = 1, \dots, N$ (as approximations to $M^{*,l+1}$ and $Y^{*,l+1}$, respectively) such that (4.4) holds for $l + 1$. To this end, we carry out the following backward procedure, or “backward subroutine”, at level $l + 1$:

- As initialization, set $\bar{Y}_T^{l+1} = H_T$, $\bar{c}_T^{l+1} = 0$. (We also set $\bar{c}_T^0 = 0$.)
- Suppose that, for $0 < j + 1 \leq T$, the values $\bar{Y}_{j+1}^{l+1,n}$, $n = 1, \dots, N$, the set of ρ -martingale increments $\left(\bar{M}_r^{l+1} - \bar{M}_{j+1}^{l+1} \right)_{j+1 < r \leq T}$ (which is empty if $j + 1 = T$), and the continuation function \bar{c}_{j+1}^{l+1} have been constructed.
- Then construct, according to the regression subroutine in Section 4.1 below, a continuation function \bar{c}_j^{l+1} , a ρ -martingale increment $\bar{\mathbf{m}}_{j+1}^{l+1} \in \mathcal{F}_{j+1}$ with $\rho_j(\bar{\mathbf{m}}_{j+1}^{l+1}) = 0$, and set $\left(\bar{M}_r^{l+1} - \bar{M}_j^{l+1} \right) = \left(\bar{M}_r^{l+1} - \bar{M}_{j+1}^{l+1} + \bar{\mathbf{m}}_{j+1}^{l+1} \right)$, for

$$\begin{aligned} \bar{Y}_j^{l+1,n} &= \max \left[\bar{U}_j^{l+1,n}, \bar{c}_j^{l+1}(X_j^n) \right], \quad \text{with} \\ \bar{U}_j^{l+1,n} &= f_j(X_j^n) + \bar{c}_j^l(X_j^n), \quad \text{and } n = 1, \dots, N. \end{aligned} \quad (4.6)$$

With these three steps explained, we now summarize our general primal-dual pseudo algorithm:

General Primal-Dual Pseudo Algorithm.

- (a.) Working forward from $l = 0, \dots, L$ yields a family of continuation functions \bar{c}^l and a family of (true) ρ -martingales \bar{M}^l , respectively:

$$\bar{c}_j^l(\cdot), \quad \text{and} \quad \bar{M}_j^l := \sum_{r=1}^j \bar{m}_r^l, \quad l = 1, \dots, L, \quad j = 0, \dots, T.$$

- (b.) An upper bound for the solution to the robust multiple stopping problem at $t = 0$ due to L exercise rights is now given by (cf. Theorem 3.1, (ii)):

$$Y_0^{\text{upp},L} := \rho_0 \left(\max_{0 \leq j_1 < j_2 < \dots < j_L} \sum_{l=1}^L \left(f_{j_l}(X_{j_l}) - \bar{M}_{j_l}^{L-l+1} + \bar{M}_{j_{l-1}}^{L-l+1} \right) \right), \quad (4.7)$$

which needs to be estimated by a separate (Monte Carlo) procedure.

- (c.) A lower bound for the solution to the robust multiple stopping problem at $t = 0$ due to L exercise rights may next be obtained from the family of stopping times

$$\tau^l := \min \left\{ j : \tau^{l-1} < j \leq T, f_j(X_j) + \bar{c}_j^{l-1}(X_j) \geq \bar{c}_j^l(X_j) \right\}, \quad (4.8)$$

via a (Monte Carlo) estimation of:

$$Y_0^{\text{low},L} := \rho_0 \left(\sum_{l=1}^L f_{\tau^l}(X_{\tau^l}) \right). \quad (4.9)$$

4.1 Regression Subroutine

In this subsection, we describe some technical details pertaining to Step 3. of the three-step-procedure introduced above. Let there be given, for each $j = 0, \dots, T-1$, some parametric collection of ‘elementary’ \mathcal{F}_{j+1} -measurable ρ_j -martingale increments $\mathcal{E}_{j+1}^\beta = \mathcal{E}_{j+1}^\beta(\omega)$ with $\rho_j(\mathcal{E}_{j+1}^\beta) = 0$, $\beta = (\beta_1, \dots, \beta_{K'}) \in \mathbb{R}^{K'}$, $K' \in \mathbb{N}$. We assume that the set

$$L_{j+1,0}^2 := \left\{ \mathcal{E} \in L^2(\Omega, \mathcal{F}_{j+1}, \mathbb{P}) : \rho_j(\mathcal{E}) = 0 \right\},$$

for $j = 0, \dots, T-1$, has a countable dense subset, say $\mathcal{D}_{j+1,0} := \left\{ \mathcal{E}_{j+1}^{(k)} : k \in \mathbb{N} \right\}$. Proposition 4.3 below provides a sufficient condition for the existence of such a dense subset. We then consider for each $j = 0, \dots, T-1$, $\beta = (\beta_1, \dots, \beta_{K'}) \in \mathbb{R}^{K'}$, $K' \in \mathbb{N}$, the ρ_j -martingale increments

$$\mathcal{E}_{j+1}^\beta := \sum_{k=1}^{K'} \beta_k \mathcal{E}_{j+1}^{(k)} - \rho_j \left(\sum_{k=1}^{K'} \beta_k \mathcal{E}_{j+1}^{(k)} \right),$$

i.e., $\mathcal{E}_{j+1}^\beta = \mathcal{E}_{j+1}^\beta(\omega) \in \mathcal{F}_{j+1}$ with $\rho_j \left(\mathcal{E}_{j+1}^\beta \right) = 0$. Since we obviously have $\mathcal{E}_{j+1}^{e^{(k)}} = \mathcal{E}_{j+1}^{(k)}$ with $e^{(k)} \in \mathbb{R}^{K'}$ being the standard K' -dimensional unit vector in $\mathbb{R}^{K'}$, the family

$$\left\{ \mathcal{E}_{j+1}^{(\beta_1, \dots, \beta_{K'})} : \beta \in \mathbb{R}^{K'}, K' \in \mathbb{N} \right\}, \quad (4.10)$$

for $j = 0, \dots, T-1$, is dense in $L_{j+1,0}^2$.

Proposition 4.3 *If $L^2(\Omega, \mathcal{F}_{j+1}, \mathbb{P})$ is separable and ρ_j satisfies (3.12) with $p = 2$, then $L_{j+1,0}^2$ has a countable dense subset.*

Furthermore, we consider a ‘rich enough’ collection of basis functions $\psi_1, \dots, \psi_{K''} : \mathbb{R}^d \rightarrow \mathbb{R}$ and then solve, for fixed N, K' , and K'' , the minimization problem

$$\begin{aligned} \text{MSE} &:= \sum_{n=1}^N \left(Y_{j+1}^{l+1,n} - \mathcal{E}_{j+1}^\beta(\omega^n) - \gamma \psi(X_j^n) \right)^2 \\ &\longrightarrow \arg \min_{\beta \in \mathbb{R}^{K'}, \gamma \in \mathbb{R}^{K''}} =: \left[\beta^{l+1,j,K',N}, \gamma^{l+1,j,K'',N} \right], \end{aligned} \quad (4.11)$$

where we use vector notation $\gamma = (\gamma_1, \dots, \gamma_{K''})$ and $\psi = (\psi_1, \dots, \psi_{K''})^\top$. We will later prove that the MSE above converges to zero as $K', K'' \rightarrow \infty$ for our choice of β and γ , which is actually sufficient for the algorithm to converge. We will suppress the superscripts K', K'' and N whenever there is no ambiguity. We set

$$\bar{\mathbf{m}}_{j+1}^{l+1} := \bar{\mathbf{m}}_{j+1}^{l+1}(X) := \mathcal{E}_{j+1}^{\beta^{l+1,j}}, \quad \bar{c}_j^{l+1}(\cdot) := \sum_{k=1}^{K''} \gamma_k^{l+1,j} \psi_k(\cdot).$$

4.2 Convergence Theorem

We state the following theorem:

Theorem 4.4 *Let ρ be subadditive (P1) and Lipschitz continuous in the sense of (3.12) for $p = 2$. We set $K = \min \{K', K''\}$ and denote by $\bar{M}_j^{l,K,N} := \bar{M}_j^l, \bar{c}_j^{l,K,N} := \bar{c}_j^{l,K,N}(X_j) := \bar{c}_j^l$ and $\bar{Y}_j^{l,K,N} := \bar{Y}_j^{l,K,N}(X_j) := \bar{Y}_j^{l,N}$ the functions constructed in the algorithm above. Then,*

$$\lim_{K \rightarrow \infty} \lim_{N \rightarrow \infty} \bar{M}_j^{l,K,N} = M_j^{*,l} \text{ in } L^2, \quad (4.12)$$

$$\lim_{K \rightarrow \infty} \lim_{N \rightarrow \infty} \bar{c}_j^{l,K,N} = c_j^{*,l} \text{ in } L^2, \quad (4.13)$$

$$\lim_{K \rightarrow \infty} \lim_{N \rightarrow \infty} \bar{Y}_j^{l,K,N} = Y_j^{*,l} \text{ in } L^2, \quad (4.14)$$

for all $j = T, T-1, \dots, 0$ and $l = 1, \dots, L$. Furthermore,

$$\begin{aligned} \left| Y_j^{*,l} - \rho_j \left(\max_{j \leq r \leq T} \left(\bar{U}_r^l - \bar{M}_r^l \right) \right) \right| &\leq \rho_j \left(\max_{j \leq r \leq T} \left| M_r^{*,l} - \bar{M}_r^l \right| \right) + \rho_j \left(\max_{j \leq r \leq T} \left| \bar{c}_r^{l-1} - c_r^{*,l-1} \right| \right) \\ &\longrightarrow_{K \rightarrow \infty, N \rightarrow \infty} 0. \end{aligned}$$

4.3 Complexity

At first sight, the pathwise maximum in (4.7) would require the evaluation of $T!/(L!(T-L)!) terms. Fortunately, due to following proposition, it only requires $O(LT)$ evaluations.$

Proposition 4.5 *Define, for $1 \leq q \leq L$ and $0 \leq i \leq T$,*

$$\Theta_i^q := \max_{i \leq j_1 < j_2 < \dots < j_q} \sum_{l=1}^q \left(f_{j_l}(X_{j_l}) - \overline{M}_{j_l}^{q-l+1} + \overline{M}_{j_{l-1}}^{q-l+1} \right), \quad \text{with } j_0 = i,$$

and naturally $\Theta_i^q = 0$, $i > T$. Then,

$$\Theta_i^q = \max \left[f_i(X_i) + \overline{M}_i^{q-1} - \overline{M}_{i+1}^{q-1} + \Theta_{i+1}^{q-1}, \overline{M}_i^q - \overline{M}_{i+1}^q + \Theta_{i+1}^q \right]. \quad (4.15)$$

Thus, the evaluation of (4.7) may be described as follows:

Recursive evaluation of (4.7)

1. Initialize $\Theta_i^0 = 0$, for $i = 0, \dots, T$;
2. Suppose that, for $0 \leq q-1 < L$ and for $i = 0, \dots, T$, the construction of Θ_i^{q-1} has been conducted;
3. Backward subroutine: Initialize $\Theta_T^q = f_T(X_T)$. When Θ_{i+1}^q has been constructed for $i+1 \leq T$, compute Θ_i^q via (4.15).

The workload of computing the coefficients $\beta^{K'}$ and $\gamma^{K''}$ for all exercise rights l and times j via the minimization problem (4.11) is of order $O(LTN \max(K', K'')^2)$, provided (4.11) can be cast into a least squares problem. The latter applies in our numerical settings in which the functionals ρ_j are constructed via (jump)-BSDEs on a fine time grid. Once these coefficients have been determined, one may construct an exercise policy and a dual estimator, and compute a lower and upper bound (in expectation), respectively. If N_{low} and N_{upp} simulations are used for the lower bound and upper bound, respectively, the computational costs required will be $O(LTN_{\text{low}} \max(K', K''))$ and $O(LTN_{\text{upp}} \max(K', K''))$, respectively. The cost estimate for the upper bound is based on the recursion (4.15). Taking all together, we thus arrive at a workload of order

$$O(LT(N \max(K', K'')^2 + N_{\text{low}} \max(K', K'') + N_{\text{upp}} \max(K', K''))))$$

for the entire primal-dual procedure.

5 Explicit Construction in a Brownian-Poisson Filtration

In the sequel, we assume that we have a completed continuous-time filtration $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ generated by a d_1 -dimensional standard (i.e., zero mean and unit variance) Brownian motion $W = (W_1, \dots, W_{d_1})^\top$ and a d_2 -dimensional Poisson process $N = (N_1, \dots, N_{d_2})^\top$ with arrival intensity $\lambda_{\mathbb{P}} = (\lambda_{\mathbb{P}}^1, \dots, \lambda_{\mathbb{P}}^{d_2})^\top$. As usual, we define the compensated counterpart of N as $\tilde{N}_t = N_t - \lambda_{\mathbb{P}} t$. The components of the processes W and

N are assumed to be independent. The stochastic drivers W and N generate the underlying Markovian adapted reward process $(X_t)_{t \in [0, T]}$ with state space \mathbb{R}^d of Section 4.

Furthermore, we assume that ρ satisfies (the continuous-time analogs of) (C1)–(C4) and (P1)–(P3). This means, in particular, that ρ is a coherent risk measure. By classical duality results (e.g., Föllmer and Schied [40]), the robust multiple stopping problem at time t is then given by

$$Y_t^{*,L} = \sup_{\substack{t \leq \tau_1 < \dots < \tau_L \\ (\tau_1, \dots, \tau_L) \in \mathcal{T}_t(L)}} \rho_t \left(\sum_{l=1}^L H_{\tau_l} \right) = \sup_{\substack{t \leq \tau_1 < \dots < \tau_L \\ (\tau_1, \dots, \tau_L) \in \mathcal{T}_t(L)}} \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}_{\mathbb{Q}} \left[\sum_{l=1}^L H_{\tau_l} \middle| \mathcal{F}_t \right], \quad 0 \leq t \leq T, \quad (5.1)$$

with $\mathcal{T}_t(L)$ our family of stopping vectors, and \mathcal{Q} a closed convex set of probability measures absolutely continuous with respect to \mathbb{P} and satisfying a stability assumption. In such a continuous-time setting, it is known that every recursive coherent risk measure can be identified with a solution to a backwards stochastic differential equation (BSDE) also called a g -expectation, modulo a compactness assumption; see Section 5.1 for the precise definitions and results. Exploiting our algorithm presented in Section 4, this section constructs explicit upper and lower bounds to $Y^{*,L}$ with desirable properties.

5.1 Bellman’s Principle, the Set of Priors, and BSDE drivers

A probability measure change from \mathbb{P} to an absolutely continuous measure $\mathbb{Q} \in \mathcal{Q}$ admits an explicit representation in our Brownian-Poisson setting. Consider the Radon-Nikodym derivative

$$D_t := \mathbb{E} \left[\frac{d\mathbb{Q}}{d\mathbb{P}} \middle| \mathcal{F}_t \right], \quad t \in [0, T].$$

As is well-known, D_t has the Doléans-Dade exponential form

$$D_t = \exp \left(\int_0^t q_s dW_s + \int_0^t \log \left(\frac{\lambda_s}{\lambda_{\mathbb{P}}} \right) dN_s - \int_0^t \left(\frac{|q_s|^2}{2} + \lambda_s - \lambda_{\mathbb{P}} \right) ds \right), \quad (5.2)$$

where q is a predictable, \mathbb{R}^{d_1} -valued, stochastic drift, and λ is a positive, predictable, \mathbb{R}^{d_2} -valued process with $\frac{\lambda_s}{\lambda_{\mathbb{P}}} := \left(\frac{\lambda_s^1}{\lambda_{\mathbb{P}}^1}, \dots, \frac{\lambda_s^{d_2}}{\lambda_{\mathbb{P}}^{d_2}} \right)^\top$, which jointly uniquely characterize \mathbb{Q} . From Girsanov’s theorem, we know that $W_t^{\mathbb{Q}} := W_t - \int_0^t q_s ds$ is a standard Brownian motion under \mathbb{Q} while the process N_t has arrival intensity λ_t . In particular, the reference model \mathbb{P} corresponds to $q \equiv 0$ and $\lambda \equiv \lambda_{\mathbb{P}}$. The stochastic drift q may be given the interpretation of a drift in the diffusive component that is misspecified to be absent by the reference model \mathbb{P} . Similarly, $\lambda_s - \lambda_{\mathbb{P}}$ represents a deviation from the misspecified arrival intensity $\lambda_{\mathbb{P}}$ under \mathbb{P} .

In our dynamic setting, with DMU evaluations satisfying the continuous-time analogs of (C1)–(C4) and (P1)–(P3), time-consistency of choice under uncertainty is satisfied as it is equivalent to recursiveness or Bellman’s dynamic programming principle. Time-consistency of a dynamic evaluation $(\rho_t(H))_{t \in [0, T]}$ requires—according to its usual definition, also referred to as ‘strong’ time-consistency—that $\rho_s(H_1) \geq \rho_s(H_2)$ whenever $\rho_t(H_1) \geq \rho_t(H_2)$, $t \geq s$. That is, if H_2 is preferred over H_1 , in each state of nature at time t , then the same preference necessarily applies prior to time t ; see e.g., Riedel [69], Ruzsarczyński and Shapiro [74], Shapiro, Dentcheva and Ruzsarczyński [79], Chapter 6, Ruzsarczyński [75] and Shapiro [80]. Indeed, requiring recursiveness or Bellman’s dynamic programming principle is equivalent to requiring time-consistency

for $\rho_t(H) = \sup_{\{\mathbb{Q} \sim \mathbb{P} | \mathbb{Q} = \mathbb{P} \text{ on } \mathcal{F}_t\}} \mathbb{E}_{\mathbb{Q}}[H | \mathcal{F}_t]$, $t \in [0, T]$, which is, in turn, equivalent to the set of priors \mathcal{Q} being m -stable; see Delbaen [35]. More formally, the following statements are equivalent (see Lemma 11.11 of Föllmer and Schied [40] for the equivalence (i)–(ii), Delbaen [35] and Delbaen, Peng and Rosazza Gianin [36] for (ii)–(iii) in a Brownian setting, and Tang and Wei [83] and Laeven and Stadjé [56] for (ii)–(iii) in a general semi-martingale setting):

- (i) ρ is recursive, i.e., ρ satisfies Bellman’s dynamic programming principle $\rho_0(\rho_t(H)I_A) = \rho_0(HI_A)$ for every $t \in [0, T]$, $A \in \mathcal{F}_t$, and bounded H .
- (ii) ρ is time-consistent over bounded rewards.
- (iii) There exists a closed, convex, set-valued predictable mapping C taking values in $\mathbb{R}^{d_1} \times (-\lambda_{\mathbb{P}}^1, \infty) \times \cdots \times (-\lambda_{\mathbb{P}}^{d_2}, \infty)$ such that

$$\rho_t(H) = \sup_{(q, \lambda) \in C} \mathbb{E}_{\mathbb{Q}}[H | \mathcal{F}_t], \quad t \in [0, T].$$

As in our continuous-time Brownian-Poisson setting recursiveness (C3) is equivalent to (iii), we assume henceforth:

(A1) $C = (C_t)_{t \in [0, T]} \subset [0, T] \times \mathbb{R}^{d_1} \times (-\lambda_{\mathbb{P}}^1 + \varepsilon, \infty) \times \cdots \times (-\lambda_{\mathbb{P}}^{d_2} + \varepsilon, \infty)$ with $\varepsilon > 0$ is compact.

We note that (A1) also implies that ρ is recursive and time-consistent over square-integrable rewards.

For $t \in [0, T]$, $z \in \mathbb{R}^{1 \times d_1}$ and $\tilde{z} \in \mathbb{R}^{1 \times d_2}$, and C satisfying Assumption (A1), let us define a function g via Fenchel’s duality:

$$g(t, z, \tilde{z}) := \sup_{(q, \lambda - \lambda_{\mathbb{P}}) \in C_t} \{zq + \tilde{z}(\lambda - \lambda_{\mathbb{P}})\}. \quad (5.3)$$

One easily verifies that g is convex, positively homogeneous and Lipschitz continuous. Then, from Krätschmer *et al.* [54], we have the following statement, which is essentially (with (A1)) equivalent to (i)–(iii) above:

- (iv) For every $H \in L^2(\mathcal{F}_{j+1})$, there exists a unique square-integrable predictable (Z, \tilde{Z}) such that

$$d\rho_t(H_{j+1}) = -g(t, Z_t, \tilde{Z}_t)dt + Z_t dW_t + \tilde{Z}_t d\tilde{N}_t, \text{ for } t \in [j, j+1], j \in \{0, \dots, T-1\}.$$

In particular, there exists a unique square-integrable predictable (Z^*, \tilde{Z}^*) such that

$$d\rho_t(Y_{j+1}^*) = -g(t, Z_t^*, \tilde{Z}_t^*)dt + Z_t^* dW_t + \tilde{Z}_t^* d\tilde{N}_t, \text{ for } t \in [j, j+1], j \in \{0, \dots, T-1\}. \quad (5.4)$$

Furthermore, for $t \in [0, T]$, the (Z^*, \tilde{Z}^*) in (5.4) recover—and later allow to practically compute—the ρ -Doob martingale as follows (cf. Eqns. (2.4)–(2.5)):

$$M_t^* = \rho_t(M_T^*) = - \int_0^t g(s, Z_s^*, \tilde{Z}_s^*) ds + \int_0^t Z_s^* dW_s + \int_0^t \tilde{Z}_s^* d\tilde{N}_s. \quad (5.5)$$

Here, $j = 0, \dots, L-1$ should be interpreted as exercise dates and $t \in [j, j+1]$ as the continuous embedding.

Because $\rho_{j+1}(Y_{j+1}^*) = Y_{j+1}^*$, by (iv), for $t \in [j, j+1]$,

$$\rho_t(Y_{j+1}^*) = Y_{j+1}^* + \int_t^{j+1} g(s, Z_s^*, \tilde{Z}_s^*) ds - \int_t^{j+1} Z_s^* dW_s - \int_t^{j+1} \tilde{Z}_s^* d\tilde{N}_s. \quad (5.6)$$

Eqn. (5.4) is referred to as a backward stochastic differential equation (BSDE). Formally, given a terminal payoff $H \in L^2$ and a function $g : [0, T] \times \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$, referred to as a driver, the solution to the corresponding BSDE is a triple of square-integrable and suitably measurable processes (Y, Z, \tilde{Z}) that satisfies

$$dY_t = -g(t, Z_t, \tilde{Z}_t) dt + Z_t dW_t + \tilde{Z}_t d\tilde{N}_t, \quad \text{and} \quad Y_T = H.$$

The solution is often referred to as a (conditional) g -expectation; see, e.g., Peng [66].

As a means of illustrating the generality of our setup given by (5.1) with (A1) we provide a few examples.

Example 5.1 (1.) *Ball scenarios: Consider a decision-maker endowed with a set of priors constituting a small ball environment surrounding \mathbb{P} all deemed equally plausible. Then,*

$$\mathcal{Q} = \left\{ \mathbb{Q}^{(q, \lambda)} \ll \mathbb{P} \mid |q_t| \leq \delta_1, \quad |\lambda_t - \lambda_{\mathbb{P}}| \leq \delta_2, \text{ for Lebesgue-a.s. all } t \right\}, \quad \delta_1, \delta_2 > 0,$$

and $C_t = \{(q, \lambda) \mid |q| \leq \delta_1, |\lambda - \lambda_{\mathbb{P}}| \leq \delta_2\}$. Suppose without losing generality that $|\lambda_{\mathbb{P}}| \geq \delta_2$. Then, from (5.3), in explicit form, $g(t, z, \tilde{z}) = \delta_1 |z| + \delta_2 |\tilde{z}|$.

(2.) *Discrete scenarios: Imagine a decision-maker who considers, at each time $t > 0$, finite-dimensional families $\{q_{1,t}, \dots, q_{m,t}\}$ and $\{\lambda_{1,t}, \dots, \lambda_{m,t}\}$, $m \in \mathbb{N}$, with all elements deemed equally plausible. Then,*

$$\mathcal{Q} = \left\{ \mathbb{Q}^{(q, \lambda)} \ll \mathbb{P} \mid (q_t, \lambda_t) \in \{(q_{i,t}, \lambda_{j,t}), i, j \in \{1, \dots, m\}\}, \text{ for Lebesgue-a.s. all } t \right\},$$

and $C_t = \{(q, \lambda) \mid (q, \lambda) \in \text{conv}(\{(q_{i,t}, \lambda_{j,t}), i, j \in \{1, \dots, m\}\})\}$, with $\text{conv}(\cdot)$ the convex hull. We can assume that $0 \in \text{conv}(\{(q_{i,t}, \lambda_{j,t}), i, j \in \{1, \dots, m\}\})$ without losing generality, upon redefining the reference measure. Furthermore, $g(t, z, \tilde{z}) = \max_{i=1, \dots, m} q_{i,t} z + \max_{j=1, \dots, m} \lambda_{j,t} \tilde{z}$.

To obtain genuine upper and lower bounds to the optimal solution of the stopping problem (5.1), we henceforth impose the following additional assumption:

(A2) $H_t = f_t(X_t)$ and we can simulate i.i.d. copies of $(X_t)_{t \in [0, T]}$.

Once we have constructed a ‘good’ family of ρ -martingales \overline{M} , we are faced with the computation of $Y_0^{\text{upp}, L}$ in (4.7). An important advantage of this dual approach for numerical stability is that we have, in fact, the pathwise dual representation

$$Y_0^{*, L} := \max_{0 \leq j_1 < j_2 < \dots < j_L} \sum_{l=1}^L \left(f_{j_l}(X_{j_l}) - M_{j_l}^{*, L-l+1} + M_{j_{l-1}}^{*, L-l+1} \right), \quad \text{almost surely.}$$

Indeed, we will obtain an estimate with ‘low’ variance provided the ρ -martingale \overline{M} is ‘good’. Furthermore, to obtain a lower bound we employ $Y_0^{\text{low},L}$ in (4.9). By the results discussed in this subsection, for a square-integrable payoff U , ρ has a representation of the form

$$\rho_t(U) = \sup_{(q,\lambda) \in \mathcal{C}} \mathbb{E}_{\mathbb{Q}(q,\lambda)} [U | \mathcal{F}_t] \quad (5.7)$$

$$= U + \int_t^T g(s, \mathcal{Z}_s, \tilde{\mathcal{Z}}_s) ds - \int_t^T \mathcal{Z}_s dW_s - \int_t^T \tilde{\mathcal{Z}}_s d\tilde{N}_s. \quad (5.8)$$

Let us now first consider the question of how to explicitly obtain a ‘good’ family of ρ -martingales \overline{M} in our Brownian-Poisson filtration with (C1)–(C4), (P1)–(P3) and (A1)–(A2). Next, we develop in Sections 5.3 and 5.4 explicit genuine lower and upper bounds.

5.2 Parameterization of the ρ -Martingale Increments

As in Section 4, we assume that $(X_j)_{0 \leq j \leq T}$ is an \mathcal{F}_j -adapted d -dimensional underlying Markovian process, now in a Brownian-Poisson filtration, and that our cash-flow process has a structure of the form $H_j = f_j(X_j)$, $j = 0, \dots, T$ such that $H \in \mathcal{H}$. We further assume that the resulting random variables H_j are square integrable.

We are going to construct a ρ -martingale backwardly. To this end, we consider, between two exercise dates j and $j+1$, the (fine) grid $\pi_j = \{s_{(j-1)n_0} = j, s_{j1}, \dots, s_{jn_0} = j+1\}$, where $s_{jp} = j + p\Delta$ with $\Delta = n_0^{-1}$. We also define

$$\Pi := \{s_{00} = 0, s_{01}, \dots, s_{0n_0} = 1, s_{11}, \dots, s_{(T-2)n_0} = T-1, s_{(T-1)1}, \dots, s_{(T-1)n_0} = T\},$$

and sometimes use the notation $\Pi = \{t_0 = 0, t_1, t_2, \dots, t_{n_1} = T\}$, where the t_i are simply the enumerated s_{jp} . For our numerical schemes we always assume that Assumption (A2) is in place, next to (A1). In particular, we can also simulate i.i.d. copies of the $(Z_{s_{jp}}, \tilde{Z}_{s_{jp}})_{s_{jp} \in \Pi}$.

We formally initialize $(\overline{M}_j^{l+1} - \overline{M}_T^{l+1})_{j \geq T}$ as a vector of zeros. Suppose that for some (fixed) $j < T$, an approximation Y_{j+1}^{l+1} to the upper Snell envelope $Y_{j+1}^{*,l+1}$ and the set of ρ -martingale increments $(\overline{M}_q^{l+1} - \overline{M}_{j+1}^{l+1})_{j+1 \leq q \leq T}$ have been constructed. Then, we carry out the following loop. For $p = n_0$, we initialize $0 \leq U_{n_0} = Y_{j+1}^{l+1} \approx \rho_{j+1}(Y_{j+1}^{*,l+1})$. Now, if $0 \leq U_{p+1}$, $p < n_0$, has been constructed, we solve the piecewise linear minimization problem

$$\begin{aligned} \left[\gamma_{s_{jp}}^{N_1}, \beta_{s_{jp}}^{N_1}, \tilde{\beta}_{s_{jp}}^{N_1} \right] = \arg \min_{\gamma, \beta, \tilde{\beta} \in \mathbb{R}^{K_1}} \frac{1}{N_1} \sum_{n=1}^{N_1} \left(U_{p+1}^n - \sum_{k=1}^{K_1} \gamma_k \psi_k(s_{jp}, X_{s_{jp}}^n) \right. \\ \left. + g\left(s_{jp}, \sum_{k=1}^{K_1} \beta_k \varphi_k(s_{jp}, X_{s_{jp}}^n), \sum_{k=1}^{K_1} \tilde{\beta}_k \tilde{\varphi}_k(s_{jp}, X_{s_{jp}}^n)\right) (s_{j(p+1)} - s_{jp}) \right. \\ \left. - \sum_{k=1}^{K_1} \beta_k \varphi_k(s_{jp}, X_{s_{jp}}^n) \Delta W_{s_{jp}}^n - \sum_{k=1}^{K_1} \tilde{\beta}_k \tilde{\varphi}_k(s_{jp}, X_{s_{jp}}^n) \Delta \tilde{N}_{s_{jp}}^n \right)^2, \quad (5.9) \end{aligned}$$

for certain basis functions (ψ_k) , (φ_k) , and $(\tilde{\varphi}_k)$, $K_1 \in \mathbb{N}$, and N_1 trajectories. Alternatively to

(5.9), we can solve

$$\begin{aligned} \left[\gamma_{s_{jp}}^{N_1}, \beta_{s_{jp}}^{N_1}, \tilde{\beta}_{s_{jp}}^{N_1} \right] = \arg \min_{\gamma, \beta, \tilde{\beta} \in \mathbb{R}^{K_1}} \frac{1}{N_1} \sum_{n=1}^{N_1} \left(U_{p+1}^n - \sum_{k=1}^{K_1} \gamma_k \psi_k(s_{jp}, X_{s_{jp}}^n) \right. \\ \left. - \sum_{k=1}^{K_1} \beta_k \varphi_k(s_{jp}, X_{s_{jp}}^n) \Delta W_{s_{jp}}^n - \sum_{k=1}^{K_1} \tilde{\beta}_k \tilde{\varphi}_k(s_{jp}, X_{s_{jp}}^n) \Delta \tilde{N}_{s_{jp}}^n \right)^2, \end{aligned} \quad (5.10)$$

which has a closed-form solution.

Remark 5.2 We note that in the case the filtration is generated by a one-dimensional process (either a Brownian motion or a Poisson process), the minimization problem (5.9) corresponds to a linear programming problem. This is seen as follows. As ρ is a coherent risk measure, it follows that g is positively homogeneous. Thus, there exist $f_{s_{jp}}^+, f_{s_{jp}}^- \geq 0$ such that $g(s_{jp}, z) = f_{s_{jp}}^+ z^+ + f_{s_{jp}}^- z^-$. Hence, the function

$$z \mapsto h(s_{j(p+1)}, z)^2 := \left(U_{p+1}^n + g(s_{jp}, z)(s_{j(p+1)} - s_{jp}) - (W_{s_{j(p+1)}}^n - W_{s_{jp}}^n)z \right)^2$$

is convex as $U_{p+1}^n \geq 0$. The reason is that $h(s_{jp}, \cdot)$ is linear on its negative part. Thus, the minimization problem (5.9) is convex. Because any piecewise linear function that is convex can be written as a supremum of finitely many linear functions, the minimization problem can be expressed as a linear programming problem.

Next, define

$$C_p(X_{s_{jp}}^n) := \sum_{k=1}^{K_1} \gamma_{s_{jp}k}^{N_1} \psi_k(s_{jp}, X_{s_{jp}}^n), \quad \mathcal{Z}_{s_{jp}}^{l+1, N_1}(X_{s_{jp}}^n) := \sum_{k=1}^{K_1} \beta_{s_{jp}k}^{N_1} \varphi_k(s_{jp}, X_{s_{jp}}^n),$$

and $\tilde{\mathcal{Z}}_{s_{jp}}^{l+1, N_1}(X_{s_{jp}}^n)$ similarly. We set

$$U_p(X_{s_{jp}}^n) := \max \left(C_p(X_{s_{jp}}^n) + g(s_{jp}, \mathcal{Z}_{s_{jp}}^{l+1, N_1}(X_{s_{jp}}^n), \tilde{\mathcal{Z}}_{s_{jp}}^{l+1, N_1}(X_{s_{jp}}^n))(s_{j(p+1)} - s_{jp}), 0 \right).$$

We then obtain the desired ρ -martingale increments $\bar{M}_{s_{jp}}^{l+1, K_1, N_1} - \bar{M}_j^{l+1, K_1, N_1}$ by defining

$$\begin{aligned} \bar{M}_{s_{jp}}^{l+1, \Delta, K_1, N_1}(X_{s_{jp}}^n) - \bar{M}_j^{l+1, \Delta, K_1, N_1}(X_{s_{jp}}^n) &\equiv \bar{M}_{s_{jp}}^{l+1, K_1, N_1}(X_{s_{jp}}^n) - \bar{M}_j^{l+1, K_1, N_1}(X_{s_{jp}}^n) \\ &:= - \sum_{u=0}^{p-1} \int_{s_{ju}}^{s_{j(u+1)}} g(u, \mathcal{Z}_{s_{ju}}^{l+1, N_1}(X_{s_{ju}}^n), \tilde{\mathcal{Z}}_{s_{ju}}^{l+1, N_1}(X_{s_{ju}}^n)) du \\ &\quad + \sum_{u=0}^{p-1} \mathcal{Z}_{s_{ju}}^{l+1, N_1}(X_{s_{ju}}^n) \Delta W_{s_{ju}}^n + \sum_{u=0}^{p-1} \tilde{\mathcal{Z}}_{s_{ju}}^{l+1, N_1}(X_{s_{ju}}^n) \Delta \tilde{N}_{s_{ju}}^n. \end{aligned} \quad (5.11)$$

In the end, when we have arrived at $p = 0$, we define $\bar{c}_j^{l+1, K_1, N_1}(\cdot) := C_0(\cdot)$ and $\bar{Y}_j^{l+1, K_1, N_1}(\cdot)$ according to Eqn. (4.6). This way, we have recursively constructed the parameterized space of ρ -martingale increments. The following proposition establishes convergence of our construction.

Proposition 5.3 *The ρ -martingale increments constructed in (5.11) are dense in L^2 and, in particular,*

$$\lim_{\Delta \rightarrow 0} \lim_{K_1 \rightarrow \infty} \lim_{N_1 \rightarrow \infty} \overline{M}_t^{l+1, \Delta, K_1, N_1} = M_t^{*, l+1}.$$

We finally note that (5.11) gives rise to a true discrete-time ρ -martingale $(\overline{M}_j^{l+1})_{j \in \{0, 1, 2, \dots, T\}}$. The thus constructed $(\overline{M}_j^{l+1})_{j \in \{0, 1, 2, \dots, T\}}$ will be exploited to establish an upper bound to the upper Snell envelope via (4.7), while $(\overline{M}_{s_{jp}}^{l+1})_{j,p}$ (living on the finer grid Π) is needed for the numerical approximation.

5.3 Converging Genuine Lower Bound

This subsection develops an explicit lower bound that converges to the upper Snell envelope asymptotically and constitutes a genuine (biased low) lower bound at the pre-limiting level. Consider (5.8) with $U = \sum_{l=1}^L f_{\tau^l}(X_{\tau^l})$ and τ^l constructed by (4.8). From e.g., Barrieu and El Karoui [4], we know that the supremum in (5.7) is attained in

$$\frac{d\mathbb{Q}^g}{d\mathbb{P}} = \mathcal{D} \left(\int_0^T H_s dW_s + \int_0^T \tilde{H}_s d\tilde{N}_s \right), \quad \text{with } (H_s, \tilde{H}_s) \in \partial g(s, \mathcal{Z}_s, \tilde{\mathcal{Z}}_s), \quad (5.12)$$

where \mathcal{D} denotes the Doleans-Dade exponential, and $\partial g(s, \cdot)$ denotes the mapping of sub-differentials of the convex driver $g(s, \cdot)$. We shall exploit this to compute the lower bound numerically. For simplicity, assume that $g(s, \cdot)$ is differentiable. If that is not satisfied, then our approach may still be applied by considering elements in the subgradient.

Let $N_2 \in \mathbb{N}$, simulate paths $(W_{s_{jp}}^n)$, $(N_{s_{jp}}^n)$ and $(X_{s_{jp}}^n)$ for $n = 1, \dots, N_2$, and also consider (the true, non-simulated) $(W_{s_{jp}})$, $(N_{s_{jp}})$ and $(X_{s_{jp}})$. Then, define and construct

$$\begin{aligned} \mathcal{Z}_t^{N_2} &:= z_t^{N_2}(s_{jp}, X_{s_{jp}}) := \beta_{s_{jp}}^{N_2} \varphi(s_{jp}, X_{s_{jp}}) \text{ for } s_{jp} \leq t < s_{j(p+1)}, \\ \tilde{\mathcal{Z}}_t^{N_2} &:= \tilde{z}_t^{N_2}(s_{jp}, X_{s_{jp}}) := \tilde{\beta}_{s_{jp}}^{N_2} \tilde{\varphi}(s_{jp}, X_{s_{jp}}) \text{ for } s_{jp} \leq t < s_{j(p+1)}, \end{aligned}$$

using least squares Monte Carlo regression as described in Section 5.2 with K_2 basis functions and terminal condition given by $U = \sum_{l=1}^L f_{\tau^l}(X_{\tau^l})$. Henceforth, we suppress K_2 in the notation. Furthermore, define and construct, using $N_3 \in \mathbb{N}$ new i.i.d. simulations,

$$\begin{aligned} \mathcal{Z}_t^{N_2, n} &:= z_t^{N_2}(s_{jp}, X_{s_{jp}}^n) := \beta_{s_{jp}}^{N_2} \varphi(s_{jp}, X_{s_{jp}}^n) \text{ for } s_{jp} \leq t < s_{j(p+1)}, \quad n = 1, \dots, N_3, \\ \tilde{\mathcal{Z}}_t^{N_2, n} &:= \tilde{z}_t^{N_2}(s_{jp}, X_{s_{jp}}^n) := \tilde{\beta}_{s_{jp}}^{N_2} \tilde{\varphi}(s_{jp}, X_{s_{jp}}^n) \text{ for } s_{jp} \leq t < s_{j(p+1)}, \quad n = 1, \dots, N_3, \end{aligned}$$

and moreover the partial derivatives

$$\begin{aligned} q_{s_{jp}}^n &:= g_{\mathcal{Z}}(s_{jp}, \mathcal{Z}_{s_{jp}}^{N_2, n}, \tilde{\mathcal{Z}}_{s_{jp}}^{N_2, n}), \quad n = 1, \dots, N_3, \\ \lambda_{s_{jp}}^n - \lambda_{\mathbb{P}} &:= g_{\tilde{\mathcal{Z}}}(s_{jp}, \mathcal{Z}_{s_{jp}}^{N_2, n}, \tilde{\mathcal{Z}}_{s_{jp}}^{N_2, n}), \quad n = 1, \dots, N_3. \end{aligned}$$

Next, define N_3 i.i.d. simulations of the measure $\frac{d\mathbb{Q}^{\text{approx}}}{d\mathbb{P}}$ via the Radon-Nikodym derivative

$$\begin{aligned} D^n &:= \exp \left(\sum_{0 \leq s_{jp}} q_{s_{jp}}^n \Delta W_{jp}^n + \sum_{0 \leq s_{jp}} \log \left(\frac{\lambda_{s_{jp}}^n}{\lambda_{\mathbb{P}}} \right) \Delta N_{jp}^n \right. \\ &\quad \left. - \sum_{0 \leq s_{jp}} \left(\frac{1}{2} |q_{s_{jp}}^n|^2 + \lambda_{s_{jp}}^n - \lambda_{\mathbb{P}} \right) \Delta_{jp} \right), \quad n = 1, \dots, N_3. \end{aligned}$$

Finally, set

$$\tilde{Y}_0^{\text{low},L} := \frac{1}{N_3} \sum_{n=1}^{N_3} D^n \sum_{l=1}^L f_{\tau^l,n}(X_{\tau^l,n}^n), \quad (5.13)$$

with $f_{\tau^l,n}(X_{\tau^l,n}^n)$, $n = 1, \dots, N_3$, simulated copies of $f_{\tau^l}(X_{\tau^l})$ constructed by applying the numerical scheme of Section 5.2.

Recall from (4.9) that $\rho_0(U)$ gives a lower bound to the upper Snell envelope. Thus, it follows from (5.7), (5.12), and the definition of U that

$$\begin{aligned} \mathbb{E} \left[\tilde{Y}_0^{\text{low},L} \right] &= \mathbb{E} \left[\frac{1}{N_3} \sum_{n=1}^{N_3} D^n \sum_{l=1}^L f_{\tau^l,n}(X_{\tau^l,n}^n) \right] \\ &= \mathbb{E}_{\mathbb{Q}^{\text{approx}}} \left[\sum_{l=1}^L f_{\tau^l}(X_{\tau^l}) \right] \leq \rho_0 \left(\sum_{l=1}^L f_{\tau^l}(X_{\tau^l}) \right) \leq Y_0^{*,L}. \end{aligned} \quad (5.14)$$

That is, our estimator (5.13) constitutes a genuine lower bound. This means that, on average, we indeed obtain a lower bound to the optimal solution given by the upper Snell envelope. Furthermore, as a consequence of Proposition 5.3 and Theorem 4.4, the lower bound converges to the optimal solution.

Our lower bound algorithm can be summarized as follows:

Lower Bound Algorithm.

Given a time grid Π , our explicit numerical construction of the lower bound consists of the following steps:

- (1.) Select K_1 basis functions and run N_1 Monte Carlo simulations to determine \overline{M}^{K_1, N_1} and \overline{c}^{K_1, N_1} . To describe their evolution, it is sufficient to store the corresponding $(\gamma_{s_{jp}}^{N_1})_{j,p}$, $(\beta_{s_{jp}}^{N_1})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_1})_{j,p}$.
- (2.) Use $(\gamma_{s_{jp}}^{N_1})_{j,p}$, $(\beta_{s_{jp}}^{N_1})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_1})_{j,p}$ to estimate $\sum_{l=1}^L f_{\tau^l}(X_{\tau^l})$. Select K_2 basis functions and run N_2 Monte Carlo simulations to determine $(\mathcal{Z}^{N_2}, \tilde{\mathcal{Z}}^{N_2})$ using the estimated $\sum_{l=1}^L f_{\tau^l}(X_{\tau^l})$ as terminal condition. To describe the evolution of this process it is sufficient to store the corresponding $(\beta_{s_{jp}}^{N_2})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_2})_{j,p}$.
- (3.) With $(\beta_{s_{jp}}^{N_2})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_2})_{j,p}$ at hand from Step (2.), simulate N_3 copies of $\frac{d\mathbb{Q}^{\text{approx}}}{d\mathbb{P}}$. Furthermore, with $(\gamma_{s_{jp}}^{N_1})_{j,p}$, $(\beta_{s_{jp}}^{N_1})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_1})_{j,p}$ at hand from Step (1.), simulate N_3 copies of $\sum_{l=1}^L f_{\tau^l,n}(X_{\tau^l,n}^n)$. Using (5.13), a genuine lower bound to the upper Snell envelope is then obtained.

In this algorithm, Step (1.) corresponds to Step (a.) in our General Primal-Dual Pseudo Algorithm of Section 4, whereas Steps (2.) and (3.) fill in details on how to implement the ‘primal’ Step (c.) of that algorithm. Note that the Monte Carlo simulations are done consecutively and are not nested, so that the total computation time depends only on $N_1 + N_2 + N_3$. We summarize the results of this subsection in the following theorem:

Theorem 5.4 *The estimator $\tilde{Y}_0^{\text{low},L}$ defined in (5.13) is a genuine lower bound to the upper Snell envelope, i.e., $\mathbb{E} \left[\tilde{Y}_0^{\text{low},L} \right] \leq Y_0^{*,L}$. Furthermore, $\tilde{Y}_0^{\text{low},L}$ converges to $Y_0^{*,L}$ as N_1, N_2, N_3, K_1 and K_2 tend to infinity and Δ tends to zero, i.e.,*

$$\lim_{\Delta \rightarrow 0} \lim_{K_i \rightarrow \infty, i=1,2} \lim_{N_i \rightarrow \infty, i=1,2,3} \tilde{Y}_0^{\text{low},L} = Y_0^{*,L}.$$

5.3.1 Subtracting the associated ρ -martingale in Eqn. (5.13): Reducing the variance while not inducing a bias

Let $(\mathcal{Z}^N, \tilde{\mathcal{Z}}^N)$ be approximations of the ‘true’ solution to the BSDE with terminal condition U . Denote by \mathbb{Q} a probability measure defined by

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \mathcal{D} \left(\int_0^T H_s dW_s + \int_0^T \tilde{H}_s d\tilde{N}_s \right), \quad \text{with } (H_s, \tilde{H}_s) \in \partial g(s, \mathcal{Z}_s^N, \tilde{\mathcal{Z}}_s^N). \quad (5.15)$$

By (5.7), $\mathbb{E}_{\mathbb{Q}}[U]$ yields a lower bound to $\rho_0(U)$. If the \mathcal{Z}^N and $\tilde{\mathcal{Z}}^N$ were exact, then $\mathbb{Q} = \mathbb{Q}^g$ and $\mathbb{E}_{\mathbb{Q}}[U] = \mathbb{E}_{\mathbb{Q}^g}[U] = \rho_0(U)$.

The following proposition will help to reduce the variation in our numerical scheme.

Proposition 5.5 *Let \overline{M}^N be the ρ -martingale defined by*

$$\overline{M}_t^N := - \int_0^t g(s, \mathcal{Z}_s^N, \tilde{\mathcal{Z}}_s^N) ds + \int_0^t \mathcal{Z}_s^N dW_s + \int_0^t \tilde{\mathcal{Z}}_s^N d\tilde{N}_s.$$

Then, subtracting the ρ -martingale from the terminal condition does not induce a bias, i.e.,

$$\mathbb{E}_{\mathbb{Q}} \left[U - \overline{M}_T^N \right] = \mathbb{E}_{\mathbb{Q}}[U].$$

We finally note that if we would have that $\mathbb{Q} = \mathbb{Q}^g$ from Eqn. (5.12), then, by the definition of \overline{M}_t^N , $U - \overline{M}_T^N = \mathbb{E}_{\mathbb{Q}^g}[U] = \rho_0(U)$ were constant a.s. Hence, if \mathbb{Q} is approximately \mathbb{Q}^g , then $U - \overline{M}_T^N$ is approximately constant. More formally, if \overline{M}_T^N converges to \overline{M}_T in L^2 , we have that

$$\lim_{N \rightarrow \infty} U - \overline{M}_T^N = U - \overline{M}_T = \text{constant},$$

where the convergence should be understood in L^2 . In particular, if $\overline{M}_T^N \rightarrow \overline{M}_T$ in L^2 , then $\text{Var}(U - \overline{M}_T^N) \rightarrow 0$ as $N \rightarrow \infty$.

Inspired by this theoretical result, which may be referred to as an almost sure property of a second kind to distinguish it from the additive dual representation’s almost sure property, we will in our numerical analysis subtract the associated ρ -martingale from the right-hand side of Eqn. (5.13) when computing the genuine lower bound: it will reduce the variance without inducing a bias.

5.4 Converging Approximate and Genuine Upper Bounds

This subsection develops an explicit *approximate* upper bound that converges to the upper Snell envelope asymptotically, and an explicit *genuine* upper bound that not only converges to the upper Snell envelope but is also biased high at the pre-limiting level. To obtain an upper bound with Theorem 3.1, we set

$$U := \max_{0 \leq j_1 < j_2 < \dots < j_L} \sum_{l=1}^L \left(f_{j_l}(X_{j_l}) - M_{j_l}^{*,L-l+1} + M_{j_{l-1}}^{*,L-l+1} \right).$$

Since we cannot compute (M^*) exactly, we approximate, in view of (4.7), the terminal condition U by the ρ -martingale (\bar{M}^{K_1, N_1}) constructed in Section 5.2, i.e., we set

$$U^{N_1} := \max_{0 \leq j_1 < j_2 < \dots < j_L} \sum_{l=1}^L \left(f_{j_l}(X_{j_l}) - \bar{M}_{j_l}^{L-l+1, K_1, N_1} + \bar{M}_{j_{l-1}}^{L-l+1, K_1, N_1} \right).$$

Next, we define

$$\mathcal{X}_{t_i}^{N_1} := \left(X_{t_i}, \bar{M}_{t_i}^{K_1, N_1}, \max_{k \leq L: 0 \leq j_1 < j_2 < \dots < j_k \leq \lfloor t_i \rfloor} \sum_{l=1}^k \left(f_{j_l}(X_{j_l}) - \bar{M}_{j_l}^{L-l+1, K_1, N_1} + \bar{M}_{j_{l-1}}^{L-l+1, K_1, N_1} \right) \right).$$

Clearly, the terminal condition U^{N_1} depends only on $\mathcal{X}_T^{N_1}$. Next, note that, for every t_i , we have that H_{t_i} is a function of X_{t_i} , and, for every $r > i$, $\bar{M}_{t_r}^{K_1, N_1}$ only depends on $\bar{M}_{t_i}^{K_1, N_1}$, $(X_{t_l})_{i \leq l \leq r}$, $(W_{t_l} - W_{t_i})_{i \leq l \leq r}$ and $(N_{t_l} - N_{t_i})_{i \leq l \leq r}$ where both of the latter are independent of \mathcal{F}_{t_r} . From this we may conclude that \mathcal{X}^{N_1} is a Markov process on the time grid Π .

Next, we solve numerically the BSDE (5.8) with U^{N_1} . To do so we will consider an approximation scheme. We can simulate paths of the adapted process

$$\begin{aligned} (\mathcal{X}_{t_i}^{N_1, n})_i &= \left(X_{t_i}^n, \bar{M}_{t_i}^{K_1, N_1, n}, \right. \\ &\quad \left. \max_{k \leq L: 0 \leq j_1 < j_2 < \dots < j_k \leq \lfloor t_i \rfloor} \sum_{l=1}^k \left(f_{j_l}(X_{j_l}^n) - \bar{M}_{j_l}^{L-l+1, K_1, N_1, n} + \bar{M}_{j_{l-1}}^{L-l+1, K_1, N_1, n} \right) \right)_i, \end{aligned}$$

for $n = 1, \dots, N_4$. To compute the BSDE with terminal condition U^{N_1} , let K_4 be the number of basis functions in the least squares Monte Carlo regression. Employing the algorithm described in Section 5.2, we can construct the coefficients $\gamma_{s_{jp}}^{N_4}$, $\beta_{s_{jp}}^{N_4}$ and $\tilde{\beta}_{s_{jp}}^{N_4}$, and processes

$$(\bar{Y}^{N_4, L}, \mathcal{Z}^{N_4}, \tilde{\mathcal{Z}}^{N_4}). \quad (5.16)$$

Note that by applying Proposition 5.3 and Theorem 4.4 twice we may conclude that, in L^2 ,

$$\lim_{\Delta \rightarrow 0} \lim_{K_i \rightarrow \infty, i=1,4} \lim_{N_i \rightarrow \infty, i=1,4} (\bar{Y}^{N_4, L}, \mathcal{Z}^{N_4}, \tilde{\mathcal{Z}}^{N_4}) = (Y^{*, L}, \mathcal{Z}^*, \tilde{\mathcal{Z}}^*). \quad (5.17)$$

In particular, $\bar{Y}^{N_4, L}$ constitutes a converging approximate upper bound to the upper Snell envelope. From now on we assume that we have already estimated $\bar{Y}_0^{N_4}$, $(\beta_{s_{jp}}^{N_4})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_4})_{j,p}$.

Next, let us develop a genuine upper bound. It is well-known that under assumption (A1), the functional ρ is Lipschitz continuous (cf. Peng [65]). Define

$$\begin{aligned}\hat{U}^{N_4} &:= \bar{Y}_0^{N_4, L} - \int_0^T g(s, \mathcal{Z}_s^{N_4}, \tilde{\mathcal{Z}}_s^{N_4}) ds + \int_0^T \mathcal{Z}_s^{N_4} dW_s + \int_0^T \tilde{\mathcal{Z}}_s^{N_4} d\tilde{N}_s \\ &= \bar{Y}_0^{N_4, L} - \sum_{j,p} g(s_{jp}, \beta_{s_{jp}}^{N_4} \psi(s_{jp}, \mathcal{X}_{s_{jp}}^{N_1}), \tilde{\beta}_{s_{jp}}^{N_4} \tilde{\psi}(s_{jp}, \mathcal{X}_{s_{jp}}^{N_1})) \Delta_{s_{jp}} \\ &\quad + \sum_{j,p} \beta_{s_{jp}}^{N_4} \psi(s_{jp}, \mathcal{X}_{s_{jp}}^{N_1}) \Delta W_{s_{jp}} + \sum_{j,p} \tilde{\beta}_{s_{jp}}^{N_4} \tilde{\psi}(s_{jp}, \mathcal{X}_{s_{jp}}^{N_1}) \Delta \tilde{N}_{s_{jp}}.\end{aligned}\quad (5.18)$$

Then, by Theorem 3.1 and the Lipschitz continuity of ρ , we have that

$$\begin{aligned}Y_0^{*, L} \leq \rho_0(U^{N_1}) &\leq \rho_0(\hat{U}^{N_4}) + \mathcal{K} \|\hat{U}^{N_4} - U^{N_1}\|_2 \\ &= \bar{Y}_0^{N_4, L} + \mathcal{K} \|\hat{U}^{N_4} - U^{N_1}\|_2,\end{aligned}\quad (5.19)$$

for a Lipschitz constant \mathcal{K} analyzed later. We will exploit inequality (5.19) to develop our genuine upper bound. By Proposition 5.3 and Theorem 4.4, $\bar{Y}_0^{N_4, L}$ converges to $Y_0^{*, L}$, $(\mathcal{Z}^{N_4}, \tilde{\mathcal{Z}}^{N_4})$ converges to $(\mathcal{Z}^*, \tilde{\mathcal{Z}}^*)$ in $L^2(d\mathbb{P} \times ds)$ and $\|\hat{U}^{N_4} - U^{N_1}\|_2$ converges to zero, as N_1, N_4, K_1 and K_4 tend to infinity and Δ tends to zero.

For $N_5 \in \mathbb{N}$, simulate i.i.d. copies of U^{N_1} through

$$U_n^{N_1} = \max_{0 \leq j_1 < j_2 < \dots < j_L} \sum_{l=1}^L \left(f_{j_l}(X_{j_l}^n) - \bar{M}_{j_l}^{L-l+1, K_1, N_1, n} + \bar{M}_{j_{l-1}}^{L-l+1, K_1, N_1, n} \right), \quad n = 1, \dots, N_5.$$

Next, using (5.18), simulate i.i.d. copies $\hat{U}_1^{N_4}, \hat{U}_2^{N_4}, \dots, \hat{U}_{N_5}^{N_4}$ of \hat{U}^{N_4} . Recall that $\bar{Y}_0^{N_4, L}$, $(\beta_{s_{jp}}^{N_4})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_4})_{j,p}$ are already available. Then, (5.19) suggests to estimate the upper bound $\rho_0(U^{N_1})$ by

$$\bar{Y}_0^{N_4, L} + \mathcal{K} \sqrt{\frac{1}{N_5} \sum_{n=1}^{N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2}.$$

Note that $\frac{1}{N_5} \sum_{i=1}^{N_5} |\hat{U}_i^{N_4} - U_i^{N_1}|^2$ is an unbiased estimator of $\mathbb{E} \left[|U^{N_1} - \hat{U}^{N_4}|^2 \right]$. However, taking the square root of an estimator gives rise to a possible downward bias. If we wish to eliminate the downward bias, we need to develop one step further, as follows. We simulate independent $(\hat{U}_i^{N_4})_{i=1, \dots, 2N_5}$. Then we set

$$\tilde{Y}_0^{\text{upp}, L} := \bar{Y}_0^{N_4, L} + \mathcal{K} \frac{\frac{1}{N_5} \sum_{n=1}^{N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2}{\sqrt{\frac{1}{N_5} \sum_{n=N_5+1}^{2N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2}}.\quad (5.20)$$

Since

$$\begin{aligned}
& \mathbb{E} \left[\frac{\frac{1}{N_5} \sum_{n=1}^{N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2}{\sqrt{\frac{1}{N_5} \sum_{n=N_5+1}^{2N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2}} \right] \\
&= \mathbb{E} \left[\frac{1}{N_5} \sum_{n=1}^{N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2 \right] \mathbb{E} \left[\frac{1}{\sqrt{\frac{1}{N_5} \sum_{n=N_5+1}^{2N_5} |\hat{U}_n^{N_4} - U_n^{N_1}|^2}} \right] \\
&\geq \|\hat{U}_n^{N_4} - U_n^{N_1}\|^2 \frac{1}{\sqrt{\frac{1}{N_5} \sum_{n=N_5+1}^{2N_5} \mathbb{E} [|\hat{U}_n^{N_4} - U_n^{N_1}|^2]}} \\
&= \|\hat{U}_n^{N_4} - U_n^{N_1}\|^2 \frac{1}{\sqrt{\mathbb{E} [|\hat{U}_n^{N_4} - U_n^{N_1}|^2]}} = \|\hat{U}_n^{N_4} - U_n^{N_1}\|,
\end{aligned}$$

$\tilde{Y}_0^{\text{upp},L}$ thus defined is biased high. Note that the first equality follows from independence and the inequality is due to a suitable (reversed) application of Jensen's inequality. As before, as a consequence of Proposition 5.3 and Theorem 4.4, we have that $\tilde{Y}_0^{\text{upp},L}$ converges to $Y_0^{*,L}$ as N_1, N_4, N_5, K_1 and K_4 tend to infinity and Δ tends to zero.

Our upper bound algorithm can be summarized as follows:

Upper Bound Algorithm.

Given a time grid Π , our explicit numerical construction of the upper bound consists of the following steps:

- (1.) Select K_1 basis functions and run N_1 Monte Carlo simulations to determine \bar{M}^{K_1, N_1} and \bar{c}^{K_1, N_1} . To describe their evolution, it is sufficient to store the corresponding $(\gamma_{s_{jp}}^{N_1})_{j,p}$, $(\beta_{s_{jp}}^{N_1})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_1})_{j,p}$.
- (2.) $(\gamma_{s_{jp}}^{N_1})_{j,p}$, $(\beta_{s_{jp}}^{N_1})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_1})_{j,p}$ give rise to a terminal condition U^{N_1} and a Markov process \mathcal{X}^{N_1} defined above. Select K_2 basis functions and run N_4 Monte Carlo simulations to calculate $(\bar{Y}^{N_4, L}, \mathcal{Z}^{N_4}, \tilde{\mathcal{Z}}^{N_4})$ as the solution to the corresponding BSDEs with the Markov process \mathcal{X}^{N_1} and terminal condition U^{N_1} . Store the corresponding $\bar{Y}_0^{N_4, L}$, $(\beta_{s_{jp}}^{N_4})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_4})_{j,p}$.
- (3.) With $\bar{Y}_0^{N_4, L}$, $(\beta_{s_{jp}}^{N_4})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_4})_{j,p}$ at hand from Step (2.), simulate N_5 copies of \hat{U}^{N_4} defined by (5.18). Furthermore, with $(\gamma_{s_{jp}}^{N_1})_{j,p}$, $(\beta_{s_{jp}}^{N_1})_{j,p}$ and $(\tilde{\beta}_{s_{jp}}^{N_1})_{j,p}$ at hand from Step (1.), simulate N_5 copies of U^{N_1} . Using (5.20), a genuine upper bound to the upper Snell envelope is then obtained.

In this algorithm, Step (1.) corresponds to Step (a.) in our General Primal-Dual Pseudo Algorithm of Section 4, whereas Steps (2.) and (3.) fill in details on how to implement the ‘dual’ Step (b.) of that algorithm. When both an upper bound and a lower bound are computed, using this algorithm and the Lower Bound Algorithm described above, then Step (1.) needs to be executed only once as the two algorithms have it in common. Note again that the Monte

Carlo simulations are done consecutively and are not nested, so that the total computation time depends only on $N_1 + N_4 + N_5$. We summarize the results of this subsection in the following theorem:

Theorem 5.6 *The estimator $\tilde{Y}_0^{\text{upp},L}$ defined in (5.20) is a genuine upper bound to the upper Snell envelope, i.e., $\mathbb{E} \left[\tilde{Y}_0^{\text{upp},L} \right] \geq Y_0^{*,L}$. Furthermore, $\tilde{Y}_0^{\text{upp},L}$ converges to $Y_0^{*,L}$ as N_1, N_4, N_5, K_1 and K_4 tend to infinity and Δ tends to zero, i.e.,*

$$\lim_{\Delta \rightarrow 0} \lim_{K_i \rightarrow \infty, i=1,4} \lim_{N_i \rightarrow \infty, i=1,4,5} \tilde{Y}_0^{\text{upp},L} = Y_0^{*,L}.$$

Moreover, the estimator $\bar{Y}_0^{N_4,L}$ defined in (5.16) also converges to $Y_0^{*,L}$ as N_1, N_4, K_1 and K_4 tend to infinity and Δ tends to zero.

We finally state the following proposition on the precise Lipschitz constant \mathcal{K} appearing in (5.19).

Proposition 5.7 *Let ξ and ξ' be square-integrable terminal conditions and denote by (Y, Z, \tilde{Z}) and (Y', Z', \tilde{Z}') the associated BSDE solutions. Then,*

$$|Y_0 - Y'_0|^2 \leq \exp(\mathcal{L}^2 T) \mathbb{E} [|\delta\xi|^2], \quad (5.21)$$

where $\delta\xi := \xi - \xi'$ and with \mathcal{L} the Lipschitz constant of the driver g .

6 Numerical Examples

In this section we analyze our approach in numerical examples, including single and multiple stopping, univariate and multivariate stochastic drivers, increasing and decreasing reward functions, and pure diffusion and jump-diffusion models. As a general observation, we recall that the computational complexity of the numerically implementable method proposed in this paper is linear in the number of exercise rights and that it does not require nested simulation.

6.1 Single Stopping: Bermudan Option in a Diffusion Model

The first example studied in this subsection is the pricing problem for a Bermudan-style option in a single risky asset Black-Scholes model with dividends in the presence of ambiguity. This example goes back to Andersen and Broadie [1] in a setting without ambiguity, which can serve as a benchmark case. Throughout this subsection, we consider the dynamics $\frac{dX_t^i}{X_t^i} = \mu^i dt + \sigma^i dW_t^i$, $i = 1, \dots, d$, with $\mu^i \in \mathbb{R}$ and $\sigma^i \in \mathbb{R}_{>0}$, where X_t^i is the price of asset i at time t . Following this literature, we assume that there is a risk-free interest rate of $\rho = 0.05$, and that the option's underlying is a single dividend-paying stock X with constant volatility $\sigma = 0.2$ and dividend rate $\delta = 0.1$, resulting in a risk-neutral drift of $\mu = \rho - \delta = -0.05$. The first product we study is a call option with strike price $K = 100$ and maturity $T = 3$. The stock price at time 0 is varied between $x_0 \in \{90, 100, 110\}$. Exercise dates are specified as $t_j = \frac{jT}{10}$, $j = 0, 1, \dots, 10$, i.e., there are 9 intermediate exercise dates, and the trivial ones at time $t = 0$ and at maturity. We allow for ambiguity in the drift and consider $g(t, z) = \delta_1 |z|$, $\delta_1 > 0$; cf. Example 5.1.

The three BSDEs— ρ -martingale, lower, and upper bounds—are solved with two sets of simulations. We consider one set of simulations with 100,000 trajectories and 1,000 time steps for the initial ρ -martingale BSDE, and a second set of simulations with 100,000 trajectories and 1,000 time steps for the BSDEs associated with the lower and upper bounds. The number of basis functions is always 52 for Y and 52 for Z . These are $1, x, (x - q_{it})^+$ where the q_{it} are 1, 3, 5, \dots , 99 percent quantiles of X_t estimated from the trajectories in the initial run of least squares Monte Carlo. Thus, we approximate all unknown functions by linear splines. Notice that this function basis is not problem-dependent and that its precision can be controlled by the chosen grid of quantiles. Numerically, this basis works much better in our experiments than a locally linear approximation, which would also include the discontinuous terms $1_{\{x > q_{it}\}}$. In all our numerical experiments, the implementation of the regression is based on (5.10).

For the evaluation of the lower bound, we draw a new sample with a larger number of 400,000 trajectories. For the evaluation of the upper bound, we need a very fine time discretization to obtain a small ‘tracking error’, defined as $\sqrt{\frac{1}{N_5} \sum_{i=1}^{N_5} |\hat{U}_i^{N_4} - U_i^{N_1}|^2}$; cf. (5.20). We need, however, fewer trajectories, because the pathwise dual representation leads to a very small variance, as expected from the theoretical results in Section 3.2. We thus choose the number of trajectories to be 1,000 and increase the number of time steps by a factor 100 to 100,000. Here, we do not run new regressions but simply repeat each set of coefficients 100 times. This device of increasing the time discretization and extrapolating regression coefficients is due to Belomestny, Bender, Schoenmakers [10], in a standard stopping setting without ambiguity. However, due to the additional non-linearity from the BSDE, it is not sufficient in our setting to just run the regressions at exercise times only; we need them at our fine grid Π .

Tables 1–3 summarize lower and upper bounds for different degrees of ambiguity (δ_1) and different values of the initial stock price (x_0). The first four columns display lower bounds along with their standard errors. Here, “LB without M” is the lower bound in Eqn. (5.13) without subtraction of the ρ -martingale, while “LB” is the definitive lower bound that subtracts the ρ -martingale, as discussed in Section 5.3.1. As the results confirm, subtraction of the ρ -martingale leads to a substantial variance reduction. The next three columns correspond to the upper bound. Here, $\bar{Y}_0^{N_4}$ is the solution to the upper bound BSDE, i.e., the approximate upper bound, “TE” is the tracking error defined above, and “UB” combines the two to obtain the genuine upper bound in Eqn. (5.20). The final two columns display the mean of $\bar{Y}_T^{N_4} = U^{N_1}$ together with its standard error. The point here is to illustrate that having only 1,000 trajectories in the upper bound simulation is already sufficient, since the terminal condition has a (very) small variance. In general, we observe that the gaps between LB on the one hand and $\bar{Y}_0^{N_4}$ and UB on the other hand are (very) small. When $1/\delta_1 \equiv \infty$ (i.e., in the no-ambiguity case), we can compare our results to benchmark values, e.g., from Andersen and Broadie [1]. With $x_0 = 100$, the true value is 7.98. This should be compared to our genuine lower bound of 7.98, our approximate upper bound of 8.00, and our genuine upper bound of 8.07, corresponding to gaps of 0.2% and 1.0% of the option value, respectively.

In Figure 1 we analyze exercise boundaries. These boundaries are theoretically independent of the starting values of x_0 . In order to obtain accurate exercise boundaries already for early time points, we need to make sure that the simulated trajectories are sufficiently spread out over the entire time span from 0 to T and are not concentrated around a fixed value of x_0 close to time 0. To this end, we start simulating trajectories at time -1 and set the drift of X to zero over the interval $[-1, 0]$. Moreover, we double the number of basis functions by choosing a finer

$1/\delta_1$	LB without M	s.e.	LB	s.e.	$\bar{Y}_0^{N_4}$	TE	UB	$\mathbb{E}[\bar{Y}_T^{N_4}]$	s.e.
10	5.4901	0.0197	5.4635	0.0107	5.4833	0.0843	5.5689	5.4679	0.0026
30	4.7384	0.0158	4.7161	0.0096	4.7326	0.0836	4.8163	4.7184	0.0025
100	4.5005	0.0147	4.4795	0.0093	4.4962	0.0812	4.5775	4.4830	0.0025
300	4.4349	0.0144	4.4143	0.0092	4.4310	0.0806	4.5116	4.4180	0.0024
1,000	4.4121	0.0143	4.3917	0.0092	4.4084	0.0804	4.4888	4.3955	0.0024
10,000	4.4030	0.0142	4.3827	0.0092	4.3997	0.0803	4.4800	4.3868	0.0024
∞	4.4024	0.0142	4.3821	0.0092	4.3987	0.0803	4.4790	4.3859	0.0024

Table 1: Bounds for $x_0 = 90$

$1/\delta_1$	LB without M	s.e.	LB	s.e.	$\bar{Y}_0^{N_4}$	TE	UB	$\mathbb{E}[\bar{Y}_T^{N_4}]$	s.e.
10	9.4387	0.0244	9.4069	0.0143	9.4331	0.0721	9.5063	9.4110	0.0021
30	8.4419	0.0199	8.4161	0.0132	8.4445	0.0700	8.5146	8.4232	0.0020
100	8.1305	0.0186	8.1068	0.0129	8.1316	0.0678	8.1994	8.1113	0.0019
300	8.0426	0.0183	8.0195	0.0128	8.0450	0.0678	8.1127	8.0249	0.0019
1,000	8.0152	0.0182	7.9923	0.0128	8.0150	0.0678	8.0828	7.9949	0.0019
10,000	8.0039	0.0181	7.9811	0.0127	8.0034	0.0678	8.0712	7.9833	0.0019
∞	8.0030	0.0181	7.9802	0.0127	8.0022	0.0678	8.0699	7.9821	0.0019

Table 2: Bounds for $x_0 = 100$

$1/\delta_1$	LB without M	s.e.	LB	s.e.	$\bar{Y}_0^{N_4}$	TE	UB	$\mathbb{E}[\bar{Y}_T^{N_4}]$	s.e.
10	14.7691	0.0279	14.7380	0.0179	14.7718	0.0682	14.8411	14.7460	0.0020
30	13.6662	0.0228	13.6414	0.0166	13.6814	0.0669	13.7484	13.6570	0.0019
100	13.3221	0.0214	13.2994	0.0163	13.3400	0.0658	13.4058	13.3164	0.0019
300	13.2257	0.0211	13.2037	0.0162	13.2459	0.0658	13.3116	13.2225	0.0019
1,000	13.1924	0.0209	13.1706	0.0162	13.2133	0.0657	13.2790	13.1900	0.0019
10,000	13.1791	0.0209	13.1574	0.0162	13.2007	0.0657	13.2664	13.1775	0.0019
∞	13.1774	0.0209	13.1556	0.0162	13.1994	0.0657	13.2650	13.1761	0.0019

Table 3: Bounds for $x_0 = 110$

grid of quantiles, 0.005, 0.0015, \dots , 0.995, and double the number of trajectories to 200,000, to account for the fact that we now have a wider space over which we approximate.

In the left panel of Figure 1, we plot the threshold value that the stock price has to exceed to make stopping optimal (i.e., the continuation value), for different degrees of ambiguity (δ_1), as a function of the exercise dates. The solid line depicts exercise boundaries without ambiguity ($1/\delta_1 = \infty$), while the dashed line corresponds to $1/\delta_1 = 100$ and the dotted line corresponds to $1/\delta_1 = 10$. In all cases, continuation values increase as we move away from the no-ambiguity case ($1/\delta_1 = \infty$). Thus, with more ambiguity, the decision-maker stops later. Intuitively, for a call option's payoff function, the possibility of a (prosperous) deviation from the reference model which accumulates over time makes the option more valuable at later dates. Hence, the decision-maker will stop later.

A similar pattern emerges when we replace the call option payoff function by a put option, with strike price $K = 100$ and $x_0 = 100$. To make exercise decisions non-trivial, we set the dividend rate equal to zero in this case, all else equal. Table 4 and Figure 1 (right panel) display the corresponding price bounds and exercise boundaries. For the put option, contrary to for the call option, exercise becomes optimal when the stock price falls *below* the exercise




Figure 1: Exercise boundaries (left panel: call; right panel: put; solid: $1/\delta_1 = \infty$; dashed: $1/\delta_1 = 100$; and dotted: $1/\delta_1 = 10$).

boundaries in the right panel of Figure 1. Similar to the call option, with more ambiguity, the decision-maker stops later. In this case, the possibility of an unfavorable deviation from the reference model that accumulates over time makes the option more valuable at later dates.

$1/\delta_1$	LB without M	s.e.	LB	s.e.	$\bar{Y}_0^{N_4}$	TE	UB	$\mathbb{E}[\bar{Y}_T^{N_4}]$	s.e.
10	9.8647	0.0202	9.8810	0.0105	9.8767	0.0482	9.9256	9.8678	0.0015
30	8.9584	0.0169	8.9749	0.0101	8.9723	0.0480	9.0204	8.9622	0.0015
100	8.6643	0.0160	8.6808	0.0100	8.6785	0.0478	8.7264	8.6684	0.0014
300	8.5832	0.0157	8.5997	0.0100	8.5967	0.0476	8.6443	8.5865	0.0014
1,000	8.5540	0.0156	8.5705	0.0100	8.5683	0.0475	8.6157	8.5581	0.0014
10,000	8.5437	0.0156	8.5603	0.0100	8.5573	0.0475	8.6048	8.5471	0.0014
∞	8.5427	0.0156	8.5593	0.0100	8.5561	0.0475	8.6036	8.5459	0.0014

Table 4: Bounds for $x_0 = 100$ (put option)

Next, we consider a two-dimensional version of this example. We suppose that there are two risky assets X^1 and X^2 , which are assumed to be independent and identically distributed with the same dynamics as the single dividend-paying stock in the univariate case. The function g is now chosen to be $g(t, z) = \delta_1|z_1| + \delta_1|z_2|$. For the processes q_t that determine our uncertainty set, this implies that $|q_t^i| \leq \delta_1$ for $i = 1, 2$. The payoff function we consider is a max-call, that is, stopping at time t yields a reward of $(\max(X_t^1, X_t^2) - K)^+$. We set $X_0^1 = X_0^2 = 100$ and $K = 100$, and allow for eleven equidistant exercise opportunities including 0 and T , as before. All other problem parameters and specifications remain the same.

Regarding the numbers of trajectories and time steps in the different stages of the algorithm, we maintain the same specifications as in the univariate case. The function basis is constructed as follows. We always use the same set of 441 basis functions for Y and for (Z, \tilde{Z}) . These consist of the constant 1, 20 univariate basis functions $\phi_i^{(1)}(x^1)$, $i = 1, \dots, 20$, which only depend on X^1 , 20 univariate basis functions $\phi_j^{(2)}(x^2)$, $j = 1, \dots, 20$, which only depend on X^2 , and all products $\phi_i^{(1)}(x^1)\phi_j^{(2)}(x^2)$, $i, j = 1, \dots, 20$. The 20 univariate basis functions are constructed

as in the univariate case but with a slightly coarser grid, i.e., for $d = 1, 2$, we choose x^d and $(x^d - q_{it}^{(d)})^+$ where the $q_{it}^{(d)}$ are the 5, 10, \dots , 95 percent quantiles of X_t^d estimated from the trajectories in the initial run of least squares Monte Carlo. This bivariate basis of linear splines is relatively large but fairly generic, i.e., it exploits additional knowledge about the problem far less than, e.g., the $2d$ -implementations in Belomestny, Bender and Schoenmakers [10] or Krättschmer *et al.* [54], which rely on prices of European max-call options. In principle, one could increase efficiency by including a variable selection step in the first regression.

We observe from Table 5 that the gaps between the genuine lower bound and the approximate upper bound are fairly small, corresponding to about only 0.4% of the option value. The presence of ambiguity amplifies in the multivariate setting and its impact is more pronounced than in the univariate case.

$1/\delta_1$	LB without M	s.e.	LB	s.e.	$\bar{Y}_0^{N_4}$	$\mathbb{E}[\bar{Y}_T^{N_4}]$	s.e.
10	16.5275	0.0352	16.5257	0.0196	16.5958	16.5722	0.0231
30	14.7579	0.0268	14.7519	0.0172	14.8068	14.7998	0.0217
100	14.1999	0.0246	14.1923	0.0167	14.2421	14.2367	0.0212
300	14.0435	0.0240	14.0350	0.0166	14.0862	14.0808	0.0210
1,000	13.9916	0.0238	13.9830	0.0165	14.0322	14.0268	0.0210
10,000	13.9700	0.0238	13.9618	0.0165	14.0115	14.0061	0.0210
∞	13.9679	0.0237	13.9597	0.0165	14.0092	14.0038	0.0210

Table 5: Bounds for $x_0^1 = x_0^2 = 100$ (bivariate case)

6.2 Multiple Stopping: Swing Option in a Two-Factor Jump-Diffusion Model

Supported by the accuracy and stability of our pathwise duality approach for optimal single stopping, we now proceed to multiple stopping. In this subsection, we analyze a canonical multiple stopping problem, that of swing option pricing in electricity markets, in the presence of ambiguity. For this purpose, we consider a two-factor jump-diffusion model for the electricity log-price process, which has been suggested by Hambly, Howison and Kluge [45] to be a more realistic extension of the one-factor Gaussian model proposed by Lucia and Schwartz [60] and implemented e.g., by Bender, Schoenmakers and Zhang [15].

Specifically, we assume that the electricity price X_t at time $t > 0$ is given by $X_t = X_0 \exp(f(t) + u_t + v_t)$, where the two stochastic factors u_t and v_t are mutually independent and $f(t)$ is a deterministic function of time that can be used to calibrate the model. The factor u is Gaussian and follows the SDE

$$du_t = -\kappa_u u_t + \sigma_u dW_t,$$

with $u_0 = 0$, $\kappa_u, \sigma_u > 0$, and W_t a standard Brownian motion. The jump component v follows the SDE

$$dv_t = -\kappa_v v_{t-} + J dN_t,$$

where $v_0 = 0$, N is a (non-compensated) Poisson process with arrival rate $\lambda_{\mathbb{P}}$, and κ_v and the (deterministic) jump size J are positive constants. In the special case $\kappa_u \equiv \kappa_v$, this model reduces to a mean-reverting one-factor jump-diffusion model for the log-price process.

For $\kappa_v \gg \kappa_u$, the model combines a mean-reverting Gaussian component, like in the Lucia-Schwartz model, with occasional highly transitory spikes; see Hambley, Howison and Kluge [45] for further discussion. Using their formula (3), the processes u and v can be simulated forward in time without discretization error.

In our illustration, we consider a swing option contract that gives the owner the right to purchase electricity at a strike price K , and consider L exercise rights, in the time interval $[0, T]$. We assume a fixed number of equidistant exercise opportunities. We set the parameters of the price process as $S_0 = 10$, $f \equiv 0$, $\kappa_u = 10$, $\sigma_u = 0.25$, $\kappa_v = 50$, $\lambda_{\mathbb{P}} = 1$ and $J \in \{0, 0.06\}$. Furthermore, we set the contract parameters as $K = 10$, $T = 5$ and consider 21 equidistant exercise opportunities (including one at time 0 and one at time T). We allow for varying degrees of ambiguity towards the Gaussian and the jump components of the price process and consider $g(t, z, \tilde{z}) = \delta_1|z| + \delta_2|\tilde{z}|$, $\delta_1, \delta_2 > 0$; cf. Example 5.1.

The overall numerical implementation is very similar to the previous optimal single stopping example. The three BSDEs— ρ -martingale, lower, and upper bounds—are solved with two sets of simulations. We consider one set of simulations with 100,000 trajectories and 1,000 time steps for the initial ρ -martingale BSDE and a second set of simulations with 100,000 trajectories and 1,000 time steps for the BSDEs associated with the lower and upper bounds. We choose the same basis functions as in the single stopping example and thus obtain 52 basis functions for Y and 52 for Z , and now also 52 basis functions for \tilde{Z} . Note that these functions depend only on X but not on u and v . For the evaluation of the lower bound, we draw a new sample of 100,000 trajectories and 1,000 time steps. For the evaluation of the upper bound, we again artificially create a finer time discretization, by repeating each set of coefficients 100 times, and reduce the number of trajectories to 1,000.

Tables 6–7 and E.1–E.6 (in the Appendix) summarize lower and upper bounds for different numbers of exercise rights, different values of δ_1 and δ_2 , and for different values of J (i.e., without and with jump component). Upon comparing the results for one exercise right to those for multiple exercise rights, we readily observe that the impact of ambiguity is even more pronounced in the multiple stopping case.

L	1	2	3	4	5
LB	0.9526	1.7020	2.3178	2.8290	3.2539
s.e.	0.0012	0.0018	0.0023	0.0027	0.0030
$\bar{Y}_0^{N_4, L}$	0.9599	1.7138	2.3319	2.8453	3.2717
TE	0.0315	0.0407	0.0484	0.0526	0.0568
UB	0.9914	1.7546	2.3803	2.8979	3.3284

Table 6: Bounds for $\delta_1 = 0$ and $J = 0$

L	1	2	3	4	5
LB	0.9796	1.7572	2.4040	2.9293	3.3946
s.e.	0.0027	0.0048	0.0065	0.0079	0.0092
$\bar{Y}_0^{N_4, L}$	1.0129	1.8174	2.4851	3.0470	3.5205
TE	0.0331	0.0426	0.0469	0.0520	0.0562
UB	1.0495	1.8644	2.5370	3.1045	3.5826

Table 7: Bounds for $\delta_1 = 0.2$ and $J = 0$

Acknowledgements. We are very grateful to an Area Editor and two anonymous referees for

their careful reading, comments and suggestions that have significantly improved the paper. We are also grateful to Christian Bender and seminar and conference participants at the University of Waterloo and the Lorentz Center for their comments and suggestions. This research was funded in part by the Netherlands Organization for Scientific Research (NWO) under grants NWO-Vidi and NWO-Vici (Laeven) and by the DFG Excellence Cluster Math+ Berlin, project AA4-2 (Schoenmakers).

References

- [1] ANDERSEN, L., AND BROADIE, M. (2004). A primal-dual simulation algorithm for pricing multi-dimensional American options. *Management Science* 50, 1222–1234.
- [2] ARTZNER, P., DELBAEN, F., EBER, J., AND HEATH, D. (1999). Coherent measures of risk. *Mathematical Finance* 9, 203–228.
- [3] BALDER, S., MAHAYNI, A., AND SCHOENMAKERS, J. (2013). Primal-dual linear Monte Carlo algorithm for multiple stopping — An application to flexible caps. *Quantitative Finance* 13, 1003–1013.
- [4] BARRIEU, P., AND EL KAROUI, N. (2009). Pricing, hedging and optimally designing derivatives via minimization of risk measures. In: Carmona, R. (ed.). *Indifference Pricing: Theory and Applications*. Princeton University Press, Princeton, 77–146.
- [5] BAYRAKTAR, E., KARATZAS, I., AND YAO, S. (2010). Optimal stopping for dynamic convex risk measures. *Illinois Journal of Mathematics* 54, 1025–1067.
- [6] BAYRAKTAR, E., AND YAO, S. (2011). Optimal stopping for non-linear expectations—Part I. *Stochastic Processes and their Applications* 121, 185–211.
- [7] BAYRAKTAR, E., AND YAO, S. (2014). On the robust optimal stopping problem. *SIAM Journal on Control and Optimization* 52, 3135–3175.
- [8] BAYRAKTAR, E., AND YAO, S. (2017a). On the robust Dynkin game. *The Annals of Applied Probability* 27, 1702–1755.
- [9] BAYRAKTAR, E., AND YAO, S. (2017b). Optimal stopping with random maturity under nonlinear expectations. *Stochastic Processes and their Applications* 127, 2586–2629.
- [10] BELOMESTNY, D., BENDER, C., AND SCHOENMAKERS, J. (2009). True upper bounds for Bermudan products via non-nested Monte Carlo. *Mathematical Finance* 19, 53–71.
- [11] BELOMESTNY, D., AND KRÄTSCHMER, V. (2016). Optimal stopping under model uncertainty: A randomized stopping times approach. *The Annals of Applied Probability* 26, 1260–1295.
- [12] BELOMESTNY, D., AND KRÄTSCHMER, V. (2017). Optimal stopping under probability distortions and law invariant coherent risk measures. *Mathematics of Operations Research* 42, 806–833.
- [13] BELOMESTNY, D., KALEDIN, M., AND SCHOENMAKERS, J. (2020). Semi-tractability of optimal stopping problems via a weighted stochastic mesh algorithm. *Mathematical Finance* 30, 1591–1616.
- [14] BENDER, C., AND STEINER, J. (2013). A posteriori estimates for Backward SDEs. *SIAM Journal of Uncertainty Quantification* 1, 139–163.
- [15] BENDER, C., SCHOENMAKERS, J., AND ZHANG, J. (2015). Dual representations for general multiple stopping problems. *Mathematical Finance* 25, 339–370.
- [16] BENDER, C., SCHWEIZER, N., AND ZHUO, J. (2017). A primal-dual algorithm for BSDEs. *Mathematical Finance* 27, 866–901.
- [17] BENDER, C., GÄRTNER, C., AND SCHWEIZER, N. (2018). Pathwise dynamic programming. *Mathematics of Operations Research* 43, 965–995.
- [18] BEN-TAL, A., AND NEMIROVSKI, A. (1998). Robust convex optimization. *Mathematics of Operations Research* 23, 769–805.
- [19] BERTSIMAS, D., AND BROWN, D.B. (2009). Constructing uncertainty sets for robust linear optimization. *Operations Research* 57, 1483–1495.
- [20] BJÖRK, T., AND SLINKO, I. (2006). Towards a general theory of good-deal bounds. *Review of Finance* 10, 221–260.
- [21] BRIAND, P., DELYON, B., AND MÉMIN, J. (2002). On the robustness of backward stochastic differential equations. *Stochastic Processes and their Applications* 97, 229–253.
- [22] BROADIE, M., AND GLASSERMAN, P. (2004). A stochastic mesh method for pricing high-dimensional American options. *Journal of Computational Finance* 7, 35–72.

- [23] BROWN, D.B., SMITH, J.E., AND SUN, P. (2010). Information relaxations and duality in stochastic dynamic programs. *Operations Research* 58, 785–801.
- [24] CARMONA, R., AND DAYANIK, S. (2008). Optimal multiple stopping of linear diffusions. *Mathematics of Operations Research* 33, 446–460.
- [25] CARMONA, R., AND TOUZI, N. (2008). Optimal multiple stopping and valuation of swing options. *Mathematical Finance* 18, 239–268.
- [26] CARMONA, R. (ed.) (2009). *Indifference Pricing: Theory and Applications*. Princeton University Press, Princeton.
- [27] CARR, P., GEMAN, H., AND MADAN, D. (2001). Pricing and hedging in incomplete markets. *Journal of Financial Economics* 62, 131–167.
- [28] CARRIERE, J. (1996). Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19, 19–30.
- [29] CHENG, X., AND RIEDEL, F. (2013). Optimal stopping under ambiguity in continuous time. *Mathematics and Financial Economics* 7, 29–68.
- [30] CHERIDITO, P., DELBAEN, F., AND KUPPER, M. (2006). Dynamic monetary risk measures for bounded discrete-time processes. *Electronic Journal of Probability* 11, 57–106.
- [31] CHOI, S., RUSZCZYŃSKI, A., AND ZHAO, Y. (2011). A multiproduct risk-averse newsvendor with law-invariant coherent measures of risk. *Operations Research* 59, 346–364.
- [32] CLÉMENT, 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.
- [33] COCHRANE, J. AND SAÁ-REQUEJO, J. (2000). Beyond arbitrage: Good-deal asset price bounds in incomplete markets. *Journal of Political Economy* 108, 79–119.
- [34] 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.
- [35] DELBAEN, F. (2006). The structure of m -stable sets and in particular of the set of risk neutral measures. In: Émery, M., and Yor, M. (eds.). *In Memoriam Paul-André Meyer*. Springer, Berlin, 215–258.
- [36] DELBAEN, F., PENG, S., AND ROSAZZA GIANIN, E. (2010). Representation of the penalty term of dynamic concave utilities. *Finance and Stochastics* 14, 449–472.
- [37] EKREN, I., TOUZI, N., AND ZHANG, J. (2014). Optimal stopping under nonlinear expectation. *Stochastic Processes and their Applications* 124, 3277–3311.
- [38] ELLSBERG, D. (1961). Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics* 75, 643–669.
- [39] FÖLLMER, H., AND SCHIED, A. (2002). Convex measures of risk and trading constraints. *Finance and Stochastics* 6, 429–447.
- [40] FÖLLMER, H., AND SCHIED, A. (2011). *Stochastic Finance*. Walter de Gruyter & Co., Berlin.
- [41] FRITTELLI, M., AND ROSAZZA GIANIN, E. (2002). Putting order in risk measures. *Journal of Banking & Finance* 26, 1473–1486.
- [42] GILBOA, I., AND SCHMEIDLER, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 141–153.
- [43] GLASSERMAN, P. (2004). *Monte Carlo Methods in Financial Engineering*. Springer-Verlag, New York.
- [44] GRAND-CLÉMENT, J., AND KROER, C. (2021). First-order methods for Wasserstein distributionally robust MDPs. *Proceedings of the 38th International Conference on Machine Learning*, PMLR 139, 2010–2019.
- [45] HAMBLY, B., HOWISON, S., AND KLUGE, T. (2009). Modelling spikes and pricing swing options in electricity markets. *Quantitative Finance* 9, 937–949.
- [46] HANSEN, L., AND JAGANNATHAN, R. (1991). Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–262.
- [47] HANSEN, L., AND SARGENT, T. (2007). *Robustness*. Princeton University Press, Princeton.
- [48] HAUGH, M., AND KOGAN, L. (2004). Pricing American options: A duality approach. *Operations Research* 52, 258–270.
- [49] HUBER, P. (1981). *Robust Statistics*. Wiley, New York.
- [50] JAMSHIDIAN, F. (2007). The duality of optimal exercise and domineering claims: A Doob-Meyer decomposition approach to the Snell envelope. *Stochastics* 79, 27–60.

- [51] JIANG, L. (2008). Convexity, translation invariance and subadditivity for g -expectations and related risk measures. *The Annals of Applied Probability* 18, 245–258.
- [52] KOBYLANSKI, M., QUENEZ, M.-C., AND ROUY-MIRONESCU, E. (2011). Optimal multiple stopping time problem. *The Annals of Applied Probability* 21, 1365–1399.
- [53] KRÄTSCHEMER, V., AND SCHOENMAKERS, J. (2010). Representations for optimal stopping under dynamic monetary utility functionals. *SIAM Journal on Financial Mathematics* 1, 811–832.
- [54] KRÄTSCHEMER, 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, 1177–1209.
- [55] LAEVEN, R.J.A., AND STADJE, M. (2013). Entropy coherent and entropy convex measures of risk. *Mathematics of Operations Research* 38, 265–293.
- [56] LAEVEN, R.J.A., AND STADJE, M. (2014). Robust portfolio choice and indifference valuation. *Mathematics of Operations Research* 39, 1109–1141.
- [57] LESNEVSKI, V., NELSON, B. AND STAUM, J. (2007). Simulation of coherent risk measures based on generalized scenarios. *Management Science* 53, 1756–1769.
- [58] LIU, Z., BAI, Q., BLANCHET, J., DONG, P., XU, W., ZHOU, Z., AND ZHOU, Z. (2022). Distributionally robust Q -learning. *Proceedings of the 39th International Conference on Machine Learning*, PMLR 162, 13623–13643.
- [59] LONGSTAFF, F., AND SCHWARTZ, E. (2001). Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies* 14, 113–147.
- [60] LUCIA, J.J., AND SCHWARTZ, E.S. (2002). Electricity prices and power derivatives: Evidence from the Nordic power exchange. *Review of Derivatives Research* 5, 5–50.
- [61] MATOUSSI, A., PIOZIN, L., AND POSSAMAÏ, D. (2014). Second-order BSDEs with general reflection and game options under uncertainty. *Stochastic Processes and their Applications* 124, 2281–2321.
- [62] MATOUSSI, A., POSSAMAÏ, D., AND ZHOU, C. (2013). Second order reflected backward stochastic differential equations. *The Annals of Applied Probability* 23, 2420–2457.
- [63] NUTZ, M., AND ZHANG, J. (2015). Optimal stopping under adverse nonlinear expectation and related games. *The Annals of Applied Probability* 25, 2503–2534.
- [64] ØKSENDAL, B., SULEM, A., AND ZHANG, T. (2014). Singular control and optimal stopping of SPDEs, and backward SPDEs with reflection. *Mathematics of Operations Research* 39, 464–486.
- [65] PENG, S. (1997). Backward SDE and related g -expectations. In: El Karoui, N., and Mazliak, L. (eds.). *Backward Stochastic Differential Equations*. Pitman Res. Notes Math. Ser. Vol. 364, Longman, Harlow, 141–159.
- [66] PENG, S. (2004). *Nonlinear Expectations, Nonlinear Evaluations and Risk Measures*. Vol. 1856 of *Lecture Notes in Mathematics*. Springer, Berlin.
- [67] PHILPOTT, A., DE MATOS, V., AND FINARDI, E. (2013). On solving multistage stochastic programs with coherent risk measures. *Operations Research* 61, 957–970.
- [68] RIEDEL, F. (2009). Optimal stopping with multiple priors. *Econometrica* 77, 857–908.
- [69] RIEDEL, F. (2004). Dynamic coherent risk measures. *Stochastic Processes and their Applications* 112, 185–200.
- [70] ROGERS, L.C.G. (2002). Monte Carlo valuation of American options. *Mathematical Finance* 12, 271–286.
- [71] ROGERS, L.C.G. (2007). Pathwise stochastic optimal control. *SIAM Journal on Control and Optimization* 46, 1116–1132.
- [72] ROGERS, L.C.G. (2010). Dual valuation and hedging of Bermudan options. *SIAM Journal on Financial Mathematics* 1, 604–608.
- [73] RUSZCZYŃSKI, A., AND SHAPIRO, A. (2006). Optimization of convex risk functions. *Mathematics of Operations Research* 31, 433–452.
- [74] RUSZCZYŃSKI, A., AND SHAPIRO, A. (2006). Conditional risk mappings. *Mathematics of Operations Research* 31, 544–561.
- [75] RUSZCZYŃSKI, A. (2010). Risk-averse dynamic programming for Markov decision processes. *Mathematical Programming Series B* 125, 235–261.
- [76] SAVAGE, L.J. (1954). *The Foundations of Statistics*. Wiley, New York (2nd ed. 1972, Dover, New York).
- [77] SCHOENMAKERS, J. (2012). A pure martingale dual for multiple stopping. *Finance and Stochastics* 16, 319–334.

- [78] SCHOENMAKERS, J., ZHANG, J., AND HUANG, J. (2013). Optimal dual martingales, their analysis and application to new algorithms for Bermudan products. *SIAM Journal on Financial Mathematics* 4, 86–116.
- [79] SHAPIRO, A., DENTCHEVA, D., AND RUSZCZYŃSKI, A. (2009). *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia.
- [80] SHAPIRO, A. (2016). Rectangular sets of probability measures. *Operations Research* 64, 528–541.
- [81] SNELL, L.J. (1952). Applications of martingale system theorems. *Trans. Amer. Math. Soc.* 73, 293–312.
- [82] SONER, M., TOUZI, N., AND ZHANG, J. (2011). Quasi-sure stochastic analysis through aggregation. *Electronic Journal of Probability* 16, 1844–1879.
- [83] TANG, S., AND WEI, W. (2012). Representation of dynamic time-consistent convex risk measures with jumps. *Risk and Decision Analysis* 3, 167–190.
- [84] TEKAYA, W., SHAPIRO, A.S., SOARES, M.P., AND DA COSTA, J.P. (2013). Worst-case-expectation approach to optimization under uncertainty. *Operations Research* 61, 1435–1449.
- [85] TSITSIKLIS, J. N., AND VAN ROY, B. (2000). Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks* 12, 694–703.
- [86] WALD, A. (1950). *Statistical Decision Functions*. Wiley, New York.
- [87] WIESEMANN, W., KUHN, D., AND RUSTEM, B. (2013). Robust Markov decision processes. *Mathematics of Operations Research* 38, 153–183.
- [88] XU, H., AND MANNOR, S. (2012). Distributionally robust Markov decision processes. *Mathematics of Operations Research* 37, 288–300.