

A nonparametric algorithm for optimal stopping based on robust optimization

Bradley Sturt

Department of Information and Decision Sciences
University of Illinois at Chicago, bsturt@uic.edu

Optimal stopping is a fundamental class of stochastic dynamic optimization problems with numerous applications in finance and operations management. We introduce a new approach for solving computationally-demanding stochastic optimal stopping problems with known probability distributions. The approach uses simulation to construct a robust optimization problem that approximates the stochastic optimal stopping problem to any arbitrary accuracy; we then solve the robust optimization problem to obtain near-optimal Markovian stopping rules for the stochastic optimal stopping problem.

In this paper, we focus on designing algorithms for solving the robust optimization problems that approximate the stochastic optimal stopping problems. These robust optimization problems are challenging to solve because they require optimizing over the infinite-dimensional space of all Markovian stopping rules. We overcome this challenge by characterizing the structure of optimal Markovian stopping rules for the robust optimization problems. In particular, we show that optimal Markovian stopping rules for the robust optimization problems have a structure that is surprisingly simple and finite-dimensional. We leverage this structure to develop an exact reformulation of the robust optimization problem as a zero-one bilinear program over totally unimodular constraints. We show that the bilinear program can be solved in polynomial time in special cases, establish computational complexity results for general cases, and develop polynomial-time heuristics by relating the bilinear program to the maximal closure problem from graph theory. Numerical experiments demonstrate that our algorithms for solving the robust optimization problems are practical and can outperform state-of-the-art simulation-based algorithms in the context of widely-studied stochastic optimal stopping problems from high-dimensional option pricing.

Key words: robust optimization; optimal stopping; options pricing.

History: First version: March 4, 2021. Revisions submitted on June 17, 2022 and December 21, 2022.

Accepted for publication on March 16, 2023.

1. Introduction

Consider the following class of stochastic dynamic optimization problems: A sequence of random states are incrementally revealed to a decision maker. After observing the state in each period, the decision maker chooses whether to continue to the next period or stop and receive a reward that depends on the current state. The problem is to find a control policy, called a stopping rule, for selecting when to stop the process to maximize the expected reward.

Such optimal stopping problems are widely studied and arise in a variety of domains like finance, promotion planning (Feng and Gallego 1995), and organ transplantation (David and Yechiali 1985). In particular, optimal stopping has considerable importance to industry for the pricing of financial derivatives. With a record trading volume that exceeded seven billion contracts in 2020, equity options are among the most widely-traded type of financial derivative (Reuters 2021), and financial firms depend on solving optimal stopping problems to determine accurate prices for American-style options, the most common type of equity option.

In this paper, we study a general class of optimal stopping problems in which the sequence of random states is driven by a non-Markovian probability distribution. We recall that a sequence of random states is non-Markovian if the state in the next time period (*e.g.*, a stock’s price tomorrow) has a probability distribution which depends both on the state in the current time period (*e.g.*, the stock’s price today) as well as the states in the past periods (*e.g.*, the stock’s price yesterday). This class of optimal stopping problems has witnessed a surge of interest as financial firms increasingly use non-Markovian probability distributions to accurately model the volatility patterns of stocks (Gatheral et al. 2018, Leão et al. 2019, Becker et al. 2019, Bezerra et al. 2020, Goudenège et al. 2020, Bayer et al. 2020). Optimal stopping problems with non-Markovian probability distributions also occur when using popular dimensionality-reduction techniques for pricing high-dimensional basket options (Bayer et al. 2019, p. 372) and pricing options when the probability distribution of underlying assets is accessed via a black-box simulator constructed from historical data (Ciocan and Mišić 2020, §5.5).

Despite their importance in practice, non-Markovian optimal stopping (NMOS) problems are “not easy to solve” (Leão et al. 2019, p. 982). The difficulty of these problems arises because the optimal decision in each period may depend on the entire history of the state process. In principle, an NMOS problem can be transformed into an equivalent Markovian optimal stopping problem by converting the original sequence of random states $x_1, \dots, x_T \in \mathcal{X}$ into a Markovian stochastic process $X_1, \dots, X_T \in \mathcal{X}^T$, where the new state $X_t := (x_1, \dots, x_t, 0, \dots, 0) \in \mathcal{X}^T$ in each period t includes the entire state history of the original process. Unfortunately, the enlarged state space \mathcal{X}^T will be high-dimensional when the optimal stopping problem has many periods, and the difficulty of solving a Markovian optimal stopping problem explodes in the dimensionality of the state space.

To contend with the curse-of-dimensionality that arises in NMOS problems, a natural approximation technique is to search only for stopping rules that are Markovian. Rather than depending on the entire history of the original sequence, a Markovian stopping rule makes a decision in each period t based only on the current original state x_t and knowledge that the sequence of random states was not stopped in any of the previous time periods. In general, the best Markovian stopping rule for an NMOS problem is not guaranteed to be an optimal stopping rule for the NMOS problem. However, recent numerical evidence demonstrates that Markovian stopping rules can lead to highly accurate approximations of optimal stopping rules in NMOS problems from options pricing; see Goudenège et al. (2020, §9.2), Ciocan and Mišić (2020, §5.5), Bayer et al. (2020, §3.5).

As far as we are aware, only two papers until now have suggested methods that are theoretically capable of finding the best Markovian stopping rules to NMOS problems. Belomestny (2011a) analyzes simulation-based methods which optimize directly over spaces of Markovian stopping rules and suggests a nonparametric space of Markovian stopping rules based on ϵ -nets; however, he does not propose any concrete algorithms for optimizing over this nonparametric space of Markovian stopping rules. Ciocan and Mišić (2020) propose optimizing over Markovian stopping rules which are restricted to decision trees with fixed depth. Due to the computational intractability of optimizing over all decision trees of fixed depth, the authors develop greedy heuristics which are shown to limit the range of attainable decision trees, and overcoming this limitation of the heuristics “is not obvious, especially in light of the structure of the optimal stopping problem that is leveraged to efficiently optimize split points in our construction” (Ciocan and Mišić 2020, p. 22).

We take a different approach to the aforementioned literature, and in doing so make our contributions to optimal stopping, by drawing on the traditionally unrelated field of *robust optimization*. Over the past two decades, robust optimization has emerged as a leading tool in operations research for dynamic decision-making when uncertainty is driven by unknown or ambiguous probability distributions (Ben-Tal et al. 2009, Delage and Iancu 2015). In this paper, we show that robust optimization can be combined with simulation to develop algorithms for finding Markovian stopping

rules to NMOS problems with known probability distributions. Compared to Belomestny (2011a) and Ciocan and Mišić (2020), our approach for optimal stopping does not restrict the space of Markovian stopping rules to any parametric class, and we develop concrete algorithms that are guaranteed to yield ϵ -optimal Markovian stopping rules for general classes of NMOS problems.

In greater detail, this paper introduces a new approach for computing Markovian stopping rules for NMOS problems with known probability distributions that is based on a combination of simulation and robust optimization. At a high level, our approach is comprised of the following steps:

Step 1: We use Monte-Carlo simulation to generate sample paths of the sequence of random states.

Step 2: From those sample paths, we construct a robust optimization problem that approximates the NMOS problem.

Step 3: We solve the robust optimization problem to obtain stopping rules for the NMOS problem.

The robust optimization problem constructed in Step 2 can be interpreted as a *proxy* or *surrogate* for the NMOS problem. Indeed, if the number of simulated sample paths in Step 1 is sufficiently large, then every Markovian stopping rule that is optimal for the robust optimization problem constructed in Step 2 is guaranteed with high probability to be an ϵ -optimal Markovian stopping rule for the NMOS problem; see §2.4. Thus, the approach comprised of the above steps enables the task of computing near-optimal Markovian stopping rules for an NMOS problem to be reduced to the task of solving a robust optimization problem.

With respect to the robust optimization literature, Step 2 of our approach follows a recent paper by Bertsimas, Shtern, and Sturt (2023, henceforth referred to as BSS22). In that paper, the authors showed that a general class of stochastic dynamic optimization problems with unknown probability distributions can in principle be approximated to arbitrary accuracy by a robust optimization problem constructed from historical data. In particular, our work draws on BSS22 in two specific ways. First, our robust optimization problem constructed in Step 2 is a variant of a robust optimization formulation that is proposed in BSS22. Second, by focusing on the application of optimal stopping, we strengthen theoretical developments from BSS22 to prove under relatively mild and verifiable assumptions that the optimal objective value and optimal Markovian stopping rules of the robust optimization problem from Step 2 converge almost surely to those of the NMOS problem as the number of simulated sample paths in Step 1 grows to infinity (Theorems 1-3 in §2.4). For the interested reader, a discussion of our improvements to the theoretical convergence guarantees from BSS22 is provided at the beginning of Appendix C.

In contrast to BSS22, the key novelty of the present paper lies *not* in showing that robust optimization can be used to construct approximations of stochastic dynamic optimization problems; rather, the main contributions of the present paper are *algorithmic*. Indeed, in order for a combination of simulation and robust optimization to yield near-optimal algorithms for a class of stochastic dynamic optimization problems with known probability distributions, it is not sufficient to construct a robust optimization problem that approximates the stochastic problem to arbitrary accuracy: one must also have exact or provably near-optimal algorithms for solving the robust optimization problem. However, up to this point, there have been no exact algorithms in the literature for solving the type of robust optimization problems proposed by BSS22 in any class of stochastic dynamic optimization problems where uncertainty unfolds over two or more periods. In particular, because solving the type of robust optimization problems from BSS22 requires optimizing over infinite-dimensional spaces of control policies, it has been unknown in any application whether optimal control policies for these robust optimization problems even *exist*, let alone whether they can ever be tractably computed. These hurdles have, practically speaking, prevented the robust optimization techniques from BSS22 from being combined with simulation to develop algorithms for

computing near-optimal control policies for any class of stochastic dynamic optimization problems with known probability distributions until now.

In this paper, we resolve the aforementioned gaps in the robust optimization literature by establishing the first characterization of optimal control policies for the type of robust optimization problems proposed by BSS22 in an application where uncertainty unfolds over two or more periods (Theorem 4 in §3). Specifically, we consider the task of solving the robust optimization problems constructed in Step 2 over the infinite-dimensional space of all Markovian stopping rules. Using a novel pruning technique, we prove in Theorem 4 that optimal Markovian stopping rules for these robust optimization problems not only exist, but also have a structure that is simple and finite-dimensional. In fact, our characterization reveals that optimal Markovian stopping rules for these robust optimization problems can be compactly parameterized by integer variables $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$, where N is the number of simulated sample paths chosen in Step 1 and T is the number of periods in the NMOS problem.

Leveraging the structure of optimal Markovian stopping rules for the robust optimization problems constructed in Step 2, we develop exact and heuristic algorithms for solving the robust optimization problem in Step 3. Specifically, we make the following algorithmic contributions:

- (a) We use our characterization of optimal Markovian stopping rules for the robust optimization problems constructed in Step 2 to transform the robust optimization problem into a finite-dimensional optimization problem over integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ (Theorem 5 in §3.3).
- (b) For NMOS problems with two periods, we show that optimal Markovian stopping rules for the robust optimization problem can be computed in $\mathcal{O}(N^3)$ time (Theorem 6 in §4.1).
- (c) For NMOS problems with three or more periods, we prove that finding optimal Markovian stopping rules for the robust optimization problem is NP-hard (Theorem 7 in §4.1).
- (d) We develop a general-purpose exact algorithm for computing optimal Markovian stopping rules for the robust optimization problem that consists of solving a zero-one bilinear program over totally unimodular constraints (Theorem 8 in §4.1).
- (e) We design and analyze a polynomial-time heuristic algorithm for approximately solving the robust optimization problem by relating the bilinear program to the maximal closure problem from graph theory (Propositions 1-5 in §4.2).

In summary, our exact and heuristic algorithms allow us to solve the robust optimization problem constructed in Step 2. Since the robust optimization problem constructed via Steps 1 and 2 serves as an approximation of the NMOS problem to any arbitrary accuracy, our algorithms for solving the robust optimization problem in turn allow us to compute near-optimal Markovian stopping rules for the NMOS problem.

We conclude with numerical experiments that demonstrate the value of our robust optimization-based algorithms in several settings. First, we consider a simple one-dimensional non-Markovian optimal stopping problem with fifty periods, and we compare the robust optimization algorithm to existing methods based on approximate dynamic programming (Longstaff and Schwartz 2001) and parametric stopping rules (Ciocan and Mišić 2020). The experiments show that our method can find stopping rules that significantly outperform those found by the other techniques, while maintaining a comparable computational cost. In particular, the experiments reveal that our method can strictly outperform alternative algorithms for finding Markovian stopping rules to NMOS problems that are based on backwards recursion. Second, we consider a widely-studied and important problem of pricing high-dimensional Bermudan barrier options with over fifty periods. Across several variants of this problem, we demonstrate that our combination of robust optimization and simulation can find stopping rules that match, and in some cases significantly outperform, those from state-of-the-art algorithms by Longstaff and Schwartz (2001), Ciocan and Mišić (2020) as well as the duality-based pathwise optimization method of Desai et al. (2012).

The rest of our paper has the following organization. §1.1 provides a review of methods for solving optimal stopping problems with Markovian and non-Markovian probability distributions. §2 formalizes the problem setting and introduces our robust optimization-based method. §3 characterizes the structure of optimal policies for the robust optimization problem. §4 develops tractable algorithms and computational complexity results for the robust optimization problem. §5 illustrates the performance of our algorithms in numerical experiments. Unless stated otherwise, all technical proofs can be found in the appendices.

1.1. Other Related Literature

Many methods based on approximate dynamic programming (ADP) have been developed for optimal stopping problems with high-dimensional Markovian stochastic processes. The most popular ADP methods for these optimal stopping problems are based on Monte-Carlo simulation and regression, which originate with Carriere (1996), Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (2001). Given sample paths of the entire stochastic process, these methods use backwards recursion and regression to obtain approximations of the value function, and exercise policies are then obtained by proceeding greedily with respect to the approximate value functions. The efficacy of regression-based methods hinges on selecting a parametrization of basis functions for the value function that strikes a balance between approximation quality and sample complexity. Nonparametric choices for the basis functions, *e.g.*, Laguerre polynomials, are discussed in the aforementioned works and subsequently analyzed in works such as Clément et al. (2002), Glasserman and Yu (2004), Egloff (2005), Belomestny (2011b) and Zanger (2020). In §5, we provide numerical comparisons of our proposed algorithms to ADP techniques in the context of NMOS problems.

A variety of other nonparametric methods have been developed for solving optimal stopping problems with Markovian probability distributions, such as quantization-based approximations of value functions (Bally and Pages 2003) and scenario tree discretizations of the sequence of random states (Broadie and Glasserman 1997). Recent works have also considered using deep learning to learn the continuation function, including Becker et al. (2019) and Goudenège et al. (2020). The efficacy of deep learning algorithms for optimal stopping hinges on carefully selecting the topology of the neural network and choosing the right tuning algorithm, and performing these tasks effectively in the context of optimal stopping is an ongoing area of research (Fathan and Delage 2021). Methods to compute upper bounds on optimal stopping problems grew in interest due to the independent works of Haugh and Kogan (2004) and Rogers (2002), and duality-based algorithms to obtain upper bounds which combine simulation and suboptimal stopping rules were first proposed by Andersen and Broadie (2004). Other works that harness dual representations to solve optimal stopping and other stochastic dynamic optimization problems include Brown et al. (2010), Desai et al. (2012), Belomestny (2013), and Goldberg and Chen (2018), among many others. In §5, we provide numerical comparisons of our proposed algorithms to the duality-based pathwise optimization method of Desai et al. (2012).

In the context of non-Markovian optimal stopping, methods have been developed which address settings that are different from ours in non-trivial ways. Leão et al. (2019) and Bezerra et al. (2020) develop discretization schemes for NMOS problems over continuous time and restrict the class of probability distributions to those based on the Brownian motion. In contrast to these works, our paper develops algorithms for finding Markovian stopping rules for discrete-time optimal stopping problems, and we do not require any parametric assumptions on the probability distributions of the underlying stochastic processes. NMOS problems can also be addressed by the scenario tree method of Broadie and Glasserman (1997) and the recursive-dual algorithm of Goldberg and Chen (2018),

provided that one can perform Monte-Carlo simulation on the conditional probability distribution of the stochastic process in each time period. In contrast to these methods, the algorithms in this paper require only the ability to simulate sample paths of the entire stochastic process and are shown in numerical experiments to be practically tractable in low-dimensional NMOS problems with dozens of time periods. Within the optimal stopping literature, our method is most closely related to a class of simulation-based methods which optimize directly over spaces of deterministic stopping rules, as explored by Garcia (2003), Andersen (1999), Belomestny (2011a), Gemmrich (2012), Ciocan and Mišić (2020), and Glasserman (2013, §8.2). We discuss connections between our approach and this stream of literature in §2.3, and a discussion of the challenges of using dynamic programming to find the best Markovian stopping rules for NMOS problems can be found in Appendix A.

Finally, we note that prior research in operations research and economics have studied robust optimal stopping problems in which the goal is to find stopping rules that perform well under worst-case probability distributions (Bayraktar and Yao 2014, Riedel 2009) or under worst-case state trajectories (Iancu et al. 2021). That stream of research differs significantly from ours, as that stream of research does not consider robust optimization problems that are approximations of stochastic optimal stopping problems with known probability distributions.

2. Robust Optimization for Stochastic Optimal Stopping

2.1. Problem Setting

We consider stochastic optimal stopping problems defined by the following components:

States: Let $x \equiv (x_1, \dots, x_T)$ denote a sequence of random states, where the state $x_t \in \mathcal{X} \equiv \mathbb{R}^d$ in each period t is a random vector of dimension d . For example, the state in each period may represent the prices of multiple assets at that point in time. The joint probability distribution of this stochastic process is assumed to be known and accessible through a simulator which generates independent sample paths of the entire stochastic process.

Policies: Let $\mu \equiv (\mu_1, \dots, \mu_T)$ represent a collection of exercise policies, where the exercise policy in each period t is a measurable function of the form $\mu_t : \mathcal{X} \rightarrow \{\text{STOP}, \text{CONTINUE}\}$. Speaking intuitively, each exercise policy is a partitioning of the state space into regions for stopping and continuing. From the exercise policies, the corresponding Markovian stopping rule $\tau_\mu : \mathcal{X}^T \rightarrow \{1, \dots, T\} \cup \{\infty\}$ is a function that maps a realization of the stochastic process to a stopping period:

$$\tau_\mu(x) \triangleq \min\{t \in \{1, \dots, T\} : \mu_t(x_t) = \text{STOP}\}.$$

Throughout this paper, a minimization problem with no feasible solutions is defined equal to ∞ .¹

Rewards: Let $g : \{1, \dots, T\} \cup \{\infty\} \times \mathcal{X}^T \rightarrow \mathbb{R}_+$ be a known and deterministic function that maps a stopping period and a realization of the entire stochastic process to a reward. The assumption that the reward function is nonnegative is common in many applications of optimal stopping,

¹ We remark that the Markovian stopping rule τ_μ is a non-anticipative control policy, meaning that the event $\{\tau_\mu(x) = t\}$ does not depend on the future states x_{t+1}, \dots, x_T for each period $t \in \{1, \dots, T\}$. To see why this is the case, consider any two realizations of the sequence of random states $x \equiv (x_1, \dots, x_T)$ and $x' \equiv (x'_1, \dots, x'_T)$, and suppose for a given period $t \in \{1, \dots, T\}$ that the two realizations satisfy $x_s = x'_s$ for all $s \in \{1, \dots, t\}$. Then we readily observe from our definition of Markovian stopping rules and from algebra that the Markovian stopping rule τ_μ satisfies $\tau_\mu(x) = t$ if and only if $\tau_\mu(x') = t$ is satisfied.

and we assume throughout the paper that a stochastic process that is never stopped yields a reward of zero: $g(\infty, x) \equiv 0$. It follows from the definition of the reward function that the reward from stopping on any period t may in general depend on the states of the stochastic process in previous or future time periods.

Problem: With the above notation and inputs, the goal of this paper is to solve stochastic optimal stopping problems of the form

$$\sup_{\mu} \mathbb{E}[g(\tau_{\mu}(x), x)], \quad (\text{OPT})$$

where the optimization is taken over the space of all Markovian stopping rules. In the following sections, we introduce and analyze a new simulation-based method for solving this class of stochastic dynamic optimization problems.

2.2. The Robust Optimization Approach

Our proposed approach for solving stochastic optimal stopping problems of the form (OPT) consists of the following steps. We first simulate sample paths of the stochastic process $x \equiv (x_1, \dots, x_T)$. Let N denote the number of sample paths, and let the values of the sample paths be denoted by

$$x^i \equiv (x_1^i, \dots, x_T^i) \text{ for each } i = 1, \dots, N.$$

We assume that the sample paths are independent and identically distributed realizations of the entire (possibly non-Markovian) stochastic process. We next choose the following robustness parameter:

$$\epsilon \geq 0.$$

The purpose of the robustness parameter will become clear momentarily, and a discussion on how to choose the number of sample paths and the robustness parameter is deferred until §2.5. With these parameters, let the uncertainty set around sample path i on period t be defined as

$$\mathcal{U}_t^i \triangleq \{y_t \in \mathcal{X} : \|y_t - x_t^i\|_{\infty} \leq \epsilon\}.$$

For notational convenience, denote the uncertainty set around sample path i across all periods by

$$\mathcal{U}^i \triangleq \mathcal{U}_1^i \times \dots \times \mathcal{U}_T^i.$$

Hence, we observe that the role of the robustness parameter is to control the size of these sets. Given the sample paths and choice of the robustness parameter, our approach obtains an approximation of (OPT) by solving the following robust optimization problem:

$$\sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i). \quad (\text{RO})$$

By solving the above robust optimization problem, we obtain exercise policies $\hat{\mu} \equiv (\hat{\mu}_1, \dots, \hat{\mu}_T)$. These exercise policies constitute our approximate solution to the stochastic optimal stopping problem (OPT).

REMARK 1. In the above formulation of the robust optimization problem (RO), the sample paths in the objective function are allowed to be perturbed by an adversary in the evaluation of the Markovian stopping rule, $\tau_{\mu}(y)$, but not in the evaluation of the reward, $g(\cdot, x^i)$. This is a deviation

from the robust optimization formulation from BSS22, presented below as (RO'), in which the worst-case reward over each uncertainty set has the form $\inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y)$:

$$\sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y). \quad (\text{RO}')$$

The difference between the objective functions of (RO) and (RO') turns out to be inconsequential from the perspective of establishing convergence guarantees (see §2.4) or characterizing the structure of optimal Markovian stopping rules (see §3). However, (RO) is significantly simpler from the perspective of algorithm design. For the interested reader, an extended discussion on the similarities and differences between (RO) and the alternative robust optimization formulation (RO') can be found in Appendix B.

2.3. Background and Motivation

In contrast to traditional robust optimization or distributionally robust optimization, our motivation behind adding adversarial noise to the sample paths in (RO) is not to find stopping rules which have worst-case performance guarantees, are attractive in risk-averse settings, or perform well the presence of an ambiguous probability distribution. Rather, this paper proposes using robust optimization purely as an algorithmic tool for solving stochastic optimal stopping problems of the form (OPT) when the joint probability distributions are known. The present section elaborates on this motivation and positions our use of robust optimization within the optimal stopping literature.

For the sake of developing intuition, let us suppose for the moment that the robustness parameter of the uncertainty sets in (RO) was set equal to zero. In this case, for any fixed exercise policies $\mu = (\mu_1, \dots, \mu_T)$, the expected reward of those exercise policies,

$$J^*(\mu) \triangleq \mathbb{E}[g(\tau_\mu(x), x)],$$

would be approximated in (RO) by the sample average approximation:

$$\widehat{J}_{N,0}(\mu) \triangleq \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{X}^T: \|y - x^i\|_\infty \leq 0} g(\tau_\mu(y), x^i) = \frac{1}{N} \sum_{i=1}^N g(\tau_\mu(x^i), x^i).$$

For these fixed exercise policies, we observe that the sample average approximation is a *consistent* estimator of the expected reward. In other words, for the fixed exercise policies μ , it follows from the strong law of large numbers under relatively mild assumptions² that $\widehat{J}_{N,0}(\mu)$ will converge almost surely to $J^*(\mu)$ as the number of simulated sample paths is taken to infinity.

However, it is well known in the optimal stopping literature that these desirable asymptotic properties of $\widehat{J}_{N,0}(\mu)$ are generally not retained when considering the problem of optimizing over the space of all exercise policies. For instance, Ciocan and Mišić (2020, EC.1.2) provide simple examples in which the following statements hold almost surely:

$$\lim_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,0}(\mu) \gg \sup_{\mu} J^*(\mu); \quad \lim_{N \rightarrow \infty} J^*(\widehat{\mu}_{N,0}) \ll \sup_{\mu} J^*(\mu). \quad (1)$$

The asymptotic suboptimality of the optimal objective value and optimal policies for the problem $\sup_{\mu} \widehat{J}_{N,0}(\mu)$ can be intuitively understood as a type of overfitting. To see why line (1) occurs, we

² For example, Assumptions 2 and 4 in §2.4.

recall for any fixed choice of exercise policies μ that the sample average approximation $\widehat{J}_{N,0}(\mu)$ is an unbiased estimate of the expected reward $J^*(\mu)$. However, when simultaneously considering the space of all exercise policies, there exists for each $N \in \mathbb{N}$ with high probability a collection of exercise policies that satisfies $\widehat{J}_{N,0}(\mu) \gg J^*(\mu)$. The problem $\sup_{\mu} \widehat{J}_{N,0}(\mu)$ will thus be biased towards choosing those exercise policies, which in general will be suboptimal for the problem $\sup_{\mu} J^*(\mu)$. Because the set of all μ is an infinite-dimensional space, the gap between the objective values $\widehat{J}_{N,0}(\mu)$ and $J^*(\mu)$ does not converge to zero uniformly over the set of all μ as the number of sample paths tends to infinity.

To circumvent this overfitting in the context of optimal stopping in line (1), a vast literature has focused on restricting the functional form of exercise policies to a finite-dimensional space, such as Garcia (2003), Andersen (1999), Belomestny (2011a), Gemmrich (2012), Ciocan and Mišić (2020). In this approach, the choice of the parameterization for the space of exercise policies must be made very carefully. On one hand, the effective dimension of the restricted space of exercise policies must be small relative to the number of simulated sample paths to ensure that the sample average approximation problem finds the parametric exercise policies that are ‘best-in-class’ with respect to the stochastic optimal stopping problem (Belomestny 2011a, §3). On the other hand, the parameterization must be chosen appropriately in order for the sample average approximation problem to obtain a good approximation of (OPT). Choosing such an appropriate parameterization “may be counterfactual in some cases”, as explained by Garcia (2003, p. 1859), “since we may not have a good understanding of what the early exercise rule should depend on.”

Our approach, in view of the above discussion, provides an alternative means to circumvent overfitting. The proposed robust optimization problem allows the space of exercise policies to remain general, and thus relieves the decision maker from the need to select and impose a parametric structure on the exercise policies. Moreover, we show in the following section that our use of robust optimization provably overcomes the asymptotic overfitting described in line (1).

2.4. Optimality Guarantees

In this subsection, we establish theoretical justification for our combination of robust optimization and simulation that is presented in §2.2. Specifically, we strengthen convergence guarantees from BSS22 to the specific problem of optimal stopping to prove that the optimal objective value and optimal exercise policies of (RO) converge almost surely to those of the stochastic optimal stopping problem (OPT) under mild and verifiable conditions. Establishing these convergence guarantees in the context of optimal stopping is necessary to ensure that (RO) will provide a high-quality approximation of the stochastic optimal stopping problem (OPT) when the robustness parameter is sufficiently small and the number of simulated sample paths is sufficiently large.

To establish our theoretical results, we make four relatively mild assumptions on the stochastic optimal stopping problem (OPT). Our first assumption, denoted below by Assumption 1, concerns the structure of the reward functions in the optimal stopping problem, and can be roughly interpreted as a requirement that the reward function changes continuously as a function of the states:

ASSUMPTION 1. $\lim_{\epsilon \rightarrow 0} \Delta_{\epsilon}(x) = 0$ almost surely, where

$$\Delta_{\epsilon}(x) \triangleq \min_{t \in \{1, \dots, T\}} \left\{ \inf_{y \in \mathcal{X}^T: \|y-x\|_{\infty} \leq \epsilon} g(t, y) - g(t, x) \right\}.$$

From a practical standpoint, it is easy to see that Assumption 1 holds whenever the functions $g(1, \cdot), \dots, g(T, \cdot)$ are continuous, and it can also hold in important stochastic optimal stopping problems with discontinuous reward functions.³

Our second and third assumptions concern the structure of the probability distribution in (OPT). Specifically, Assumption 2 enforces that the stochastic process has a light tail, and Assumption 3 says the stochastic process is drawn from a continuous probability distribution.

ASSUMPTION 2. *The stochastic process satisfies $\mathbb{E}[\exp(\|x\|_\infty^a)] < \infty$ for some $a > 1$.*

ASSUMPTION 3. *The stochastic process $x \equiv (x_1, \dots, x_T)$ has a probability density function.*

Let us reflect on the practical restrictiveness of these two assumptions. The second assumption, Assumption 2, is a standard light-tail assumption on the stochastic process which is satisfied, for example, if the stochastic process is bounded or has a multivariate normal distribution. This assumption greatly simplifies our analysis, as it allows us to invoke a convergence result by BSS22 in our proofs (see Appendix C). We impose the third assumption, Assumption 3, to ensure that there exist arbitrarily near-optimal Markovian stopping rules for (OPT) that satisfy a certain technical continuity structure that we can exploit in our proof. As far as we can tell, these assumptions on the probability distribution are relatively mild and routinely satisfied in applications of optimal stopping in the context of the options pricing literature. Nonetheless, we do not preclude the possibility that these assumptions on the probability distribution can be weakened while still establishing convergence guarantees.

Our fourth and final assumption imposes boundedness on the reward function. This assumption, presented below as Assumption 4, leads to a considerably simpler proof and statement of the results, but can generally be relaxed to reward functions bounded above by an integrable, Lipschitz-continuous function.

ASSUMPTION 4. *The reward function satisfies $0 \leq g(t, y) \leq U$ for all $t \in \{1, \dots, T\}$ and $y \in \mathcal{X}^T$.*

We emphasize that each of the aforementioned four assumptions on the stochastic optimal stopping problem (OPT) can be verified *a priori*. In particular, they do not require any knowledge of the structure of optimal exercise policies for the stochastic optimal stopping problem (OPT). As a result, each of these assumptions can be verified using the information typically available in practice. We note that these assumptions are considerably weaker than those in BSS22, which require knowledge of the structure of optimal control policies to establish convergence results.

Under the above conditions, the following theorems provide justification for using the robust optimization problem as a proxy for the stochastic optimal stopping problem. In a nutshell, the following Theorems 1-3 show that (RO) will, for all sufficiently small choices of the robustness parameter and all sufficiently large choices of the number of simulated sample paths, yield a near-optimal approximation of (OPT). Stated another way, the following theorems show that our use of robust optimization provably overcomes the asymptotic overfitting described in line (1) of §2.3. While the following theorems do not specify how to choose the robustness parameter and number of simulated sample paths for any particular optimal stopping problem, we provide guidance (§2.5) and numerical evidence (§5) which suggest that these parameters can be found

³ To illustrate, consider Robbin's problem (Bruss 2005), in which the reward functions $g(t, x) = \text{rank}(x_t; x_1, \dots, x_T)$ are discontinuous and the probability distribution is $x_1, \dots, x_T \stackrel{\text{iid}}{\sim} \text{Uniform}[0, 1]$. To show that Assumption 1 is satisfied, we observe that the random variable $\bar{\epsilon} \triangleq \min_{s < t} |x_s - x_t|$ is strictly positive with probability one, which implies that $\Delta_\epsilon(x) = 0$ for all $\epsilon < \bar{\epsilon}$.

effectively in practice. A discussion of the technical innovations as well the proofs of the following theorems in this subsection can be found in Appendix C.

Our first theorem shows that the optimal objective value of the robust optimization problem (RO) will converge almost surely to that of the stochastic problem (OPT) as the robustness parameter tends to zero and the number of sample paths tends to infinity. In the following result, we use the notation $\widehat{J}_{N,\epsilon}(\mu)$ to denote the objective value of the robust optimization problem (RO) corresponding to exercise policies μ .

THEOREM 1 (Consistency of optimal objective value). *Under Assumptions 1, 2, 3, and 4,*

$$\lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) = \sup_{\mu} J^*(\mu) \quad \text{almost surely.}$$

Our second theorem shows that the expected reward of the optimal exercise policies for the robust optimization problem (RO) will converge almost surely to the optimal objective value of the stochastic problem (OPT). We let $\widehat{\mu}_{N,\epsilon}$ denote optimal exercise policies for (RO), and we remark that the existence of optimal exercise policies for the robust optimization problem will be established in §3.

THEOREM 2 (Consistency of optimal policies). *Under Assumptions 1, 2, 3, and 4,*

$$\lim_{\epsilon \rightarrow 0} \liminf_{N \rightarrow \infty} J^*(\widehat{\mu}_{N,\epsilon}) = \lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} J^*(\widehat{\mu}_{N,\epsilon}) = \sup_{\mu} J^*(\mu) \quad \text{almost surely.}$$

Because we will develop algorithms that solve the robust optimization problem approximately as well as exactly, it is imperative for us to have theoretical guarantees that hold for any Markovian stopping rule that can be found by the robust optimization problem. To this end, our third and final theorem of this section shows that the (in-sample) robust objective value will asymptotically provide a low-bias estimate of the expected reward, and this bound holds uniformly over all exercise policies. The result yields theoretical assurance, provided that the robustness parameter is sufficiently small and the number of sample paths is sufficiently large, that searching for exercise policies with high robust objective values $\widehat{J}_{N,\epsilon}(\mu)$ will typically result in exercise policies with high expected rewards $J^*(\mu)$.

THEOREM 3 (Asymptotic low-bias). *Under Assumptions 1, 2, and 4,*

$$\lim_{\epsilon \rightarrow 0} \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ J^*(\mu) - \widehat{J}_{N,\epsilon}(\mu) \right\} \geq 0 \quad \text{almost surely.}$$

2.5. Implementation Details

In anticipation of algorithmic techniques for solving the robust optimization problem (RO) in the remainder of the paper, it remains to be specified how the parameters of the robust optimization problem (the number of simulated sample paths $N \in \mathbb{N}$ and the robustness parameter $\epsilon \geq 0$) should be selected in practice. For the sake of concreteness, we conclude §2 by briefly providing guidance for choosing these parameters and applying the robust optimization approach in practice. The procedures described below are formalized in Algorithm 1 and implemented in our numerical experiments in §5.

As described previously, this paper addresses stochastic optimal stopping problems in which the probability distributions are known. Consequently, the decision-maker is granted flexibility in choosing the number of sample paths N to simulate. On one hand, we have established in the

The Robust Optimization Approach for Stochastic Optimal Stopping

Inputs:

- *Sizes of Training Sets:* A collection of integers $\mathcal{N} \triangleq \{N_1, \dots, N_K\}$ sorted in ascending order.
- *Size of Validation Set:* An integer $\bar{N} \in \mathbb{N}$.
- *Size of Testing Set:* An integer $\tilde{N} \in \mathbb{N}$.
- *Robustness Parameters:* A collection of nonnegative real numbers $\mathcal{E} \triangleq \{\epsilon_1, \dots, \epsilon_L\}$.

Outputs:

- Exercise policies $\hat{\mu} \equiv (\hat{\mu}_1, \dots, \hat{\mu}_T)$ for the stochastic optimal stopping problem (OPT).
- An (unbiased) estimate of the expected reward $J^*(\hat{\mu})$ for the exercise policies.

Procedure:

1. Simulate the validation set $\bar{\mathcal{S}} \triangleq \{\bar{x}^1, \dots, \bar{x}^{\bar{N}}\}$ and testing set $\tilde{\mathcal{S}} \triangleq \{\tilde{x}^1, \dots, \tilde{x}^{\tilde{N}}\}$.
2. For each $N \in \mathcal{N}$:
 - (a) Simulate a training set $\mathcal{S} \triangleq \{x^1, \dots, x^N\}$.
 - (b) For each $\epsilon \in \mathcal{E}$:
 - i. Obtain exercise policies $\hat{\mu}_{N,\epsilon}$ by solving the robust optimization problem (RO) constructed from the training set \mathcal{S} and robustness parameter ϵ .
 - ii. Estimate the expected reward for these exercise policies using the validation set:

$$\bar{J}(\hat{\mu}_{N,\epsilon}) \triangleq \frac{1}{\bar{N}} \sum_{i=1}^{\bar{N}} g(\tau_{\hat{\mu}_{N,\epsilon}}(\bar{x}^i), \bar{x}^i).$$

- (c) If the total computation time of steps (2a)-(2b) reached the allowed budget (or if $N = N_K$), go to step (3) with the current value of N .
3. Output the exercise policies which maximized the estimated expected reward:

$$\hat{\mu} \leftarrow \arg \max_{\epsilon \in \mathcal{E}} \bar{J}(\hat{\mu}_{N,\epsilon}),$$

and output an unbiased estimate of their expected reward using the testing set:

$$J^*(\hat{\mu}) \approx \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} g(\tau_{\hat{\mu}}(\tilde{x}^i), \tilde{x}^i).$$

Algorithm 1: A procedure for selecting the robustness parameter and number of sample paths.

previous section that larger choices of the number of simulated sample paths will generally lead to tighter approximations of the stochastic optimal stopping problem. On the other hand, larger choices of N require a greater computation cost in performing the Monte-Carlo simulation and creates a robust optimization problem of a larger size. To balance these tradeoffs in particular applications, we recommend using a straightforward procedure of starting out with a small choice of N and iteratively increasing the number of simulated sample paths until the total computational cost meets the allocated computational budget.

Given a fixed number of sample paths, the choice of the robustness parameter $\epsilon \geq 0$ can have a significant impact on the policies produced by the robust optimization problem. To this end, we recommend solving (RO) over a grid of possible choices for the robustness parameter. Because the probability distribution is known, we can generate a second set of ‘validation’ sample paths to select the best choice of the robustness parameter. Specifically, for each choice of the robustness parameter, one solves the robust optimization problem to obtain exercise policies. The expected reward of the exercise policies is then estimated using the validation set of sample paths. Finally, we select the value of the robustness parameter (and the corresponding exercise policies) which maximizes the average reward with respect to the validation set.

In summary, we have described straightforward and easy-to-implement heuristics for choosing the parameters of the robust optimization problem. Applying the heuristics and solving the robust optimization problem yields exercise policies for the stochastic optimal stopping problem, and an unbiased estimate of the expected reward of these exercise policies can similarly be obtained by simulating a set of ‘testing’ sample paths (see Algorithm 1). Because the exercise policies obtained from the robust optimization problem are feasible for the stochastic optimal stopping problem, the expected reward of these exercise policies is thus a lower bound on the optimal objective value of the stochastic optimal stopping problem. Finally, we remark that under a stronger assumption in which one has the ability to perform conditional Monte Carlo simulation, the exercise policies obtained from solving the robust optimization problem can be combined with the method of Andersen and Broadie (2004) to obtain an upper bound on the optimal objective value of the stochastic optimal stopping problem.

3. Characterization of Optimal Markovian Stopping Rules

In §2, we showed that our combination of robust optimization and simulation (§2.2) can yield an arbitrarily close approximation of the stochastic optimal stopping problem (§2.4-§2.5). In this section, we develop the key technical result of this paper, Theorem 4, which will enable us to design exact and heuristic algorithms for solving the robust optimization problem. Specifically, Theorem 4 establishes the existence and characterizes the structure of optimal Markovian stopping rules for the robust optimization problem (RO). In §3.1 and §3.2, we present the statement of Theorem 4 and provide a sketch of its proof. In §3.3, we use the characterization of optimal Markovian stopping rules to transform (RO) from an optimization problem over an infinite-dimensional space of exercise policies into a finite-dimensional optimization problem over integer decision variables $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ (Theorem 5).

3.1. Statement of Theorem 4

We begin by introducing the notation that will be used in our characterization of the structure of optimal Markovian stopping rules for the robust optimization problem (RO). Consider any instance of the robust optimization problem (RO). For any choice of integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$, we

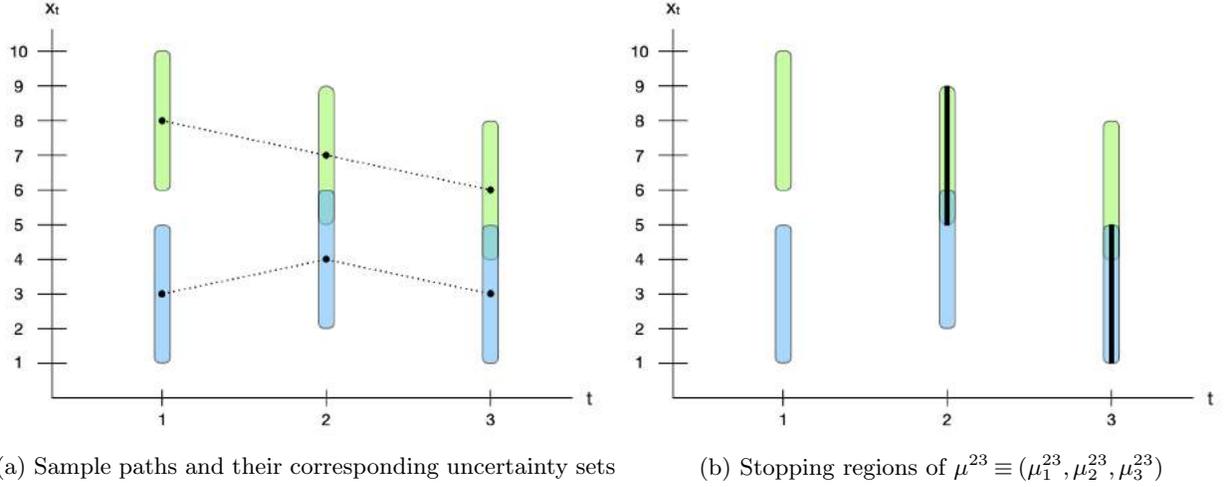


Figure 1 Visualization of a robust optimization problem with $N = 2$ sample paths and $T = 3$ periods. The values of the sample paths in the robust optimization problem are given by $x^1 = (x_1^1, x_2^1, x_3^1) = (8, 7, 6)$ and $x^2 = (x_1^2, x_2^2, x_3^2) = (3, 4, 3)$, the state space is the real numbers ($\mathcal{X} = \mathbb{R}^1$) and the robustness parameter is $\epsilon = 2.0$. The green intervals correspond to the uncertainty sets $\mathcal{U}_1^1 = [6, 10]$, $\mathcal{U}_2^1 = [5, 9]$, and $\mathcal{U}_3^1 = [4, 8]$ around the first sample path, and the blue intervals correspond to the uncertainty sets $\mathcal{U}_1^2 = [1, 5]$, $\mathcal{U}_2^2 = [2, 6]$, and $\mathcal{U}_3^2 = [1, 5]$ around the second sample path. In Figure 1a, we visualize the two sample paths using dashed lines. In Figure 1b, we visualize the stopping regions induced by the exercise policy $\mu^{\sigma^1 \sigma^2} \equiv (\mu_1^{\sigma^1 \sigma^2}, \mu_2^{\sigma^1 \sigma^2}, \mu_3^{\sigma^1 \sigma^2})$ in the case where the integers $\sigma^1, \sigma^2 \in \{1, 2, 3\}$ satisfy $\sigma^1 = 2$ and $\sigma^2 = 3$. Specifically, the black vertical lines in Figure 1b show the stopping regions $\{y_1 \in \mathbb{R} : \mu_1^{23}(y_1) = \text{STOP}\}$, $\{y_2 \in \mathbb{R} : \mu_2^{23}(y_2) = \text{STOP}\}$, and $\{y_3 \in \mathbb{R} : \mu_3^{23}(y_3) = \text{STOP}\}$ induced by $\mu^{23} \equiv (\mu_1^{23}, \mu_2^{23}, \mu_3^{23})$.

define $\mu^{\sigma^1 \dots \sigma^N} \equiv (\mu_1^{\sigma^1 \dots \sigma^N}, \dots, \mu_T^{\sigma^1 \dots \sigma^N})$ as the exercise policy that satisfies the following equality for each period $t \in \{1, \dots, T\}$ and each state $y_t \in \mathcal{X}$:

$$\mu_t^{\sigma^1 \dots \sigma^N}(y_t) \triangleq \begin{cases} \text{STOP}, & \text{if } y_t \in \bigcup_{i: \sigma^i = t} \mathcal{U}_t^i, \\ \text{CONTINUE}, & \text{if } y_t \notin \bigcup_{i: \sigma^i = t} \mathcal{U}_t^i. \end{cases} \quad (2)$$

To develop intuition of the above exercise policy, we remark for each sample path $i \in \{1, \dots, N\}$ that the integer $\sigma^i \in \{1, \dots, T\}$ can be interpreted as a selection of one of the uncertainty sets $\mathcal{U}_1^i, \dots, \mathcal{U}_T^i$. Specifically, if $\sigma^i = t$, then we observe from the above definition that the exercise policy will satisfy $\mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}$ for all $y_t \in \mathcal{U}_t^i$. More generally, we observe that the exercise policy will satisfy $\mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}$ for period $t \in \{1, \dots, T\}$ and state $y_t \in \mathcal{X}$ if and only if there exists an uncertainty set \mathcal{U}_t^i such that the state is contained in the uncertainty set, $y_t \in \mathcal{U}_t^i$, and the integer corresponding to the i th sample path is equal to the current period, $\sigma^i = t$. A visualization of the exercise policy $\mu^{\sigma^1 \dots \sigma^N}$ is found in Figure 1.

Let the set of all exercise polices generated by integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ be denoted by

$$\mathcal{M} \triangleq \left\{ \mu \equiv (\mu_1, \dots, \mu_T) : \text{there exist } \sigma^1, \dots, \sigma^N \in \{1, \dots, T\} \text{ such that } \mu = \mu^{\sigma^1 \dots \sigma^N} \right\}.$$

It is clear from the above definition that each of the exercise policies $\mu \in \mathcal{M}$ is parameterized by integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$; thus, we readily observe that the cardinality of \mathcal{M} is always finite

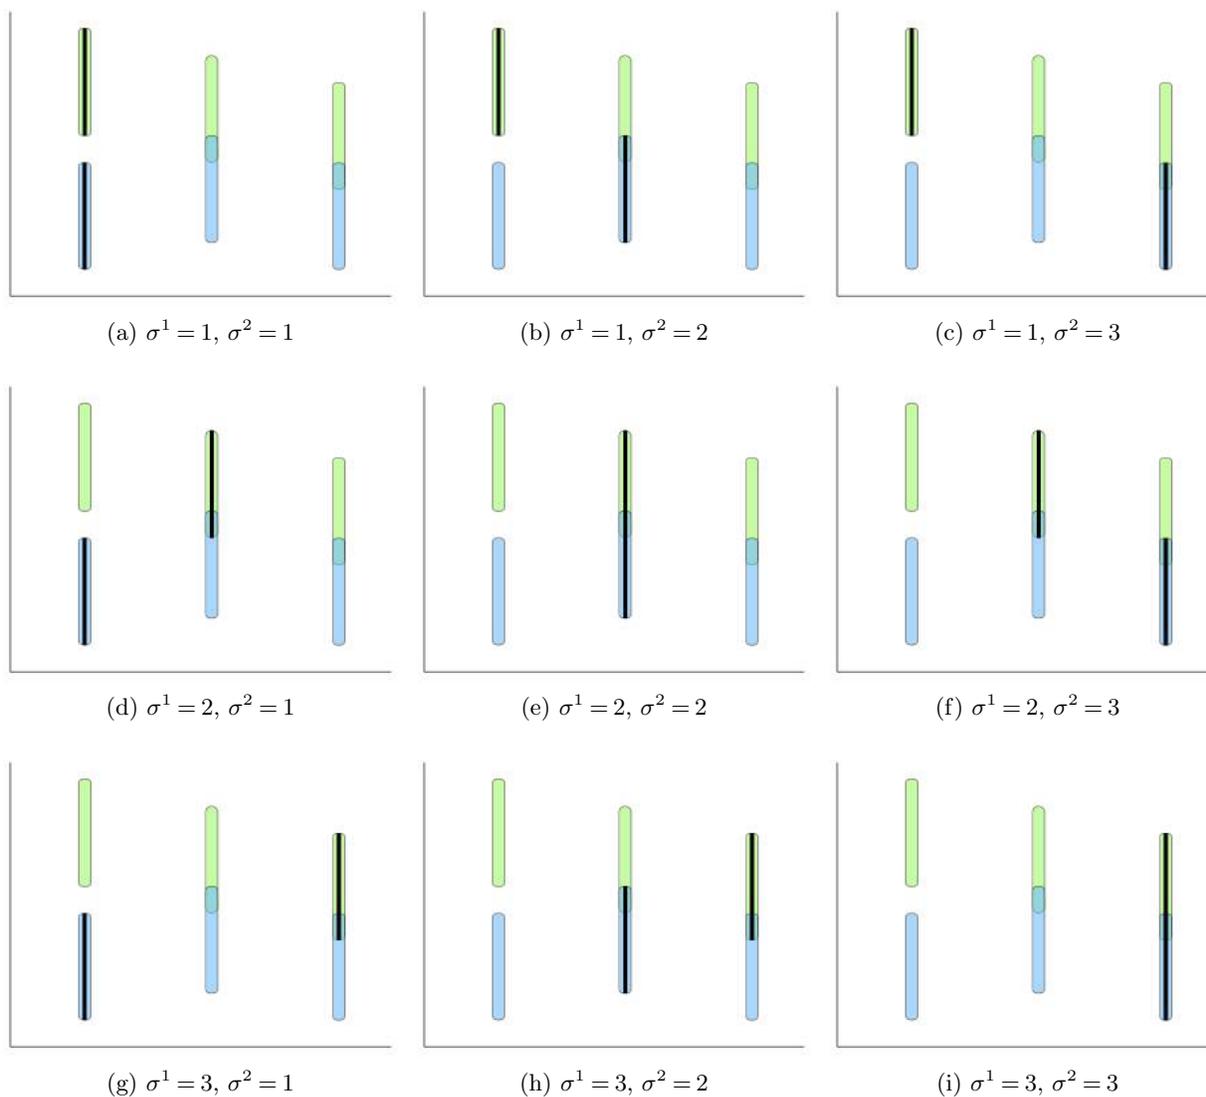


Figure 2 The set of exercise policies \mathcal{M} for the robust optimization problem that is visualized in Figure 1a. Specifically, each of the Figures 2a-2i visualizes the stopping regions induced by the exercise policy $\mu^{\sigma^1\sigma^2} \equiv (\mu_1^{\sigma^1\sigma^2}, \mu_2^{\sigma^1\sigma^2}, \mu_3^{\sigma^1\sigma^2})$ overlaid on top of the uncertainty set around the first sample path $\mathcal{U}^1 \equiv \mathcal{U}_1^1 \times \mathcal{U}_2^1 \times \mathcal{U}_3^1 \equiv [6, 10] \times [5, 9] \times [4, 8]$ and the uncertainty set around the second sample path $\mathcal{U}^2 \equiv \mathcal{U}_1^2 \times \mathcal{U}_2^2 \times \mathcal{U}_3^2 \equiv [1, 5] \times [2, 6] \times [1, 5]$. We observe that none of the exercise policies in the Figures 2a-2i induce identical stopping regions; hence, for this robust optimization problem, the cardinality of the set of exercise policies \mathcal{M} is equal to $3^2 = 9$.

and upper bounded by T^N . Moreover, we observe that the definition of the set of exercise policies is sample-path dependent, in the sense that \mathcal{M} depends on the number and realizations of the simulated sample paths $x^1 \equiv (x_1^1, \dots, x_T^1), \dots, x^N \equiv (x_1^N, \dots, x_T^N)$ and on the choice of the robustness parameter $\epsilon \geq 0$. In Figure 2, we present a visualization of the set of exercise policies \mathcal{M} .

In view of the above notation, we now present our main result:

THEOREM 4. *There exists $\mu \in \mathcal{M}$ that is optimal for (RO).*

The above theorem shows that there always exists an exercise policy in the set \mathcal{M} that is optimal for the robust optimization problem. The result is significant because it will allow us to transform (RO) from an optimization problem over an infinite-dimensional space of exercise policies into a finite-dimensional combinatorial optimization problem over the set \mathcal{M} . Moreover, Theorem 4 is important because it establishes the existence of optimal Markovian stopping rules for the robust optimization problem (RO).

3.2. Proof Sketch of Theorem 4

Our proof of Theorem 4 follows an exchange argument that is rooted in a technique that we refer to as *pruning*. The technique consists of starting with an initial exercise policy $\mu \equiv (\mu_1, \dots, \mu_T)$ and then modifying the exercise policy to reduce the size of the stopping regions $\{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}$. By carefully pruning (*i.e.*, reducing the size of) the stopping regions of the initial exercise policy $\mu \equiv (\mu_1, \dots, \mu_T)$, we will show that any initial exercise policy can be transformed into a new exercise policy $\mu' \in \mathcal{M}$ with the same or better objective value in the robust optimization problem (RO).

To discuss the proof in greater detail, we begin by stating two preliminary lemmas:

LEMMA 1. *The optimal objective value of (RO) is equal to the optimal objective value of*

$$\begin{aligned} & \sup_{\mu} \quad \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) \\ & \text{subject to} \quad \text{for each } i \in \{1, \dots, N\}, \text{ there exists } t \in \{1, \dots, T\} \\ & \quad \quad \quad \text{such that } \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i. \end{aligned} \tag{RO}_T$$

LEMMA 2. *Consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO)_T, and define*

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

Then the following equality holds for each $i \in \{1, \dots, N\}$:

$$\inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) = \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t(y_t) = \text{STOP}\}.$$

The first preliminary lemma shows that a structural constraint can be imposed onto the exercise policies in the robust optimization problem (RO) without any loss of optimality. In particular, Lemma 1 says that we can restrict to exercise policies in which the resulting Markovian stopping rule satisfies $\tau_{\mu}(y) < \infty$ for each trajectory y in each of the uncertainty sets $\mathcal{U}^1, \dots, \mathcal{U}^N$.⁴ The second

⁴ Recall that the definition of a Markovian stopping rule is $\tau_{\mu}(y) \triangleq \min\{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP}\}$. Therefore, we readily observe for each sample path i that an exercise policy $\mu \equiv (\mu_1, \dots, \mu_T)$ satisfies $\tau_{\mu}(y) < \infty$ for all $y \in \mathcal{U}^i \equiv \mathcal{U}_1^i \times \dots \times \mathcal{U}_T^i$ if and only if there exists a period $t \in \{1, \dots, T\}$ that satisfies $\mu_t(y_t) = \text{STOP}$ for all $y_t \in \mathcal{U}_t^i$.

preliminary lemma develops a convenient representation of the objective function of (RO_T) . Specifically, Lemma 2 shows for each sample path i that the quantity $\inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i)$ in the objective function of (RO_T) is equal to the minimum $g(t, x^i)$ among all periods $t \in \{1, \dots, \max_{y \in \mathcal{U}^i} \tau_\mu(y)\}$ for which the uncertainty set \mathcal{U}_t^i has a nonempty intersection with the stopping region $\{y_t \in \mathcal{X}_t : \mu_t(y_t) = \text{STOP}\}$.

Equipped with the above preliminary lemmas, we now formally define the pruning technique that underpins our proof of Theorem 4.

DEFINITION 1. *Let μ be an exercise policy that is feasible for (RO_T) . We say that μ' is a pruned version of μ if the stopping regions of μ' are a subset of the stopping regions of μ , i.e.,*

$$\{y_t \in \mathcal{X} : \mu'_t(y_t) = \text{STOP}\} \subseteq \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\} \quad \forall t \in \{1, \dots, T\},$$

and if the following equality holds for each sample path $i \in \{1, \dots, N\}$:

$$\min \{t \in \{1, \dots, T\} : \mu'_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\} = \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\}.$$

Speaking intuitively, an exercise policy μ' is a pruned version of an exercise policy μ if the stopping regions of μ' are a subset of the stopping regions of μ and if $\max_{y \in \mathcal{U}^i} \tau_\mu(y) = \max_{y \in \mathcal{U}^i} \tau_{\mu'}(y)$ for each sample path i . The significance of Definition 1 in combination with Lemma 2 is established by the following Lemma 3; specifically, the following Lemma 3 shows that if μ' is a pruned version of μ , then the robust objective value associated with μ' is greater than or equal to the robust objective value associated with μ .

LEMMA 3. *If μ' is a pruned version of μ , then $\inf_{y \in \mathcal{U}^i} g(\tau_{\mu'}(y), x^i) \geq \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i) \forall i \in \{1, \dots, N\}$.*

In summary, we have shown in Lemma 1 that the robust optimization problem (RO) is equivalent to the robust optimization problem (RO_T) . Moreover, for every initial exercise policy μ that is feasible for (RO_T) , Lemma 3 shows that the objective value associated with μ is less than or equal to the objective value associated with every exercise policy μ' that is a pruned version of μ . The final step of our proof of Theorem 4 is given by the following lemma:

LEMMA 4. *Consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO_T) , and define*

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

Then $\mu^{\sigma^1 \dots \sigma^N}$ is a pruned version of μ .

The above lemma implies that if μ is a feasible exercise policy for (RO_T) , then there always exists an exercise policy $\mu' \in \mathcal{M}$ that is a pruned version of μ . Hence, we conclude from Lemmas 1, 3, and 4 that there always exists an exercise policy $\mu \in \mathcal{M}$ that is optimal for (RO). The omitted details of the proofs of Theorem 4 and Lemmas 1, 2, 3, and 4 are found in Appendix D.

3.3. Reformulation of (RO) as a Finite-Dimensional Optimization Problem

We conclude §3 by using our characterization of optimal Markovian stopping rules (Theorem 4) to develop a reformulation of (RO) as a finite-dimensional optimization problem over integer decision variables $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. The reformulation from this subsection is important because it establishes a natural combinatorial interpretation of the robust optimization problem (RO), which will provide the foundation for our algorithmic developments in §4.1 and §4.2.

Our finite-dimensional reformulation is presented as the following optimization problem (IP):

$$\text{maximize}_{\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}} \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\}. \quad (\text{IP})$$

The decision variables in the above optimization problem are the integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. We further observe that the inner minimization problems in (IP) involves constraints that depend on whether the intersection $\mathcal{U}_t^i \cap \mathcal{U}_t^j$ is nonempty for each pair of sample paths i, j and each period t . An important insight is that these intersections can be precomputed; that is, given the construction of the uncertainty sets from §2.2, we can efficiently precompute the set of all pairs of sample paths and periods that satisfy $\mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset$.⁵

The equivalence of (IP) and (RO) is formalized by the following theorem.

THEOREM 5. *The optimal objective value of (IP) is equal to the optimal objective value of (RO). Furthermore, for any choice of integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$,*

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g\left(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), x^i\right) \\ & \geq \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\}. \end{aligned}$$

The proof of Theorem 5, which is found in Appendix D, follows readily from Lemmas 1, 2, 3, and 4. Stated in words, the above theorem shows, for any given feasible solution $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ for the optimization problem (IP), that the objective value associated with the exercise policy $\mu^{\sigma^1 \dots \sigma^N}$ in the optimization problem (RO) is greater than or equal to the objective value associated with $\sigma^1, \dots, \sigma^N$ in the optimization problem (IP). Because the above theorem also shows that optimal objective value of (RO) is equal to the optimal objective value of (IP), it follows immediately from Theorem 5 that any optimal solution for (IP) can be transformed into an optimal solution for (RO).

4. Computation of Optimal Markovian Stopping Rules

In this section, we use our characterization of optimal Markovian stopping rules from §3 to develop exact and heuristic algorithms for solving the robust optimization problem (RO). In §4.1, we use this finite-dimensional reformulation of (RO) to design exact algorithms and hardness results for solving the robust optimization problem. In §4.2, we propose and analyze an efficient heuristic algorithm for approximately solving the robust optimization problem. We emphasize that the algorithms and analysis in this section are general and do not require any of the assumptions that were made in §2.4.

⁵ We recall from §2.2 that each uncertainty set $\mathcal{U}_t^i \subseteq \mathbb{R}^d$ is a hypercube. Thus, determining whether $\mathcal{U}_t^i \cap \mathcal{U}_t^j$ is nonempty for each pair of sample paths $i, j \in \{1, \dots, N\}$ and each period $t \in \{1, \dots, T\}$ can be precomputed in a total of $\mathcal{O}(N^2 T d)$ computation time.

4.1. Exact Algorithm

In §3.3, we developed a finite-dimensional optimization problem (IP) with integer decision variables $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ that is equivalent to the infinite-dimensional robust optimization problem (RO). Specifically, we showed that the optimal objective values of the optimization problems (IP) and (RO) are always equal. Moreover, we readily observe from Theorem 5 that any optimal solution for (IP) can further be transformed into an optimal exercise policy for (RO) through the transformation described by line (2) in §3. Harnessing this optimization problem (IP), we now design exact algorithms for solving the robust optimization problem (RO).

We begin our discussion by considering the case of (IP) when there are two periods. Optimal stopping problems with two periods has been studied in the literature as a testbed for understanding the complexity of solving optimal stopping problems, *e.g.*, Glasserman and Yu (2004). We defer the proof of the following result to §4.2.4.

THEOREM 6. *If $T = 2$, then (IP) can be solved in $\mathcal{O}(N^3)$ time.*

Continuing with our discussion on the theoretical tractability of (IP), we next consider problems with three or more periods. The following negative result shows that the computational tractability of (IP) in the case of two periods established in Theorem 6 does not generally extend to optimal stopping problems with three or more periods. The proof of the following result, which is found in Appendix E, consists of a reduction from MIN-2-SAT, which is shown to be strongly NP-hard by Kohli et al. (1994).

THEOREM 7. *(IP) is strongly NP-hard for any fixed $T \geq 3$.*

Motivated by the above hardness result, we proceed to develop an exact algorithm for solving (IP) by reformulating it as a zero-one bilinear program over totally unimodular constraints. This reformulation, presented below as (BP), is valuable for three primary reasons. First, it is well known that bilinear programs can generally be transformed into mixed-integer linear optimization problems using linearization techniques; see Appendix F for details. Thus, for robust optimization problems with small values of T and N , the exact reformulation from this section can directly solved by off-the-shelf commercial optimization software. Second, the exact reformulation (BP) will provide the foundation for our heuristic algorithm for the robust optimization problem in the following §4.2. Third, the exact reformulation (BP) is relatively compact, requiring only $\mathcal{O}(NT^2)$ decision variables and $\mathcal{O}(NT(T+N))$ constraints. In particular, the mild dependence of the size of the exact reformulation on the number of simulated sample paths is attractive in practical settings where the number of simulated sample paths is much larger than the number of time periods.

Our reformulation of (IP) as a zero-one bilinear program over totally unimodular constraints (BP) requires the following additional notation. For each sample path $i \in \{1, \dots, N\}$, we define the following set:

$$\mathcal{K}^i \triangleq \{\kappa : \text{there exists } t \in \{1, \dots, T\} \text{ such that } g(t, x^i) = \kappa\} \cup \{0\}.$$

The above set can be interpreted as the set of distinct values among $g(1, x^i), \dots, g(T, x^i)$ and 0. For notational convenience, let the elements of \mathcal{K}^i be indexed in ascending order, $\kappa_1^i < \dots < \kappa_{|\mathcal{K}^i|}^i$, and let $L_t^i \in \{1, \dots, |\mathcal{K}^i|\}$ be defined for each period $t \in \{1, \dots, T\}$ as the unique index that satisfies the equality $g(t, x^i) = \kappa_{L_t^i}^i$. We readily observe that the quantities L_t^i and $\kappa_1^i, \dots, \kappa_{L_t^i}^i$ can be efficiently

precomputed for each sample path i and period t . With the above notation, our exact reformulation of (IP) is stated as follows:

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_t^i (1 - w_{t\ell}^i) \\
& \text{subject to} && w_{t\ell}^i \leq w_{t+1,\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_t^i \leq w_{t+1,1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\
& && b_t^j \leq w_{t\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\} \text{ such that } g(t, x^i) = \kappa_{\ell}^i \\
& && \text{and } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \\
& && b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\} \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{BP}$$

Let us reflect on the structure of the above optimization problem. First, we observe that (BP) is a bilinear program because its objective function is the sum of products of decision variables b_t^i and $w_{t\ell}^i$. Second, we observe that the bilinear program (BP) is comprised of binary decision variables $b_t^i \in \{0, 1\}$ as well as continuous decision variables $w_{t\ell}^i \in \mathbb{R}$. In particular, we readily observe from the structure of the constraints of (BP) and from the fact that each quantity $\kappa_{\ell+1}^i - \kappa_{\ell}^i$ is strictly positive that there always exists an optimal solution for (BP) in which each decision variable $w_{t\ell}^i$ is equal to zero or one. For this reason, (BP) will be henceforth referred to as a zero-one bilinear program without any ambiguity in terminology. Finally, we remark that the constraints of (BP) are totally unimodular, which implies that every extreme point of the polyhedron defined by the constraints of (BP) is integral (Conforti et al. 2014, §4.2). We will utilize the total unimodularity of (BP) in §4.2 when designing our heuristic algorithm for the robust optimization problem.

We next show that (BP) is indeed equivalent to (IP). In the following Theorem 8, we establish this equivalence and show that any feasible solution for (BP) can be transformed into a feasible solution for (IP) with the same or greater objective value. The proof of Theorem 8, as well as the proofs of the subsequent lemmas in §4.1, are found in Appendix G.

THEOREM 8. *The optimal objective value of (IP) is equal to the optimal objective value of (BP). Furthermore, if b, w is a feasible solution for (BP) and if $\sigma^i \triangleq \min\{\min\{t \in \{1, \dots, T\} : b_t^i = 1\}, T\}$ for each $i \in \{1, \dots, N\}$, then*

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\
& \geq \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_t^i (1 - w_{t\ell}^i).
\end{aligned}$$

The above theorem shows for each sample path i that the decision variables $b_1^i, \dots, b_T^i \in \{0, 1\}$ in (BP) can be interpreted as an encoding of an integer $\sigma^i \in \{1, \dots, T\}$ for (IP). In other words, given any feasible solution b, w for (BP), Theorem 8 shows how to construct a feasible solution $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ for (IP) such that the objective value associated with $\sigma^1, \dots, \sigma^N$ in the optimization problem (IP) is greater than or equal to the objective value associated with b, w in the optimization problem (BP). Because the above theorem also shows that optimal objective value of (IP) is equal to the optimal objective value of (BP), it follows immediately from Theorem 8 that any optimal solution for (BP) can be transformed into an optimal solution for (IP).

We conclude §4.1 by providing intuition for the role of the decision variables $w_{t\ell}^i$ in the optimization problem (BP). To do this, we outline the two key steps of our derivation of (BP). Indeed, recall that the optimization problem (IP) is comprised of integer decision variables $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. The first step in our derivation of (BP) is to introduce binary decision variables $b_1^i, \dots, b_T^i \in \{0, 1\}$ for each sample path i to represent each integer decision variable σ^i . Specifically, consider the following nonlinear binary optimization problem:

$$\begin{aligned} & \text{maximize}_b \quad \frac{1}{N} \sum_{i=1}^N \psi^i(b) \\ & \text{subject to} \quad b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\}, \end{aligned} \tag{BP-1}$$

where the function $\psi^i(b)$ is defined for each sample path $i \in \{1, \dots, N\}$ as

$$\psi^i(b) \triangleq \sum_{t=1}^T b_t^i \left(\prod_{s=1}^{t-1} (1 - b_s^i) \right) \min_{s \in \{1, \dots, t\}} \{g(s, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \text{ and } b_s^j = 1\}.$$

To make sense of the function $\psi^i(b)$, we remark that each quantity $b_t^i \prod_{s=1}^{t-1} (1 - b_s^i)$ will be equal to one if and only if t is the earliest period for sample path i that satisfies the equality $b_t^i = 1$. With this observation, the first step in our derivation of (BP) is comprised of the following lemma:

LEMMA 5. *The optimal objective value of (IP) is equal to the optimal objective value of (BP-1). Furthermore, if b is a feasible solution for (BP-1) and if $\sigma^i \triangleq \min\{\min\{t \in \{1, \dots, T\} : b_t^i = 1\}, T\}$ for each $i \in \{1, \dots, N\}$, then*

$$\frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \geq \frac{1}{N} \sum_{i=1}^N \psi^i(b).$$

The above lemma establishes the equivalence of (IP) of (BP-1), and it shows that any feasible solution for (BP-1) can be transformed into a feasible solution for (IP) with the same or greater objective value. In the second and final step in our derivation of (BP), denoted below by Lemma 6, we show that each function $\psi^i(b)$ can be represented as the optimal objective function of a linear optimization problem.

LEMMA 6. *For each sample path $i \in \{1, \dots, N\}$, $\psi^i(b)$ is equal to the optimal objective value of the following linear optimization problem:*

$$\begin{aligned} & \text{maximize}_{w^i} \quad \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) b_t^i (1 - w_{t\ell}^i) \\ & \text{subject to} \quad w_{t\ell}^i \leq w_{t+1, \ell}^i \quad \text{for all } t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\ & \quad \quad \quad w_{t\ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\ & \quad \quad \quad b_t^j \leq w_{t\ell}^i \quad \text{for all } j \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\} \\ & \quad \quad \quad \quad \quad \quad \text{such that } g(t, x^i) = \kappa_\ell^i \text{ and } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \\ & \quad \quad \quad b_t^i \leq w_{t+1, 1}^i \quad \text{for all } t \in \{1, \dots, T-1\} \\ & \quad \quad \quad w_{t\ell}^i \in \mathbb{R} \quad \text{for all } t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}. \end{aligned} \tag{3}$$

Lemma 6 thus shows that the purpose of the decision variables w_s^i and the constraints of (BP) is to encode the value of $\psi^i(b)$ for each sample path $i \in \{1, \dots, N\}$. Finally, we readily observe that the combination of Lemmas 5 and 6 immediately yields the proof of Theorem 8.

4.2. Heuristic Algorithm

In §4.1, we developed an exact reformulation of the robust optimization problem (RO) as a zero-one bilinear program over totally unimodular constraints (BP). That bilinear program can be solved by off-the-shelf software for mixed-integer linear optimization; see Appendix F. Thus, the previous subsection can be viewed as a concrete and easily implementable exact algorithm for solving the robust optimization problem.

Building upon (BP), we now turn to the task of developing efficient algorithms that can approximately solve the robust optimization problem. Specifically, the main contribution of §4.2 is a practical heuristic for solving the robust optimization problem with computation time that is polynomial in both the number of sample paths N and the number of periods T . The proposed heuristic is thus significantly more computationally tractable than solving (BP), which was shown in Theorem 7 to be NP-hard. Moreover, we will provide theoretical and empirical evidence that the proposed heuristic can perform surprisingly well, both with respect to approximation quality and computational tractability. All omitted proofs of results from §4.2 can be found in Appendix H.

This subsection is organized as follows. In §4.2.1, we discuss the high-level motivation and intuition behind our proposed heuristic for the robust optimization problem. In §4.2.2 and §4.2.3, we formalize the heuristic and offer a strongly polynomial-time algorithm for implementing it. In §4.2.4, we provide theoretical guarantees which show that the approximation gap of the heuristic cannot be arbitrarily bad and, in some cases, is guaranteed to be equal to zero. We study the empirical performance of the heuristic in the subsequent §5.

4.2.1. Preliminaries. Our heuristic for approximately solving the robust optimization problem is motivated by the structure of the exact reformulation (BP) from §4.1 for the robust optimization problem. Recall that (BP) is a zero-one optimization problem with a bilinear objective function and totally unimodular constraints. Because the constraints of (BP) are totally unimodular, the computational difficulty of solving (BP) can thus be attributed to the nonlinearity of its objective function. The proposed heuristic, which is formalized in §4.2.2, aims to contend with this difficulty by approximating the nonlinear objective function of (BP) with a linear function.

To describe the motivation behind our heuristic in greater detail, consider any function $f(b, w)$ that is linear in the vectors b, w and is a lower bound on the objective function of the optimization problem (BP) for all b, w that satisfy the constraints of (BP). That is, let $f(b, w)$ be a linear function that satisfies the following condition:

$$f(b, w) \leq \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_t^i (1 - w_{t\ell}^i) \quad \text{for all } b, w \text{ that are feasible for (BP)}. \quad (4)$$

Then it follows immediately from the condition on line (4) that a conservative, lower-bound approximation of the optimization problem (BP) is given by the following optimization problem:

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && f(b,w) \\
& \text{subject to} && w_{t\ell}^i \leq w_{t+1,\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_t^i \leq w_{t+1,1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\
& && b_t^j \leq w_{t\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\} \text{ such that } g(t, x^i) = \kappa_\ell^i \\
& && \text{and } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \\
& && b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\} \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{H}$$

Indeed, we observe that the constraints of (H) are identical to the constraints of (BP). The only difference between these two optimization problems is that the nonlinear objective function of (BP) has been replaced with the linear function $f(b,w)$ that satisfies the condition from line (4). Therefore, we conclude that every optimal solution for the optimization problem (H) will be a feasible solution for the optimization problem (BP), and the optimal objective value of (H) will always be less than or equal to the optimal objective value of (BP).

Given any linear function $f(b,w)$ that satisfies the condition from line (4), the above optimization problem (H) is useful because it provides a means to computing approximate solutions for (BP). Indeed, the fact that (H) and (BP) have the same constraints implies that any optimal solution for (H) is a feasible solution for (BP). Moreover, it is always theoretically possible to choose the function $f(b,w)$ such that every optimal solution for (H) is an optimal solution for (BP), as shown by the following Proposition 1. We thus conclude that solving the optimization problem (H) has the potential to yield high-quality approximate solutions for (BP).

PROPOSITION 1. *For any instance of (BP), there exists a linear function $f(b,w)$ satisfying the condition from line (4) such that the optimal objective value of (H) is equal to the optimal objective value of (BP), and every optimal solution for (H) is also an optimal solution for (BP).*

The optimization problem (H) is ultimately attractive compared to (BP) from the perspective of computational tractability. Indeed, we observe that (H) is a zero-one linear optimization problem over totally unimodular constraints. This implies that the integrality constraints on the decision variables can be relaxed without loss of generality; that is, (H) can be solved as a linear optimization problem in which each constraints $b_t^i \in \{0, 1\}$ are replaced with $0 \leq b_t^i \leq 1$ (Conforti et al. 2014, §4.2). The optimization problem (H) is thus particularly convenient from an implementation standpoint: as a linear optimization problem, (H) can be easily formulated and solved directly by commercial linear optimization solvers such as CPLEX or Gurobi.

In fact, the optimization problem (H) can be solved *very* efficiently by exploiting the relationship between (H) and the maximal closure problem. The maximal closure problem is a problem from combinatorial optimization that has been widely studied in the operations research literature dating back to Rhys (1970) and Picard (1976), with applications ranging from project scheduling to open-pit mining.⁶ The goal in the maximal closure problem is to find a maximum-weight closure in a directed graph, where a closure is defined as a subset of vertices without edges that leave the subset. In particular, it follows immediately from Picard (1976, §3) that the linear optimization

⁶ For further background on the maximal closure problem, we refer the interested reader to Hochbaum (2004).

problem (H) with $\mathcal{O}(NT^2)$ decision variables and $\mathcal{O}(NT^2)$ constraints is equivalent to computing the maximal closure of a directed graph with $\mathcal{O}(NT^2)$ nodes and $\mathcal{O}(NT(T+N))$ edges. An important algorithmic property is that maximal closure problems can be solved by computing the maximum flow in the directed graph (Picard 1976, §4). Thus, as we will formalize in §4.2.3, the optimization problem (H) can be solved with running time that is strongly polynomial with respect to both the number of simulated sample paths N and the number of time periods T .⁷

4.2.2. Description of Heuristic. In §4.2.1, we discussed the high-level motivation and intuition behind our proposed heuristic. In view of that motivation, we now formally describe our proposed heuristic for approximately solving the optimization problem (BP). Specifically, our proposed heuristic consists of solving the optimization problem (H) with a particular linear objective function, denoted below by $\bar{f}(b, w)$.

To formally describe our proposed heuristic, we require the following notation. Recall that T is the number of periods in the optimal stopping problem. For each sample path i , let $T^i \triangleq \arg \max_{t \in \{1, \dots, T\}} g(t, x^i)$ be defined as the period in which the sample path i achieves its maximum reward, and if there are multiple optimal solutions, we choose the optimal solution that is smallest. To simplify our notation, let us define $\mathcal{T}^i \triangleq \{T^i\} \cup \{T\}$ as the set that contains the period in which sample path i achieves its maximum reward as well as the last period of the optimal stopping period.

With this additional notation, the proposed heuristic obtains approximate solutions b, w for the optimization problem (BP) by solving the optimization problem (H) with the objective function

$$\bar{f}(b, w) \triangleq \frac{1}{N} \sum_{i=1}^N \sum_{t \in \mathcal{T}^i} \sum_{\ell=1}^{L_t^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (b_t^i - w_{t\ell}^i).$$

In other words, our heuristic for solving the optimization problem (BP) consists of solving the optimization problem (H) with objective function given by $f(b, w) = \bar{f}(b, w)$. The optimal solution for (H) with objective function $\bar{f}(b, w)$ constitutes a feasible solution for (BP), since the constraints of (H) and (BP) are identical. This solution can thus be transformed into a Markovian stopping rule using the transformations described in Theorem 5 and 8.

REMARK 2. We readily observe that the function $\bar{f}(b, w)$ is linear in b and w . Moreover, it follows from algebra that $\bar{f}(b, w)$ is less than or equal to the objective function of (BP) for all feasible solutions of (BP); that is, the linear function $\bar{f}(b, w)$ satisfies the condition from line (4).⁸ Thus, we conclude that solving the optimization problem (H) with objective function $f(b, w) = \bar{f}(b, w)$ will provide a lower-bound approximation of the optimization problem (BP), and any optimal solution for (H) will be a feasible solution for (BP).

Our motivations for using the heuristic outlined above are two-fold. First, we find that the optimization problem (H) with the objective function $\bar{f}(b, w)$ is highly tractable from both a theoretical and empirical standpoint on realistic sizes of robust optimization problems. Second,

⁷ We note that Proposition 1 does not contradict our hardness result for the robust optimization problem (RO) from Theorem 7, as the right choice of the linear function $f(b, w)$ in Proposition 1 may be difficult to identify.

⁸ To see why $\bar{f}(b, w)$ is a lower bound on the objective function of (BP), consider any arbitrary vectors b, w that satisfy the constraints of (BP). Since feasibility for the optimization problem (BP) implies that b is a binary vector, we observe that the equality $b_t^i(1 - w_{t\ell}^i) = \max\{b_t^i - w_{t\ell}^i, 0\}$ holds for each $i \in \{1, \dots, N\}$, $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, |\mathcal{K}^i|\}$. Therefore, a lower-bound approximation of the objective function of the optimization problem (BP) can be obtained by replacing each term $b_t^i(1 - w_{t\ell}^i)$ with $b_t^i - w_{t\ell}^i$ if $t \in \mathcal{T}^i$ and with 0 if $t \notin \mathcal{T}^i$.

we provide theoretical and empirical evidence that the above heuristic can find high-quality and in some cases optimal solutions for (BP). We elaborate on the theoretical aspects of these two motivations in §4.2.3 and §4.2.4, and we explore the empirical performance of the heuristic in the subsequent §5.

4.2.3. Algorithms for Implementing Heuristic. In §4.2.1, we argued generally that the optimization problem (H) with any linear objective function $f(b, w)$ can be solved in strongly polynomial time by reducing (H) to a maximal closure problem. We now formalize those earlier arguments to derive an explicit, strongly polynomial-time algorithm for solving the optimization problem (H) in the particular case where $f(b, w) = \bar{f}(b, w)$. Specifically, the main contribution of §4.2.3 is the development of an algorithm that achieves the running time that is specified in the following proposition:

PROPOSITION 2. *If $f(b, w) = \bar{f}(b, w)$, then (H) can be solved in $\mathcal{O}(N^2T(N+T))$ time.*

The above proposition establishes the computational tractability of our heuristic by demonstrating that an optimal solution for the optimization problem (H) with objective function $f(b, w) = \bar{f}(b, w)$ can be computed with running time that is polynomial in both the number of simulated sample paths as well as the number of time periods. In particular, if the number of periods T is held constant, then the running time in Proposition 2 scales cubically in the number of simulated sample paths N . Such a tractability guarantee is ultimately important from a practical perspective: indeed, in the following §5, we will present numerical experiments which show that our heuristic can run in seconds on realistic problem sizes with over fifty periods and thousands of sample paths. An outline of the proof of Proposition 2 is found throughout the rest of §4.2.3.

To develop an algorithm with the computation time that is specified in Proposition 2, we begin by deriving a compact reformulation of (H) when the objective function satisfies $f(b, w) = \bar{f}(b, w)$. The compact reformulation is denoted below as the optimization problem (\bar{H}). To derive this compact reformulation, we first show that many of the decision variables in the optimization problem (H) can, without loss of generality, be removed from (H) when the objective function satisfies $f(b, w) = \bar{f}(b, w)$. In particular, our compact reformulation will utilize the following lemma:

LEMMA 7. *If $f(b, w) = \bar{f}(b, w)$, then there exists an optimal solution b, w for (H) that satisfies $b_T^1 = \dots = b_T^N = 1$ and satisfies $b_t^i = 0$ for each sample path $i \in \{1, \dots, N\}$ and period $t \in \{1, \dots, T\} \setminus \mathcal{T}^i$.*

The above lemma demonstrates that the optimal values for many of the decision variables of (H) can be known in advance when the objective function satisfies $f(b, w) = \bar{f}(b, w)$. Hence, Lemma 7 implies that each decision variable b_t^i can be fixed to its optimal value and removed from (H) when $f(b, w) = \bar{f}(b, w)$, except for the decision variables b_t^i in which the sample path $i \in \{1, \dots, N\}$ and period $t \in \{1, \dots, T\}$ satisfy $t = T^i$ and $T^i < T$. In view of Lemma 7, we now consider the following optimization problem:

$$\begin{aligned}
& \underset{b, w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i < T} \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (b_{T^i}^i - w_{T^i \ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_T^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (1 - w_{T \ell}^i) \\
& \text{subject to} && w_{t \ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_{T^i}^i \leq w_{T^i 1}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\
& && b_{T^j}^j \leq w_{t \ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \mathcal{T}^i \text{ such that } g(T^j, x^i) = \kappa_{\ell}^i, \\
& && \quad \quad \quad T^j \leq t, T^j < T, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\
& && b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\
& && w_{t \ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{\bar{H}}$$

It can be shown using straightforward algebra that the above optimization problem (\bar{H}) is equivalent to the optimization problem (H) when the objective function of (H) satisfies $f(b, w) = \bar{f}(b, w)$ and when the decision variables in (H) are restricted to satisfy $b_T^1 = \dots = b_T^N = 1$ and $b_t^i = 0$ for all $i \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\} \setminus \mathcal{T}^i$; see proof of the following Lemma 8 for details. Hence, it follows from Lemma 7 that any optimal solution for (\bar{H}) can be transformed into an optimal solution for (H) when the objective function satisfies $f(b, w) = \bar{f}(b, w)$. We formalize the equivalence of (H) and (\bar{H}) in the following Lemma 8:

LEMMA 8. *Let $f(b, w) = \bar{f}(b, w)$. Then, the optimal objective value of (\bar{H}) is equal to the optimal objective value of (H) . Moreover, let \bar{b}, \bar{w} denote an optimal solution for (\bar{H}) . Then there exists an optimal solution b, w for (H) that satisfies the following equality for each sample path $i \in \{1, \dots, N\}$ and period $t \in \{1, \dots, T\}$:*

$$b_t^i = \begin{cases} \bar{b}_t^i, & \text{if } t = T^i \text{ and } T^i < T, \\ 1, & \text{if } t = T, \\ 0, & \text{otherwise.} \end{cases}$$

To summarize, we have shown in Lemmas 7 and 8 that solving the optimization problem (H) can be reduced to solving the optimization problem (\bar{H}) when the objective function of (H) satisfies $f(b, w) = \bar{f}(b, w)$. More precisely, we have shown that any optimal solution for (\bar{H}) can be transformed into an optimal choice for the decision variables b in (H) when the objective function of (H) satisfies $f(b, w) = \bar{f}(b, w)$. By applying the transformations described in Theorems 5 and 8, we can construct exercise policies $\mu \equiv (\mu_1, \dots, \mu_T)$ from the decision variables b whose robust objective value in the robust optimization problem $\widehat{J}_{N, \epsilon}(\mu)$ is greater than or equal to the optimal objective value of (\bar{H}) .

From a computational tractability standpoint, the optimization problem (\bar{H}) is ultimately attractive because it has significantly fewer decision variables and fewer constraints than the optimization problem (H) . Indeed, we observe from inspection that (\bar{H}) has $\mathcal{O}(NT)$ decision variables and $\mathcal{O}(NT + N^2)$ constraints, whereas (H) has $\mathcal{O}(NT^2)$ decision variables and $\mathcal{O}(NT^2 + N^2T)$ constraints. Because of the same reasoning as given in §4.2.1, the integrality constraints on the decision variables can also be relaxed without loss of generality; that is, (\bar{H}) can be solved as a linear optimization problem in which each constraints $b_t^i \in \{0, 1\}$ are replaced with $0 \leq b_t^i \leq 1$. Thus, (\bar{H}) can be easily formulated and solved directly by commercial linear optimization solvers.

We conclude §4.2.3 by using the above optimization problem (\bar{H}) to outline the proof of Proposition 2. Specifically, our proof of Proposition 2 consists of reformulating (\bar{H}) as a maximal closure problem in a directed graph. We then use the efficient maximum-flow algorithm of Orlin (2013) to solve the maximal closure problem. The remaining details for the proof of Proposition 2 are found in Appendix H.

4.2.4. Approximation Guarantees. We conclude §4.2 by performing a theoretical analysis of the approximation quality of the proposed heuristic (\bar{H}) for the optimization problem (BP) . Our motivation here is to understand whether the proposed heuristic is ever guaranteed to find optimal solutions for (BP) and, conversely, whether it is ever possible for the gap between the optimal objective values of (\bar{H}) and (BP) to be arbitrarily large. Our answers to these theoretical questions are presented below in Propositions 3, 4, and 5.

We begin by focusing on optimal stopping problems in which the number of periods is equal to two. We recall from Theorem 6 in §4.1 that any instance of the robust optimization problem with

two periods can be solved exactly in strongly polynomial-time. We will now prove Theorem 6 by establishing that the approximation gap between (\bar{H}) and (BP) is always equal to zero when the number of periods is equal to two:

PROPOSITION 3. *If $T = 2$, then the optimal objective values of (\bar{H}) and (BP) are equal.*

Our takeaways from the above proposition are two-fold. First, Proposition 3 shows that there indeed exist settings in which the proposed heuristic is guaranteed to find optimal solutions for (BP). Second, Proposition 3 in combination with Theorems 5 and 8 implies that solving the robust optimization problem (RO) can be reduced to solving the optimization problem (\bar{H}) when $T = 2$. Therefore, we observe that the proof of Theorem 6 follows immediately from Propositions 2 and 3.

We next develop a general bound on the gap between the optimal objective values of (\bar{H}) and (BP) that holds for any fixed number of periods T . For notational convenience, we let J^{BP} denote the optimal objective value of (BP) and $J^{\bar{H}}$ denote the optimal objective value of (\bar{H}) .

PROPOSITION 4. $J^{\bar{H}} \geq \frac{1}{T} J^{\text{BP}}$.

The above proposition can be viewed as valuable from a theoretical perspective, as it shows that the gap between the optimal objective values of (\bar{H}) and (BP) can never be arbitrarily large. Indeed, we recall from Remark 2 and Lemma 8 that the optimal objective value of (\bar{H}) is always less than or equal to the optimal objective value of (BP). Thus, Proposition 4 establishes that there exist bounds on the gap between the optimal objective values of (\bar{H}) and (BP) that are independent of the choice of the number of simulated sample paths N , the choice of the robustness parameter ϵ that controls the size of the uncertainty sets, the reward function, etc.

Our third theoretical guarantee in §4.2.4, stated below as Proposition 5, provides a bound on the gap between the optimal objective values of (\bar{H}) and (BP) that does not depend on the number of periods T . Our theoretical bound in Proposition 5 will require that the reward function of the NMOS problem satisfies the following regularity condition:

ASSUMPTION 5. *There exists a constant $L > 0$ such that, for each period $t \in \{1, \dots, T\}$, the reward function satisfies $|g(t, y) - g(t, y')| \leq L \|y_t - y'_t\|_\infty$ for all $y \equiv (y_1, \dots, y_T), y' \equiv (y'_1, \dots, y'_T) \in \mathcal{X}^T$.*

The above assumption stipulates for each period $t \in \{1, \dots, T\}$ that the reward function $y \mapsto g(t, y)$ is Lipschitz-continuous with respect to the state $y_t \in \mathcal{X}$. The above assumption is relatively mild in applications such as options pricing when the reward function depends only on the current state in each period.⁹ We will impose Assumption 5 in the following Proposition 5 to eliminate pathological situations in which slight perturbations of sample paths lead to drastic changes in reward. Once again, for notational convenience, we let $J^{\bar{H}}$ and J^{BP} denote the optimal objective values of the optimization problems (\bar{H}) and (BP), respectively.

PROPOSITION 5. *If Assumption 5 holds, then $J^{\bar{H}} \geq \frac{1}{\log N + 1} J^{\text{BP}} - 2\epsilon L$.*

Compared to the bound from Proposition 4, we observe that the bound from Proposition 5 is attractive when the robustness parameter ϵ is relatively small and when the number of simulated sample paths N is relatively small compared to the number of periods T . For example, we observe that $\frac{1}{\log N + 1} \ll \frac{1}{T}$ will hold whenever the number of sample paths N is subexponential in the number of periods T .

⁹ For example, we observe that Assumption 5 is satisfied by the reward function in the multi-dimensional barrier option pricing problem from §5.2 when the state space in the optimal stopping problem is augmented as $(x_t, q_t) \equiv (\max_{a \in \{1, \dots, d\}} \xi_{t,a}, 1 - \mathbb{I}\{\max_{a \in \{1, \dots, d\}} \xi_{s,a} \leq B(s) \text{ for all } s \in \{1, \dots, t\}\})$.

In summary, Propositions 3, 4, and 5 establish under mild and verifiable conditions that the approximation gap between the proposed heuristic (\bar{H}) and the robust optimization problem (RO) cannot be arbitrarily large and, in some cases, is guaranteed to be equal to zero. Ultimately, the practical value of the proposed heuristic (\bar{H}) lies in its performance in the context of optimal stopping applications. In the following §5, we provide numerical evidence that (\bar{H}) can indeed find high-quality stopping rules for realistic stochastic optimal stopping problems in practical computation times.

5. Numerical Experiments

In this section, we perform numerical experiments to compare our robust optimization approach and three state-of-the-art benchmarks from the literature (Longstaff and Schwartz 2001, Ciocan and Mišić 2020, Desai et al. 2012). The first two benchmarks serve as representatives of two classes of approximation methods for stochastic optimal stopping problems (approximate dynamic programming and parametric exercise policies) which, similarly as the robust optimization approach, only require the ability to simulate sample paths of the entire sequence of random states. The third benchmark is representative of state-of-the-art duality-based methods for stochastic optimal stopping problems. All experiments were conducted on a 2.6 GHz 6-Core Intel Core i7 processor with 16 GB of memory. All methods are implemented in the Julia programming language and solved using the JuMP library and Gurobi optimization software.

5.1. A Simple Non-Markovian Problem

To demonstrate the value of our robust optimization approach, we begin by investigating a simple, one-dimensional stochastic optimal stopping problem with a non-Markovian probability distribution. The optimal stopping problem of consideration involves a state space which is equal to the real numbers ($\mathcal{X} = \mathbb{R}^1$) and a reward function of $g(t, x) = x_t$ in each period $t \in \{1, \dots, T\}$. For any fixed duration $\Delta \in \mathbb{N}$, the joint probability distribution of the stochastic process is given by

$$x_t \sim \text{Uniform}[0, 1] + \frac{2\theta}{T} \mathbb{I}\{\theta \leq t \leq \theta + \Delta\} \quad \text{for all } t = 1, \dots, T,$$

where the random parameter $\theta \sim \text{Uniform}\{1, 2, \dots, T - \Delta\}$ is selected once per sample path and is unobserved. Simulated sample paths of this non-Markovian stochastic process $x = (x_1, \dots, x_T)$ are visualized in Figure 3. We perform numerical experiments on the following methods:

- **Robust Optimization (RO):** The robust optimization approach is used here to approximate (OPT) over the non-Markovian stochastic process $x \equiv (x_1, \dots, x_T)$, and thus aims to find the best Markovian stopping rules for the stochastic optimal stopping problem. The method was run with robustness parameters $\epsilon \in \{0, 0.01, \dots, 0.2\}$ for $N < 1000$ and $\epsilon \in \{0, 0.01, \dots, 0.1\}$ otherwise, and the robust optimization problem is solved approximately using our proposed heuristic (\bar{H}) from §4.2.3.
- **Least-Squares Regression (LS):** We implement the method of Longstaff and Schwartz, which employs least-squares regression to approximate the continuation value function (*i.e.*, the expected reward from not stopping) in each period using backwards recursion. To apply this method, we first transform each non-Markovian sample path $x^i \equiv (x_1^i, \dots, x_T^i)$ into an augmented Markovian sample path of the form $X^i \equiv (X_1^i, \dots, X_T^i)$ by adding the full history into the state in each period: $X_t^i \triangleq (x_1^i, \dots, x_t^i, 0, \dots, 0)$. The regression step requires a specification of basis functions, and we consider the following categories of basis functions.

Full-History: This category of basis functions uses the entire vector X_t^i of states observed up to that point. The basis functions that we consider in this category are ONE (the constant function 1), PRICES (the states observed up to that point, $X_t^i \in \mathbb{R}^T$), and PRICES2 (the product of each pair of states observed up to that point, $X_t^i (X_t^i)^\top \in \mathbb{R}^{T \times T}$).

Markovian: This category of basis functions uses only the current state x_t^i in each period. The purpose of considering these basis functions is to analyze the performance when, like the robust optimization approach, the method of Longstaff-Schwartz is restricted to stopping rules which depend only on the current state in each period. We consider basis functions based on the Laguerre polynomials, where LAGUERRE-K is the polynomial $\sum_{\ell=0}^k \binom{k}{\ell} \frac{(-1)^\ell}{\ell!} (x_t^i)^\ell$.

- **Tree Method (Tree):** The method of Ciocan and Mišić approximates the stochastic optimal stopping problem (OPT) by restricting the space of exercise policies to decision trees. Like the robust optimization approach, the tree approach is used to find Markovian stopping rules for the non-Markovian optimal stopping problem (OPT). We apply the method with the same information as the robust optimization approach at each time period (the current state and the time period) and with a splitting parameter of 0.005.

In our computational experiments, we consider a time horizon of $T = 50$ stopping periods with a duration parameter of $\Delta = 5$. All methods were run using simulated training datasets of sizes $N \in \{10^2, 10^{2.1}, \dots, 10^{3.9}, 10^4\}$. The robustness parameters in the “RO” method were selected using a validation set of size $\tilde{N} = 10^3$; see §2.5 for more details. All methods were evaluated on a common and independent testing dataset of $\tilde{N} = 10^5$ sample paths, and experiments were repeated over 10 replications.

Figures 3 and 4 visualize the Markovian stopping rules found by the robust optimization approach, the expected rewards of the stopping rules obtained by the various methods, and the computation times of the various methods. The results of this experiment show that the robust optimization approach outperforms the other approaches, producing stopping rules with an expected reward of approximately 1.62 from training datasets of $N = 10^3$ simulated sample paths and with a computation time of approximately 0.713 seconds. We reflect below on the main differences between our robust optimization approach and the benchmark methods.

For the least-squares regression method with full-history, selecting a good choice of basis functions is found to be a first-order challenge. For example, the basis functions (ONE, PRICES) turn out to be insufficiently rich to provide an accurate approximation of the continuation function with the full state history. When the number of sample paths is sufficiently large, we expect that the basis functions (ONE, PRICES, PRICES2) should provide a better approximation of the stochastic optimal stopping problem. However, the computational cost resulting from this rich class of basis functions precluded its practicality on sufficiently large training datasets. In contrast, the robust optimization approach only searches for Markovian stopping rules, and in doing so, has a reduction in sample complexity. In addition to achieving a significantly better expected reward in this example, the Markovian stopping rules found by the robust optimization approach are considerably more interpretable than those which have full-history dependence, as illustrated in Figure 3.

Compared to the tree method, the complex structure of the Markovian stopping rules found by the robust optimization approach demonstrates the value of algorithms that do not impose parametric restrictions on the exercise policies. In theory, decision trees with sufficient depth are capable of approximating the best Markovian stopping rules to the NMOS problem to arbitrary accuracy (Ciocan and Mišić 2020, Theorem 2). However, the greedy heuristic that is proposed by Ciocan and Mišić (2020) to efficiently optimize over decision trees is not able to find a good approximation of the optimal Markovian stopping rule in this example, even as the number of sample paths is large. We note that this issue is not unique to the decision trees *per se*; the

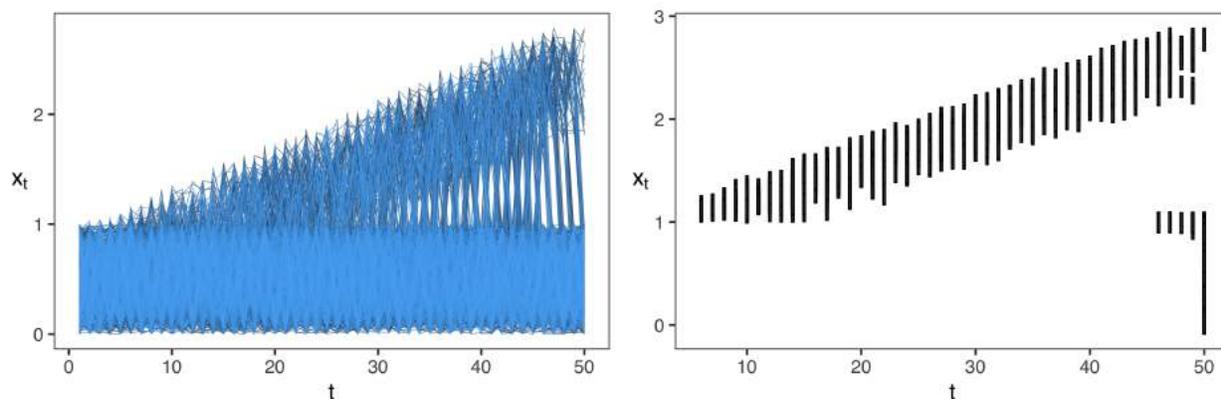


Figure 3 The left figure shows simulated sample paths of the non-Markovian stochastic process with $T = 50$ and $\Delta = 5$. The right figure shows the exercise policies obtained from solving the robust optimization problem constructed from a training dataset of $N = 10^3$ simulated sample paths and with robustness parameter $\epsilon = 0.1$.

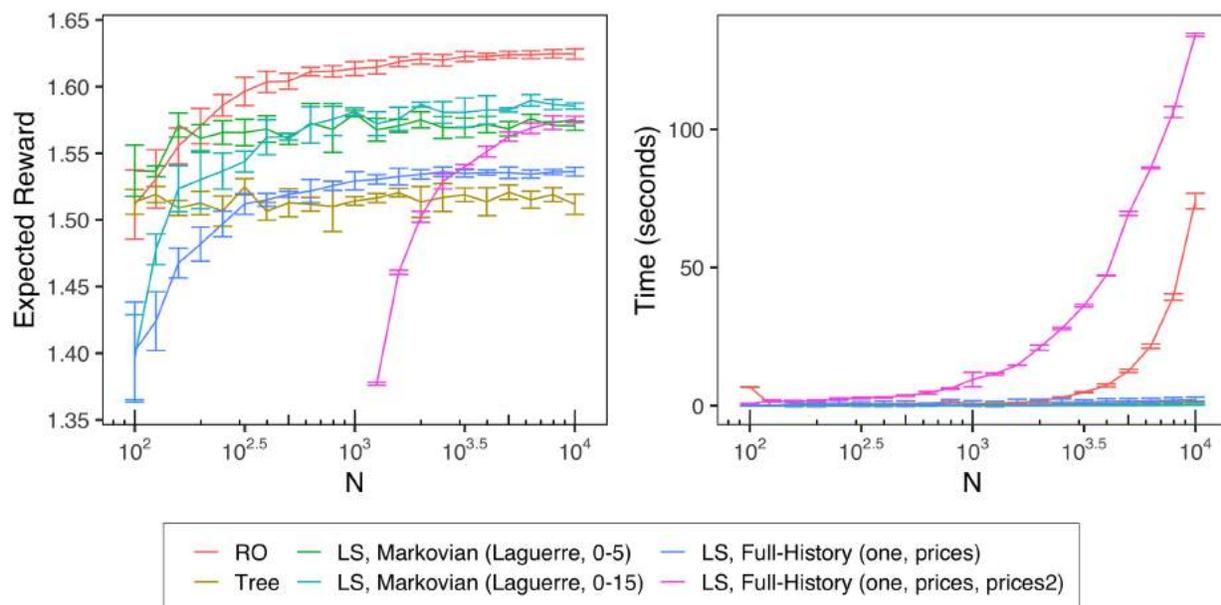


Figure 4 The left figure shows the expected reward of stopping rules obtained by the various methods, and the right figure shows the total computation time. The results for the robust optimization approach are shown using the heuristic (\bar{H}) from §4.2.3, the displayed computation time is the total time to construct and approximately solve the robust optimization problems over all choices of the robustness parameter, and the expected rewards are displayed for the robustness parameter that is chosen using a validation set of size $\bar{N} = 10^3$ for each choice of N . All results for all methods are averaged over 10 replications. The results in the left figure are truncated at 1.35 for visual clarity.

computational intractability of optimizing over parametric spaces of Markovian stopping rules has resulted in heuristics for other settings as well (Glasserman 2013, §8.2). This example provides evidence that our robust optimization approach, in conjunction with the proposed approximation from §4.2, can yield high-quality Markovian stopping rules to complex (low-dimensional) non-Markovian optimal stopping problems.

Like the robust optimization and tree methods, “LS, Markovian” aims to find Markovian stopping rules for the non-Markovian optimal stopping problem. In particular, we observe that “LS, Markovian” was implemented with basis functions of Laguerre polynomials with a rather large degree.¹⁰ The expected reward of “LS, Markovian (Laguerre 0-15)” may thus be interpreted as an estimate for the best Markovian stopping rule that one could achieve using dynamic programming. However, the optimality of backwards induction no longer holds for exercise policies that depend on the current state when the stochastic process is non-Markovian. This explains the inferior performance of this method in this example and motivates the use of methods for addressing non-Markovian optimal stopping problems that optimize directly over the exercise policies in all periods simultaneously. A further discussion on the limitations of dynamic programming in finding Markovian stopping rules for NMOS problems can be found in Appendix A.

In Figure 5, we present the relationship between the best choice of the robustness parameter for our heuristic (\bar{H}) and the number of simulated sample paths N . The results in this figure show that the best choice of the robustness parameter ϵ decreases as the number of simulated sample paths N increases. Although a formal analysis of the relationship between these two quantities is outside the scope of this paper, Figure 5 provides empirical evidence that only smaller choices of the robustness parameter need to be considered when using larger numbers of simulated sample paths.

In Figure 6, we compare the (in-sample) robust objective values and the expected rewards of stopping rules obtained by our heuristic (\bar{H}) and our exact reformulation (BP) of the robust optimization problem. The comparison in Figure 6 is performed on a smaller instance where $T = 20$, $\Delta = 2$, and the robust optimization problem is constructed from a training set of $N = 100$ simulated sample paths. We observe the gap in objective values between the two algorithms is relatively small across choices of the robustness parameter, and the gap is equal to zero when the robustness parameter is set equal to zero. Moreover, the expected rewards of stopping rules obtained from these two algorithms were similarly relatively close across choices of the robustness parameter. We view these results as promising as they suggest, at least for the present example, that the significantly more tractable optimization problem (\bar{H}) can provide a close approximation of the robust optimization problem (RO).

In summary, the experiments from this subsection reveal our first setting in which the robust optimization approach is attractive. In non-Markovian optimal stopping problems, there can exist high-quality stopping rules which are Markovian. However, finding such stopping rules is a non-trivial challenge, as the best Markovian stopping rules will not necessarily exhibit a simple structure, and cannot in general be found using dynamic programming. The robust optimization approach thus provides a practical nonparametric approach for finding them. In view of these observations, the next subsection considers a benchmark optimal stopping problem from the options pricing literature.

¹⁰ We were unable to run the “LS, Markovian (Laguerre 0-K)” for degrees $K > 15$ due to the numerical precision required to encode the coefficients in the basis functions.

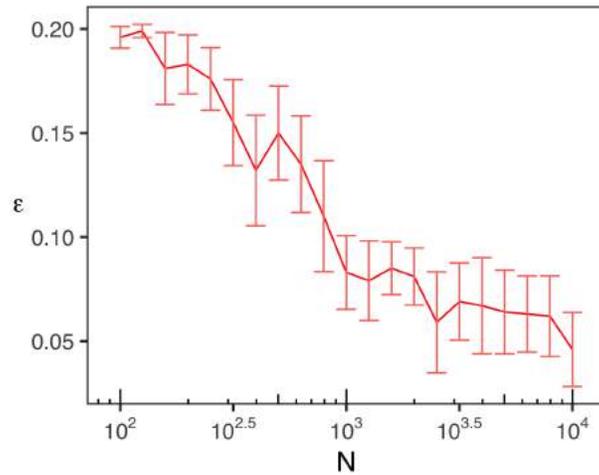


Figure 5 The figure shows the best choice of the robustness parameter ϵ for the heuristic (\bar{H}) as a function of the number of training sample paths N . The optimal robustness parameter is chosen using a validation set of size $\bar{N} = 10^3$ for each choice of N . All results are averaged over 10 replications.

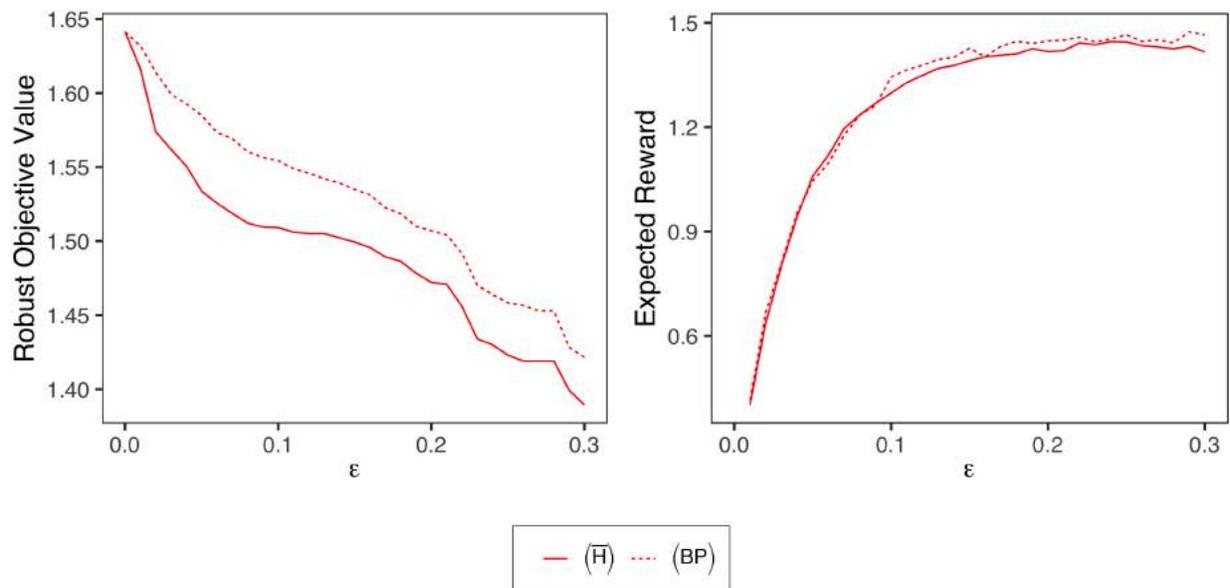


Figure 6 The left figure shows the robust objective value of stopping rules obtained from solving the heuristic (\bar{H}) and the exact reformulation (BP) of the robust optimization problem (RO) on an instance with $T = 20$ periods. The right figure shows the expected rewards of the stopping rules obtained by these two algorithms. The expected rewards are estimated by evaluating the stopping rules on a testing dataset of 10^5 simulated sample paths.

5.2. Pricing Multi-Dimensional Barrier Options

Building upon the previous section, we next consider the well-studied problem of pricing discretely-monitored barrier call options over multiple assets. These are stochastic optimal stopping problems in which the random state at each period, $\xi_t \in \Xi \equiv \mathbb{R}^d$, is comprised of d non-dividend paying assets. Provided that the option has not been ‘knocked-out’, the reward from stopping on period t is essentially an increasing function of the maximum value of the current assets, $\max_{a \in \{1, \dots, d\}} \xi_{t,a}$. However, if the maximum value of the assets exceeds a prespecified barrier, the option becomes ‘knocked-out’ and the reward becomes equal to zero for the remainder of the time horizon. In this section, we compare various methods on several instances of these stochastic optimal stopping problems in which the barrier threshold changes over time (Kunitomo and Ikeda 1992).

To formalize the problem setting, let the components of the state vector $\xi_t \in \Xi \equiv \mathbb{R}^d$ be referenced through the notation $\xi_t \equiv (\xi_{t,1}, \dots, \xi_{t,d})$, where $\xi_{t,a} \in \mathbb{R}$ represents the value of asset a at exercise opportunity $t \in \{1, \dots, T\}$. The exercise opportunities are evenly spaced over a calendar of Y years and thus, defining $\lambda \triangleq Y/T$, it follows that exercise opportunity t occurs at the calendar time λt . Let $r \in [0, 1)$ be the annualized discount rate, $K \geq 0$ be the strike price, $B(t) \triangleq B_0 e^{\delta \lambda t}$ be the barrier threshold on exercise opportunity t , and $\xi \equiv (\xi_1, \dots, \xi_T)$ be the sequence of random states. With this notation, the reward function of this stochastic optimal stopping problem is defined as

$$g(t, \xi) = \begin{cases} e^{-r\lambda t} \max \left\{ 0, \max_{a \in \{1, \dots, d\}} \xi_{t,a} - K \right\}, & \text{if } \max_{a \in \{1, \dots, d\}} \xi_{s,a} \leq B(s) \text{ for all } s \in \{1, \dots, t\}, \\ 0, & \text{otherwise.} \end{cases}$$

Under the standard Black-Scholes setup, the sequence of random states $\xi \equiv (\xi_1, \dots, \xi_T)$ obeys a multidimensional geometric Brownian motion where each asset $a \in \{1, \dots, d\}$ has an initial value of \bar{x} , drift equal to the risk-free rate r , and annualized volatility equal to σ_a . The value of asset a at exercise opportunity t is thus given by $\xi_{t,a} \triangleq \bar{x} e^{(r - \sigma_a^2/2)\lambda t + \sigma_a \lambda W_{t,a}}$, where each $W_{t,a}$ is a standard Brownian motion process and the instantaneous correlation of $W_{\cdot,a}$ and $W_{\cdot,a'}$ is equal to $\rho_{a,a'}$. We perform numerical experiments on the following methods:

- **Robust Optimization (RO):** The robust optimization approach is used here to approximate (OPT) over the projected stochastic process $x \equiv (x_1, \dots, x_T)$, where $x_t \triangleq \max_{a \in \{1, \dots, d\}} \xi_{t,a}$ denotes the maximum value of the assets in each exercise opportunity t . Note that this stochastic process does not include information on whether the option has been knocked-out. The robust optimization approach is run with robustness parameters $\epsilon \in \{0\} \cup \{0.01, \dots, 0.09\} \cup \{0.1, \dots, 0.9\} \cup \{1, \dots, 10\}$ and solved approximately using our heuristic (\bar{H}) from §4.2.3.
- **Least-Squares Regression (LS):** The method of Longstaff and Schwartz is applied to stochastic optimal stopping problem (OPT) over the Markovian stochastic process $X \equiv (X_1, \dots, X_T)$, where the state at each exercise opportunity $X_t \triangleq (\xi_{t,1}, \dots, \xi_{t,d}, q_t)$ consists of the values of the assets at the current exercise opportunity as well as an indicator variable which equals one if the option has been knocked out: $q_t \triangleq 1 - \mathbb{I}\{\max_{a \in \{1, \dots, d\}} \xi_{s,a} \leq B(s) \text{ for all } s \in \{1, \dots, t\}\}$. We consider the following basis functions:
 - ONE: The constant function, 1.
 - KOIND: An indicator variable for whether the option was knocked out, q_t .
 - PRICES: The values of the assets, $\xi_{t,1}, \dots, \xi_{t,d}$.
 - PRICESKO: The assets multiplied by the indicator variable, $\xi_{t,1}(1 - q_t), \dots, \xi_{t,d}(1 - q_t)$.
 - MAXPRICE: The maximum asset value, $x_t \triangleq \max_{a \in \{1, \dots, d\}} \xi_{t,a}$.
 - PAYOFF: The reward of exercising the option, $(1 - q_t)e^{-r\lambda t} \max\{0, x_t - K\}$.

- **Pathwise Optimization (PO):** The duality-based method of Desai et al. is applied to the stochastic optimal stopping problem (OPT) over the Markovian stochastic process $X \equiv (X_1, \dots, X_T)$, where the state at each exercise opportunity $X_t \triangleq (\xi_{t,1}, \dots, \xi_{t,d}, q_t)$ consists of the values of the assets at the current exercise opportunity as well as an indicator variable which equals one if the option has been knocked out: $q_t \triangleq 1 - \mathbb{I}\{\max_{a \in \{1, \dots, d\}} \xi_{s,a} \leq B(s) \text{ for all } s \in \{1, \dots, t\}\}$. We consider the same possible basis functions for PO as were considered for LS.
- **Tree Method (Tree):** The method of Ciocan and Mišić approximates the stochastic optimal stopping problem (OPT) by restricting the space of exercise policies to decision trees. Like the robust optimization approach, the tree approach is used to find Markovian stopping rules for the optimal stopping problem over the projected stochastic process $x \equiv (x_1, \dots, x_T)$. We apply the method with the same information as the robust optimization approach at each time period (the current state x_t and the time period t) and with a splitting parameter of 0.005.

Our experiments and parameter settings closely parallel those of Ciocan and Mišić (2020, §5), albeit with two primary differences. First, our experiments consider a barrier threshold that changes as a function of time. Second, we perform experiments in which the underlying assets have symmetrical as well as asymmetrical annualized volatilities (*i.e.*, $\sigma_a \neq \sigma_{a'}$ for $a \neq a'$). Our motivations behind these differences are to compare the various methods in settings in which the best Markovian stopping rules may not have a simple structure and in which the low-dimensional projections x_1, \dots, x_T are not sufficient statistics for the optimal stopping rules (Broadie and Detemple 1997, §7). In other words, our experiments aim to evaluate the performance of the proposed robust optimization approach in settings where the best Markovian stopping rules are not guaranteed to be optimal stopping rules for the optimal stopping problem. Similarly as Ciocan and Mišić (2020, §5), we perform experiments on various numbers of assets ($d \in \{8, 16, 32\}$) and various initial prices for the assets ($\bar{x} \in \{90, 100, 110\}$).

In our implementation, we experimented with a variety of combinations of basis functions, and, for the sake of brevity, we report results only for a subset of combinations with the best performance. The expected rewards of the stopping rules obtained by the various methods were estimated using a common and independent testing dataset of $\tilde{N} = 10^5$ sample paths, and all experiments were repeated over 10 replications. The RO method is run in each experiment with $N = 10^3$ training sample paths and $\tilde{N} = 10^3$ validation sample paths (see §2.5 for additional details), the LS and Tree methods are run in each experiment with 10^5 training sample paths, and the PO method is run in each experiment with 10^3 outer sample paths and 500 inner sample paths. We use smaller numbers of training and outer sample paths for RO and PO because the computation times of these methods scales relatively quickly with the number of sample paths.

In Table 1, we present the expected rewards of stopping rules obtained from the various methods in experiments with symmetrical annualized volatilities.¹¹ Across the parameter settings, the RO method yields stopping rules with either the best or close-to-best expected reward relative to the alternative state-of-the-art methods. The improvements of the robust optimization approach over the state-of-the-art benchmarks are viewed as particularly encouraging given the practical significance of this class of options pricing problems. In Figure 7, we present visualizations of the Markovian stopping rules produced by the RO method. These visualizations provide interpretability to the stopping rules found by the RO method, which is not possible for the LS method due to the high-dimensional state space. Further visualizations of the stopping rules obtained by the RO method under additional problem parameters are provided in Appendix I.

¹¹ Expected rewards and computation times for experiments with asymmetrical annualized volatility can be found in Tables EC.1 and EC.2 in Appendix I.

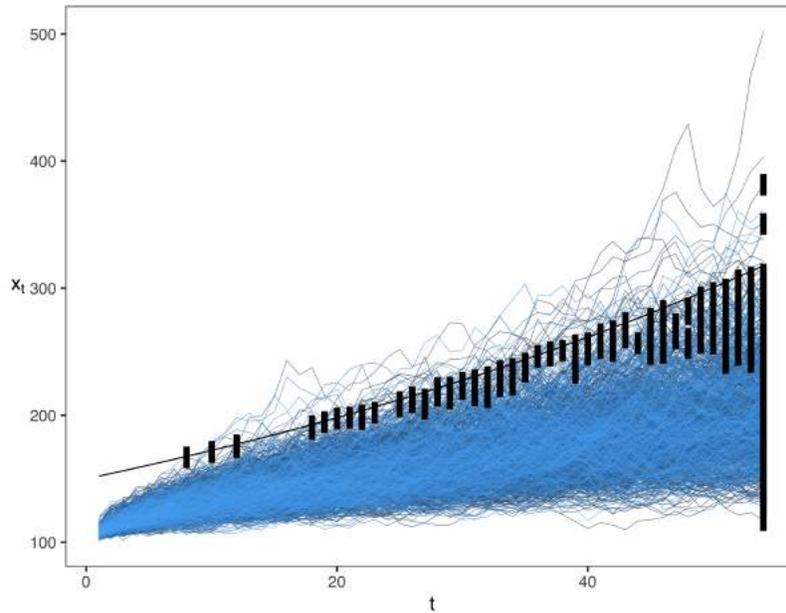


Figure 7 Plot shows exercise policies obtained from solving a robust optimization problem constructed from a training dataset of size $N = 10^3$ and with the robustness parameter selected using a validation set of size $\tilde{N} = 10^3$. The problem parameters are $\bar{x} = 100$ and $d = 16$. The remaining parameters are the same as those shown in Table 1. The thick black rectangles are the stopping regions and the thin black line shows the barrier threshold.

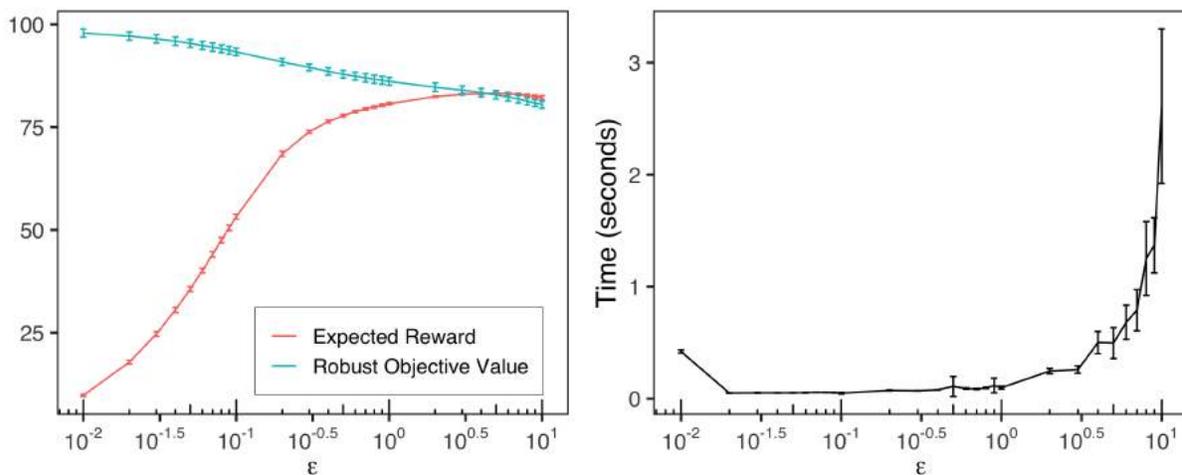


Figure 8 Each plot shows performance metrics of the robust optimization problems constructed from a training dataset of size $N = 10^3$ as a function of the robustness parameter ϵ . The left plot shows the robust objective value and expected reward of policies obtained by solving the robust optimization problem. The right plot shows the computation time for solving the robust optimization problem. In both plots, the robust optimization problem is solved approximately using our heuristic algorithm (\bar{H}) from §4.2.3. The problem parameters are $\bar{x} = 100$ and $d = 16$, and the remaining parameters are the same as those shown in Table 1.

Table 1 Barrier Option (Symmetric) - Expected Reward.

d	Method	Basis functions	Initial Price			# of Sample Paths
			$\bar{x} = 90$	$\bar{x} = 100$	$\bar{x} = 110$	
8	RO	maxprice	54.88 (0.26)	68.35 (0.13)	75.93 (0.40)	10^3 training, 10^3 validation
8	LS	one, pricesKO, KOind, payoff	54.71 (0.09)	67.22 (0.10)	73.52 (0.11)	10^5
8	LS	one, pricesKO, payoff	54.80 (0.09)	67.29 (0.11)	73.77 (0.12)	10^5
8	LS	payoff, KOind, pricesKO	54.71 (0.09)	67.22 (0.10)	73.52 (0.11)	10^5
8	LS	pricesKO, payoff	54.78 (0.09)	67.28 (0.11)	73.78 (0.12)	10^5
8	LS	one, prices, payoff	54.74 (0.09)	65.57 (0.08)	68.05 (0.14)	10^5
8	LS	one	44.96 (0.04)	57.13 (0.10)	60.41 (0.09)	10^5
8	LS	one, KOind, prices	49.38 (0.04)	62.87 (0.10)	69.86 (0.08)	10^5
8	LS	one, prices	48.90 (0.05)	59.96 (0.12)	59.53 (0.08)	10^5
8	LS	one, pricesKO	48.99 (0.05)	62.19 (0.07)	69.68 (0.08)	10^5
8	LS	maxprice, KOind, pricesKO	51.16 (0.07)	63.37 (0.11)	69.92 (0.08)	10^5
8	PO	payoff, KOind, pricesKO	54.47 (0.11)	65.73 (0.18)	67.90 (0.20)	2×10^3 outer, 500 inner
8	PO	prices	52.19 (0.17)	64.80 (0.16)	68.11 (0.33)	2×10^3 outer, 500 inner
8	Tree	payoff, time	54.73 (0.08)	66.82 (0.12)	71.30 (0.27)	10^5
8	Tree	maxprice, time	54.73 (0.08)	66.82 (0.12)	71.30 (0.27)	10^5
16	RO	maxprice	71.00 (0.21)	83.11 (0.36)	82.77 (0.36)	10^3 training, 10^3 validation
16	LS	one, pricesKO, KOind, payoff	70.65 (0.07)	81.25 (0.14)	81.38 (0.15)	10^5
16	LS	one, pricesKO, payoff	70.68 (0.08)	81.40 (0.15)	81.05 (0.16)	10^5
16	LS	payoff, KOind, pricesKO	70.65 (0.07)	81.25 (0.14)	81.38 (0.15)	10^5
16	LS	pricesKO, payoff	70.67 (0.08)	81.41 (0.15)	81.05 (0.16)	10^5
16	LS	one, prices, payoff	70.04 (0.10)	76.50 (0.17)	70.00 (0.15)	10^5
16	LS	one	61.48 (0.06)	70.28 (0.07)	57.58 (0.09)	10^5
16	LS	one, KOind, prices	65.28 (0.08)	77.36 (0.09)	74.91 (0.12)	10^5
16	LS	one, prices	64.03 (0.08)	70.34 (0.08)	59.75 (0.09)	10^5
16	LS	one, pricesKO	64.67 (0.07)	77.15 (0.09)	74.56 (0.14)	10^5
16	LS	maxprice, KOind, pricesKO	66.45 (0.08)	77.62 (0.10)	74.93 (0.14)	10^5
16	PO	payoff, KOind, pricesKO	69.85 (0.15)	77.85 (0.16)	68.30 (0.17)	2×10^3 outer, 500 inner
16	PO	prices	68.16 (0.18)	77.48 (0.16)	68.44 (0.30)	2×10^3 outer, 500 inner
16	Tree	payoff, time	70.44 (0.10)	79.88 (0.12)	74.71 (0.54)	10^5
16	Tree	maxprice, time	70.44 (0.10)	79.88 (0.12)	74.71 (0.54)	10^5
32	RO	maxprice	86.15 (0.27)	93.17 (0.29)	78.63 (0.48)	10^3 training, 10^3 validation
32	LS	one, pricesKO, KOind, payoff	85.00 (0.10)	90.89 (0.15)	79.72 (0.16)	10^5
32	LS	one, pricesKO, payoff	85.03 (0.09)	91.03 (0.16)	78.49 (0.14)	10^5
32	LS	payoff, KOind, pricesKO	85.00 (0.10)	90.89 (0.15)	79.72 (0.16)	10^5
32	LS	pricesKO, payoff	85.03 (0.10)	91.04 (0.15)	78.48 (0.14)	10^5
32	LS	one, prices, payoff	82.21 (0.08)	81.15 (0.15)	60.94 (0.16)	10^5
32	LS	one	76.08 (0.08)	72.07 (0.19)	52.26 (0.08)	10^5
32	LS	one, KOind, prices	80.33 (0.07)	86.28 (0.12)	68.40 (0.13)	10^5
32	LS	one, prices	77.16 (0.11)	71.90 (0.17)	54.46 (0.05)	10^5
32	LS	one, pricesKO	80.00 (0.06)	86.22 (0.12)	65.40 (0.20)	10^5
32	LS	maxprice, KOind, pricesKO	80.91 (0.07)	86.33 (0.11)	68.81 (0.15)	10^5
32	PO	payoff, KOind, pricesKO	83.06 (0.20)	82.22 (0.32)	56.67 (0.44)	2×10^3 outer, 500 inner
32	PO	prices	82.12 (0.29)	82.50 (0.27)	56.85 (0.17)	2×10^3 outer, 500 inner
32	Tree	payoff, time	84.36 (0.16)	86.68 (0.13)	61.09 (0.13)	10^5
32	Tree	maxprice, time	84.36 (0.16)	86.68 (0.13)	61.09 (0.13)	10^5

Optimal is indicated in bold for each number of assets $d \in \{8, 16, 32\}$ and initial price $\bar{x} \in \{90, 100, 110\}$. Problem parameters are $T = 54$, $Y = 3$, $r = 0.05$, $K = 100$, $B_0 = 150$, $\delta = 0.25$, $\sigma_a = 0.2$, $\rho_{a,a'} = 0$ for all $a \neq a'$.

Table 2 Barrier Option (Symmetric) - Computation Times.

d	Method	Basis functions	Initial Price			# of Sample Paths
			$\bar{x} = 90$	$\bar{x} = 100$	$\bar{x} = 110$	
8	RO	maxprice	7.73 (0.33)	8.63 (1.36)	13.93 (2.74)	10^3 training, 10^3 validation
8	LS	one, pricesKO, KOind, payoff	3.35 (0.21)	3.53 (0.5)	3.14 (0.22)	10^5
8	LS	one, pricesKO, payoff	3.24 (0.2)	2.97 (0.29)	2.94 (0.23)	10^5
8	LS	payoff, KOind, pricesKO	3.52 (0.83)	2.97 (0.28)	3.03 (0.44)	10^5
8	LS	pricesKO, payoff	3.04 (0.25)	2.81 (0.22)	2.84 (0.19)	10^5
8	LS	one, prices, payoff	3.08 (0.18)	2.97 (0.25)	2.99 (0.1)	10^5
8	LS	one	0.86 (0.07)	0.88 (0.03)	0.93 (0.05)	10^5
8	LS	one, KOind, prices	3.13 (0.29)	2.99 (0.32)	2.92 (0.28)	10^5
8	LS	one, prices	2.64 (0.19)	2.63 (0.15)	2.62 (0.13)	10^5
8	LS	one, pricesKO	2.72 (0.34)	2.77 (0.17)	2.57 (0.17)	10^5
8	LS	maxprice, KOind, pricesKO	3.02 (0.15)	2.99 (0.32)	2.87 (0.08)	10^5
8	PO	payoff, KOind, pricesKO	20.40 (0.54)	19.29 (0.69)	17.58 (0.92)	2×10^3 outer, 500 inner
8	PO	prices	32.61 (7.86)	29.01 (8.21)	23.70 (0.58)	2×10^3 outer, 500 inner
8	Tree	payoff, time	14.18 (0.43)	14.21 (0.18)	19.06 (3.12)	10^5
8	Tree	maxprice, time	13.72 (0.3)	13.84 (0.3)	18.28 (3.13)	10^5
16	RO	maxprice	6.82 (0.19)	10.07 (1.17)	28.96 (8.34)	10^3 training, 10^3 validation
16	LS	one, pricesKO, KOind, payoff	5.49 (0.62)	5.15 (0.17)	5.20 (0.24)	10^5
16	LS	one, pricesKO, payoff	5.29 (1.02)	4.84 (0.43)	5.10 (0.33)	10^5
16	LS	payoff, KOind, pricesKO	4.71 (0.44)	4.64 (0.3)	4.61 (0.32)	10^5
16	LS	pricesKO, payoff	5.17 (0.71)	4.82 (0.15)	4.74 (0.25)	10^5
16	LS	one, prices, payoff	4.86 (0.81)	4.69 (0.41)	4.97 (0.1)	10^5
16	LS	one	1.27 (0.15)	1.21 (0.08)	1.22 (0.08)	10^5
16	LS	one, KOind, prices	6.24 (1.34)	5.69 (1.28)	5.50 (0.69)	10^5
16	LS	one, prices	5.97 (1.17)	5.51 (1.55)	5.01 (0.87)	10^5
16	LS	one, pricesKO	4.30 (0.3)	4.17 (0.29)	4.45 (0.31)	10^5
16	LS	maxprice, KOind, pricesKO	5.12 (0.61)	4.68 (0.27)	4.63 (0.28)	10^5
16	PO	payoff, KOind, pricesKO	40.18 (1.59)	32.82 (1.38)	28.44 (1.41)	2×10^3 outer, 500 inner
16	PO	prices	57.17 (24.52)	57.76 (25.68)	44.87 (0.49)	2×10^3 outer, 500 inner
16	Tree	payoff, time	14.80 (0.28)	14.19 (0.48)	25.00 (7.02)	10^5
16	Tree	maxprice, time	13.88 (0.17)	13.74 (0.21)	23.81 (6.46)	10^5
32	RO	maxprice	8.48 (1.28)	14.93 (3.58)	91.80 (15.54)	10^3 training, 10^3 validation
32	LS	one, pricesKO, KOind, payoff	13.73 (1.39)	13.32 (2.22)	12.00 (1.35)	10^5
32	LS	one, pricesKO, payoff	15.42 (3.9)	15.33 (2.99)	12.89 (2.17)	10^5
32	LS	payoff, KOind, pricesKO	13.05 (3.69)	11.33 (1.27)	11.66 (2.39)	10^5
32	LS	pricesKO, payoff	14.63 (2.31)	14.01 (3.91)	11.35 (1.54)	10^5
32	LS	one, prices, payoff	15.18 (5.56)	14.17 (3.53)	11.73 (1.72)	10^5
32	LS	one	2.13 (1.23)	1.78 (0.93)	1.67 (0.73)	10^5
32	LS	one, KOind, prices	13.25 (2.33)	12.38 (2.44)	11.95 (1.75)	10^5
32	LS	one, prices	12.23 (2.71)	10.83 (1.02)	10.91 (0.8)	10^5
32	LS	one, pricesKO	14.55 (4.76)	11.12 (1.1)	10.56 (0.6)	10^5
32	LS	maxprice, KOind, pricesKO	12.05 (1.72)	10.92 (0.83)	10.93 (1.42)	10^5
32	PO	payoff, KOind, pricesKO	91.66 (5.6)	62.49 (5.23)	53.57 (1.55)	2×10^3 outer, 500 inner
32	PO	prices	120.46 (3.27)	120.76 (1.98)	113.86 (2.31)	2×10^3 outer, 500 inner
32	Tree	payoff, time	17.53 (0.23)	17.42 (0.28)	10.42 (0.18)	10^5
32	Tree	maxprice, time	17.01 (0.18)	16.97 (0.27)	10.90 (0.31)	10^5

Problem parameters are $T = 54$, $Y = 3$, $r = 0.05$, $K = 100$, $B_0 = 150$, $\delta = 0.25$, $\sigma_a = 0.2$, $\rho_{a,a'} = 0$ for all $a \neq a'$.

Table 3 Barrier Option (Symmetric) - Best Choice of Robustness Parameter.

d	Initial Price		
	$\bar{x} = 90$	$\bar{x} = 100$	$\bar{x} = 110$
8	8.0 (2.45)	6.7 (1.25)	3.7 (1.06)
16	8.3 (1.89)	5.5 (1.27)	2.8 (1.03)
32	6.8 (1.32)	3.4 (0.84)	2.6 (0.70)

Best choice of robustness parameter ϵ found using the validation method from §2.5 in the robust optimization problems constructed from a training dataset of size $N = 10^3$ and validation set of size $\bar{N} = 10^3$. The remaining parameters are the same as those shown in Table 1.

In Table 2, we present the computation times of the various methods. We observe from Table 2 that the computation times of the RO method are highly competitive compared with the computation times of the alternative methods when the initial price is $\bar{x} = 90$ or $\bar{x} = 100$. Interestingly, the computation time of the RO method increases significantly when the initial price is $\bar{x} = 110$. To understand why this is the case, we recall that the computation time for the RO method is equal to the sum of the computation times for solving the heuristic (\bar{H}) over all considered choices of the robustness parameter. Moreover, in Figure 8, we show that the computation time of solving (\bar{H}) is highly dependent on the robustness parameter; indeed, most of the computation time for the RO method comes from solving the robust optimization problem with unnecessarily large choices of the robustness parameter. Finally, we show in Table 3 that the best choice of the robustness parameter is smallest in the numerical experiments with $\bar{x} = 110$, even though the set of robustness parameters remains constant across all of the numerical experiments. From these results, we conclude that a dynamic search over the space of robustness parameters has the potential to significantly reduce the computation times of the RO method, particularly the case where $\bar{x} = 110$. Additional numerical experiments which show the relationship between the robustness parameter and the robust optimization approach are provided in Appendix I.

6. Conclusion

Over the past two decades, dynamic robust optimization has experienced a surge of algorithmic advances. Until now, these advances from robust optimization have not been harnessed to develop algorithms for stochastic dynamic optimization problems with known probability distributions. In this paper, we showed the value of bridging these traditionally separate fields of research, in application to the classical and widely-studied problem of optimal stopping. In this context, we devised new and theoretically-justified algorithms for non-Markovian optimal stopping problems and highlighted the performance of these new algorithms on stylized and well-studied problems from options pricing. Along the way, we also developed novel theoretical and computational results for solving dynamic robust optimization problems which average over multiple uncertainty sets. We believe this work takes a meaningful step towards broadening the impact of robust optimization to address stochastic dynamic optimization problems of importance to industry.

Acknowledgments

The author thanks Vivek Farias for his feedback during the early stages of this project, and Velibor Mišić and Selvaprabu Nadarajah for comments on early versions of the paper. The author also thanks the anonymous associate editor and two referees for constructive comments and suggestions that helped improve

the paper. Finally, the author is grateful to the authors of the paper Ciocan and Mišić (2020) for making their high-quality code publicly available.

References

- Adams WP, Sherali HD (1986) A tight linearization and an algorithm for zero-one quadratic programming problems. *Management Science* 32(10):1274–1290.
- Andersen L (1999) A simple approach to the pricing of bermudan swaptions in the multi-factor libor market model. *Journal of Computational Finance* 3(2):5–32.
- Andersen L, Broadie M (2004) Primal-dual simulation algorithm for pricing multidimensional american options. *Management Science* 50(9):1222–1234.
- Bally V, Pages G (2003) A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. *Bernoulli* 9(6):1003–1049.
- Bayer C, Häppölä J, Tempone R (2019) Implied stopping rules for american basket options from markovian projection. *Quantitative Finance* 19(3):371–390.
- Bayer C, Tempone R, Wolfers S (2020) Pricing american options by exercise rate optimization. *Quantitative Finance* 1–12.
- Bayraktar E, Yao S (2014) On the robust optimal stopping problem. *SIAM Journal on Control and Optimization* 52(5):3135–3175.
- Becker S, Cheridito P, Jentzen A (2019) Deep optimal stopping. *Journal of Machine Learning Research* 20:74.
- Belomestny D (2011a) On the rates of convergence of simulation-based optimization algorithms for optimal stopping problems. *The Annals of Applied Probability* 21(1):215–239.
- Belomestny D (2011b) Pricing bermudan options by nonparametric regression: optimal rates of convergence for lower estimates. *Finance and Stochastics* 15(4):655–683.
- Belomestny D (2013) Solving optimal stopping problems via empirical dual optimization. *The Annals of Applied Probability* 23(5):1988–2019.
- Ben-Tal A, El Ghaoui L, Nemirovski A (2009) *Robust Optimization*, volume 28 (Princeton University Press).
- Bertsimas D, Shtern S, Sturt B (2023) A data-driven approach to multistage stochastic linear optimization. *Management Science* 69(1):51–74.
- Bezerra SC, Ohashi A, Russo F, de Souza F (2020) Discrete-type approximations for non-markovian optimal stopping problems: Part ii. *Methodology and Computing in Applied Probability* 1–35.
- Broadie M, Detemple J (1997) The valuation of american options on multiple assets. *Mathematical Finance* 7(3):241–286.
- Broadie M, Glasserman P (1997) Pricing american-style securities using simulation. *Journal of Economic Dynamics and Control* 21(8-9):1323–1352.
- Brown DB, Smith JE, Sun P (2010) Information relaxations and duality in stochastic dynamic programs. *Operations Research* 58(4-part-1):785–801.
- Bruss FT (2005) What is known about robbins’ problem? *Journal of Applied Probability* 42(1):108–120.
- Carriere JF (1996) Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics* 19(1):19–30.
- Ciocan DF, Mišić VV (2020) Interpretable optimal stopping. *Management Science (forthcoming)* .
- Clément E, Lamberton D, Protter P (2002) An analysis of a least squares regression method for american option pricing. *Finance and Stochastics* 6(4):449–471.
- Conforti M, Cornuéjols G, Zambelli G (2014) *Integer Programming*, volume 271 (Springer).
- David I, Yechiali U (1985) A time-dependent stopping problem with application to live organ transplants. *Operations Research* 33(3):491–504.

- Delage E, Iancu DA (2015) Robust multistage decision making. *The Operations Research Revolution*, 20–46 (INFORMS).
- Desai VV, Farias VF, Moallemi CC (2012) Pathwise optimization for optimal stopping problems. *Management Science* 58(12):2292–2308.
- Egloff D (2005) Monte carlo algorithms for optimal stopping and statistical learning. *The Annals of Applied Probability* 15(2):1396–1432.
- Fathan A, Delage E (2021) Deep reinforcement learning for optimal stopping with application in financial engineering. *arXiv preprint arXiv:2105.08877* .
- Feng Y, Gallego G (1995) Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science* 41(8):1371–1391.
- Garcia D (2003) Convergence and biases of monte carlo estimates of american option prices using a parametric exercise rule. *Journal of Economic Dynamics and Control* 27(10):1855–1879.
- Gatheral J, Jaisson T, Rosenbaum M (2018) Volatility is rough. *Quantitative Finance* 18(6):933–949.
- Gemrich S (2012) *Multilevel Monte Carlo methods for American options*. Master’s thesis, University of Oxford.
- Glasserman P (2013) *Monte Carlo Methods in Financial Engineering*, volume 53 (Springer Science & Business Media).
- Glasserman P, Yu B (2004) Number of paths versus number of basis functions in american option pricing. *The Annals of Applied Probability* 14(4):2090–2119.
- Goldberg DA, Chen Y (2018) Beating the curse of dimensionality in options pricing and optimal stopping. *arXiv preprint arXiv:1807.02227* .
- Goudenège L, Molent A, Zanette A (2020) Machine learning for pricing american options in high-dimensional markovian and non-markovian models. *Quantitative Finance* 20(4):573–591.
- Haugh MB, Kogan L (2004) Pricing american options: a duality approach. *Operations Research* 52(2):258–270.
- Hochbaum DS (2004) 50th anniversary article: Selection, provisioning, shared fixed costs, maximum closure, and implications on algorithmic methods today. *Management Science* 50(6):709–723.
- Iancu DA, Trichakis N, Yoon DY (2021) Monitoring with limited information. *Management Science* 67(7):4233–4251.
- Kohli R, Krishnamurti R, Mirchandani P (1994) The minimum satisfiability problem. *SIAM Journal on Discrete Mathematics* 7(2):275–283.
- Kunitomo N, Ikeda M (1992) Pricing options with curved boundaries. *Mathematical Finance* 2(4):275–298.
- Leão D, Ohashi A, Russo F (2019) Discrete-type approximations for non-markovian optimal stopping problems: Part i. *Journal of Applied Probability* 56(4):981–1005.
- Longstaff FA, Schwartz ES (2001) Valuing american options by simulation: a simple least-squares approach. *The Review of Financial Studies* 14(1):113–147.
- Orlin JB (2013) Max flows in $o(nm)$ time, or better. *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, 765–774.
- Picard JC (1976) Maximal closure of a graph and applications to combinatorial problems. *Management Science* 22(11):1268–1272.
- Reuters (2021) U.s. equity options set new volume record URL <https://www.reuters.com/article/us-usa-stocks-options-idUSKBN29K20I>.
- Rhys JM (1970) A selection problem of shared fixed costs and network flows. *Management Science* 17(3):200–207.
- Riedel F (2009) Optimal stopping with multiple priors. *Econometrica* 77(3):857–908.
- Rogers LC (2002) Monte carlo valuation of american options. *Mathematical Finance* 12(3):271–286.

- Rudin W (1964) *Principles of Mathematical Analysis*, volume 3 (McGraw-Hill New York).
- Shapiro A, Dentcheva D, Ruszczyński A (2014) *Lectures on stochastic programming: modeling and theory* (SIAM).
- Sturt B (2020) *Dynamic Optimization in the Age of Big Data*. Ph.D. thesis, Massachusetts Institute of Technology.
- Tsitsiklis JN, Van Roy B (2001) Regression methods for pricing complex american-style options. *IEEE Transactions on Neural Networks* 12(4):694–703.
- Zanger DZ (2020) General error estimates for the longstaff–schwartz least-squares monte carlo algorithm. *Mathematics of Operations Research* 45(3):923–946.

Technical Proofs and Additional Results

Appendix A: Limitations of DP for Finding Markovian Stopping Rules in Non-Markovian Optimal Stopping Problems

In the context of non-Markovian optimal stopping problems, a subtle but important challenge is that the exercise policies which define the best Markovian stopping rule cannot be found in general using backwards recursion. Intuitively, this problem arises because Bellman's dynamic programming equations no longer hold when the stochastic process is non-Markovian. This fact is illustrated numerically in §5.1 and motivates the development of methods which optimize over the exercise policies in all time periods simultaneously. For the sake of completeness, we provide the following example in which there is a Markovian stopping rule that is optimal for the non-Markovian stopping problem but is not obtained using backwards recursion.

EXAMPLE EC.1. Consider a three-period optimal stopping problem with a one-dimensional non-Markovian stopping process that obeys the following probability distribution:

$$\mathbb{P}(x_1 = 3, x_2 = 2, x_3 = 1) = \frac{2}{3}; \quad \mathbb{P}(x_1 = 1, x_2 = 2, x_3 = 3) = \frac{1}{3}.$$

Let the reward of stopping on each period t be equal to the current state x_t , and recall that our goal is to find a stopping rule which maximizes the expected reward. We observe that there is an optimal stopping rule for this non-Markovian stopping problem that is a Markovian stopping rule, defined by exercise policies $\mu_1^*(x_1) = \text{STOP}$ if and only if $x_1 = 3$, $\mu_2^*(x_2) = \text{CONTINUE}$ for all x_2 , and $\mu_3^*(x_3) = \text{STOP}$ for all x_3 . Applying this stopping rule yields an expected reward of $\frac{2}{3} \times 3 + \frac{1}{3} \times 3 = 3$.

We now show that the best Markovian stopping rule which is obtained using backwards recursion will have a strictly lower expected reward. Indeed, assume that the reward from not stopping on any period is equal to zero. Then, starting on the last period, it is clear from the dynamic programming principle that the optimal exercise policy for the last period is $\mu_3^{\text{DP}}(x_3) = \text{STOP}$. Hence, conditioned on $x_2 = 2$, the expected reward from stopping on the third period is $1 \times \mathbb{P}(x_3 = 1 \mid x_2 = 2) + 3 \times \mathbb{P}(x_3 = 3 \mid x_2 = 2) = 1 \times \frac{2}{3} + 3 \times \frac{1}{3} = \frac{5}{3}$. Because the reward from stopping on the second period (2) is greater than the conditional expected reward of not stopping on the second period ($\frac{5}{3}$), the dynamic programming principle says that the exercise policy in the second period should be chosen to satisfy $\mu_2^{\text{DP}}(2) = \text{STOP}$. Finally, unfolding to the first period, we conclude that the exercise policy obtained from dynamic programming is $\mu_1^{\text{DP}}(x_1) = \text{STOP}$ if and only if $x_1 = 3$. All together, the Markovian stopping rule for the non-Markovian optimal stopping problem that is obtained using backwards recursion yields an expected reward of $\frac{2}{3} \times 3 + \frac{1}{3} \times 2 = 2.6\bar{6}$. \square

Appendix B: Comparison of Robust Optimization Formulations

As described in §2.2, our formulation of the robust optimization problem (RO) deviates from that of (RO'), which followed from BSS22. In this section, we show that all of the main results from §2.4 and our characterization of optimal Markovian stopping rules in §3 also hold for formulation (RO') with only minor modifications to the proofs. We will then discuss the advantages of using formulation (RO) in the context of our algorithmic techniques that are developed in §4.

B.1. The Relationship Between (RO) and (RO')

Before proceeding further, let us develop intuition for the relationship between the two formulations (RO) and (RO') by comparing them in a simple example. Speaking informally, the following Example EC.2 shows that if the reward functions $g(1, \cdot), \dots, g(T, \cdot)$ are well behaved functions of stochastic process, and if the radius of the uncertainty sets is small, then the two robust optimization formulations are essentially equivalent. As we will see afterwards, this intuition will extend to general classes of reward functions.

EXAMPLE EC.2. Consider an optimal stopping problem in which the state space is one-dimensional and the reward functions satisfy $g(t, x) = x_t$ for all periods t . In this case, we observe that the following equalities hold:

$$\begin{aligned}
 (\text{RO}') &= \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) \\
 &= \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} \sum_{t=1}^T y_t \mathbb{I}\{\tau_{\mu}(y) = t\} \\
 &= \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} \sum_{t=1}^T ((y_t - x_t^i) + x_t^i) \mathbb{I}\{\tau_{\mu}(y) = t\}. \tag{EC.1}
 \end{aligned}$$

Indeed, the first equality follows from the definition of (RO'), the second equality follows from the fact that $g(t, x) = x_t$ for all periods t , and the third equality follows from algebra.

For notational convenience, let $\widehat{J}_{N,\epsilon}(\mu)$ denote the objective value of the robust optimization formulation (RO) corresponding to exercise policies μ . Moreover, we recall from the definition of the uncertainty sets in §2.2 that the inequality $|y_t - x_t^i| \leq \epsilon$ is satisfied for all sample paths $i \in \{1, \dots, N\}$, periods $t \in \{1, \dots, T\}$, and $y \equiv (y_1, \dots, y_T) \in \mathcal{U}^i$. Therefore, it follows from line (EC.1) and algebra that

$$\begin{aligned}
 (\text{RO}') &\leq \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} \sum_{t=1}^T (\epsilon + x_t^i) \mathbb{I}\{\tau_{\mu}(y) = t\} = T\epsilon + \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) = T\epsilon + (\text{RO}), \text{ and} \\
 (\text{RO}') &\geq \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} \sum_{t=1}^T (-\epsilon + x_t^i) \mathbb{I}\{\tau_{\mu}(y) = t\} = -T\epsilon + \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) = -T\epsilon + (\text{RO}).
 \end{aligned}$$

We have thus shown that if the robustness parameter $\epsilon \geq 0$ is small, then the two formulations will be close to one another with respect to optimal objective value. Moreover, it follows by identical reasoning that for all exercise policies μ ,

$$\left| \widehat{J}_{N,\epsilon}(\mu) - \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) \right| \leq T\epsilon.$$

Therefore, we conclude that if the robustness parameter $\epsilon \geq 0$ is small, then the objective values of the two formulations will be close to one another, uniformly over the space of all exercise policies. \square

The above example is insightful because it reveals, at least for problems with linear reward functions, that the robust optimization formulations (RO) and (RO') become essentially equivalent in the asymptotic regime in which the robustness parameter converges to zero. We will see shortly that the same intuition from Example EC.2 will extend to a broad class of reward functions.

We emphasize that the above intuition, of course, does *not imply* that the two formulations (RO) and (RO') are guaranteed to have identical performance on any fixed collection of simulated sample paths. Indeed, we do not preclude the possibility that one of the two formulations (RO) and (RO') may have better finite-sample performance than the other formulation in finding Markovian stopping rules which perform well with respect to (OPT). Fortunately, in the particular setting of the present paper in which the joint probability distribution of the underlying stochastic problem is *known*, any potential differences in finite-sample performance between the two formulations can be decreased arbitrarily by simulating larger number of sample paths when constructing the robust optimization problem and then choosing a smaller robustness parameter.

B.2. Optimality Guarantees for (RO')

In view of the above intuition, we proceed to prove that the convergence guarantees from §2.4 will also hold if we opted instead to use formulation (RO'). Our convergence guarantees in the following Theorem EC.1 will be developed by extending the intuition from Example EC.2, that is, by showing that the gap between the objective values for formulations (RO') and (RO) converges to zero, almost surely, uniformly over the space of all exercise policies. We establish these convergence guarantees for formulation (RO') when the reward function in the optimal stopping problem satisfies the following assumption:

ASSUMPTION EC.1. $\lim_{\epsilon \rightarrow 0} \Omega_\epsilon(x) = 0$ almost surely, where

$$\Omega_\epsilon(x) \triangleq \max_{t \in \{1, \dots, T\}} \left\{ \sup_{y \in \mathcal{X}^T: \|y-x\|_\infty \leq \epsilon} g(t, y) - g(t, x) \right\}.$$

We readily observe that above assumption, in conjunction with Assumption 1, is equivalent to requiring that the reward functions $g(1, \cdot), \dots, g(T, \cdot)$ are continuous functions of the stochastic process almost surely. Hence, Assumption EC.1 can be viewed as a mild assumption that is frequently satisfied in the applications of optimal stopping to options pricing. In the following theorem, as in §2.4, we let $\widehat{J}_{N, \epsilon}(\mu)$ denote the objective value of the robust optimization problem (RO) corresponding to exercise policies μ .

THEOREM EC.1 (Uniform convergence for robust optimization formulations). *If Assumptions 4 and EC.1 hold, then the gap in objective values between formulations (RO) and (RO') converges to zero, almost surely, uniformly over the space of all exercise policies:*

$$\limsup_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \left| \widehat{J}_{N, \epsilon}(\mu) - \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y) \right| = 0 \quad \text{almost surely.}$$

Proof. Consider any arbitrary sample path $x^i = (x_1^i, \dots, x_T^i)$ and robustness parameter $\epsilon \geq 0$. We first observe from algebra that

$$\begin{aligned} \left| \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y) - \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i) \right| &\leq \sup_{y \in \mathcal{U}^i} |g(\tau_\mu(y), y) - g(\tau_\mu(y), x^i)| \\ &\leq \max_{t \in \{1, \dots, T\}} \sup_{y \in \mathcal{U}^i} |g(t, y) - g(t, x^i)| \\ &\leq \max_{t \in \{1, \dots, T\}} \left\{ \sup_{y \in \mathcal{X}^T: \|y-x^i\|_\infty \leq \epsilon} g(t, y) - g(t, x^i) \right\} \\ &= \Omega_\epsilon(x^i), \end{aligned} \tag{EC.2}$$

where the last inequality follows from the definition of the uncertainty sets. Therefore,

$$\begin{aligned}
& \limsup_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \left| \widehat{J}_{N,\epsilon}(\mu) - \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) \right| \\
& \leq \limsup_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \left| \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) - \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) \right| \\
& \leq \limsup_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \Omega_{\epsilon}(x^i) \\
& = \limsup_{\epsilon \rightarrow 0} \mathbb{E}[\Omega_{\epsilon}(x)] \quad \text{almost surely} \\
& = 0.
\end{aligned}$$

where the first inequality follows from triangle inequality, the second inequality follows from (EC.2), the first equality follows from the strong law of large numbers, and the second equality follows from the dominated convergence theorem and Assumption EC.1. We note that the strong law of large numbers and dominated convergence theorem can be applied in both cases because of the boundedness of the reward function (Assumption 4). This concludes the proof of Theorem EC.1. \square

Using standard proof techniques from stochastic programming (see, for example, Shapiro et al. (2014, §5.1.1)), Theorem EC.1 readily implies that the alternative formulation (RO') enjoys identical convergence guarantees as those in Theorems 1-3 under Assumptions 1-4 and EC.1.

B.3. Characterization of Optimal Markovian Stopping Rules for (RO')

Next, we show that our characterization of the structure of optimal Markovian stopping rules for formulation (RO) from §3 can be readily extended to the alternative formulation (RO') using similar proof techniques to those used in Theorem 4. Our following analysis in Theorem EC.2 uses the following additional but relatively mild assumption on the reward functions in the optimal stopping problem:

ASSUMPTION EC.2. *For each period t , the reward function satisfies $g(t, x) = h(t, x_t)$.*

The above assumption says that the reward function depends in each period only on the current state. This is a common assumption in the optimal stopping and options pricing literature and is often without loss of generality. With this assumption, the following theorem establishes the structure of optimal Markovian stopping rules for formulation (RO').

THEOREM EC.2. *Under Assumption EC.2, there exists $\mu \in \mathcal{M}$ that is optimal for (RO').*

Our proof of Theorem EC.2 follows the same pruning technique that was used in the proof of Theorem 4 in §3.2. Specifically, our proof of Theorem EC.2 makes use of the following four lemmas, which are essentially restatements of Lemmas 1-4 from §3.2.

LEMMA EC.1. *The optimal objective value of (RO') is equal to the optimal objective value of*

$$\begin{aligned}
& \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) \\
& \text{subject to} \quad \text{for each } i \in \{1, \dots, N\}, \text{ there exists } t \in \{1, \dots, T\} \\
& \quad \text{such that } \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i
\end{aligned} \tag{RO'_T}$$

Proof. The proof of Lemma EC.1 is identical to the proof of Lemma 1. \square

LEMMA EC.2. *Let Assumption EC.2 hold, consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO_T') , and define*

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

Then the following equality holds for each $i \in \{1, \dots, N\}$:

$$\inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i) = \min_{t \in \{1, \dots, \sigma^i\}} \inf_{y_t \in \mathcal{U}_t^i} \{h(t, y_t) : \mu_t(y_t) = \text{STOP}\}.$$

Proof. Let Assumption EC.2 hold, consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO_T') , and define

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

It follows from identical reasoning as in the proof of Lemma 2 that the following equality holds for each sample path $i \in \{1, \dots, N\}$:

$$\{\tau_\mu(y) : y \in \mathcal{U}^i\} = \{t \in \{1, \dots, \sigma^i\} : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t(y_t) = \text{STOP}\}. \quad (\text{EC.3})$$

We thus conclude for each sample path $i \in \{1, \dots, N\}$ that

$$\inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y) = \min_{t \in \{1, \dots, \sigma^i\}} \inf_{y_t \in \mathcal{U}_t^i} \{h(t, y_t) : \mu_t(y_t) = \text{STOP}\},$$

where the equality follows from Assumption EC.2 and line (EC.3). This completes our proof of Lemma EC.2. \square

LEMMA EC.3. *If Assumption EC.2 holds and if μ' is a pruned version of μ , then $\inf_{y \in \mathcal{U}^i} g(\tau_{\mu'}(y), y) \geq \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y)$ for all $i \in \{1, \dots, N\}$.*

REMARK EC.1. We observe that (RO_T') and (RO_T) have identical constraints. Thus, we use the same definition of pruning in Appendix B.3 as given by Definition 1.

Proof of Lemma EC.3. Let Assumption EC.2 hold, and let μ' be a pruned version of μ . Let $\sigma^1, \dots, \sigma^N$ satisfy the following equalities for each $i \in \{1, \dots, N\}$:

$$\sigma^i = \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\} = \min \{t \in \{1, \dots, T\} : \mu'_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\}. \quad (\text{EC.4})$$

Then it follows from the fact that μ is feasible for (RO_T') that $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. Therefore, for each $i \in \{1, \dots, N\}$,

$$\begin{aligned} \inf_{y \in \mathcal{U}^i} g(\tau_{\mu'}(y), y) &= \min_{t \in \{1, \dots, \sigma^i\}} \inf_{y_t \in \mathcal{U}_t^i} \{h(t, y_t) : \mu'_t(y_t) = \text{STOP}\} \\ &\geq \min_{t \in \{1, \dots, \sigma^i\}} \inf_{y_t \in \mathcal{U}_t^i} \{h(t, y_t) : \mu_t(y_t) = \text{STOP}\} \\ &= \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y). \end{aligned}$$

Indeed, the two equalities follow from Lemma EC.2 and line (EC.4). The inequality follows from the fact that μ' is a pruned version of μ , which implies that $\{y_t \in \mathcal{X} : \mu'_t(y_t) = \text{STOP}\} \subseteq \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}$ for all $t \in \{1, \dots, T\}$. Our proof of Lemma EC.3 is thus complete. \square

LEMMA EC.4. Consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO_T') , and define

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

Then $\mu^{\sigma^1 \dots \sigma^N}$ is a pruned version of μ .

Proof. The proof of Lemma EC.4 is identical to the proof of Lemma 4. \square

In view of the above Lemmas EC.1-EC.4, we now present the proof of Theorem EC.2.

Proof of Theorem EC.2. Lemma EC.1 shows that the robust optimization problem (RO') is equivalent to the robust optimization problem (RO_T') . Moreover, for any arbitrary exercise policy μ that is feasible for (RO_T') , Lemmas EC.3 and EC.4 together show that there exists an exercise policy $\mu' \in \mathcal{M}$ such that the objective value associated with μ is less than or equal to the objective value associated with μ' . Since μ was chosen arbitrarily, our proof of Theorem EC.2 is complete. \square

B.4. Computational Tractability of (RO')

Finally, we now discuss our primary motivation for using formulation (RO) instead of formulation (RO') . As we have shown up to this point, these two robust optimization formulations are essentially equivalent with respect to convergence guarantees and characterization of optimal Markovian stopping rules. However, as we will show momentarily, we find that formulation (RO) is significantly more amenable than (RO') from an algorithmic perspective. Our subsequent discussion on the computational tractability of formulation (RO') makes use of the following Theorem EC.3, which, analogously to Theorem 5 from §3.3, provides a reformulation of (RO') as a finite-dimensional optimization problem over integer decision variables.

THEOREM EC.3. Under Assumption EC.2, the alternative formulation (RO') is equivalent to

$$\underset{\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}}{\text{maximize}} \quad \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \min_{j: \sigma^j = t} \inf_{y_t \in \mathcal{U}_t^i \cap \mathcal{U}_t^j} h(t, y_t). \quad (\text{IP}')$$

Proof. Our proof of Theorem EC.3 is essentially identical to the proof of Theorem 5. Indeed, let Assumption EC.2 hold. With this assumption, our proof of Theorem EC.3 is split into the following two intermediary claims:

CLAIM EC.1. The optimal objective value of (RO') is less than or equal to the optimal objective value of (IP') .

Proof of Claim EC.1. Consider any arbitrary exercise policy $\mu \in \mathcal{M}$, and define the integers

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP } \forall y_t \in \mathcal{U}_t^i\}, \quad \forall i \in \{1, \dots, N\}. \quad (\text{EC.5})$$

It follows from the definition of \mathcal{M} that μ satisfies the constraints of (RO_T') , and so Lemma EC.4 implies that $\mu^{\sigma^1 \dots \sigma^N}$ is a pruned version of μ . Therefore, it follows from Definition 1 and line (EC.5) that the following equalities hold:

$$\sigma^i = \min \{t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}. \quad (\text{EC.6})$$

Therefore, we observe that

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), y) \\
& \leq \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g\left(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), y\right) \\
& = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \inf_{y_t \in \mathcal{U}_t^i} \left\{ h(t, y_t) : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP} \right\} \\
& = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \min_{j: \sigma^j = t} \inf_{y_t \in \mathcal{U}_t^i \cap \mathcal{U}_t^j} h(t, y_t), \tag{EC.7}
\end{aligned}$$

where the first inequality follows from Lemma EC.3, the first equality follows from Lemma EC.2 and line (EC.6), and the second equality follows from the fact that $\mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}$ if and only if there exists a sample path j such that $y_t \in \mathcal{U}_t^j$ and $\sigma^j = t$. Because $\mu \in \mathcal{M}$ was chosen arbitrarily, we conclude from line (EC.7) and Theorem EC.2 that the optimal objective value of (RO') is less than or equal to the optimal objective value of (IP'). Our proof of Claim EC.1 is thus complete. \square

CLAIM EC.2. *The optimal objective value of (RO') is greater than or equal to the optimal objective value of (IP'). Furthermore, for any choice of integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$, the corresponding exercise policy $\mu^{\sigma^1 \dots \sigma^N} \equiv (\mu_1^{\sigma^1 \dots \sigma^N}, \dots, \mu_T^{\sigma^1 \dots \sigma^N})$ satisfies*

$$\frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g\left(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), y\right) \geq \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \min_{j: \sigma^j = t} \inf_{y_t \in \mathcal{U}_t^i \cap \mathcal{U}_t^j} h(t, y_t).$$

Proof of Claim EC.2. Consider any integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. For each sample path $i \in \{1, \dots, N\}$, we observe that

$$\begin{aligned}
& \min \left\{ t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i \right\} \\
& = \min \left\{ t \in \{1, \dots, \sigma^i\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i \right\} \\
& \leq \sigma^i, \tag{EC.8}
\end{aligned}$$

where the equality follows from the fact that $\mu^{\sigma^1 \dots \sigma^N} \equiv (\mu_1^{\sigma^1 \dots \sigma^N}, \dots, \mu_T^{\sigma^1 \dots \sigma^N})$ by construction satisfies $\mu_{\sigma^i}^{\sigma^1 \dots \sigma^N}(y_{\sigma^i}) = \text{STOP}$ for all $y_{\sigma^i} \in \mathcal{U}_{\sigma^i}^i$, and the inequality follows from algebra. Therefore,

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g\left(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), y\right) \\
& \geq \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \inf_{y_t \in \mathcal{U}_t^i} \left\{ h(t, y_t) : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP} \right\} \\
& = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \min_{j: \sigma^j = t} \inf_{y_t \in \mathcal{U}_t^i \cap \mathcal{U}_t^j} h(t, y_t), \tag{EC.9}
\end{aligned}$$

where the inequality follows from Lemma EC.2 and line (EC.8), and the equality follows from the fact that $\mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}$ if and only if there exists a sample path j such that $y_t \in \mathcal{U}_t^j$ and $\sigma^j = t$. Because $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ were chosen arbitrarily, we conclude from line (EC.9) that the optimal objective value of (RO') is greater than or equal to the optimal objective value of (IP'), which concludes our proof of Claim EC.2. \square

Combining Claims EC.1 and EC.2, our proof of Theorem EC.3 is thus complete.

□

Equipped with the above Theorem EC.3, we now explain why formulation (RO) is more amenable from an tractability perspective than (RO'). Indeed, we recall from Theorem 5 in §3.3 that the robust optimization problem (RO) is equivalent to

$$\underset{\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}}{\text{maximize}} \quad \frac{1}{N} \sum_{i=1}^N \underbrace{\min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\}}_{\nu_i(\sigma)}, \quad (\text{IP})$$

where we observe for each sample path $i \in \{1, \dots, N\}$ that the quantity $\nu_i(\sigma)$ satisfies

$$\nu_i(\sigma) \in \bigcup_{t=1}^T \{g(t, x^i)\}.$$

In contrast, we showed in the above Theorem EC.3 that (RO') is equivalent to

$$\underset{\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}}{\text{maximize}} \quad \frac{1}{N} \sum_{i=1}^N \underbrace{\min_{t \in \{1, \dots, \sigma^i\}} \min_{j: \sigma^j = t} \inf_{y_t \in \mathcal{U}_t^i \cap \mathcal{U}_t^j} h(t, y_t)}_{\nu'_i(\sigma)}, \quad (\text{IP}')$$

where we observe for each sample path $i \in \{1, \dots, N\}$ that the quantity $\nu'_i(\sigma)$ satisfies

$$\nu'_i(\sigma) \in \bigcup_{t=1}^T \bigcup_{j: \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset} \left\{ \inf_{y_t \in \mathcal{U}_t^i \cap \mathcal{U}_t^j} h(t, y_t) \right\}.$$

Hence, for each sample path $i \in \{1, \dots, N\}$, we observe that the number of possible values for $\nu'_i(\sigma)$ can be a factor of $\mathcal{O}(N)$ greater than the number of possible values for $\nu_i(\sigma)$. This demonstrates that (RO) is more amenable to compact reformulations than (RO') and concludes our motivation for using formulation (RO) instead of formulation (RO') throughout the paper.

Appendix C: Proofs of Theorems 1, 2, and 3

Establishing the convergence guarantees from §2.4 for the robust optimization problem (RO) can be organized into two high-level steps. The first high-level step, comprised of Theorem 3, consists of showing that the objective function of the robust optimization problem (RO) converges almost surely to a (conservative) lower bound approximation of the objective function of the stochastic optimal stopping problem (OPT), uniformly over the space of all exercise policies. The second high-level step, comprised of Theorems 1 and 2, consists of showing that the conservativeness of the robust optimization problem (RO) disappears almost surely as the robustness parameter tends to zero and the number of sample paths tends to infinity. We present the first high-level step in Appendix C.1, and the second high-level step can be found in Appendix C.2.

Let us reflect on the novelty of the results in the present Appendix C. The first high-level step in Appendix C.1 (i.e., the proof of Theorem 3) is not novel and follows immediately from a uniform convergence result established by BSS22. Rather, the significant novelty of Appendix C is found in the second high-level step in Appendix C.2 (i.e., the proofs of Theorems 1 and 2). Stated succinctly, the second high-level step is challenging to establish because it requires showing

that (a) the limit of the optimal objective value of the robust optimization problem (RO) exists almost surely as $\epsilon \rightarrow \infty$ and $N \rightarrow \infty$, and showing that (b) there always exists an arbitrarily near-optimal Markovian stopping rule μ for the stochastic optimal stopping problem (OPT) that almost surely satisfies $\tau_\mu(y) = \tau_\mu(x)$ for all sufficiently close realizations of $y \equiv (y_1, \dots, y_T)$ to the stochastic process $x \equiv (x_1, \dots, x_T)$.

For the general classes of stochastic dynamic optimization problems considered by BSS22 and Sturt (2020), the authors were unable to identify simple and verifiable conditions under which the aforementioned properties (a) and (b) are guaranteed to hold. Consequently, the authors only presented upper bounds on the gap between the optimal objective value of the robust optimization problem and the optimal objective value of the stochastic dynamic optimization (BSS22, Theorem 1). In certain cases, the authors alternatively assumed that the stochastic dynamic optimization problem happened to have arbitrarily near-optimal control policies with convenient structure (Sturt 2020, Assumption 15). The authors of BSS22 reflect on the weaknesses of these results in the literature at the end of their §4.3, saying that “future work may identify subclasses of [stochastic dynamic optimization problems] where the equality of the bounds can be ensured.”

In Appendix C.2 of this paper, we resolve this gap in the literature by showing that the aforementioned properties (a) and (b) are guaranteed to hold for the specific class of stochastic dynamic optimization problems (OPT) under the mild assumptions that the stochastic problem has a bounded objective function (Assumption 4) and that the random variables in the stochastic problem have a continuous joint probability distribution (Assumption 3). Specifically, we establish in Lemma EC.6 of Appendix C.2 that property (a) holds by combining Assumption 4 with McDiarmid’s inequality, and we establish in Lemma EC.7 of Appendix C.2 that property (b) holds by combining Assumptions 3 and 4 with elementary techniques from topology and measure theory. These lemmas allow us to establish the proofs of Theorems 1 and 2 at the end of Appendix C.2.

C.1. Proof of Theorem 3

In Appendix C.1, we establish the proof of Theorem 3 by using a uniform convergence result from BSS22. We begin by presenting some preliminary notation. Recall that $x \equiv (x_1, \dots, x_T)$, $x^1 \equiv (x_1^1, \dots, x_T^1)$, $x^2 \equiv (x_1^2, \dots, x_T^2), \dots \in \mathcal{X}^T \equiv \mathbb{R}^{Td}$ are sample paths drawn independently from an identical joint probability distribution. We will make use of the following additional notation:

$$\begin{aligned} \delta_N &\triangleq N^{-\frac{1}{\max\{3, Td+1\}}}; & \mathcal{U}^i(\delta_N) &\triangleq \{y \in \mathcal{X}^T : \|y - x^i\|_\infty \leq \delta_N\}; \\ \mathcal{U}^i &\triangleq \{y \in \mathcal{X}^T : \|y - x^i\|_\infty \leq \epsilon\}; & M_N &\triangleq N^{-\frac{1}{(Td+1)(Td+2)}} \log N. \end{aligned}$$

We now state the uniform convergence result from BSS22, which has been adapted to the notation of the present paper.

LEMMA EC.5 (Theorem 2 of BSS22). *Let Assumption 2 hold. Then there exists a finite $\bar{N} \in \mathbb{N}$, almost surely, such that the following inequality holds for all $N \geq \bar{N}$ and all measurable functions $f : \mathbb{R}^{Td} \rightarrow \mathbb{R}$:*

$$\mathbb{E} [f(x) \mathbb{I} \{x \in \cup_{i=1}^N \mathcal{U}^i(\delta_N)\}] \geq \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i(\delta_N)} f(y) - M_N \sup_{y \in \cup_{i=1}^N \mathcal{U}^i(\delta_N)} |f(y)|.$$

In view of the above notation and lemma, we are now ready to present the proof of Theorem 3.

Proof of Theorem 3. Consider any arbitrary choice of the robustness parameter $\epsilon > 0$, and recall that the reward function satisfies $g(\infty, y) = 0$ (see §2) and $0 \leq g(1, y), \dots, g(T, y) \leq U$ for all trajectories $y \in \mathcal{X}^T$ (Assumption 4). With this notation, we observe that

$$\begin{aligned} & \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ J^*(\mu) - \widehat{J}_{N,\epsilon}(\mu) \right\} \\ &= \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \mathbb{E}[g(\tau_{\mu}(x), x)] - \widehat{J}_{N,\epsilon}(\mu) \right\} \end{aligned} \quad (\text{EC.10})$$

$$\geq \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \mathbb{E} \left[g(\tau_{\mu}(x), x) \mathbb{I} \{ x \in \cup_{i=1}^N \mathcal{U}^i(\delta_N) \} \right] - \widehat{J}_{N,\epsilon}(\mu) \right\} \quad (\text{EC.11})$$

$$\geq \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i(\delta_N)} g(\tau_{\mu}(y), y) - M_N U - \widehat{J}_{N,\epsilon}(\mu) \right\} \text{ almost surely} \quad (\text{EC.12})$$

$$= \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i(\delta_N)} g(\tau_{\mu}(y), y) - \widehat{J}_{N,\epsilon}(\mu) \right\} \quad (\text{EC.13})$$

$$\geq \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) - \widehat{J}_{N,\epsilon}(\mu) \right\}. \quad (\text{EC.14})$$

(EC.10) follows from applying the definition of $J^*(\mu)$; (EC.11) holds because the reward function is nonnegative; (EC.12) follows from Lemma EC.5 and the boundedness of the reward function; (EC.13) holds because $M_N \rightarrow 0$; (EC.14) holds because $\delta_N \rightarrow 0$ implies, for any arbitrary $\epsilon > 0$, that the inequality $\inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), y) \leq \inf_{y \in \mathcal{U}^i(\delta_N)} g(\tau_{\mu}(y), y)$ is satisfied for all $i \in \mathbb{N}$, for all μ , and for all large $N \in \mathbb{N}$. Moreover:

$$(\text{EC.14}) = \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} \{ g(\tau_{\mu}(y), x^i) + g(\tau_{\mu}(y), y) - g(\tau_{\mu}(y), x^i) \} - \widehat{J}_{N,\epsilon}(\mu) \right\} \quad (\text{EC.15})$$

$$\geq \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \left(\inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) + \inf_{y \in \mathcal{U}^i} \{ g(\tau_{\mu}(y), y) - g(\tau_{\mu}(y), x^i) \} \right) - \widehat{J}_{N,\epsilon}(\mu) \right\} \quad (\text{EC.16})$$

$$\geq \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \left(\inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) + \min_{t \in \{1, \dots, T\}} \left\{ \inf_{y \in \mathcal{U}^i} g(t, y) - g(t, x^i) \right\} \right) - \widehat{J}_{N,\epsilon}(\mu) \right\} \quad (\text{EC.17})$$

$$\geq \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \frac{1}{N} \sum_{i=1}^N \left(\inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) + \Delta_{\epsilon}(x^i) \right) - \widehat{J}_{N,\epsilon}(\mu) \right\} \quad (\text{EC.18})$$

$$= \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ \widehat{J}_{N,\epsilon}(\mu) - \widehat{J}_{N,\epsilon}(\mu) \right\} + \liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \Delta_{\epsilon}(x^i) \quad (\text{EC.19})$$

$$= \mathbb{E}[\Delta_{\epsilon}(x)] \text{ almost surely.} \quad (\text{EC.20})$$

(EC.15), (EC.16), and (EC.17) follow from algebra; (EC.18) follows from the definition of $\Delta_{\epsilon}(x^i)$ (see Assumption 1); (EC.19) follows from the definition of $\widehat{J}_{N,\epsilon}(\mu)$; (EC.20) follows from the strong law of large numbers.

Since $\epsilon > 0$ was chosen arbitrarily, we have shown that

$$\lim_{\epsilon \rightarrow 0} \liminf_{N \rightarrow \infty} \inf_{\mu} \left\{ J^*(\mu) - \widehat{J}_{N,\epsilon}(\mu) \right\} \geq \lim_{\epsilon \rightarrow 0} \mathbb{E}[\Delta_{\epsilon}(x)] \geq 0 \text{ almost surely,}$$

where the first inequality holds almost surely from lines (EC.10)-(EC.20), and the second inequality holds almost surely due to the dominated convergence theorem and Assumption 1. Note that the

above limits exist because $\epsilon \mapsto \liminf_{N \rightarrow \infty} \inf_{\mu} \{J^*(\mu) - \widehat{J}_{N,\epsilon}(\mu)\}$ and $\epsilon \mapsto \mathbb{E}[\Delta_{\epsilon}(x)]$ are monotonic functions. This concludes the proof of Theorem 3. \square

C.2. Proofs of Theorems 1 and 2

We begin Appendix C.2 by presenting the two novel intermediary results, Lemmas EC.6 and EC.7, that were discussed at the beginning of Appendix C. In the first novel intermediary result, denoted below by Lemma EC.6, we prove that the limit of the optimal objective value of the robust optimization problem (RO) exists almost surely as $\epsilon \rightarrow \infty$ and $N \rightarrow \infty$. The proof of the following lemma is based on McDiarmid's inequality.

LEMMA EC.6. *Let Assumption 4 hold. Then for all $\epsilon > 0$,*

$$\liminf_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) = \limsup_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) \quad \text{almost surely.}$$

Proof. Consider any fixed $\epsilon > 0$, and, for notational convenience, define the following function:

$$h_{\epsilon}(x^1, \dots, x^N) \triangleq \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu).$$

We will utilize the following intermediary claim:

CLAIM EC.3. *The function $h_{\epsilon} : \mathcal{X}^T \times \dots \times \mathcal{X}^T \rightarrow \mathbb{R}$ has the following 'bounded differences' property: for all $\check{x}^1, \dots, \check{x}^N \in \mathcal{X}^T$ and $\bar{x}^1, \dots, \bar{x}^N \in \mathcal{X}^T$ that differ only on the j th coordinate ($\check{x}^i = \bar{x}^i$ for $i \neq j$),*

$$|h_{\epsilon}(\bar{x}^1, \dots, \bar{x}^N) - h_{\epsilon}(\check{x}^1, \dots, \check{x}^N)| \leq \frac{U}{N}.$$

Proof of Claim EC.3. For any arbitrary $\eta > 0$, let the exercise policies $\check{\mu}^{\eta}$ be chosen to satisfy

$$\frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{X}^T: \|y - \check{x}^i\|_{\infty} \leq \epsilon} g(\tau_{\check{\mu}^{\eta}}(y), \check{x}^i) \geq h_{\epsilon}(\check{x}^1, \dots, \check{x}^N) - \eta. \quad (\text{EC.21})$$

We observe that

$$\begin{aligned} & h_{\epsilon}(\check{x}^1, \dots, \check{x}^N) - h_{\epsilon}(\bar{x}^1, \dots, \bar{x}^N) \\ & \leq \left(\frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{X}^T: \|y - \check{x}^i\|_{\infty} \leq \epsilon} g(\tau_{\check{\mu}^{\eta}}(y), \check{x}^i) + \eta \right) - \left(\frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{X}^T: \|y - \bar{x}^i\|_{\infty} \leq \epsilon} g(\tau_{\check{\mu}^{\eta}}(y), \bar{x}^i) \right) \\ & = \frac{1}{N} \inf_{y \in \mathcal{X}^T: \|y - \check{x}^j\|_{\infty} \leq \epsilon} g(\tau_{\check{\mu}^{\eta}}(y), \check{x}^j) - \frac{1}{N} \inf_{y \in \mathcal{X}^T: \|y - \bar{x}^j\|_{\infty} \leq \epsilon} g(\tau_{\check{\mu}^{\eta}}(y), \bar{x}^j) + \eta \\ & \leq \frac{U}{N} + \eta. \end{aligned}$$

Indeed, the first inequality holds because of line (EC.21) and because $\check{\mu}^{\eta}$ is a feasible but possibly suboptimal solution to the optimization problem $\sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{X}^T: \|y - \bar{x}^i\|_{\infty} \leq \epsilon} g(\tau_{\mu}(y), \bar{x}^i) \equiv h_{\epsilon}(\bar{x}^1, \dots, \bar{x}^N)$, and the second inequality follows from Assumption 4. Because $\eta > 0$ was chosen arbitrarily, we have shown that

$$h_{\epsilon}(\check{x}^1, \dots, \check{x}^N) - h_{\epsilon}(\bar{x}^1, \dots, \bar{x}^N) \leq \frac{U}{N}.$$

It follows from symmetry that

$$h_{\epsilon}(\bar{x}^1, \dots, \bar{x}^N) - h_{\epsilon}(\check{x}^1, \dots, \check{x}^N) \leq \frac{U}{N},$$

which concludes our proof of Claim EC.3. \square

Because the above Claim EC.3 holds, it follows from McDiarmid's inequality that

$$\mathbb{P}(|h_\epsilon(x^1, \dots, x^N) - \mathbb{E}[h_\epsilon(x^1, \dots, x^N)]| > \eta) \leq 2\exp\left(-\frac{2\eta^2 N}{U^2}\right) \quad \forall \eta > 0.$$

It follows from the above line that

$$\sum_{N=1}^{\infty} \mathbb{P}(|h_\epsilon(x^1, \dots, x^N) - \mathbb{E}[h_\epsilon(x^1, \dots, x^N)]| > \eta) < \infty \quad \forall \eta > 0,$$

and so the Borel-Cantelli lemma implies that

$$\lim_{N \rightarrow \infty} |h_\epsilon(x^1, \dots, x^N) - \mathbb{E}[h_\epsilon(x^1, \dots, x^N)]| = 0 \quad \text{almost surely.} \quad (\text{EC.22})$$

We observe from identical reasoning as in the proof of Shapiro et al. (2014, Proposition 5.6) that $\mathbb{E}[h_\epsilon(x^1, \dots, x^N)]$ is monotonically decreasing with respect to $N \in \mathbb{N}$. Since the random variables $h_\epsilon(x^1, \dots, x^N)$ are also contained in the interval $[0, U]$ for all $N \in \mathbb{N}$, we conclude that $\lim_{N \rightarrow \infty} \mathbb{E}[h_\epsilon(x^1, \dots, x^N)]$ exists, and thus it follows from line (EC.22) that $\lim_{N \rightarrow \infty} h_\epsilon(x^1, \dots, x^N)$ exists almost surely. This concludes the proof of Lemma EC.6. \square

In the second novel intermediary result, denoted below by Lemma EC.7, we show that there always exists an arbitrarily near-optimal Markovian stopping rule μ for the stochastic optimal stopping problem (OPT) that almost surely satisfies $\tau_\mu(y) = \tau_\mu(x)$ for all sufficiently close realizations of $y \equiv (y_1, \dots, y_T)$ to the stochastic process $x \equiv (x_1, \dots, x_T)$. The proof of the following lemma is based on elementary techniques from topology and measure theory.

LEMMA EC.7. *Under Assumptions 3 and 4,*

$$\sup_{\mu} \mathbb{E}[g(\tau_\mu(x), x)] = \sup_{\mu} \mathbb{E}\left[\liminf_{y \rightarrow x} g(\tau_\mu(y), x)\right].$$

Proof. Our proof of Lemma EC.7 is organized into the following steps. In our first step, we will show that any feasible exercise policy to the stochastic optimal stopping problem (OPT) can be approximated to arbitrary accuracy by an exercise policy with exercise regions that are open sets. In our second step, we will show that these exercise policies from the previous step can be approximated to arbitrary accuracy by an exercise policy with exercise regions that are the union of finitely many open balls. Finally, we will then invoke Assumption 3 to conclude that the probability of a random state lying a strictly positive distance from the boundary of finitely many open balls is equal to one, which implies that the corresponding Markovian stopping rule almost surely satisfies $\tau_\mu(y) = \tau_\mu(x)$ for all sufficiently close realizations of $y \equiv (y_1, \dots, y_T)$ to the stochastic process $x \equiv (x_1, \dots, x_T)$. From this, the desired lemma will follow readily.

In view of the above organization, we now present our proof. Choose any arbitrary exercise policy $\bar{\mu} \equiv (\bar{\mu}_1, \dots, \bar{\mu}_T)$ which is feasible for the stochastic optimal stopping problem (OPT), and choose any arbitrary constant $\eta > 0$. In our first claim, we show that this exercise policy can be approximated by an exercise policy with exercise regions that are open sets.

CLAIM EC.4. *There exists an exercise policy $\tilde{\mu}^\eta \equiv (\tilde{\mu}_1^\eta, \dots, \tilde{\mu}_T^\eta)$ such that:*

- (a) $\mathbb{E}[g(\tau_{\tilde{\mu}^\eta}(x), x)] \geq \mathbb{E}[g(\tau_{\bar{\mu}}(x), x)] - \eta$;
- (b) For each period $t \in \{1, \dots, T\}$, $A_t^\eta \triangleq \{y_t \in \mathcal{X} : \tilde{\mu}^\eta(y_t) = \text{STOP}\}$ is an open set.

Proof of Claim EC.4. Because $\bar{\mu}_1, \dots, \bar{\mu}_T : \mathbb{R}^d \rightarrow \{\text{STOP}, \text{CONTINUE}\}$ are feasible for the optimization problem (OPT), it follows from §2.1 that $\bar{\mu}_1, \dots, \bar{\mu}_T : \mathbb{R}^d \rightarrow \{\text{STOP}, \text{CONTINUE}\}$ are measurable functions. Thus, it follows for each period $t \in \{1, \dots, T\}$ that the set $A_t \triangleq \{y_t \in \mathbb{R}^d : \mu_t(y_t) = \text{STOP}\}$ is a Borel set in \mathbb{R}^d . Now, for each period $t \in \{1, \dots, T\}$, let $\mathbb{P}_{x,t}(\cdot) \triangleq$

$\mathbb{P}(x_t \in \cdot)$ denote the marginal probability law of the stochastic process on period t . Since A_t is a Borel set, it is a well known result from measure theory (Rudin 1964, Remark 11.11(b)) that there exists an open set $A_t^\eta \subseteq \mathbb{R}^d$ which satisfies $A_t \subseteq A_t^\eta$ and $\mathbb{P}_{x,t}(A_t^\eta \setminus A_t) \leq \frac{\eta}{T}$. Using these sets, we define an exercise policy on each period t by

$$\tilde{\mu}_t^\eta(y_t) \triangleq \begin{cases} \text{STOP}, & \text{if } y_t \in A_t^\eta, \\ \text{CONTINUE}, & \text{otherwise.} \end{cases}$$

It follows from the above construction that the new exercise policy $\tilde{\mu}^\eta \equiv (\tilde{\mu}_1^\eta, \dots, \tilde{\mu}_T^\eta)$ satisfies property (b) of Claim EC.4. To show that property (a) of Claim EC.4 holds, let us first define the following set for notational convenience:

$$B \triangleq \{y \equiv (y_1, \dots, y_T) : y_t \notin A_t^\eta \setminus A_t \text{ for all } t \in \{1, \dots, T\}\}$$

We observe that B is a Borel set in \mathbb{R}^{Td} , since B is comprised of a finite number of complements and intersections of Borel sets. Then,

$$\begin{aligned} & \mathbb{E}[|g(\tau_{\tilde{\mu}^\eta}(x), x) - g(\tau_{\bar{\mu}}(x), x)|] \\ &= \mathbb{E}[|g(\tau_{\tilde{\mu}^\eta}(x), x) - g(\tau_{\bar{\mu}}(x), x)| \mathbb{I}\{x \notin B\}] + \mathbb{E}[|g(\tau_{\tilde{\mu}^\eta}(x), x) - g(\tau_{\bar{\mu}}(x), x)| \mathbb{I}\{x \in B\}] \\ &= \mathbb{E}[|g(\tau_{\tilde{\mu}^\eta}(x), x) - g(\tau_{\bar{\mu}}(x), x)| \mathbb{I}\{x \notin B\}] + 0 \\ &\leq U\mathbb{P}(x \notin B) \\ &\leq \sum_{t=1}^T \mathbb{P}(x_t \in A_t^\eta \setminus A_t) \\ &\leq \eta. \end{aligned}$$

The first equality follows from the law of total expectation. The second equality holds because $g(\tau_{\tilde{\mu}^\eta}(x), x) = g(\tau_{\bar{\mu}}(x), x)$ when $x \in B$. The first inequality follows from Assumption 4. The second inequality follows from the union bound and the definition of the set B . The third and final inequality holds because $\mathbb{P}(x_t \in A_t^\eta \setminus A_t) = \mathbb{P}_{x,t}(A_t^\eta \setminus A_t) \leq \frac{\eta}{T}$. This concludes our proof of Claim EC.4. \square

In the above Claim EC.4, we showed that we can construct a new exercise policy $\tilde{\mu}^\eta \equiv (\tilde{\mu}_1^\eta, \dots, \tilde{\mu}_T^\eta)$ which is close to the original exercise policy $\bar{\mu} \equiv (\bar{\mu}_1, \dots, \bar{\mu}_T)$ with respect to expected reward. However, the new exercise policy is comprised of exercise regions in each period which are open sets. We now use this open set property, along with Assumption 3, to show the following second claim.

CLAIM EC.5. *There exists an exercise policy $\check{\mu}^\eta \equiv (\check{\mu}_1^\eta, \dots, \check{\mu}_T^\eta)$ such that:*

- (a) $\mathbb{E}[g(\tau_{\check{\mu}^\eta}(x), x)] \geq \mathbb{E}[g(\tau_{\bar{\mu}^\eta}(x), x)] - \eta$;
- (b) $\lim_{y \rightarrow x} \tau_{\check{\mu}^\eta}(y) = \tau_{\bar{\mu}^\eta}(x)$ almost surely.

Proof of Claim EC.5. For each period $t \in \{1, \dots, T\}$, consider the exercise region $A_t^\eta \triangleq \{y_t \in \mathcal{X} : \tilde{\mu}^\eta(y_t) = \text{STOP}\}$ corresponding to the exercise policy $\tilde{\mu}_t^\eta$. Since A_t^η is an open set, it is a well-known result from measure theory (Rudin 1964, Remark 11.11(a)) that A_t^η is the union of a countable collection of open balls. That is, there exists a countable set of tuples $\{(y_t^{\eta, \ell}, \epsilon_t^{\eta, \ell})\}_{\ell \in \mathbb{N}}$ such that $\epsilon_t^{\eta, \ell} > 0$ for all $\ell \in \mathbb{N}$ and the following equality is satisfied:

$$A_t^\eta = \lim_{k \rightarrow \infty} \underbrace{\bigcup_{\ell=1}^k \{y_t \in \mathbb{R}^d : \|y_t - y_t^{\eta, \ell}\|_\infty \leq \epsilon_t^{\eta, \ell}\}}_{C_t^{\eta, k}}.$$

We observe from the above construction that $C_t^{\eta,k}$, $k \in \mathbb{N}$, is an increasing sequence of sets which converges to A_t^η . Consequently, we have $\lim_{k \rightarrow \infty} \mathbb{P}(x_t \in A_t^\eta \setminus C_t^{\eta,k}) = 0$, and so it follows from the definition of a limit that there exists a $k_t^\eta \in \mathbb{N}$ such that

$$\mathbb{P}(x_t \in A_t^\eta \setminus C_t^{\eta,k_t^\eta}) \leq \frac{\eta}{T}. \quad (\text{EC.23})$$

Using these sets, we define an exercise policy on each period t by

$$\check{\mu}_t^\eta(y_t) \triangleq \begin{cases} \text{STOP}, & \text{if } y_t \in C_t^{\eta,k_t^\eta}, \\ \text{CONTINUE}, & \text{otherwise.} \end{cases}$$

Since the inequality (EC.23) holds for each period t , it follows from identical reasoning as the proof of property (a) of Claim EC.4 that the new exercise policy $\check{\mu}^\eta \equiv (\check{\mu}_1^\eta, \dots, \check{\mu}_T^\eta)$ satisfies property (a) of Claim EC.5.

Moreover, since the exercise regions C_t^{η,k_t^η} are unions of finite numbers of open balls, and since the Lebesgue measure of the boundaries of a finite number of open balls is equal to zero, it follows from Assumption 3 that the random state in each period $x_t \in \mathbb{R}^d$ will be a strictly positive distance from the boundary of C_t^{η,k_t^η} with probability one. This concludes our proof of property (b) of Claim EC.5, and thus concludes the proof of Claim EC.5. \square

We now combine Claims EC.4 and EC.5 to conclude our proof of Lemma EC.7. Indeed, it follows from these claims that

$$\mathbb{E}[g(\tau_{\bar{\mu}}(x), x)] \leq \mathbb{E}[g(\tau_{\check{\mu}^\eta}(x), x)] + \eta \leq \mathbb{E}[g(\tau_{\bar{\mu}^\eta}(x), x)] + 2\eta = \mathbb{E}\left[\liminf_{y \rightarrow x} g(\tau_{\bar{\mu}^\eta}(y), x)\right] + 2\eta,$$

where the first inequality follows from property (a) of Claim EC.4, the second inequality follows from property (a) of Claim EC.5, and the third inequality follows from property (b) of Claim EC.5. Since the exercise policy $\bar{\mu} \equiv (\bar{\mu}_1, \dots, \bar{\mu}_T)$ and constant $\eta > 0$ were chosen arbitrarily, we have proven that

$$\sup_{\mu} \mathbb{E}[g(\tau_{\bar{\mu}}(x), x)] \leq \sup_{\mu} \mathbb{E}\left[\liminf_{y \rightarrow x} g(\tau_{\mu}(y), x)\right].$$

The other direction of the inequality obviously holds, and so our proof of Lemma EC.7 is complete. \square

We now combine the above novel intermediary lemmas to establish our proofs of Theorems 1 and 2.

Proof of Theorem 1. We first show that the optimal objective value of (RO) is an asymptotic lower bound on the optimal objective value of (OPT). Indeed,

$$0 \geq \lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \left\{ \widehat{J}_{N,\epsilon}(\mu) - J^*(\mu) \right\} \quad \text{almost surely} \quad (\text{EC.24})$$

$$\geq \lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \left(\sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) - \sup_{\mu} J^*(\mu) \right) \quad (\text{EC.25})$$

$$= - \sup_{\mu} J^*(\mu) + \lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) \quad (\text{EC.26})$$

$$= - \sup_{\mu} J^*(\mu) + \lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) \quad \text{almost surely,} \quad (\text{EC.27})$$

where (EC.24) follows from Theorem 3, (EC.25) and (EC.26) follow from algebra, and (EC.27) follows from Lemma EC.6. Note that all of the above limits exist because $\epsilon \mapsto \limsup_{N \rightarrow \infty} \sup_{\mu} \{ \widehat{J}_{N,\epsilon}(\mu) - J^*(\mu) \}$ and $\epsilon \mapsto \limsup_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu)$ are monotonic functions.

We next show that the optimal objective value of (RO) provides an asymptotic upper bound on the optimal objective value of (OPT). Indeed, we observe that

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \limsup_{N \rightarrow \infty} \sup_{\mu} \widehat{J}_{N,\epsilon}(\mu) &= \lim_{\epsilon \rightarrow 0} \liminf_{N \rightarrow \infty} \sup_{\mu} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) \\ &\geq \limsup_{\epsilon \rightarrow 0} \sup_{\mu} \liminf_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) \end{aligned} \quad (\text{EC.28})$$

$$= \limsup_{\epsilon \rightarrow 0} \sup_{\mu} \mathbb{E} \left[\inf_{y \in \mathcal{X}^T: \|y-x\| \leq \epsilon} g(\tau_{\mu}(y), x) \right] \quad \text{almost surely} \quad (\text{EC.29})$$

$$= \sup_{\mu} \mathbb{E} \left[\liminf_{y \rightarrow x} g(\tau_{\mu}(y), x) \right] \quad \text{almost surely.} \quad (\text{EC.30})$$

$$= \sup_{\mu} \mathbb{E} [g(\tau_{\mu}(x), x)] \quad (\text{EC.31})$$

(EC.28) follows from algebra; (EC.29) follows from the strong law of large numbers; (EC.30) follows from the dominated convergence theorem. We note that the strong law of large numbers and dominated convergence theorem can both be applied because of Assumption 4. Finally, (EC.31) follows from Lemma EC.7. Combining the above, our proof of Theorem 1 is complete. \square

Proof of Theorem 2. We observe from Theorems 1 and 3 that for every arbitrary $\eta > 0$, there exists a finite $\bar{\epsilon}(\eta) > 0$ almost surely such that the following statements hold for all $0 < \epsilon < \bar{\epsilon}(\eta)$:

$$\left| \lim_{N \rightarrow \infty} \widehat{J}_{N,\epsilon}(\hat{\mu}_{N,\epsilon}) - \sup_{\mu} J^*(\mu) \right| \leq \eta \quad \text{almost surely;} \quad (\text{EC.32})$$

$$\liminf_{N \rightarrow \infty} \left(J^*(\hat{\mu}_{N,\epsilon}) - \widehat{J}_{N,\epsilon}(\hat{\mu}_{N,\epsilon}) \right) \geq -\eta \quad \text{almost surely.} \quad (\text{EC.33})$$

Therefore,

$$\sup_{\mu} J^*(\mu) \geq \limsup_{N \rightarrow \infty} J^*(\hat{\mu}_{N,\epsilon}) \geq \liminf_{N \rightarrow \infty} J^*(\hat{\mu}_{N,\epsilon}) \geq \liminf_{N \rightarrow \infty} \widehat{J}_{N,\epsilon}(\hat{\mu}_{N,\epsilon}) - \eta \geq \sup_{\mu} J^*(\mu) - 2\eta,$$

where the first inequality holds because each $\hat{\mu}_{N,\epsilon}$ is a feasible but possibly suboptimal solution to (OPT), the second inequality is obvious, the third inequality follows from (EC.33), and the final inequality follows from (EC.32). Rearranging the above line, we have shown that the following statements hold for all $0 < \epsilon < \bar{\epsilon}(\eta)$:

$$\begin{aligned} \left| \limsup_{N \rightarrow \infty} J^*(\hat{\mu}_{N,\epsilon}) - \sup_{\mu} J^*(\mu) \right| &\leq 2\eta \quad \text{almost surely;} \\ \left| \liminf_{N \rightarrow \infty} J^*(\hat{\mu}_{N,\epsilon}) - \sup_{\mu} J^*(\mu) \right| &\leq 2\eta \quad \text{almost surely.} \end{aligned}$$

Since $\eta > 0$ was chosen arbitrarily, our proof of Theorem 2 is complete. \square

Appendix D: Proofs from §3

Proof of Lemma 1. We recall from §2.1 that a Markovian stopping rule satisfies $\tau_{\mu}(y) = \infty$ for a trajectory $y \equiv (y_1, \dots, y_T) \in \mathcal{X}^T$ if and only if $\mu_t(y_t) = \text{CONTINUE}$ for each period $t \in \{1, \dots, T\}$. Therefore, we observe the robust optimization problem (RO_T) is equivalent to the following optimization problem:

$$\begin{aligned} \sup_{\mu} \quad & \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i) \\ \text{subject to} \quad & \tau_{\mu}(y) < \infty \quad \text{for all } i \in \{1, \dots, N\} \text{ and } y \in \mathcal{U}^i \end{aligned} \quad (\text{RO''})$$

Now consider any arbitrary exercise policy $\mu \equiv (\mu_1, \dots, \mu_T)$, and let $\mu' \equiv (\mu'_1, \dots, \mu'_T)$ be an exercise policy that is defined for each period $t \in \{1, \dots, T\}$ and state $y_t \in \mathcal{X}$ as

$$\mu'_t(y_t) \triangleq \begin{cases} \mu_t(y_t), & \text{if } t \in \{1, \dots, T-1\}, \\ \text{STOP}, & \text{if } t = T. \end{cases}$$

We readily observe that $\tau_{\mu'}(y) \leq T$ for each $i \in \{1, \dots, N\}$ and $y \in \mathcal{U}^i$, which implies that μ' satisfies the constraints of (ROⁿ). Moreover, we observe for each $i \in \{1, \dots, N\}$ and $y \in \mathcal{U}^i$ that

$$\tau_{\mu'}(y) = \begin{cases} \tau_\mu(y), & \text{if } \tau_\mu(y) < \infty, \\ T, & \text{if } \tau_\mu(y) = \infty. \end{cases}$$

Since the reward function satisfies $g(1, y), \dots, g(T, y) \geq 0$ and $g(\infty, y) = 0$ for all $y \in \mathcal{X}^T$, we have shown for each $i \in \{1, \dots, N\}$ and $y \in \mathcal{U}^i$ that

$$\begin{aligned} g(\tau_\mu(y), x^i) &= g(\tau_{\mu'}(y), x^i) && \text{if } \tau_\mu(y) < \infty, \text{ and} \\ g(\tau_\mu(y), x^i) &= 0 \leq g(T, x^i) = g(\tau_{\mu'}(y), x^i) && \text{if } \tau_\mu(y) = \infty. \end{aligned}$$

We thus conclude that the objective value $\frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu'}(y), x^i)$ associated with the new exercise policy $\mu' \equiv (\mu'_1, \dots, \mu'_T)$ is always greater than or equal to the objective value $\frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i)$ associated with the original exercise policy $\mu \equiv (\mu_1, \dots, \mu_T)$. Since $\mu \equiv (\mu_1, \dots, \mu_T)$ was chosen arbitrarily, our proof of Lemma 1 is complete. \square

Proof of Lemma 2. Consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO_T), and define

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

It follows from the fact that μ is feasible for (RO_T) that $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. Moreover, for each sample path $i \in \{1, \dots, N\}$ and trajectory $y \in \mathcal{U}^i$, we observe that

$$\tau_\mu(y) = \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP}\} \leq \sigma^i,$$

where the equality is simply the definition of a Markovian stopping rule and the inequality follows from the fact that $\mu_{\sigma^i}(y_{\sigma^i}) = \text{STOP}$ for all $y_{\sigma^i} \in \mathcal{U}_{\sigma^i}^i$. Therefore, we observe for each sample path $i \in \{1, \dots, N\}$ that

$$\begin{aligned} \{\tau_\mu(y) : y \in \mathcal{U}^i\} &= \left\{ t \in \{1, \dots, \sigma^i\} : \begin{array}{l} \text{there exists } y \in \mathcal{U}^i \text{ such that } \mu_{t'}(y_{t'}) = \text{CONTINUE} \\ \text{for all } t' \in \{1, \dots, t-1\} \text{ and } \mu_t(y) = \text{STOP} \end{array} \right\} \\ &= \{t \in \{1, \dots, \sigma^i\} : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t(y_t) = \text{STOP}\}. \end{aligned}$$

Indeed, the first equality follows from the definition of a Markovian stopping rule and from the fact that $\tau_\mu(y) \leq \sigma^i$ for all $y \in \mathcal{U}^i$. The second equality follows from the fact that \mathcal{U}_t^i is not a subset of $\{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}$ for each period $t \in \{1, \dots, \sigma^i - 1\}$, which implies for each $t \in \{1, \dots, \sigma^i - 1\}$ that there exists a $y_t \in \mathcal{U}_t^i$ that satisfies $\mu_t(y_t) = \text{CONTINUE}$. We thus conclude for each sample path $i \in \{1, \dots, N\}$ that

$$\inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i) = \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t(y_t) = \text{STOP}\},$$

which completes our proof of Lemma 2. \square

Proof of Lemma 3. Let μ' be a pruned version of μ , and let $\sigma^1, \dots, \sigma^N$ satisfy the following equalities for each $i \in \{1, \dots, N\}$:

$$\sigma^i = \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\} = \min \{t \in \{1, \dots, T\} : \mu'_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\}. \quad (\text{EC.34})$$

Then it follows from the fact that μ is feasible for (RO_T) that $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. Therefore, for each $i \in \{1, \dots, N\}$,

$$\begin{aligned} \inf_{y \in \mathcal{U}^i} g(\tau_{\mu'}(y), x^i) &= \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu'_t(y_t) = \text{STOP}\} \\ &\geq \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t(y_t) = \text{STOP}\} \\ &= \inf_{y \in \mathcal{U}^i} g(\tau_{\mu}(y), x^i). \end{aligned}$$

Indeed, the two equalities follow from Lemma 2 and line (EC.34). The inequality follows from the fact that μ' is a pruned version of μ , which implies that $\{y_t \in \mathcal{X} : \mu'_t(y_t) = \text{STOP}\} \subseteq \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}$ for all $t \in \{1, \dots, T\}$. Our proof of Lemma 3 is thus complete. \square

Proof of Lemma 4. Consider any $\mu \equiv (\mu_1, \dots, \mu_T)$ that satisfies the constraints of (RO_T) , and define

$$\sigma^i \triangleq \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \text{ for all } y_t \in \mathcal{U}_t^i\} \quad \forall i \in \{1, \dots, N\}.$$

Our proof that $\mu^{\sigma^1 \dots \sigma^N}$ is a pruned version of μ is split into the following two intermediary claims.

CLAIM EC.6. $\{y_t \in \mathcal{X} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}\} \subseteq \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}$ for all $t \in \{1, \dots, T\}$.

Proof of Claim EC.6. Indeed, we observe for each period $t \in \{1, \dots, T\}$ that

$$\begin{aligned} &\{y_t \in \mathcal{X} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}\} \\ &= \left\{ y_t \in \mathcal{X} : y_t \in \bigcup_{i: \sigma^i = t} \mathcal{U}_t^i \right\} \\ &\subseteq \left\{ y_t \in \mathcal{X} : y_t \in \bigcup_{i: \sigma^i = t} \{y'_t \in \mathcal{X} : \mu_t(y'_t) = \text{STOP}\} \right\} \\ &= \begin{cases} \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}, & \text{if there exists } i \in \{1, \dots, N\} \text{ such that } \sigma^i = t, \\ \emptyset, & \text{otherwise} \end{cases} \\ &\subseteq \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}, \end{aligned}$$

where the first equality follows from the definition of $\mu_t^{\sigma^1 \dots \sigma^N}$, the first inclusion follows from the fact that $\mathcal{U}_t^i \subseteq \{y_t \in \mathcal{X} : \mu_t(y_t) = \text{STOP}\}$ for each sample path i that satisfies $\sigma^i = t$, the second equality follows from algebra, and the second inclusion follows from algebra. This concludes our proof of Claim EC.6. \square

CLAIM EC.7. For each $i \in \{1, \dots, N\}$,

$$\min \{t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\} = \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\}.$$

Proof of Claim EC.7. Consider any sample path $i \in \{1, \dots, N\}$. We first observe that

$$\begin{aligned} \min \{t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\} &\leq \sigma^i \\ &= \min \{t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i\}, \end{aligned}$$

where the inequality follows from the fact that $\mu_{\sigma^i}^{\sigma^1 \dots \sigma^N}(y_{\sigma^i}) = \text{STOP}$ for all $y_{\sigma^i} \in \mathcal{U}_{\sigma^i}^i$, and the equality follows from the definition of σ^i . Moreover, it follows immediately from Claim EC.6 that

$$\min \left\{ t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i \right\} \geq \min \left\{ t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP} \forall y_t \in \mathcal{U}_t^i \right\}.$$

Our proof of Claim EC.7 is thus complete. \square

Combining Claims EC.6 and EC.7 with Definition 1, we conclude that $\mu^{\sigma^1 \dots \sigma^N}$ is a pruned version of μ , which completes our proof of Lemma 4. \square

The following proof follows the identical reasoning as discussed in §3.2, and is stated formally here for the sake of completeness.

Proof of Theorem 4. Lemma 1 shows that the robust optimization problem (RO) is equivalent to the robust optimization problem (RO_T). Moreover, for any arbitrary exercise policy μ that is feasible for (RO_T), Lemmas 3 and 4 together show that there exists an exercise policy $\mu' \in \mathcal{M}$ such that the objective value associated with μ is less than or equal to the objective value associated with μ' . Since μ was chosen arbitrarily, our proof of Theorem 4 is complete. \square

Proof of Theorem 5. We split our proof into the following two intermediary claims.

CLAIM EC.8. *The optimal objective value of (RO) is less than or equal to the optimal objective value of (IP).*

Proof of Claim EC.8. Consider any arbitrary $\mu \in \mathcal{M}$, and define the integers

$$\sigma^i \triangleq \min \left\{ t \in \{1, \dots, T\} : \mu_t(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i \right\} \quad \forall i \in \{1, \dots, N\}. \quad (\text{EC.35})$$

It follows from the definition of \mathcal{M} that μ satisfies the constraints of (RO_T), and so Lemma 4 implies that $\mu^{\sigma^1 \dots \sigma^N}$ is a pruned version of μ . Therefore, it follows from Definition 1 and line (EC.35) that the following equalities hold:

$$\sigma^i = \min \left\{ t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i \right\} \quad \forall i \in \{1, \dots, N\}. \quad (\text{EC.36})$$

Therefore, we observe that

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_\mu(y), x^i) \\ & \leq \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), x^i) \\ & = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \left\{ g(t, x^i) : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP} \right\} \\ & = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \left\{ g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t \right\}, \end{aligned} \quad (\text{EC.37})$$

where the first inequality follows from Lemma 3, the first equality follows from Lemma 2 and line (EC.36), and the second equality follows from the fact that $\mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}$ if and only if there exists a sample path j such that $y_t \in \mathcal{U}_t^j$ and $\sigma^j = t$. Because $\mu \in \mathcal{M}$ was chosen arbitrarily, we conclude from line (EC.37) and Theorem 4 that the optimal objective value of (RO) is less than or equal to the optimal objective value of (IP). Our proof of Claim EC.8 is thus complete. \square

CLAIM EC.9. *The optimal objective value of (RO) is greater than or equal to the optimal objective value of (IP). Furthermore, for any choice of integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$, the corresponding exercise policy $\mu^{\sigma^1 \dots \sigma^N} \equiv (\mu_1^{\sigma^1 \dots \sigma^N}, \dots, \mu_T^{\sigma^1 \dots \sigma^N})$ satisfies*

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g\left(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), x^i\right) \\ & \geq \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \left\{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\right\}. \end{aligned}$$

Proof of Claim EC.9. Consider any arbitrary integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$. For each sample path $i \in \{1, \dots, N\}$, we observe that

$$\begin{aligned} & \min \left\{t \in \{1, \dots, T\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\right\} \\ & = \min \left\{t \in \{1, \dots, \sigma^i\} : \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP for all } y_t \in \mathcal{U}_t^i\right\} \\ & \leq \sigma^i, \end{aligned} \tag{EC.38}$$

where the equality follows from the fact that $\mu^{\sigma^1 \dots \sigma^N} \equiv (\mu_1^{\sigma^1 \dots \sigma^N}, \dots, \mu_T^{\sigma^1 \dots \sigma^N})$ by construction satisfies $\mu_{\sigma^i}^{\sigma^1 \dots \sigma^N}(y_{\sigma^i}) = \text{STOP}$ for all $y_{\sigma^i} \in \mathcal{U}_{\sigma^i}^i$, and the inequality follows from algebra. Therefore,

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \inf_{y \in \mathcal{U}^i} g\left(\tau_{\mu^{\sigma^1 \dots \sigma^N}}(y), x^i\right) \\ & \geq \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \left\{g(t, x^i) : \text{there exists } y_t \in \mathcal{U}_t^i \text{ such that } \mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}\right\} \\ & = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \left\{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\right\}, \end{aligned} \tag{EC.39}$$

where the inequality follows from Lemma 2 and line (EC.38), and the equality follows from the fact that $\mu_t^{\sigma^1 \dots \sigma^N}(y_t) = \text{STOP}$ if and only if there exists a sample path j such that $y_t \in \mathcal{U}_t^j$ and $\sigma^j = t$. Because $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ were chosen arbitrarily, we conclude from line (EC.39) and Theorem 4 that the optimal objective value of (RO) is greater than or equal to the optimal objective value of (IP), which concludes our proof of Claim EC.9 is complete. \square

Combining Claims EC.8 and EC.9, our proof of Theorem 5 is thus complete. \square

Appendix E: Proof of Theorem 7

Our proof of the computational complexity of (IP) consists of a reduction from MIN-2-SAT, which is shown to be strongly NP-hard by Kohli et al. (1994):

MIN-2-SAT

The optimization version of MIN-2-SAT is to compute the optimal objective value of the binary linear optimization problem

$$\begin{aligned}
v^{\text{MIN-2-SAT}} \triangleq & \underset{b,z}{\text{minimize}} && \sum_{k=1}^K z_k \\
& \text{subject to} && z_k \geq b_\ell \quad \forall k \in \{1, \dots, K\}, \forall \ell \in I_k^+ \\
& && z_k \geq 1 - b_\ell \quad \forall k \in \{1, \dots, K\}, \forall \ell \in I_k^- \\
& && b_\ell \in \{0, 1\} \quad \forall \ell \in \{1, \dots, L\},
\end{aligned}$$

where the given sets $I_k^+, I_k^- \subseteq \{1, \dots, L\}$ satisfy $|I_k^+| + |I_k^-| = 2$ for each $k \in \{1, \dots, K\}$.

Note that the following equality is obtained by replacing each decision variable z_k with $1 - z_k$:

$$\begin{aligned}
v^{\text{MIN-2-SAT}} = & N - \underset{b,z}{\text{maximize}} && \sum_{k=1}^K z_k \\
& \text{subject to} && z_k \leq 1 - b_\ell \quad \forall k \in \{1, \dots, K\}, \forall \ell \in I_k^+ \\
& && z_k \leq b_\ell \quad \forall k \in \{1, \dots, K\}, \forall \ell \in I_k^- \\
& && b_\ell \in \{0, 1\} \quad \forall \ell \in \{1, \dots, L\}.
\end{aligned}$$

We now show that any instance of the above maximization problem can be equivalently reformulated as polynomially-size instance of (IP) with $T = 3$ periods.

Proof of Theorem 7. Consider any arbitrary instance of the binary linear optimization problem

$$\begin{aligned}
& \underset{b,z}{\text{maximize}} && \sum_{k=1}^K z_k \\
& \text{subject to} && z_k \leq 1 - b_\ell \quad \forall k \in \{1, \dots, K\}, \forall \ell \in I_k^+ \\
& && z_k \leq b_\ell \quad \forall k \in \{1, \dots, K\}, \forall \ell \in I_k^- \\
& && b_\ell \in \{0, 1\} \quad \forall \ell \in \{1, \dots, L\},
\end{aligned} \tag{-MIN-2-SAT}$$

and let $e_\ell \in \mathbb{R}^{L+1}$ denote the ℓ -th column vector of the identity matrix. We construct an instance of (IP) defined as follows:

- The number of periods is $T = 3$.
- The state space is $\mathcal{X} = \mathbb{R}^{L+1}$.
- The reward function for each period $t \in \{1, 2, 3\}$ is $g(t, y) = y_t \cdot e_{L+1} + K$.
- The robustness parameter in the uncertainty sets is $\epsilon = \frac{2}{3}$.
- The number of sample paths is $N \triangleq L + K$, and the sample paths are defined as follows:
 - For each $\ell \in \{1, \dots, L\}$, let $x_1^\ell = x_2^\ell = e_\ell$ and $x_3^\ell = -K e_{L+1}$.
 - For each $k \in \{1, \dots, K\}$, let $x_1^{L+k} = \frac{1}{2} \sum_{\ell \in I_k^+} e_\ell$, $x_2^{L+k} = \frac{1}{2} \sum_{\ell \in I_k^-} e_\ell$, and $x_3^{L+k} = e_{L+1}$.

In the remainder of the proof, we show that the above instance of (IP) is equivalent to (-MIN-2-SAT). Indeed, it follows immediately from the above construction that the values of $g(t, x^i)$ for each sample path i and period t are:

$$\begin{aligned}
g(1, x^\ell) &= K; & g(2, x^\ell) &= K; & g(3, x^\ell) &= 0, & \forall \ell \in \{1, \dots, L\}, \\
g(1, x^{L+k}) &= K; & g(2, x^{L+k}) &= K; & g(3, x^{L+k}) &= K + 1, & \forall k \in \{1, \dots, K\}.
\end{aligned}$$

We require two intermediary claims:

CLAIM EC.10. *There exists an optimal solution for (IP) which satisfies*

$$\sigma^1, \dots, \sigma^L \in \{1, 2\} \text{ and } \sigma^{L+1} = \dots = \sigma^{L+K} = 3.$$

Proof of Claim EC.10. We observe that the optimal objective value of (IP) is greater than or equal to K , since this objective value would be achieved by setting $\sigma^1 = \dots = \sigma^N = 1$. Now consider any solution to (IP) where $\sigma^{\ell'} = 3$ for some $\ell' \in \{1, \dots, L\}$. For that solution,

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ &= \frac{1}{N} \sum_{\ell \in \{1, \dots, L\} : \ell \neq \ell'} \min_{t \in \{1, \dots, \sigma^\ell\}} \{g(t, x^\ell) : \text{there exists } j \text{ such that } \mathcal{U}_t^\ell \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ &+ \frac{1}{N} \min_{t \in \{1, \dots, \sigma^{\ell'}\}} \{g(t, x^{\ell'}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{\ell'} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ &+ \frac{1}{N} \sum_{k=1}^K \min_{t \in \{1, \dots, \sigma^{L+k}\}} \{g(t, x^{L+k}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{L+k} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ &\leq \frac{1}{N} \sum_{\ell \in \{1, \dots, L\} : \ell \neq \ell'} K + 0 + \sum_{k=1}^K (K + 1) \\ &= K. \end{aligned}$$

Because the objective value associated with this solution is never better than the objective value obtained by the solution $\sigma^1 = \dots = \sigma^N = 1$, we have shown that there exists an optimal solution for (IP) that satisfies $\sigma^1, \dots, \sigma^L \in \{1, 2\}$.

Consider any arbitrary solution $\sigma^1, \dots, \sigma^N \in \{1, 2, 3\}$ that satisfies $\sigma^1, \dots, \sigma^L \in \{1, 2\}$, and suppose that $\sigma^{L+k} \in \{1, 2\}$ for some $k \in \{1, \dots, K\}$. To perform an exchange argument, we construct an alternative solution $\bar{\sigma}^1, \dots, \bar{\sigma}^N \in \{1, 2, 3\}$ defined as

$$\bar{\sigma}^i \triangleq \begin{cases} \sigma^i, & \text{if } i \neq L+k, \\ 3, & \text{if } i = L+k. \end{cases}$$

We observe that the inclusion $\{j : \sigma^j \leq \sigma^i\} \supseteq \{j : \bar{\sigma}^j \leq \bar{\sigma}^i\}$ holds for all $i \in \{1, \dots, N\} \setminus \{L+k\}$. Moreover, we observe that

$$\begin{aligned} & \min_{t \in \{1, \dots, \sigma^{L+k}\}} \{g(t, x^{L+k}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{L+k} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} = K, \text{ and} \\ & \min_{t \in \{1, \dots, \bar{\sigma}^{L+k}\}} \{g(t, x^{L+k}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{L+k} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \bar{\sigma}^j = t\} \in \{K, K+1\}. \end{aligned}$$

Therefore,

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ & \leq \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \bar{\sigma}^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \bar{\sigma}^j = t\} \end{aligned}$$

Because $\sigma^1, \dots, \sigma^N$ was chosen arbitrarily, we conclude that there exists an optimal solution for (IP) which satisfies $\sigma^1, \dots, \sigma^L \in \{1, 2\}$ and $\sigma^{L+1} = \dots = \sigma^{L+K} = 3$. This concludes our proof of Claim EC.10. \square

CLAIM EC.11. If $\sigma^{L+1} = \dots = \sigma^{L+K} = 3$, then the following equality holds for each $k \in \{1, \dots, K\}$:

$$\begin{aligned} & \min_{t \in \{1, \dots, \sigma^{L+k}\}} \{g(t, x^{L+k}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{L+k} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ &= K + \mathbb{I} \{ \sigma^\ell = 2 \text{ for all } \ell \in I_k^+ \text{ and } \sigma^\ell = 1 \text{ for all } \ell \in I_k^- \}. \end{aligned}$$

Proof of Claim EC.11. For each $\ell \in \{1, \dots, L\}$ and $k \in \{1, \dots, K\}$, we observe that

$$\begin{aligned} \|x_1^\ell - x_1^{L+k}\|_\infty &= \left\| e_\ell - \frac{1}{2} \sum_{\ell' \in I_k^+} e_{\ell'} \right\|_\infty = \begin{cases} \frac{1}{2}, & \text{if } \ell \in I_k^+, \\ 1, & \text{otherwise.} \end{cases} \\ \|x_2^\ell - x_2^{L+k}\|_\infty &= \left\| e_\ell - \frac{1}{2} \sum_{\ell' \in I_k^-} e_{\ell'} \right\|_\infty = \begin{cases} \frac{1}{2}, & \text{if } \ell \in I_k^-, \\ 1, & \text{otherwise.} \end{cases} \end{aligned}$$

This implies that the set

$$\mathcal{U}_1^\ell \cap \mathcal{U}_1^{L+k} = \left\{ y_1 \in \mathbb{R}^{L+1} : \|y_1 - x_1^\ell\|_\infty \leq \frac{2}{3} \right\} \cap \left\{ y_1 \in \mathbb{R}^{L+1} : \|y_1 - x_1^{L+k}\|_\infty \leq \frac{2}{3} \right\}$$

is nonempty if and only if $\ell \in I_k^+$, and the set

$$\mathcal{U}_2^\ell \cap \mathcal{U}_2^{L+k} = \left\{ y_2 \in \mathbb{R}^{L+1} : \|y_2 - x_2^\ell\|_\infty \leq \frac{2}{3} \right\} \cap \left\{ y_2 \in \mathbb{R}^{L+1} : \|y_2 - x_2^{L+k}\|_\infty \leq \frac{2}{3} \right\}$$

is nonempty if and only if $\ell \in I_k^-$. Consequently, since $\sigma^{L+1} = \dots = \sigma^{L+K} = 3$, we conclude that the following equalities hold for all $k \in \{1, \dots, K\}$:

$$\begin{aligned} & \min_{t \in \{1, \dots, \sigma^{L+k}\}} \{g(t, x^{L+k}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{L+k} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \\ &= \min \left\{ \min_{\ell \in I_k^+ : \sigma^\ell = 1} K, \min_{\ell \in I_k^- : \sigma^\ell = 2} K, K + 1 \right\} \\ &= K + \mathbb{I} \{ \sigma^\ell = 2 \text{ for all } \ell \in I_k^+ \text{ and } \sigma^\ell = 1 \text{ for all } \ell \in I_k^- \} \end{aligned}$$

This concludes the proof of Claim EC.11. \square

We now combine Claims EC.10 and EC.11 to complete our proof of Theorem 7:

$$(IP) \tag{EC.40}$$

$$= \max_{\substack{\sigma^1, \dots, \sigma^L \in \{1, 2\} \\ \sigma^{L+1} = \dots = \sigma^{L+K} = 3}} \frac{1}{L+K} \sum_{i=1}^{L+K} \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \tag{EC.41}$$

$$\begin{aligned} &= \max_{\substack{\sigma^1, \dots, \sigma^L \in \{1, 2\} \\ \sigma^{L+1} = \dots = \sigma^{L+K} = 3}} \left\{ \frac{1}{L+K} \sum_{\ell=1}^L K \right. \\ & \left. + \frac{1}{L+K} \sum_{k=1}^K \min_{t \in \{1, \dots, \sigma^{L+k}\}} \{g(t, x^{L+k}) : \text{there exists } j \text{ such that } \mathcal{U}_t^{L+k} \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \right\} \tag{EC.42} \end{aligned}$$

$$= \max_{\sigma^1, \dots, \sigma^L \in \{1,2\}} \left\{ \frac{LK}{L+K} + \frac{1}{L+K} \sum_{k=1}^K (K + \mathbb{I}\{\sigma^\ell = 2 \text{ for all } \ell \in I_k^+ \text{ and } \sigma^\ell = 1 \text{ for all } \ell \in I_k^-\}) \right\} \quad (\text{EC.43})$$

$$= K + \left(\frac{1}{L+K} \right) \max_{\sigma^1, \dots, \sigma^L \in \{1,2\}} \left\{ \sum_{k=1}^K \mathbb{I}\{\sigma^\ell = 2 \text{ for all } \ell \in I_k^+ \text{ and } \sigma^\ell = 1 \text{ for all } \ell \in I_k^-\} \right\}$$

$$= K + \left(\frac{1}{L+K} \right) * (\neg\text{MIN-2-SAT}). \quad (\text{EC.44})$$

Indeed, (EC.41) follows from Claim EC.10; (EC.42) holds because $g(1, x^\ell) = g(2, x^\ell) = K$ for all $\ell \in \{1, \dots, L\}$; (EC.43) follows from Claim EC.11; (EC.44) follows from algebra and setting $b^\ell = 0$ if and only if $\sigma^\ell = 2$. We have thus shown that any instance of MIN-2-SAT can be reduced to solving a polynomially-sized instance of (IP) with $T = 3$, which concludes our proof of Theorem 7. \square

Appendix F: Reformulation of (BP) as Mixed-Integer Linear Optimization Problem

Zero-one bilinear programs can be transformed into equivalent mixed-integer linear optimization problems by introducing auxiliary decision variables (Adams and Sherali 1986). In numerical experiments in §5, we perform such a linearization of the bilinear program (BP) by introducing auxiliary continuous decision variables $f_{t\ell}^i$ which obey the constraints

$$\begin{aligned} f_{t\ell}^i &\leq b_t^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, L_t^i - 1\} \\ f_{t\ell}^i &\leq 1 - w_{t\ell}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, L_t^i - 1\} \end{aligned}$$

and replacing the objective function of (BP) with

$$\underset{b, w, f}{\text{maximize}} \quad \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) f_{t\ell}^i.$$

To strengthen this linear relaxation of (BP), we also add the valid constraints:

$$\begin{aligned} w_{1,1}^i &= 0 && \text{for all } i \in \{1, \dots, N\} \\ b_t^i + w_{t,1}^i &= w_{t+1,1}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\ b_T^i + w_{T,1}^i &= 1 && \text{for all } i \in \{1, \dots, N\}. \end{aligned}$$

Indeed, the validity of the above constraints for (BP) follows from the fact that there is an optimal solution to this zero-one bilinear program which satisfies $\sum_{t=1}^T b_t^i = 1$ for each sample path i . In summary, this linearization procedure transforms (BP) into the following equivalent mixed-integer

linear optimization problem:

$$\begin{aligned}
& \underset{b,w,f}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) f_{t\ell}^i \\
& \text{subject to} && f_{t\ell}^i \leq b_t^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, L_t^i - 1\} \\
& && f_{t\ell}^i \leq 1 - w_{t\ell}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, L_t^i - 1\} \\
& && w_{t,\ell}^i \leq w_{t+1,\ell}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && w_{1,1}^i = 0 && \text{for all } i \in \{1, \dots, N\} \\
& && b_t^i + w_{t,1}^i = w_{t+1,1}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\
& && b_T^i + w_{T,1}^i = 1 && \text{for all } i \in \{1, \dots, N\} \\
& && b_t^i \leq w_{t+1,1}^i && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\
& && b_t^j \leq w_{t\ell}^i && \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\} \text{ such that } g(t, x^i) = \kappa_{\ell}^i \\
& && && \text{and } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \\
& && b_t^i \in \{0, 1\} && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\} \\
& && w_{t\ell}^i \in \mathbb{R} && \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

Appendix G: Proofs of Theorem 8, Lemma 5, and Lemma 6

Proof of Lemma 5. Our proof of Lemma 5 is split into two intermediary steps.

In the first intermediary step of our proof of Lemma 5, we show that every feasible solution for (IP) can be transformed into a feasible solution for the (BP-1) with the same objective value. Indeed, consider any feasible solution $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ for the optimization problem (IP). From these integers, we can define binary variables $b_t^i \triangleq \mathbb{I}\{\sigma^i = t\}$ for each sample path $i \in \{1, \dots, N\}$ and period $t \in \{1, \dots, T\}$. With this definition of a binary vector b , we observe for each sample path $i \in \{1, \dots, N\}$ that

$$\begin{aligned}
& \psi^i(b) \\
& = \sum_{t=1}^T b_t^i \left(\prod_{s=1}^{t-1} (1 - b_s^i) \right) \min_{s \in \{1, \dots, t\}} \{g(s, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \text{ and } b_s^j = 1\} \quad (\text{EC.45})
\end{aligned}$$

$$= \min_{s \in \{1, \dots, \sigma^i\}} \{g(s, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \text{ and } b_s^j = 1\} \quad (\text{EC.46})$$

$$= \min_{s \in \{1, \dots, \sigma^i\}} \{g(s, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \text{ and } \sigma^j = s\}, \quad (\text{EC.47})$$

where line (EC.45) follows from the definition of $\psi^i(b)$, and lines (EC.46) and (EC.47) follow from the fact that we have defined the binary variables to satisfy $b_t^j = \mathbb{I}\{\sigma^j = t\}$ for each sample path $j \in \{1, \dots, N\}$ and period $t \in \{1, \dots, T\}$. Because the above reasoning holds for each sample path $i \in \{1, \dots, N\}$, we have shown that

$$\frac{1}{N} \sum_{i=1}^N \psi^i(b) = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\}. \quad (\text{EC.48})$$

Therefore, we conclude that every feasible solution for (IP) can be transformed into a feasible solution for (BP-1) with the same objective value.

In the second intermediary step of the proof of Lemma 5, we show that if b is a feasible solution for (BP-1) and if $\sigma^i \triangleq \min\{\min\{t \in \{1, \dots, T\} : b_t^i = 1\}, T\}$ for each $i \in \{1, \dots, N\}$, then

$$\frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\} \geq \frac{1}{N} \sum_{i=1}^N \psi^i(b). \quad (\text{EC.49})$$

Indeed, consider any feasible solution b for (BP-1), and let us define a new binary vector \bar{b} as

$$\bar{b}_t^i \triangleq \begin{cases} b_t^i, & \text{if } t \in \{1, \dots, T-1\}, \\ 1, & \text{if } t = T. \end{cases}$$

It follows immediately from the above definition of \bar{b} and from the definition of the function $\psi^i(\cdot)$ for each sample path $i \in \{1, \dots, N\}$ that

$$\begin{aligned} & \psi^i(\bar{b}) - \psi^i(b) \\ &= \begin{cases} \min_{s \in \{1, \dots, T\}} \{g(s, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \text{ and } b_s^j = 1\} - 0, & \text{if } b_1^i = \dots = b_T^i = 0, \\ 0, & \text{otherwise} \end{cases} \\ &\geq 0. \end{aligned}$$

where the inequality holds because the reward function is nonnegative, that is, $g(t, x^i) \geq 0$ for each sample path $i \in \{1, \dots, N\}$ and period $s \in \{1, \dots, T\}$ (see §2.1). Therefore, we have proven that \bar{b} has the same or better objective value as b in the optimization problem (BP-1), i.e.,

$$\frac{1}{N} \sum_{i=1}^N \psi^i(b) \leq \frac{1}{N} \sum_{i=1}^N \psi^i(\bar{b}). \quad (\text{EC.50})$$

Now, for each sample path $i \in \{1, \dots, N\}$, let us define the following integer:

$$\sigma^i \triangleq \min \left\{ \min_{t \in \{1, \dots, T\}} \{t : b_t^i = 1\}, T \right\} = \min_{t \in \{1, \dots, T\}} \{t : \bar{b}_t^i = 1\}.$$

Given the integers $\sigma^1, \dots, \sigma^N \in \{1, \dots, T\}$ defined above, it follows from identical reasoning as in lines (EC.45), (EC.46), (EC.47), and (EC.48) that

$$\frac{1}{N} \sum_{i=1}^N \psi^i(\bar{b}) = \frac{1}{N} \sum_{i=1}^N \min_{t \in \{1, \dots, \sigma^i\}} \{g(t, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \text{ and } \sigma^j = t\}.$$

Combining the above equality with (EC.50), we have proven that line (EC.49) holds, thereby completing our proof of the second intermediary step.

In conclusion, we have shown through the above two intermediary steps that every feasible solution for either (IP) or (BP-1) can be transformed into a feasible solution for the other problem with the same or better objective value, which implies that the optimal objective values of (IP) and (BP-1) are equal. Moreover, the inequality in Lemma 5 follows immediately from the second intermediary step. Thus, our proof of Lemma 5 is complete. \square

Proof of Lemma 6. Consider any sample path $i \in \{1, \dots, N\}$ and binary vector b . For the sake of convenience, let us repeat below the linear optimization problem that is found in the statement of Lemma 6 for the given sample path i and binary vector b :

$$\begin{aligned}
& \underset{w^i}{\text{maximize}} && \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) b_t^i (1 - w_{t\ell}^i) \\
& \text{subject to} && w_{t\ell}^i \leq w_{t+1,\ell}^i \quad \text{for all } t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_t^j \leq w_{t\ell}^i \quad \text{for all } j \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\} \\
& && \text{such that } g(t, x^i) = \kappa_\ell^i \text{ and } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \\
& && b_t^i \leq w_{t+1,1}^i \quad \text{for all } t \in \{1, \dots, T-1\} \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{3}$$

In the remainder of this proof, we show that the optimal objective value of (3) is equal to $\psi^i(b)$.

To this end, we begin by characterizing the decision variables w^i in (3) at optimality. Indeed, it follows from the construction of the constants κ_ℓ^i that each quantity $\kappa_{\ell+1}^i - \kappa_\ell^i$ is strictly positive. Therefore, since each b_t^j is binary, we readily observe that there always exists an optimal solution w^i for (3) where the following equality holds for each period $t \in \{1, \dots, T\}$ and each $\ell \in \{1, \dots, |\mathcal{K}^i|\}$:

$$w_{t\ell}^i = \begin{cases} 0, & \text{if } [b_s^j = 0 \text{ for all } j \in \{1, \dots, N\} \text{ and } s \in \{1, \dots, t\} \text{ such that } g(s, x^i) \leq \kappa_\ell^i \text{ and } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset] \\ & \text{and } [b_s^i = 0 \text{ for all } s \in \{1, \dots, t-1\}], \\ 1, & \text{otherwise.} \end{cases} \tag{EC.51}$$

Now consider any optimal solution w^i for the optimization problem (3) for which the equality in (EC.51) is satisfied for each period $t \in \{1, \dots, T\}$ and each $\ell \in \{1, \dots, |\mathcal{K}^i|\}$. For each period $t \in \{1, \dots, T\}$, let $\tilde{\ell}_t^i$ be defined as the smallest integer such that there exists a sample path $j \in \{1, \dots, N\}$ and a period $s \in \{1, \dots, t\}$ that satisfy $b_s^j = 1$, $g(s, x^i) = \kappa_{\tilde{\ell}_t^i}^i$, and $\mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset$.¹² In particular, we observe from (EC.51) that the quantity $\tilde{\ell}_t^i$ is equal to the smallest integer such that $w_{t\tilde{\ell}_t^i}^i = 1$. Hence, the optimal objective value of the optimization problem (3) is equal to:

$$\begin{aligned}
& \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) b_t^i (1 - w_{t\ell}^i) \\
& = \sum_{t \in \{1, \dots, T\}: b_t^i=1} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (1 - w_{t\ell}^i)
\end{aligned} \tag{EC.52}$$

$$= \begin{cases} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (1 - w_{t\ell}^i), & \text{if } b_1^i = \dots = b_{t-1}^i = 0 \text{ and } b_t^i = 1 \text{ for } t \in \{1, \dots, T\}, \\ 0, & \text{otherwise} \end{cases} \tag{EC.53}$$

$$= \begin{cases} \sum_{\ell=1}^{\tilde{\ell}_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i), & \text{if } b_1^i = \dots = b_{t-1}^i = 0 \text{ and } b_t^i = 1 \text{ for } t \in \{1, \dots, T\}, \\ 0, & \text{otherwise} \end{cases} \tag{EC.54}$$

$$= \begin{cases} \kappa_{\tilde{\ell}_t^i}^i, & \text{if } b_1^i = \dots = b_{t-1}^i = 0 \text{ and } b_t^i = 1 \text{ for } t \in \{1, \dots, T\}, \\ 0, & \text{otherwise} \end{cases} \tag{EC.55}$$

¹² If no such integer exists, then we assign $\tilde{\ell}_t^i$ to be equal to L_t^i .

$$= \sum_{t=1}^T b_t^i \left(\prod_{s=1}^{t-1} (1 - b_s^i) \right) \kappa_{\tilde{\ell}_t^i}^i \quad (\text{EC.56})$$

$$= \sum_{t=1}^T b_t^i \left(\prod_{s=1}^{t-1} (1 - b_s^i) \right) \min_{s \in \{1, \dots, t\}} \{g(s, x^i) : \text{there exists } j \text{ such that } \mathcal{U}_s^i \cap \mathcal{U}_j^i \neq \emptyset \text{ and } b_s^j = 1\} \quad (\text{EC.57})$$

$$= \psi^i(b).$$

Indeed, we observe that the optimal objective value of (3) is equal to $\sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_t^i (1 - w_{t\ell}^i)$ because w^i is an optimal solution for (3). Line (EC.52) follows from algebra. Line (EC.53) follows from (EC.51), which implies that if $b_t^i = 1$, then $w_{s\ell}^i = 1$ for all $s \in \{t+1, \dots, T\}$ and all $\ell \in \{1, \dots, |\mathcal{K}^i|\}$. Line (EC.54) follows from the definition of $\tilde{\ell}_t^i$ and from (EC.51). Line (EC.55) follows from the fact that $\kappa_1^i = 0$. Line (EC.56) follows from algebra. Line (EC.57) follows from the definition of $\tilde{\ell}_t^i$ and the definition of the constants $\kappa_1^i, \dots, \kappa_{|\mathcal{K}^i|}^i$. The final equality follows from the definition of $\psi^i(b)$. This completes our proof of Lemma 6. \square

Proof of Theorem 8. The proof of Theorem 8 follows immediately from Lemmas 5 and 6. \square

Appendix H: Proofs from §4.2

H.1. Proofs from §4.2.1

Proof of Proposition 1. Let \hat{b}, \hat{w} denote an optimal solution for (BP), and let the set $\hat{\mathcal{T}}^i \triangleq \{t : \hat{b}_t^i = 1\}$ be defined for each sample path i . We will henceforth consider the optimization problem (H) in the case where the linear objective function $f(b, w)$ is defined equal to the following function:

$$\hat{f}(b, w) \triangleq \frac{1}{N} \sum_{i=1}^N \sum_{t \in \hat{\mathcal{T}}^i} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (b_t^i - w_{t\ell}^i).$$

We readily observe that the function $\hat{f}(b, w)$ is linear in b and w . Moreover, it follows from algebra that $\hat{f}(b, w)$ is less than or equal to the objective function of (BP) for all feasible solutions of (BP); that is, the linear function $\hat{f}(b, w)$ satisfies the condition from line (4).¹³ Thus, we conclude that solving the optimization problem (H) with objective function $f(b, w) = \hat{f}(b, w)$ will provide a lower-bound approximation of the optimization problem (BP), and any optimal solution for (H) will be a feasible solution for (BP). For notational convenience, we henceforth let J^H denote the optimal objective value of (H) with objective function $f(b, w) = \hat{f}(b, w)$, and we let J^{BP} denote the optimal objective value of (BP).

We first show that the optimal objective value of (H) with objective function $f(b, w) = \hat{f}(b, w)$ is equal to the optimal objective value of (BP). Indeed, it follows from the fact that (H) and (BP) have the same constraints that the optimal solution \hat{b}, \hat{w} for (BP) is a feasible solution for (H). Therefore,

$$J^H \geq \frac{1}{N} \sum_{i=1}^N \sum_{t \in \hat{\mathcal{T}}^i} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (\hat{b}_t^i - \hat{w}_{t\ell}^i)$$

¹³ To see why $\hat{f}(b, w)$ is a lower bound on the objective function of (BP), consider any arbitrary vectors b, w that satisfy the constraints of (BP). Since feasibility for the optimization problem (BP) implies that b is a binary vector, we observe that the equality $b_t^i (1 - w_{t\ell}^i) = \max\{b_t^i - w_{t\ell}^i, 0\}$ holds for each $i \in \{1, \dots, N\}$, $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, |\mathcal{K}^i|\}$.

$$\begin{aligned}
&= \frac{1}{N} \sum_{i=1}^N \sum_{t: \hat{b}_t^i=1} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i)(1 - \hat{w}_{t\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{t: \hat{b}_t^i=0} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) 0 \\
&= \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) \hat{b}_t^i (1 - \hat{w}_{t\ell}^i) \\
&= J^{\text{BP}},
\end{aligned}$$

where the first inequality holds because \hat{b}, \hat{w} is a feasible but possibly suboptimal solution for (H), the first equality follows from the construction of the sets $\hat{\mathcal{T}}^1, \dots, \hat{\mathcal{T}}^N \subseteq \{1, \dots, T\}$, the second equality follows from algebra, and the final equality follows from the fact that \hat{b}, \hat{w} is an optimal solution for (BP). Since the optimal objective value of (H) is always less than or equal to the optimal objective value of (BP), we have proved that the optimal objective value of (H) is equal to the optimal objective value of (BP).

We conclude the proof of Proposition 1 by showing that every optimal solution for (H) with objective function $f(b, w) = \hat{f}(b, w)$ is an optimal solution for (BP). Indeed, let \bar{b}, \bar{w} denote an optimal solution for (H). Then,

$$J^{\text{H}} = \hat{f}(\bar{b}, \bar{w}) \leq \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) \bar{b}_t^i (1 - \bar{w}_{t\ell}^i) \leq J^{\text{BP}},$$

where the first equality follows from the fact that \bar{b}, \bar{w} is an optimal solution for (H) and from the fact that the objective function of (H) is $f(b, w) = \hat{f}(b, w)$, the first inequality follows from the fact that the linear function $\hat{f}(b, w)$ satisfies the condition from line (4), and the final inequality follows from the fact that \bar{b}, \bar{w} is a feasible but possibly suboptimal solution for (BP). Since we have previously shown that the optimal objective value J^{H} of (H) is equal to the optimal objective value J^{BP} of (BP), we have thus proven that \bar{b}, \bar{w} is an optimal solution for (BP). This concludes our proof of Proposition 1. \square

H.2. Proofs from §4.2.3

Proof of Lemma 7. Let $f(b, w) = \bar{f}(b, w)$, and consider any binary vector b that is optimal for the optimization problem (H). Since b is binary, we readily observe that there exists an optimal choice for the remaining decision variables w in the optimization problem (H) in which the following equality holds for each sample path $i \in \{1, \dots, N\}$, period $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, |\mathcal{K}^i|\}$:

$$w_{t\ell}^i = \begin{cases} 0, & \text{if } [b_s^j = 0 \text{ for all } j \in \{1, \dots, N\} \text{ and } s \in \{1, \dots, t\} \text{ such that } g(s, x^i) \leq \kappa_\ell^i \text{ and } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \\ & \text{and } [b_s^i = 0 \text{ for all } s \in \{1, \dots, t-1\}], \\ 1, & \text{otherwise.} \end{cases} \quad (\text{EC.58})$$

We will now construct a new solution for (H) that has the same or greater objective value than b, w . Indeed, let \bar{b}, \bar{w} be a solution for (H) defined by the following equalities for each sample path $i \in \{1, \dots, N\}$, period $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, |\mathcal{K}^i|\}$:

$$\begin{aligned}
\bar{b}_t^i &\triangleq \begin{cases} 1, & \text{if } [b_t^i = 1 \text{ and } t = T^i] \text{ or } [t = T], \\ 0, & \text{otherwise;} \end{cases} \\
\bar{w}_{t\ell}^i &\triangleq \begin{cases} 0, & \text{if } [\bar{b}_s^j = 0 \text{ for all } j \in \{1, \dots, N\} \text{ and } s \in \{1, \dots, t\} \text{ such that } g(s, x^i) \leq \kappa_\ell^i \text{ and } \mathcal{U}_s^i \cap \mathcal{U}_s^j \neq \emptyset \\ & \text{and } [\bar{b}_s^i = 0 \text{ for all } s \in \{1, \dots, t-1\}], \\ 1, & \text{otherwise.} \end{cases}
\end{aligned}$$

We observe for each sample path $i \in \{1, \dots, N\}$ that $\bar{b}_T^i = 1$ and $\bar{b}_t^i = 0$ for all $t \in \{1, \dots, T\} \setminus \mathcal{T}^i$. Moreover, we readily observe from inspection that the solution \bar{b}, \bar{w} is feasible for (H). Therefore, it remains for us to prove that \bar{b}, \bar{w} is an optimal solution for (H).

To show that \bar{b}, \bar{w} is an optimal solution for (H), we observe for each sample path $i \in \{1, \dots, N\}$, period $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, L_t^i - 1\}$,

$$\bar{w}_{t\ell}^i = \begin{cases} 0, & \text{if } [b_{T^j}^j = 0 \text{ for all } j \in \{1, \dots, N\} \text{ such that } T^j \leq t, g(T^j, x^i) \leq \kappa_\ell^i, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset] \\ & \text{and } [b_{T^i}^i = 0 \text{ if } T^i \in \{1, \dots, t-1\}], \\ 1, & \text{otherwise} \end{cases} \leq w_{t\ell}^i \quad (\text{EC.59})$$

where the equality follows from the definition of \bar{b}, \bar{w} , and the inequality follows from line (EC.58). Therefore,

$$\begin{aligned} \bar{f}(\bar{b}, \bar{w}) &= \frac{1}{N} \sum_{i=1}^N \sum_{t \in \mathcal{T}^i} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (\bar{b}_t^i - \bar{w}_{t\ell}^i) \\ &\geq \frac{1}{N} \sum_{i=1}^N \sum_{t \in \mathcal{T}^i} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_t^i - w_{t\ell}^i) \\ &= \bar{f}(b, w). \end{aligned}$$

Indeed, the first equality follows from the definition of $\bar{f}(\cdot, \cdot)$. The inequality follows from line (EC.59), from the facts that $\bar{b}_{T^i}^i \geq b_{T^i}^i$ and $\bar{b}_T^i \geq b_T^i$, and from the fact that each quantity $\kappa_{\ell+1}^i - \kappa_\ell^i$ is strictly positive. The final equality follows from the definition of $\bar{f}(\cdot, \cdot)$. Since b, w was an optimal solution for (H), our proof of Lemma 7 is complete. \square

Proof of Lemma 8. Let $f(b, w) = \bar{f}(b, w)$. In this case, we recall from Lemma 7 that there exists an optimal solution b, w for (H) that satisfies $b_T^1 = \dots = b_T^N = 1$ and $b_t^i = 0$ for each sample path $i \in \{1, \dots, N\}$ and period $t \in \{1, \dots, T\} \setminus \mathcal{T}^i$. Therefore, we can without loss of generality impose those equality constraints into the optimization problem (H). That is, (H) can be rewritten equivalently as

$$\begin{aligned} &\underset{b, w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{t \in \mathcal{T}^i} \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_t^i - w_{t\ell}^i) \\ &\text{subject to} && w_{t\ell}^i \leq w_{t+1, \ell}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\ &&& w_{t\ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\ &&& b_t^i \leq w_{t+1, 1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\ &&& b_t^j \leq w_{t\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \{1, \dots, T\} \text{ such that } g(t, x^i) = \kappa_\ell^i \quad (\text{H-1}) \\ &&& \text{and } \mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset \\ &&& b_T^i = 1 \quad \text{for all } i \in \{1, \dots, N\} \\ &&& b_t^i = 0 \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\} \setminus \mathcal{T}^i \\ &&& b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\} \\ &&& w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}. \end{aligned}$$

After substituting out the decision variables b_t^i that have been set to zero or one, it follows from algebra that the above optimization problem can be rewritten equivalently as

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i < T} \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (b_{T^i}^i - w_{T^i \ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_T^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (1 - w_{T \ell}^i) \\
& \text{subject to} && w_{t \ell}^i \leq w_{t+1, \ell}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t \ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_{T^i}^i \leq w_{T^i+1, 1}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\
& && b_{T^j}^j \leq w_{T^j \ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } T^j < T, g(T^j, x^i) = \kappa_{\ell}^i, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\
& && w_{T, L_T^i}^i = 1 \quad \text{for all } i \in \{1, \dots, N\} \\
& && b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\
& && w_{t \ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{H-2}$$

Let us make three observations about the above optimization problem. First, we observe from inspection that the constraints in the above optimization problem of the form $w_{T, L_T^i}^i = 1$ for all $i \in \{1, \dots, N\}$ can be removed from the above optimization problem without affecting its optimal objective value. Second, we observe that the decision variables $w_{t \ell}^i$ for each $t \in \{1, \dots, T\} \setminus \mathcal{T}^i$ do not appear in the objective function of (H-2). Third, since each term $\kappa_{\ell+1}^i - \kappa_{\ell}^i$ is strictly positive, we observe that there exists an optimal solution for (H-2) that satisfies the following equality for each sample path $i \in \{1, \dots, N\}$, period $t \in \mathcal{T}^i$, and $\ell \in \{1, \dots, |\mathcal{K}^i|\}$:

$$w_{t \ell}^i = \begin{cases} 0, & \text{if } [b_{T^j}^j = 0 \text{ for all } j \in \{1, \dots, N\} \text{ such that } T^j \leq t, T^j < T, g(T^j, x^i) \leq \kappa_{\ell}^i, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset] \\ & \text{and [if } T^i < t, \text{ then } b_{T^i}^i = 0], \\ 1, & \text{otherwise.} \end{cases}$$

It follows from the aforementioned three observations that (H-2) can be rewritten equivalently as

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i < T} \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (b_{T^i}^i - w_{T^i \ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_T^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (1 - w_{T \ell}^i) \\
& \text{subject to} && w_{t \ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_{T^i}^i \leq w_{T^i 1}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\
& && b_{T^j}^j \leq w_{t \ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \mathcal{T}^i \text{ such that } g(T^j, x^i) = \kappa_{\ell}^i, \\
& && \quad \quad \quad T^j \leq t, T^j < T, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\
& && b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\
& && w_{t \ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{\bar{H}}$$

We have thus shown that the optimal objective value of (\bar{H}) is equal to the optimal objective value of (H) when $f(b, w) = \bar{f}(b, w)$. Moreover, it follows from the above reasoning that any optimal solution \bar{b}, \bar{w} for (\bar{H}) can be transformed into an optimal solution for (H) using the following equality:

$$b_t^i = \begin{cases} \bar{b}_t^i, & \text{if } t = T^i \text{ and } T^i < T, \\ 1, & \text{if } t = T, \\ 0, & \text{otherwise.} \end{cases}$$

This concludes our proof of Lemma 8. \square

Proof of Proposition 2. We observe that (\bar{H}) is equivalent to a binary linear optimization problem with $\mathcal{O}(NT)$ binary decision variables and $\mathcal{O}(N^2 + NT)$ constraints, with each constraint of the form $\lambda_i \geq \lambda_j$. It thus follows from Picard (1976, §3) that the optimization problem (\bar{H}) is equivalent to a problem of computing the maximal closure of a directed graph with $\mathcal{O}(NT)$ nodes and $\mathcal{O}(N^2 + NT)$ edges. Furthermore, Picard (1976, §4) shows that any maximal closure problem can be solved by computing the maximum flow in an augmented graph of identical size. Applying the algorithm of Orlin (2013) to compute the maximum flow in this augmented graph, we obtain an $\mathcal{O}(N^2T(N + T))$ algorithm for solving (\bar{H}) . Finally, the output of the maximal closure problem can be transformed into an optimal solution for (H) using Lemma 8. This concludes our proof of Proposition 2. \square

H.3. Proofs from §4.2.4

Proof of Proposition 3. Suppose that the number of periods is $T = 2$. We recall from Remark 2 that the optimal objective value of (\bar{H}) is less than or equal to the optimal objective value of (BP) . Therefore, it remains for us to show that the optimal objective value of (BP) is less than or equal to the optimal objective value of (\bar{H}) .

We begin by restating the optimization problem (BP) when $T = 2$ for the sake of convenience:

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_1^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_1^i (1 - w_{1\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_2^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_2^i (1 - w_{2\ell}^i) \\
& \text{subject to} && w_{1\ell}^i \leq w_{2\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_1^i \leq w_{21}^i \quad \text{for all } i \in \{1, \dots, N\} \\
& && b_1^j \leq w_{1\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(1, x^i) = \kappa_{\ell}^i \text{ and } \mathcal{U}_1^i \cap \mathcal{U}_1^j \neq \emptyset \\
& && b_2^j \leq w_{2\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(2, x^i) = \kappa_{\ell}^i \text{ and } \mathcal{U}_2^i \cap \mathcal{U}_2^j \neq \emptyset \\
& && b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\} \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned} \tag{BP}$$

We readily observe from inspection of the optimization problem (BP) that there exists an optimal solution b, w for (BP) that satisfies the equalities $b_2^i = \dots = b_2^N = 1$ and $b_1^i = 0$ for all $i \in \{1, \dots, N\}$ such that $T^i = 2$. Therefore, we observe that those equality constraints can be added to the optimization problem (BP) without changing its optimal objective value. That is, the optimization problem (BP) can be rewritten equivalently as

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_1^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_1^i (1 - w_{1\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_2^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_2^i (1 - w_{2\ell}^i) \\
& \text{subject to} && w_{1\ell}^i \leq w_{2\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_1^i \leq w_{21}^i \quad \text{for all } i \in \{1, \dots, N\} \\
& && b_1^j \leq w_{1\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(1, x^i) = \kappa_{\ell}^i \text{ and } \mathcal{U}_1^i \cap \mathcal{U}_1^j \neq \emptyset \\
& && b_2^j \leq w_{2\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(2, x^i) = \kappa_{\ell}^i \text{ and } \mathcal{U}_2^i \cap \mathcal{U}_2^j \neq \emptyset \\
& && b_2^i = 1 \quad \text{for all } i \in \{1, \dots, N\} \\
& && b_1^i = 0 \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 2 \\
& && b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\} \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

After eliminating the decision variables b_t^i that have been constrained to be equal to zero or one, the above optimization problem can be rewritten equivalently as

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i=1} \sum_{\ell=1}^{L_1^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_1^i (1 - w_{1\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_2^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (1 - w_{2\ell}^i) \\
& \text{subject to} && w_{1\ell}^i \leq w_{2\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_1^i \leq w_{21}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && b_1^j \leq w_{1\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(1, x^i) = \kappa_{\ell}^i, T^j = 1, \text{ and } \mathcal{U}_1^i \cap \mathcal{U}_1^j \neq \emptyset \\
& && w_{2, L_2^i}^i = 1 \quad \text{for all } i \in \{1, \dots, N\} \\
& && b_1^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

Let us make two observations about the above optimization problem. First, we observe from inspection that the constraints in the above optimization problem of the form $w_{2, L_2^i}^i = 1$ for all $i \in \{1, \dots, N\}$ can be removed from the above optimization problem without affecting its optimal objective value. Second, since each term $\kappa_{\ell+1}^i - \kappa_{\ell}^i$ is strictly positive, we observe that there exists an optimal solution for the above optimization problem in which the equalities $w_{11}^i = \dots = w_{1, L_1^i-1}^i = 0$ are satisfied for each sample path $i \in \{1, \dots, N\}$. Therefore, the above optimization problem can be rewritten equivalently as

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i=1} \sum_{\ell=1}^{L_1^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) b_1^i (1 - w_{1\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_2^i-1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (1 - w_{2\ell}^i) \\
& \text{subject to} && w_{1\ell}^i \leq w_{2\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_1^i \leq w_{21}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && b_1^j \leq w_{1\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(1, x^i) = \kappa_{\ell}^i, T^j = 1, \text{ and } \mathcal{U}_1^i \cap \mathcal{U}_1^j \neq \emptyset \\
& && w_{1\ell}^i = 0 \quad \text{for all } i \in \{1, \dots, N\} \text{ and } \ell \in \{1, \dots, L_1^i - 1\} \\
& && b_1^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

Because the above optimization problem has the constraint that $w_{1\ell}^i = 0$ for all $i \in \{1, \dots, N\}$ and $\ell \in \{1, \dots, L_1^i - 1\}$, we can without loss of generality replace each term $b_1^i (1 - w_{1\ell}^i)$ in the objective function of the above optimization problem with $b_1^i - w_{1\ell}^i$. That is, the above optimization problem

is equivalent to

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i=1} \sum_{\ell=1}^{L_1^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i)(b_1^i - w_{1\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_2^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i)(1 - w_{2\ell}^i) \\
& \text{subject to} && w_{1\ell}^i \leq w_{2\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_1^i \leq w_{21}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && b_1^j \leq w_{1\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(1, x^i) = \kappa_\ell^i, T^j = 1, \text{ and } \mathcal{U}_1^i \cap \mathcal{U}_1^j \neq \emptyset \\
& && w_{1\ell}^i = 0 \quad \text{for all } i \in \{1, \dots, N\} \text{ and } \ell \in \{1, \dots, L_1^i - 1\} \\
& && b_1^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

Finally, we can relax the above optimization problem by removing the constraint that $w_{1\ell}^i = 0$ for all $i \in \{1, \dots, N\}$ and $\ell \in \{1, \dots, L_1^i - 1\}$. That is, the optimal objective value of the above optimization problem is less than or equal to the optimal objective value of the following optimization problem:

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i=1} \sum_{\ell=1}^{L_1^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i)(b_1^i - w_{1\ell}^i) + \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_2^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i)(1 - w_{2\ell}^i) \\
& \text{subject to} && w_{1\ell}^i \leq w_{2\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_1^i \leq w_{21}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && b_1^j \leq w_{1\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(1, x^i) = \kappa_\ell^i, T^j = 1, \text{ and } \mathcal{U}_1^i \cap \mathcal{U}_1^j \neq \emptyset \\
& && b_1^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i = 1 \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, 2\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

We observe that the above optimization problem is equivalent to $(\bar{\text{H}})$ for the case where $T = 2$. Therefore, we have shown that the optimal objective value of (BP) is less than or equal to the optimal objective value of $(\bar{\text{H}})$, which concludes our proof of Proposition 3. \square

Our proofs of Propositions 4 and 5 will make use of two intermediary results, denoted below by Lemmas EC.8 and EC.9. These intermediary lemmas establish an upper bound on the optimal objective value of (BP) and a lower bound on the optimal objective value of $(\bar{\text{H}})$, respectively. Throughout the proofs, we let J^{BP} denote the optimal objective value of (BP) and $J^{\bar{\text{H}}}$ denote the optimal objective value of $(\bar{\text{H}})$.

LEMMA EC.8. $J^{\text{BP}} \leq \frac{1}{N} \sum_{i=1}^N g(T^i, x^i)$.

Proof of Lemma EC.8. By removing constraints from the optimization problem (BP), we obtain the following optimization problem:

$$\begin{aligned}
& \underset{b,w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) b_t^i (1 - w_{t\ell}^i) \\
& \text{subject to} && w_{t\ell}^i \leq w_{t+1,\ell}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\}, \ell \in \{1, \dots, |\mathcal{K}^i|\} \\
& && w_{t\ell}^i \leq w_{t,\ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_t^i \leq w_{t+1,1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T-1\} \\
& && b_t^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\} \\
& && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \{1, \dots, T\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}.
\end{aligned}$$

In greater detail, we observe that the above optimization problem is identical to (BP), with the exception that the above optimization problem does not include the constraints of the form $b_t^j \leq w_{t\ell}^i$ for all $i, j \in \{1, \dots, N\}$ and $t \in \{1, \dots, T\}$ such that $g(t, x^i) = \kappa_\ell^i$ and $\mathcal{U}_t^i \cap \mathcal{U}_t^j \neq \emptyset$. Consequently, the optimal objective value of the above optimization problem is greater than or equal to the optimal objective value of (BP). Moreover, it follows from inspection and from the fact that each term $\kappa_{\ell+1}^i - \kappa_\ell^i$ is strictly positive that there exists an optimal solution for the above optimization problem that satisfies $b_t^i = \mathbb{I}\{t = T^i\}$ for each sample path $i \in \{1, \dots, N\}$. Therefore, the optimal objective value of the above optimization problem is equal to

$$\begin{aligned} \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^T \sum_{\ell=1}^{L_t^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) \mathbb{I}\{t = T^i\} (1 - w_{t\ell}^i) &= \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (1 - w_{T^i\ell}^i) \\ &\leq \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) \\ &= \frac{1}{N} \sum_{i=1}^N g(T^i, x^i), \end{aligned}$$

where the first equality follows from algebra, the inequality follows from the fact that each decision variable $w_{T^i\ell}^i$ must be greater than or equal to zero, and the final equality follows from the definitions of the constants κ_ℓ^i and L_t^i . This concludes our proof of Lemma EC.8. \square

LEMMA EC.9. $J^{\bar{\text{H}}}$ is greater than or equal to the optimal objective value of the following optimization problem:

$$\begin{aligned} &\text{maximize}_{b,w} \quad \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_{T^i}^i - w_{T^i\ell}^i) \\ &\text{subject to} \quad w_{T^i\ell}^i \leq w_{T^i, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\ &\quad \quad \quad b_{T^j}^j \leq w_{T^i\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(T^j, x^i) = \kappa_\ell^i, \\ &\quad \quad \quad \kappa_\ell^i < g(T^i, x^i), T^j < T^i, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\ &\quad \quad \quad b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \\ &\quad \quad \quad w_{T^i\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\}. \end{aligned} \tag{\bar{\text{H}}\text{-LB}}$$

Proof of Lemma EC.9. We begin by rewriting $(\bar{\text{H}})$ as the following equivalent optimization problem:

$$\begin{aligned} &\text{maximize}_{b,w} \quad \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_{T^i}^i - w_{T^i\ell}^i) + \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i < T} \sum_{\ell=1}^{L_{T^i}^i-1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (1 - w_{T\ell}^i) \\ &\text{subject to} \quad w_{t\ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\ &\quad \quad \quad b_{T^i}^i \leq w_{T^i 1}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\ &\quad \quad \quad b_{T^j}^j \leq w_{t\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \mathcal{T}^i \text{ such that } g(T^j, x^i) = \kappa_\ell^i, \\ &\quad \quad \quad T^j \leq t, T^j < T, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\ &\quad \quad \quad b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \\ &\quad \quad \quad w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i|\}. \end{aligned} \tag{\bar{\text{H}}'}$$

We claim that the optimal objective value of the above optimization problem $(\bar{\text{H}}')$ is equal to the optimal objective value of $(\bar{\text{H}})$. To see why this is true, we first observe that the above optimization

problem is identical to (\bar{H}) , except for the fact that the term $(1 - w_{T\ell}^i)$ in the objective function of (\bar{H}) has been replaced with $(b_{T^i}^i - w_{T\ell}^i)$ in the objective function of (\bar{H}') for each $i \in \{1, \dots, N\}$ such that $T^i = T$. Because the new decision variables $b_{T^i}^i \in \{0, 1\}$ for each $i \in \{1, \dots, N\}$ such that $T^i = T$ do not appear in any inequality constraints, and because each term $\kappa_{\ell+1}^i - \kappa_\ell^i$ is strictly positive, we observe that there exists an optimal solution for the above optimization problem in which $b_{T^i}^i = 1$ for each $i \in \{1, \dots, N\}$ such that $T^i = T$. Hence, we conclude that the optimal objective value of (\bar{H}') is equal to the optimal objective value of (\bar{H}) .

We now construct a lower bound approximation of (\bar{H}') by modifying its objective function and adding constraints. Indeed, we first observe for every feasible solution for (\bar{H}') that the objective function of (\bar{H}') satisfies

$$\begin{aligned} & \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_{T^i}^i - w_{T\ell}^i) + \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i < T} \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (1 - w_{T\ell}^i) \\ & \geq \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_{T^i}^i - w_{T\ell}^i), \end{aligned}$$

where the inequality holds because each term $\kappa_{\ell+1}^i - \kappa_\ell^i$ is strictly positive and because the decision variables $w_{T\ell}^i$ are always nonnegative. Therefore, a lower bound on the optimal objective value of (\bar{H}') is given by the optimal objective value of the following optimization problem:

$$\begin{aligned} & \underset{b, w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_{T^i}^i - w_{T\ell}^i) \\ & \text{subject to} && w_{t\ell}^i \leq w_{t, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\ & && b_{T^i}^i \leq w_{T^i}^i \quad \text{for all } i \in \{1, \dots, N\} \text{ such that } T^i < T \\ & && b_{T^j}^j \leq w_{t\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ and } t \in \mathcal{T}^i \text{ such that } g(T^j, x^i) = \kappa_\ell^i, \\ & && \quad T^j \leq t, T^j < T, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\ & && b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \\ & && w_{t\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, t \in \mathcal{T}^i, \ell \in \{1, \dots, |\mathcal{K}^i|\}. \end{aligned} \tag{\bar{H}-LB-1}$$

Moreover, we observe from inspection that the optimal objective value of the above optimization problem would not change if we added the constraints $w_{T\ell}^i = 1$ to the above optimization problem for each sample path $i \in \{1, \dots, N\}$ and $\ell \in \{1, \dots, |\mathcal{K}^i|\}$ that satisfies $T^i < T$. By eliminating these decision variables $w_{T\ell}^i$ that can be constrained without loss of generality to be equal to one, it follows from algebra that the optimal objective value of the optimization problem $(\bar{H}-LB-1)$ is equal to

$$\begin{aligned} & \underset{b, w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) (b_{T^i}^i - w_{T\ell}^i) \\ & \text{subject to} && w_{T^i\ell}^i \leq w_{T^i, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\ & && b_{T^j}^j \leq w_{T^i\ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(T^j, x^i) = \kappa_\ell^i, \\ & && \quad T^j \leq T^i, T^j < T, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\ & && b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \\ & && w_{T^i\ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i|\}. \end{aligned} \tag{\bar{H}-LB-2}$$

We further observe from inspection that the optimal objective value of the above optimization problem would not change if we added the constraints $w_{T^i, |\mathcal{K}^i|}^i = 1$ to the above optimization

problem for each sample path $i \in \{1, \dots, N\}$.¹⁴ By eliminating these decision variables $w_{T^i, |\mathcal{K}^i|}^i$ that can be constrained without loss of generality to be equal to one, and by observing that the above optimization problem has a constraint of the form $b_t^i \leq w_{T^i, \ell}^i$ if and only if $\ell = |\mathcal{K}^i|$ and $t = T^i = T^j$, we conclude that the optimal objective value of the optimization problem ($\bar{\text{H-LB-2}}$) is equal to

$$\begin{aligned}
& \underset{b, w}{\text{maximize}} && \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (b_{T^i}^i - w_{T^i, \ell}^i) \\
& \text{subject to} && w_{T^i, \ell}^i \leq w_{T^i, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\} \\
& && b_{T^j}^j \leq w_{T^i, \ell}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(T^j, x^i) = \kappa_{\ell}^i, \\
& && \quad \kappa_{\ell}^i < g(T^i, x^i), T^j < T^i, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset \\
& && b_{T^i}^i \in \{0, 1\} \quad \text{for all } i \in \{1, \dots, N\} \\
& && w_{T^i, \ell}^i \in \mathbb{R} \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\}.
\end{aligned} \tag{\bar{\text{H-LB}}}$$

This concludes our proof of Lemma EC.9. \square

In view of the above Lemmas EC.8 and EC.9, we now present the proofs of Propositions 4 and 5.

Proof of Proposition 4. In the above Lemma EC.8, we showed that $J^{\text{BP}} \leq \frac{1}{N} \sum_{i=1}^N g(T^i, x^i)$. Therefore, it remains for us to show that $J^{\bar{\text{H}}} \geq \frac{1}{T} (\frac{1}{N} \sum_{i=1}^N g(T^i, x^i))$. Indeed, let

$$t^* \triangleq \arg \max_{t \in \{1, \dots, T\}} \sum_{i \in \{1, \dots, N\}: T^i = t} g(T^i, x^i)$$

be defined as any period that maximizes the sum of the rewards $g(T^i, x^i)$ over all of the sample paths $i \in \{1, \dots, N\}$ that satisfy $T^i = t^*$. We observe from algebra and from our construction of t^* that the following inequality must hold:

$$\frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i = t^*} g(T^i, x^i) \geq \frac{1}{T} \left(\frac{1}{N} \sum_{i=1}^N g(T^i, x^i) \right). \tag{EC.60}$$

Moreover, we recall from the above Lemma EC.9 that $J^{\bar{\text{H}}}$ is greater than or equal to the optimal objective value of the optimization problem ($\bar{\text{H-LB}}$). Consider the solution b, w for the optimization problem ($\bar{\text{H-LB}}$) that satisfies the following equalities for each sample path $i \in \{1, \dots, N\}$, period $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, |\mathcal{K}^i| - 1\}$:

$$b_{T^i}^i = \begin{cases} 1, & \text{if } T^i \in \{t^*, \dots, T\}, \\ 0, & \text{if } T^i \in \{1, \dots, t^* - 1\}; \end{cases} \quad w_{T^i, \ell}^i = \begin{cases} 1, & \text{if } T^i \in \{t^* + 1, \dots, T\}, \\ 0, & \text{if } T^i \in \{1, \dots, t^*\}. \end{cases}$$

We observe from inspection that the solution b, w defined by the above equalities is a feasible solution for the optimization problem ($\bar{\text{H-LB}}$). Therefore, the optimal objective value of ($\bar{\text{H-LB}}$) is greater than or equal to

$$\begin{aligned}
& \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) (\mathbb{I}\{T^i \in \{t^*, \dots, T\}\} - \mathbb{I}\{T^i \in \{t^* + 1, \dots, T\}\}) \\
& = \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) \mathbb{I}\{T^i = t^*\}
\end{aligned}$$

¹⁴ We note that the equality $L_{T^i}^i = |\mathcal{K}^i|$ holds for each sample path $i \in \{1, \dots, N\}$. This equality can be verified by simply applying the definitions of $L_{T^i}^i$, \mathcal{K}^i , and T^i .

$$\begin{aligned}
&= \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i = t^*} \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_{\ell}^i) \\
&= \frac{1}{N} \sum_{i \in \{1, \dots, N\}: T^i = t^*} g(T^i, x^i),
\end{aligned}$$

where the last equality follows from the definition of the constants κ_{ℓ}^i . Combining the above equalities with line (EC.60), our proof of Proposition 4 is complete. \square

Proof of Proposition 5. Let Assumption 5 hold. Our proof of Proposition 5 is split into the following two intermediary claims.

CLAIM EC.12. *Assume without loss of generality that $g(T^1, x^1) \geq \dots \geq g(T^N, x^N)$. Then,*

$$J^{\bar{H}} \geq \max_{i \in \{1, \dots, N\}} \frac{i}{N} g(T^i, x^i) - 2\epsilon L,$$

where ϵ is the radius of the uncertainty sets (see §2.2) and L is the constant from Assumption 5.

Proof of Claim EC.12. Assume without loss of generality that $g(T^1, x^1) \geq \dots \geq g(T^N, x^N)$, and consider any arbitrary $i^* \in \{1, \dots, N\}$. In the remainder of the proof, we will show that

$$J^{\bar{H}} \geq \frac{i^*}{N} g(T^{i^*}, x^{i^*}) - 2\epsilon L.$$

To begin, we recall from the above Lemma EC.9 that $J^{\bar{H}}$ is greater than or equal to the optimal objective value of the optimization problem $(\bar{H}\text{-LB})$. Next, consider the solution b, w for the optimization problem $(\bar{H}\text{-LB})$ that satisfies the following equalities for each sample path $i \in \{1, \dots, N\}$, period $t \in \{1, \dots, T\}$, and $\ell \in \{1, \dots, |\mathcal{K}^i| - 1\}$:

$$b_{T^i}^i = \begin{cases} 1, & \text{if } i \in \{1, \dots, i^*\}, \\ 0, & \text{if } i \in \{i^* + 1, \dots, N\}; \end{cases} \quad w_{T^i \ell}^i = \begin{cases} 1, & \text{if } \kappa_{\ell}^i \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon, \\ 0, & \text{if } T^i \in \{1, \dots, t^*\}. \end{cases}$$

We will now prove that the solution b, w defined by the above equalities is a feasible solution for the optimization problem $(\bar{H}\text{-LB})$. Indeed, we readily observe from inspection that the solution b, w defined by the above equalities satisfies the first set of constraints in the optimization problem $(\bar{H}\text{-LB})$, i.e.,

$$w_{T^i \ell}^i \leq w_{T^i, \ell+1}^i \quad \text{for all } i \in \{1, \dots, N\}, \ell \in \{1, \dots, |\mathcal{K}^i| - 1\}.$$

To show that the solution b, w satisfies the second set of constraints in the optimization problem $(\bar{H}\text{-LB})$, consider any sample paths $i, j \in \{1, \dots, N\}$ such that $g(T^j, x^i) = \kappa_{\ell}^i$, $\kappa_{\ell}^i < g(T^i, x^i)$, $T^j < T^i$, and $\mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset$. We have two cases to consider. First, suppose that the sample path j satisfies $j \in \{i^* + 1, \dots, N\}$. In this case, it follows from the fact that $b_{T^j}^j = 0$ that the inequality $b_{T^j}^j \leq w_{T^i \ell}^i$ is satisfied. Second, suppose that sample path j satisfies $j \in \{1, \dots, i^*\}$. In this case, we observe that

$$g(T^j, x^i) = \kappa_{\ell}^i \geq g(T^j, x^j) - 2\epsilon L \geq g(T^{i^*}, x^{i^*}) - 2\epsilon L,$$

where the first inequality follows from Assumption 5 and from the fact that $\mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset$, and the second inequality follows from the fact that $j \in \{1, \dots, i^*\}$. Therefore, we conclude that from our construction of the solution b, w that $w_{T^i \ell}^i = 1$, which implies that the inequality $b_{T^j}^j \leq w_{T^i \ell}^i$ is

satisfied. We have thus shown that the solution b, w satisfies the second set of constraints in the optimization problem $(\bar{\text{H}}\text{-LB})$, i.e.,

$$b_{T^j}^j \leq w_{T^i}^i \quad \text{for all } i, j \in \{1, \dots, N\} \text{ such that } g(T^j, x^i) = \kappa_\ell^i, \\ \kappa_\ell^i < g(T^i, x^i), T^j < T^i, \text{ and } \mathcal{U}_{T^j}^i \cap \mathcal{U}_{T^j}^j \neq \emptyset.$$

This concludes our proof that the solution b, w is a feasible solution for the optimization problem $(\bar{\text{H}}\text{-LB})$.

Since b, w is a feasible solution for the optimization problem $(\bar{\text{H}}\text{-LB})$, we have the following:

$$J^{\bar{\text{H}}} \geq \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) \left(\mathbb{I}\{i \in \{1, \dots, i^*\}\} - \mathbb{I}\{\kappa_\ell^i \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon\} \right) \quad (\text{EC.61})$$

$$= \frac{1}{N} \sum_{i=1}^{i^*} \sum_{\ell=1}^{L_{T^i}^i - 1} (\kappa_{\ell+1}^i - \kappa_\ell^i) - \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} \mathbb{I}\{\kappa_\ell^i \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon\} (\kappa_{\ell+1}^i - \kappa_\ell^i) \quad (\text{EC.62})$$

$$= \frac{1}{N} \sum_{i=1}^{i^*} g(T^i, x^i) - \frac{1}{N} \sum_{i=1}^N \sum_{\ell=1}^{L_{T^i}^i - 1} \mathbb{I}\{\kappa_\ell^i \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon\} (\kappa_{\ell+1}^i - \kappa_\ell^i) \quad (\text{EC.63})$$

$$\geq \frac{1}{N} \sum_{i=1}^{i^*} g(T^i, x^i) \\ - \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{\exists \ell \in \{1, \dots, L_{T^i}^i - 1\} \text{ such that } \kappa_\ell^i \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon\} \left(g(T^i, x^i) - \left(g(T^{i^*}, x^{i^*}) - 2L\epsilon \right) \right) \quad (\text{EC.64})$$

$$\geq \frac{1}{N} \sum_{i=1}^{i^*} g(T^i, x^i) - \frac{1}{N} \sum_{i=1}^N \mathbb{I}\{g(T^i, x^i) \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon\} \left(g(T^i, x^i) - g(T^{i^*}, x^{i^*}) + 2L\epsilon \right) \quad (\text{EC.65})$$

$$= \frac{1}{N} \sum_{i=1}^{i^*} \left(g(T^{i^*}, x^{i^*}) - 2L\epsilon \right) - \frac{1}{N} \sum_{i \in \{i^*+1, \dots, N\}: g(T^i, x^i) \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon} \left(g(T^i, x^i) - g(T^{i^*}, x^{i^*}) + 2L\epsilon \right) \quad (\text{EC.66})$$

$$\geq \frac{1}{N} \sum_{i=1}^{i^*} \left(g(T^{i^*}, x^{i^*}) - 2L\epsilon \right) - \frac{1}{N} \sum_{i \in \{i^*+1, \dots, N\}: g(T^i, x^i) \geq g(T^{i^*}, x^{i^*}) - 2L\epsilon} 2L\epsilon \quad (\text{EC.67})$$

$$\geq \frac{1}{N} \sum_{i=1}^{i^*} \left(g(T^{i^*}, x^{i^*}) - 2L\epsilon \right) - \frac{N - i^*}{N} 2L\epsilon \quad (\text{EC.68})$$

$$= \frac{i^*}{N} g(T^{i^*}, x^{i^*}) - 2\epsilon L, \quad (\text{EC.69})$$

Indeed, (EC.61) holds because $J^{\bar{\text{H}}}$ is greater than or equal to the optimal objective value of the optimization problem $(\bar{\text{H}}\text{-LB})$ and because b, w is a feasible but possibly suboptimal solution for the optimization problem $(\bar{\text{H}}\text{-LB})$. Lines (EC.62), (EC.63), and (EC.64) follow from algebra. Line (EC.65) follows from the fact that $\kappa_0^i < \dots < \kappa_{L_{T^i}^i}^i$. Line (EC.66) follows from rearranging terms. Line (EC.67) follows from the fact that $g(T^i, x^i) - g(T^{i^*}, x^{i^*}) \leq 0$ for all $i \in \{i^* + 1, \dots, N\}$. Lines (EC.68) and (EC.69) follow from algebra.

Since $i^* \in \{1, \dots, N\}$ was chosen arbitrarily, our proof of Claim EC.12 is complete. \square

CLAIM EC.13. $J^{\bar{H}} \geq \frac{1}{\log N + 1} \left(\frac{1}{N} \sum_{i=1}^N g(T^i, x^i) \right) - 2\epsilon L$.

Proof of Claim EC.13. Assume without loss of generality that $g(T^1, x^1) \geq \dots \geq g(T^N, x^N)$. In Claim EC.12, we showed that

$$J^{\bar{H}} \geq \max_{i \in \{1, \dots, N\}} \frac{i}{N} g(T^i, x^i) - 2\epsilon L.$$

For notational convenience, let us define $\gamma^i \triangleq \frac{i}{N} g(T^i, x^i)$ for each sample path $i \in \{1, \dots, N\}$. With this notation, we observe from algebra that

$$\begin{aligned} \frac{1}{\log N + 1} \left(\frac{1}{N} \sum_{i=1}^N g(T^i, x^i) \right) - 2\epsilon L &= \frac{1}{\log N + 1} \left(\sum_{i=1}^N \frac{1}{i} \gamma^i \right) - 2\epsilon L \\ &\leq \frac{1}{\log N + 1} \left(\max_{i \in \{1, \dots, N\}} \gamma^i \right) \sum_{i=1}^N \frac{1}{i} - 2\epsilon L \\ &\leq \max_{i \in \{1, \dots, N\}} \gamma^i - 2\epsilon L \\ &= \max_{i \in \{1, \dots, N\}} \frac{i}{N} g(T^i, x^i) - 2\epsilon L \\ &\leq J^{\bar{H}}, \end{aligned}$$

where the last line follows from Claim EC.12. This completes the proof of Claim EC.13. \square

Combining Claim EC.13 with Lemma EC.8 completes our proof of Proposition 5. \square

Appendix I: Additional Numerical Results

In this appendix, we present numerical results for additional parameter settings which were omitted from §5.2 due to length considerations.

Table EC.1 Barrier Option (Asymmetric) - Expected Reward.

d	Method	Basis functions	Initial Price			# of Sample Paths
			$\bar{x} = 90$	$\bar{x} = 100$	$\bar{x} = 110$	
8	RO	maxprice	57.53 (0.16)	67.25 (0.27)	69.10 (0.42)	10^3 training, 10^3 validation
8	LS	one, pricesKO, KOind, payoff	56.44 (0.11)	65.52 (0.16)	68.34 (0.09)	10^5
8	LS	one, pricesKO, payoff	56.54 (0.12)	65.74 (0.17)	68.31 (0.08)	10^5
8	LS	payoff, KOind, pricesKO	56.44 (0.11)	65.52 (0.16)	68.34 (0.09)	10^5
8	LS	pricesKO, payoff	56.51 (0.12)	65.75 (0.18)	68.31 (0.08)	10^5
8	LS	one, prices, payoff	55.56 (0.09)	61.92 (0.09)	60.07 (0.12)	10^5
8	LS	one	45.80 (0.06)	52.84 (0.07)	48.56 (0.08)	10^5
8	LS	one, KOind, prices	52.33 (0.09)	61.85 (0.10)	63.88 (0.08)	10^5
8	LS	one, prices	48.49 (0.11)	53.03 (0.09)	50.83 (0.06)	10^5
8	LS	one, pricesKO	51.67 (0.08)	61.49 (0.10)	63.14 (0.09)	10^5
8	LS	maxprice, KOind, pricesKO	52.63 (0.09)	61.97 (0.10)	64.14 (0.12)	10^5
8	PO	payoff, KOind, pricesKO	53.61 (0.18)	58.29 (0.36)	52.96 (0.30)	2×10^3 outer, 500 inner
8	PO	prices	53.90 (0.16)	60.66 (0.28)	58.14 (0.12)	2×10^3 outer, 500 inner
8	Tree	payoff, time	55.75 (0.13)	62.78 (0.31)	63.95 (0.52)	10^5
8	Tree	maxprice, time	55.75 (0.13)	62.78 (0.31)	63.95 (0.52)	10^5
16	RO	maxprice	72.56 (0.15)	77.44 (0.29)	69.74 (0.30)	10^3 training, 10^3 validation
16	LS	one, pricesKO, KOind, payoff	70.94 (0.08)	76.10 (0.10)	70.59 (0.13)	10^5
16	LS	one, pricesKO, payoff	71.06 (0.08)	76.14 (0.08)	70.05 (0.12)	10^5
16	LS	payoff, KOind, pricesKO	70.94 (0.08)	76.10 (0.10)	70.59 (0.13)	10^5
16	LS	pricesKO, payoff	71.06 (0.08)	76.14 (0.08)	70.04 (0.12)	10^5
16	LS	one, prices, payoff	67.57 (0.08)	67.07 (0.16)	56.36 (0.14)	10^5
16	LS	one	59.09 (0.05)	56.20 (0.11)	45.55 (0.05)	10^5
16	LS	one, KOind, prices	66.78 (0.06)	71.06 (0.14)	63.54 (0.20)	10^5
16	LS	one, prices	59.35 (0.07)	57.68 (0.08)	49.67 (0.07)	10^5
16	LS	one, pricesKO	66.51 (0.05)	70.70 (0.13)	62.53 (0.16)	10^5
16	LS	maxprice, KOind, pricesKO	67.00 (0.06)	71.27 (0.13)	64.04 (0.19)	10^5
16	PO	payoff, KOind, pricesKO	65.14 (0.31)	62.12 (0.34)	47.16 (0.37)	2×10^3 outer, 500 inner
16	PO	prices	66.69 (0.18)	66.23 (0.23)	53.84 (0.26)	2×10^3 outer, 500 inner
16	Tree	payoff, time	68.88 (0.15)	71.11 (0.19)	55.05 (0.11)	10^5
16	Tree	maxprice, time	68.88 (0.15)	71.11 (0.19)	55.05 (0.11)	10^5
32	RO	maxprice	84.12 (0.28)	79.14 (0.40)	60.66 (0.57)	10^3 training, 10^3 validation
32	LS	one, pricesKO, KOind, payoff	82.38 (0.09)	79.17 (0.10)	62.92 (0.15)	10^5
32	LS	one, pricesKO, payoff	82.45 (0.10)	78.91 (0.10)	62.27 (0.14)	10^5
32	LS	payoff, KOind, pricesKO	82.38 (0.09)	79.17 (0.10)	62.92 (0.15)	10^5
32	LS	pricesKO, payoff	82.45 (0.10)	78.91 (0.10)	62.27 (0.14)	10^5
32	LS	one, prices, payoff	73.46 (0.19)	62.81 (0.17)	48.79 (0.17)	10^5
32	LS	one	64.26 (0.12)	50.78 (0.09)	44.32 (0.05)	10^5
32	LS	one, KOind, prices	77.53 (0.11)	71.06 (0.12)	56.06 (0.21)	10^5
32	LS	one, prices	64.63 (0.11)	55.21 (0.12)	46.58 (0.05)	10^5
32	LS	one, pricesKO	77.36 (0.10)	70.44 (0.16)	55.11 (0.24)	10^5
32	LS	maxprice, KOind, pricesKO	77.67 (0.11)	71.53 (0.13)	56.30 (0.22)	10^5
32	PO	payoff, KOind, pricesKO	70.79 (0.49)	55.67 (0.31)	35.36 (1.21)	2×10^3 outer, 500 inner
32	PO	prices	73.86 (0.32)	62.09 (0.22)	44.23 (0.29)	2×10^3 outer, 500 inner
32	Tree	payoff, time	77.64 (0.44)	63.91 (0.13)	50.95 (0.04)	10^5
32	Tree	maxprice, time	77.64 (0.44)	63.91 (0.13)	50.95 (0.04)	10^5

Optimal is indicated in bold for each number of assets $d \in \{8, 16, 32\}$ and initial price $\bar{x} \in \{90, 100, 110\}$. Problem parameters are $T = 54$, $Y = 3$, $r = 0.05$, $K = 100$, $B_0 = 150$, $\delta = 0.25$, $\sigma_a = 0.1 + s \frac{\alpha \times d}{5}$, $\rho_{a,a'} = 0$ for all $a \neq a'$.

Table EC.2 Barrier Option (Asymmetric) - Computation Times.

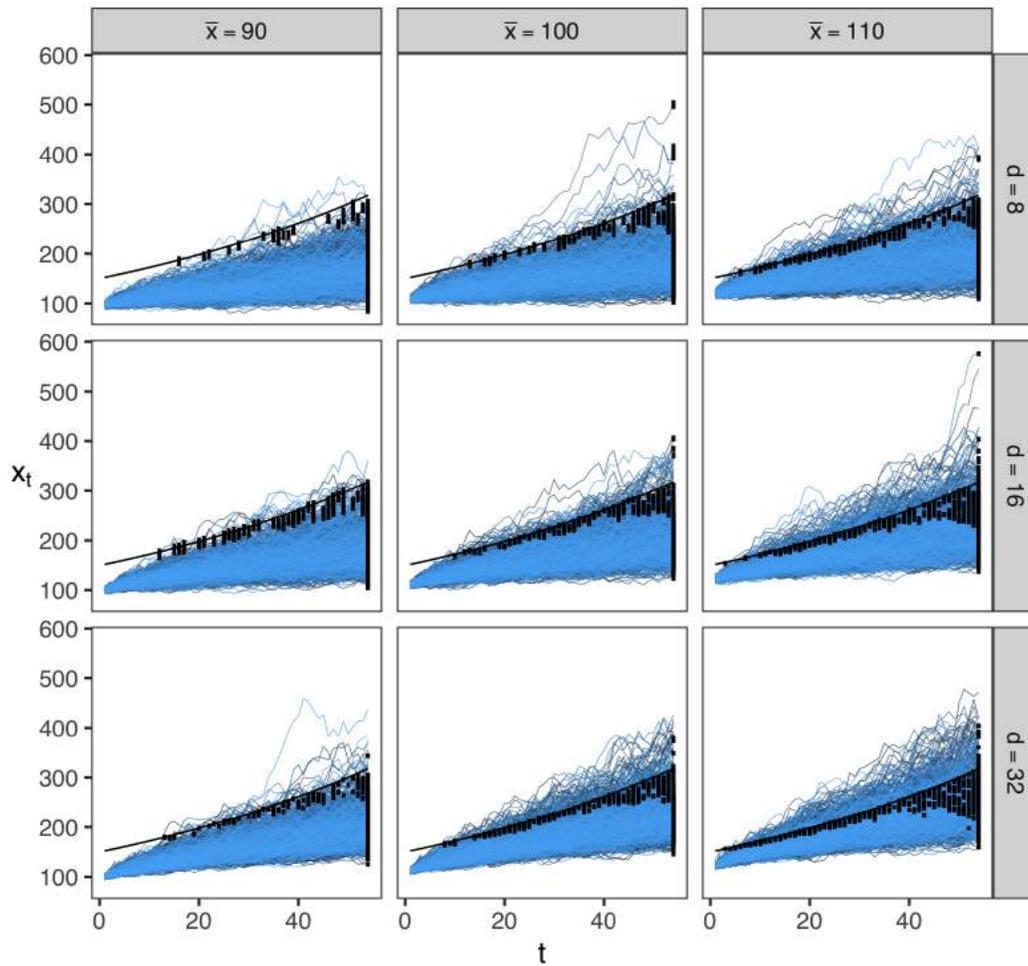
d	Method	Basis functions	Initial Price			# of Sample Paths
			$\bar{x} = 90$	$\bar{x} = 100$	$\bar{x} = 110$	
8	RO	maxprice	7.19 (0.37)	8.62 (1.09)	18.38 (4.59)	10^3 training, 10^3 validation
8	LS	one, pricesKO, KOind, payoff	3.29 (0.2)	3.20 (0.31)	3.24 (0.15)	10^5
8	LS	one, pricesKO, payoff	3.16 (0.2)	3.05 (0.25)	3.13 (0.21)	10^5
8	LS	payoff, KOind, pricesKO	3.76 (0.63)	3.29 (0.58)	3.05 (0.44)	10^5
8	LS	pricesKO, payoff	3.03 (0.14)	2.89 (0.23)	2.87 (0.12)	10^5
8	LS	one, prices, payoff	3.31 (0.56)	2.97 (0.29)	2.90 (0.24)	10^5
8	LS	one	0.89 (0.07)	0.93 (0.06)	0.94 (0.08)	10^5
8	LS	one, KOind, prices	3.39 (0.55)	2.99 (0.26)	2.98 (0.21)	10^5
8	LS	one, prices	2.63 (0.19)	2.71 (0.22)	2.72 (0.19)	10^5
8	LS	one, pricesKO	2.78 (0.35)	2.62 (0.18)	2.53 (0.21)	10^5
8	LS	maxprice, KOind, pricesKO	3.01 (0.13)	2.85 (0.26)	2.87 (0.08)	10^5
8	PO	payoff, KOind, pricesKO	18.75 (0.71)	18.12 (0.41)	16.23 (0.71)	2×10^3 outer, 500 inner
8	PO	prices	23.31 (0.45)	23.57 (0.36)	22.69 (0.29)	2×10^3 outer, 500 inner
8	Tree	payoff, time	13.80 (0.28)	15.72 (5.51)	57.70 (10.81)	10^5
8	Tree	maxprice, time	13.47 (0.19)	15.04 (5.42)	55.45 (10.47)	10^5
16	RO	maxprice	6.80 (1)	10.15 (1.4)	51.58 (9.06)	10^3 training, 10^3 validation
16	LS	one, pricesKO, KOind, payoff	5.55 (0.61)	5.30 (0.34)	5.01 (0.42)	10^5
16	LS	one, pricesKO, payoff	5.50 (0.78)	5.23 (0.33)	5.24 (0.33)	10^5
16	LS	payoff, KOind, pricesKO	4.74 (0.72)	4.62 (0.54)	4.55 (0.41)	10^5
16	LS	pricesKO, payoff	5.29 (0.37)	5.18 (0.38)	4.87 (0.26)	10^5
16	LS	one, prices, payoff	5.57 (0.55)	4.93 (0.29)	5.10 (0.32)	10^5
16	LS	one	1.17 (0.11)	1.26 (0.08)	1.24 (0.05)	10^5
16	LS	one, KOind, prices	6.21 (1.46)	5.72 (0.76)	5.16 (0.51)	10^5
16	LS	one, prices	6.02 (1.63)	5.11 (1.25)	4.86 (0.73)	10^5
16	LS	one, pricesKO	4.61 (0.62)	4.58 (0.52)	4.67 (0.58)	10^5
16	LS	maxprice, KOind, pricesKO	5.11 (0.64)	4.90 (0.6)	4.73 (0.27)	10^5
16	PO	payoff, KOind, pricesKO	32.13 (1.23)	28.51 (0.96)	25.56 (0.48)	2×10^3 outer, 500 inner
16	PO	prices	43.73 (0.61)	43.75 (0.67)	41.84 (0.65)	2×10^3 outer, 500 inner
16	Tree	payoff, time	14.24 (0.4)	48.39 (0.94)	7.89 (0.09)	10^5
16	Tree	maxprice, time	13.35 (0.3)	47.49 (0.74)	8.31 (0.19)	10^5
32	RO	maxprice	7.88 (1.21)	22.32 (7.25)	154.12 (29.91)	10^3 training, 10^3 validation
32	LS	one, pricesKO, KOind, payoff	15.07 (2.97)	11.55 (1.37)	13.68 (2.69)	10^5
32	LS	one, pricesKO, payoff	15.29 (2.15)	12.20 (1.68)	13.04 (3.25)	10^5
32	LS	payoff, KOind, pricesKO	13.13 (2.59)	10.96 (1.81)	11.30 (1.81)	10^5
32	LS	pricesKO, payoff	15.11 (3.27)	11.39 (1.91)	11.62 (1.44)	10^5
32	LS	one, prices, payoff	16.78 (4.89)	10.68 (1.52)	11.92 (3.11)	10^5
32	LS	one	1.62 (0.6)	1.51 (0.16)	1.69 (0.74)	10^5
32	LS	one, KOind, prices	14.61 (2.44)	10.77 (0.79)	11.48 (1.94)	10^5
32	LS	one, prices	12.49 (3.23)	10.42 (0.57)	10.47 (0.84)	10^5
32	LS	one, pricesKO	13.61 (3.66)	10.55 (0.43)	11.02 (1.31)	10^5
32	LS	maxprice, KOind, pricesKO	12.63 (0.86)	10.82 (1.36)	11.60 (1.09)	10^5
32	PO	payoff, KOind, pricesKO	61.15 (1.43)	52.57 (4.38)	52.56 (1.89)	2×10^3 outer, 500 inner
32	PO	prices	108.40 (2.23)	108.40 (3.02)	97.29 (2.01)	2×10^3 outer, 500 inner
32	Tree	payoff, time	44.24 (19.71)	15.35 (0.31)	9.41 (0.32)	10^5
32	Tree	maxprice, time	43.92 (20.08)	16.22 (0.62)	10.92 (0.4)	10^5

Problem parameters are $T = 54$, $Y = 3$, $r = 0.05$, $K = 100$, $B_0 = 150$, $\delta = 0.25$, $\sigma_a = 0.1 + s \frac{\alpha \times d}{5}$, $\rho_{a,a'} = 0$ for all $a \neq a'$.

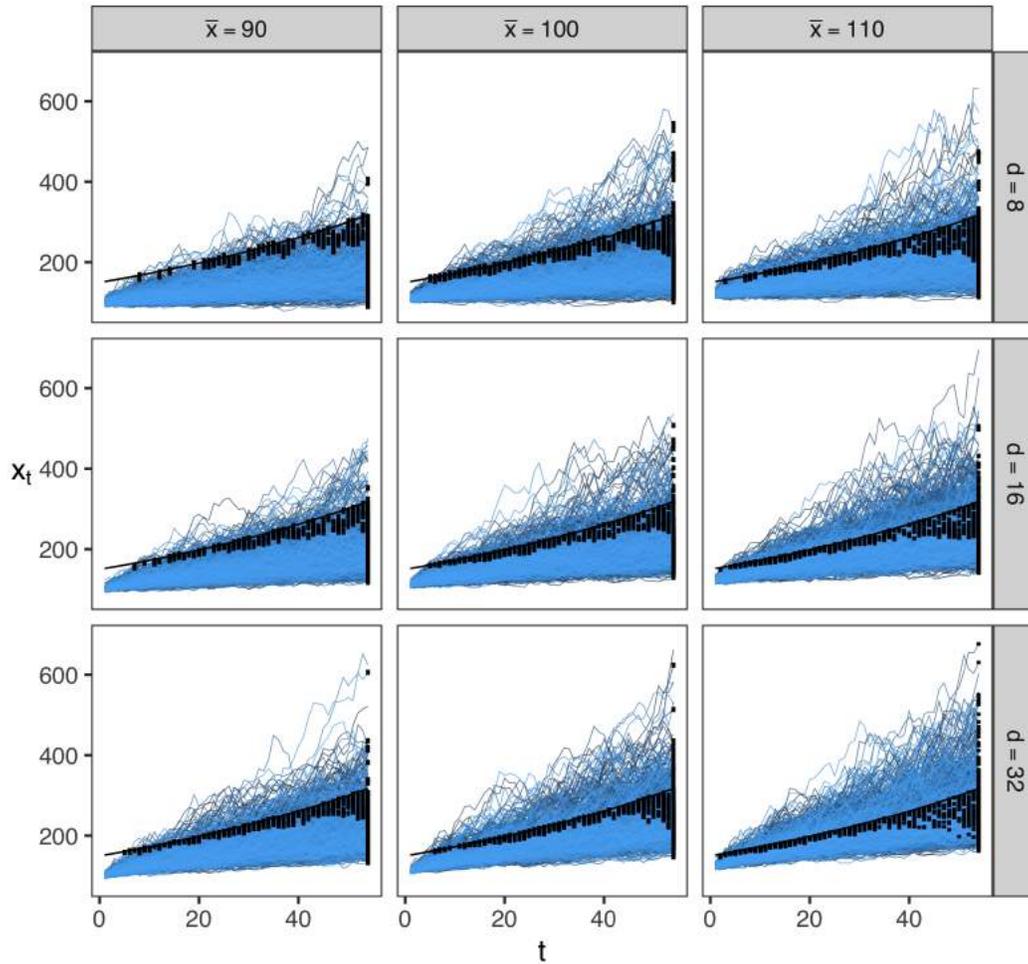
Table EC.3 Barrier Option
(Asymmetric) - Best Choice of Robustness
Parameter.

d	Initial Price		
	$\bar{x} = 90$	$\bar{x} = 100$	$\bar{x} = 110$
8	8.2 (1.75)	5.9 (1.73)	4.0 (1.33)
16	6.2 (1.69)	4.6 (1.26)	3.4 (1.17)
32	5.5 (1.27)	4.0 (1.05)	2.3 (0.95)

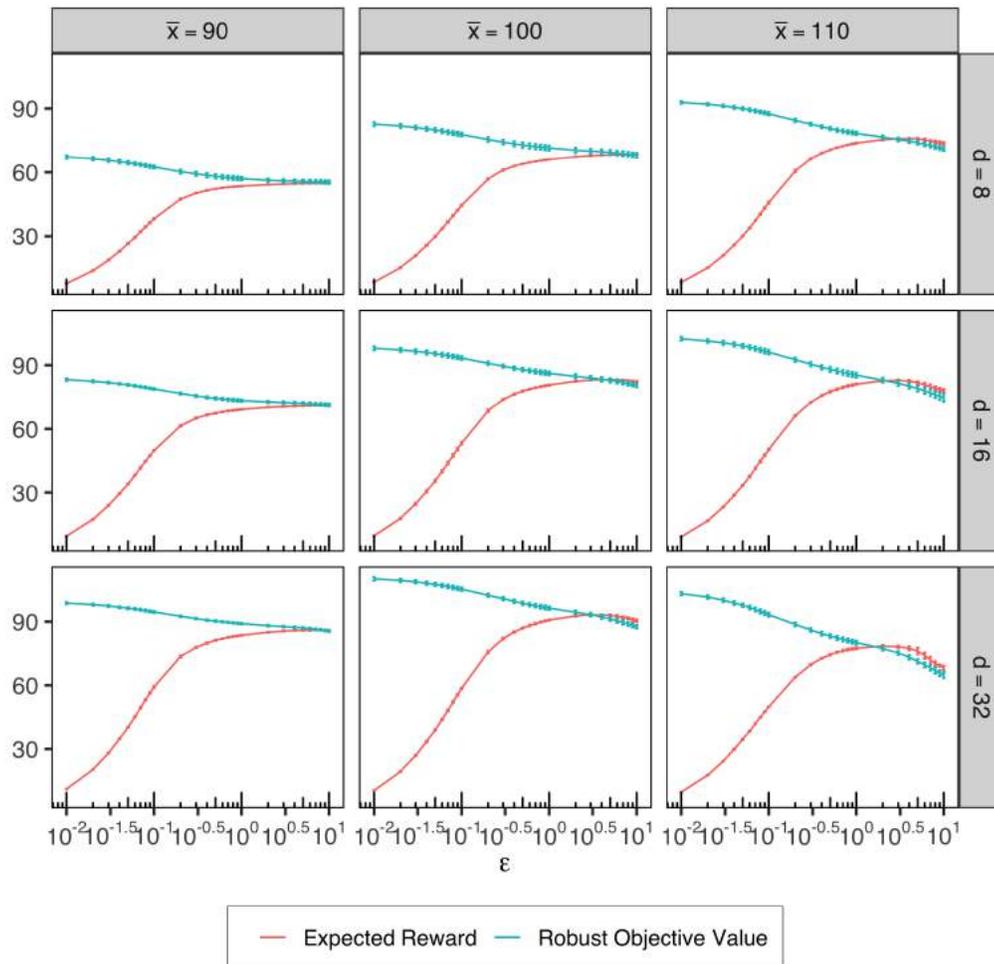
Best choice of robustness parameter ϵ found using the validation method from §2.5 in the robust optimization problems constructed from a training dataset of size $N = 10^3$ and validation set of size $\bar{N} = 10^3$. The remaining parameters are the same as those shown in Table EC.1.

Figure EC.1 Barrier Option (Symmetric) - Visualization of Robust Optimization Stopping Rules.

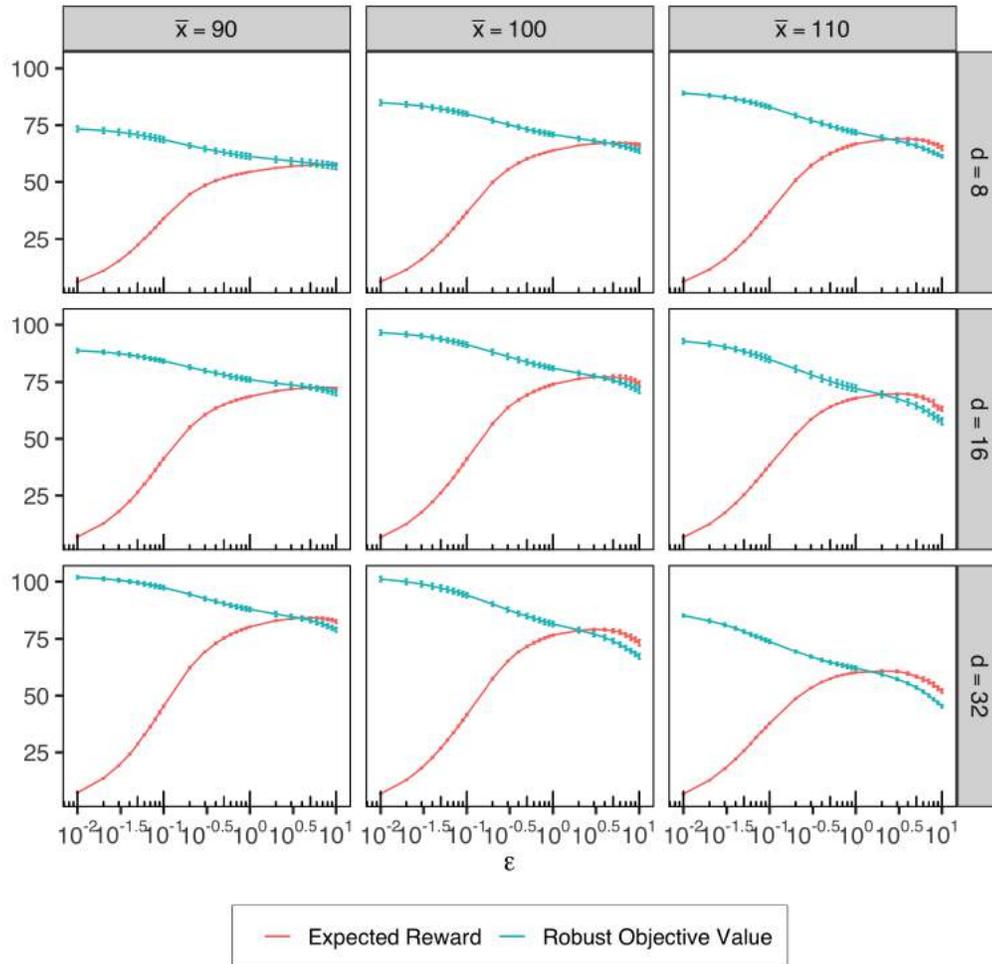
Note. Each plot shows the exercise policies obtained from solving the heuristic (\bar{H}) constructed from a training dataset of size $N = 10^3$ and with the robustness parameter selected using a validation set of size $\bar{N} = 10^3$. The problem parameters are the same as those shown in Table 1.

Figure EC.2 Barrier Option (Asymmetric) - Visualization of Robust Optimization Stopping Rules.

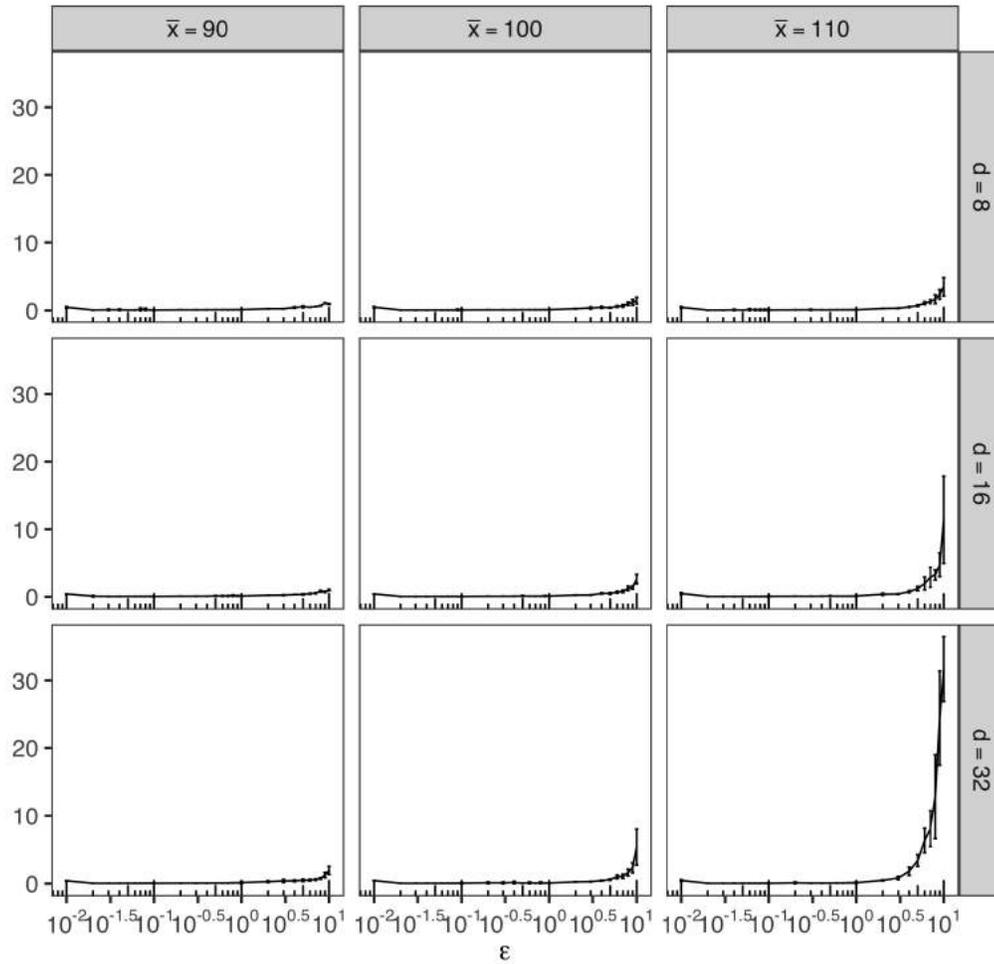
Note. Each plot shows the exercise policies obtained from solving the heuristic (\bar{H}) constructed from a training dataset of size $N = 10^3$ and with the robustness parameter selected using a validation set of size $\bar{N} = 10^3$. The problem parameters are the same as those shown in Table EC.1.

Figure EC.3 Barrier Option (Symmetric) - Impact of Robustness Parameter on Reward.

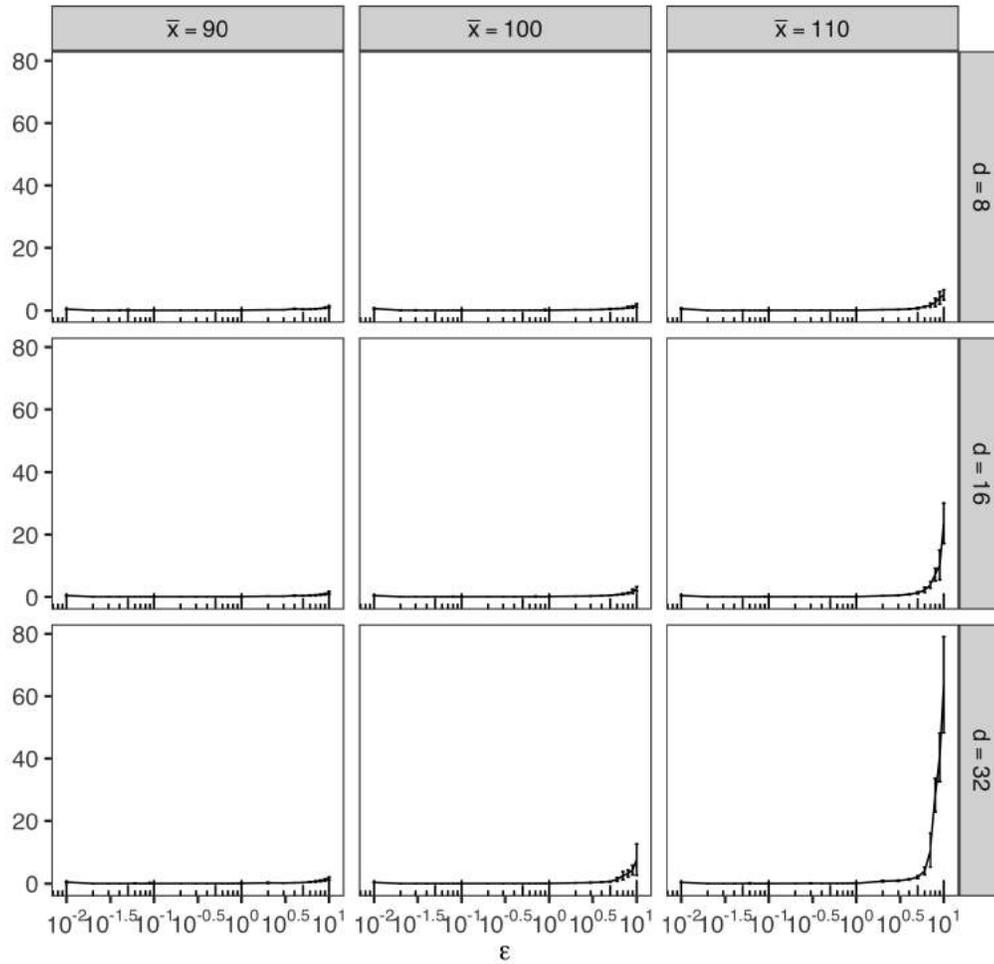
Note. Each plot shows the robust objective value and expected reward of policies obtained by solving the heuristic (\bar{H}) constructed from training datasets of size $N = 10^3$. The problem parameters are the same as those shown in Table 1.

Figure EC.4 Barrier Option (Asymmetric) - Impact of Robustness Parameter on Reward.

Note. Each plot shows the robust objective value and expected reward of policies obtained by solving the heuristic (\bar{H}) constructed from training datasets of size $N = 10^3$. The problem parameters are the same as those shown in Table EC.1.

Figure EC.5 Barrier Option (Symmetric) - Impact of Robustness Parameter on Computation Time.

Note. Each plot shows the computation times from solving the heuristic (\bar{H}) constructed from training datasets of size $N = 10^3$. The problem parameters are the same as those shown in Table 1.

Figure EC.6 Barrier Option (Asymmetric) - Impact of Robustness Parameter on Computation Time.

Note. Each plot shows the computation times from solving the heuristic (\bar{H}) constructed from training datasets of size $N = 10^3$. The problem parameters are the same as those shown in Table EC.1.