

Using Stochastic Games for Learning to Control and Stop Optimally in Worst-Case Scenarios

David Mguni
PROWLER.io
Cambridge, UK
davidmg@proowler.io

ABSTRACT

Recently, there has been a surge in interest in safe and robust techniques within reinforcement learning (RL). Current notions of risk in RL fail to capture the potential for systemic failures such as abrupt stoppages from system failures or surpassing of safety thresholds and the appropriate responsive controls in such instances. We propose a novel approach to risk minimisation within RL in which, in addition to taking actions that maximise its expected return, the controller learns a policy that is robust against stoppages due to an adverse event such as an abrupt failure. The results of the paper cover fault-tolerant control in *worst-case scenarios* under random stopping and optimal stopping, all in unknown environments. By demonstrating that the class of problems is represented by a variant of stochastic games, we prove the existence of a solution which is a unique fixed point equilibrium of the game and characterise the optimal controller behaviour. We then introduce a value function approximation algorithm that converges to the solution through simulation in unknown environments.

KEYWORDS

Robust optimal stopping, reinforcement learning, stochastic game

1 INTRODUCTION

A significant amount of focus within reinforcement learning (RL) is now being placed on safe execution, robust control, fault-tolerance and risk-minimisation [12]. Driving this interest is an increase in application of RL in real-world environments and industrial applications such as traffic light control [2], robotics [9], autonomous vehicles [24] and healthcare [13]. Applying RL in various environments requires safe operation of autonomous agents is ensured. At present however, such frameworks within RL are restricted to models in which the agent modifies the state process using an expectation measure which is altered to accommodate a predefined notion of risk e.g. H_∞ control [21]. Other notions of risk include coherent risk, conditional value at risk (CVar) [27].

There are numerous instances in which controllers are required to act in systems which suffer the potential for random stoppages or failures that produce catastrophic outcomes [12]. Examples include, in finance, optimal trading under random counterparty risk [16], and in control systems, optimal robotic control with random sensor failure and helicopter control under engine failure [1]. Consequently, when using RL in environments that present potential issues of safety, the important question of how to control the system in a way that is robust against faults that lead to catastrophic events arises. An additional issue in matters of safety is when to optimally stop the system with concern for risk of adverse events (e.g. when

to sell all asset holdings with concern for financial ruin). Despite its importance however, current notions of risk do not offer a method of mitigating risk by selectively stopping the system.

To this end, we for the first time construct a method that enables an RL controller to determine an optimal sequence of actions that is robust against failures that lead to adverse events. In order to find the optimal control policy, it is necessary to determine a stopping criterion that stops the system which produces a worst-case scenario. Secondly, we construct a method that enables an RL agent to determine when to stop the system in order to maximise its expected payoff in the presence of adverse risk. As we show, each problem admits a two-player stochastic game (SG) representation in which one of the players is delegated the task of modifying the system dynamics through its actions and the other player has the task of stopping the system under an adversarial criterion.

We perform a formal analysis of an SG between a ‘controller’ and a ‘stopper’. Under this interpretation, the outcome is determined by a controller that affects the state process through its actions whilst playing against an adversary that has the right to choose when to stop the game. This produces a framework that finds an optimal sequence of actions that is robust against stoppages at times that pose adverse risk. The notion of risk is defined in the *worst-case scenario* sense — given the complete set of probability distributions, the agent considers the worst-case in assessing the expected payoff.

These results tackle optimal stopping problems (OSPs) under worst-case scenarios. OSPs are a subclass of optimal stochastic control (OSC) problems in which the goal is to determine a criterion for stopping the system at a time that maximises some state-dependent payoff [22]. Despite the fundamental relevance of risk in RL, current iterative methods in OSPs in unknown environments are restricted to risk-neutral settings [28] and do not permit the inclusion of a controller that modifies the dynamics. Introducing a notion of risk (generated adversarially) adds considerable difficulty as the solution concept is now an *SG saddle point equilibrium*, the existence of which must be established.

As we show, our framework provides an iterative method of solving worst-case scenario OSPs in unknown environments. The framework is developed through a series of theoretical results: first, we establish the existence of a value of the game which characterises the payoff for the (saddle point) equilibrium. Second, we prove a contraction mapping property of a Bellman operator of the game and that the value is a unique fixed point of the operator. Third, we prove the existence and characterise the optimal stopping time. We then prove an equivalence between the game of control and stopping and worst-case OSPs and show that the fixed point solution of the game solves the OSP.

Finally, using an approximate dynamic programming method, we develop a simulation-based iterative scheme that computes the optimal controls. The method applies in settings in which neither the system dynamics nor the reward function are known. Hence, the agent need only observe its realised rewards by interacting with the environment.

1.1 Related Work

Presently, the coverage of FT within RL is extremely limited. In [32] RL is *applied* to tackle systems in which faults might occur with the occurrence of a fault incurring a large cost. Similarly, RL is applied to a problem in [30] in which an RL method for Bayesian discrimination which is used to segment the state and action spaces. Unlike these methods in which infrequent faults from the environment generate negative feedback, our method uses a game-theoretic framework in order to simulate faults leading to an FT trained RL policy.

Our main results are centered around a minimax proof that establishes the existence of a value of the game. This is necessary for simulating the stopping action to induce fault-tolerance. Although minimax proofs are well-known in game theory [11, 19, 25], replacing a player’s action set with stopping times necessitates a minimax proof (which now relies on a construction of open sets) which markedly differs to the standard methods within game theory. Additionally, crucial to our analysis is the characterisation of the adversary optimal stopping time (Theorem 5.15).

A relevant framework is a two-player optimal stopping game (Dynkin game) in which each player chooses one of two actions; to stop the game or continue [10]. Dynkin games have generated a vast literature since the setting requires a markedly different analysis from standard stochastic game theory. In the case with one stopper and one controller such as we are concerned with, the minimax proof requires a novel construction using open sets to cope with the stopping problem for the minimax result.

Presently, the study of optimal control that combines control and stopping is limited to a few studies e.g. [8]. Similarly, games of control and stopping have been analysed in continuous-time in specific contexts e.g. linear diffusions [17], geometric Brownian motion [4] and jump-diffusions [3, 20]. In these analyses, all aspects of the environment are known and the controller affects the dynamics of a *continuous diffusion process*. In general, under these methods, solving these problems requires computing analytic solutions to non-linear partial differential equations which are typically insoluble.

There is a plethora of work on OSPs in continuous and discrete-time [22]. [28] use approximate dynamic programming methods to construct an iterative scheme to compute the solution of an OSP. Our results generalise existing analyses to strategic settings with both a controller and an adversarial stopper which tackles risk in OSPs.

1.2 Organisation

The paper is organised as follows: in Section 2, we introduce some relevant mathematical preliminaries and give a canonical description of both the fault-tolerant RL problem and the OSP under worst-case scenarios. In Section 3, we provide illustrative examples for each problem within the context of finance and RL. In Section 4,

we introduce the underlying SG framework which we use within the main theoretical analysis which we perform in Section 5. Lastly in Section 6, we develop an approximate dynamic programming approach that enables the optimal controls to be computed through simulation, followed by some concluding remarks.

2 CANONICAL DESCRIPTION

In this setting, the state of the system is determined by a stochastic process $\{s_t | t = 0, 1, 2, \dots\}$ whose values are drawn from a state space $\mathcal{S} \subseteq \mathbb{R}^p$ for some $p \in \mathbb{N}$. The state space is defined on a probability space (Ω, \mathcal{B}, P) , where Ω is the sample space, \mathcal{B} is the set of events and P is a map from events to probabilities. We denote by $\mathcal{F} = (\mathcal{F}_n)_{n \geq 0}$ the filtration over (Ω, \mathcal{B}, P) which is an increasing family of σ -algebras generated by the random variables s_1, s_2, \dots

We operate in a Hilbert space \mathcal{V} of real-valued functions on \mathbb{L}_2 , i.e. a complete¹ vector space which we equip with a norm $\|\cdot\| : \mathcal{V} \rightarrow \mathbb{R}_{>0} \times \{0\}$ given by $\|f\|_\mu := \sqrt{\mathbb{E}_\mu[f^2(s)]}$ and its inner product $\langle f, f^T \rangle_\mu := \mathbb{E}_\mu[f(s)f^T(s)]$ where $\mu : \mathcal{B}(\mathbb{R}^n) \rightarrow [0, 1]$ is a probability measure. The problem occurs over a time interval $[0, K]$ where $K \in \mathbb{N} \times \{\infty\}$ is the time horizon. A **stopping time** is defined as a random variable $\tau \in \{0, 1, 2, \dots\}$ for which $\{\omega \in \Omega | \tau(\omega) \leq t\} \in \mathcal{F}_t$ for any $t \in [0, K]$ — this says that given the information generated by the state process, we can determine if the stopping criterion has occurred.

We now describe the two problems with which we are concerned that is, FT RL and OSPs under worst-case scenarios. We later prove an equivalence between the two problems and characterise the solution of each problem.

2.1 Fault-Tolerant Reinforcement Learning

We concern ourselves with finding control policy that copes with abrupt system stoppages and failures at the worst times in problems. In the current setting, the reward and transition functions are assumed a priori unknown. Unlike standard methods in RL and game theory that have fixed time horizons (or purely random exit times) in the following, the process is stopped by a fictitious adversary that uses a stopping strategy to decide when to stop given its observations of the state. In order to generate an FT control, we simulate the adversary’s action whilst the controller seeks to determine its optimal policy. This as we show, induces a form of control that is an FT best-response control.

A formal description is as follows: an agent exercises actions that influence the sequence of states visited by the system. At each state, the agent receives a reward which is dependent on the state. The agent’s actions are selected by a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which is a map from the set of states \mathcal{S} and the set of actions \mathcal{A} to a probability. We assume that the action set is a discrete compact set and that the agent’s policy π is drawn from a compact policy set Π . The horizon of the problem is at most T but the process may be terminated earlier at some (\mathcal{F} -measurable) stopping time at which point the agent receives a terminal reward.

¹A vector space is complete if it contains the limit points of all its Cauchy sequences.

The agent’s **performance function** is given by:

$$J^{k,\pi}[s] = \mathbb{E} \left[\sum_{t=0}^{k \wedge T} \gamma^t R(s_t, a_t) + \gamma^{k \wedge T} G(s_{k \wedge T}) \middle| s_0 = s \right], \quad (1)$$

where $a \wedge b := \min\{a, b\}$, \mathbb{E} is taken w.r.t. the transition function P and the controller’s policy $\pi \in \Pi$. The performance function (1) consists of a **reward function** $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ which quantifies the agent’s immediate reward when the system transitions from one state to the next, a **bequest function** $G : \mathcal{S} \rightarrow \mathbb{R}$ which quantifies the reward received by the agent when the system is stopped and $\gamma \in [0, 1[$, a discount factor. We assume R and G are bounded and measurable.

The problem we consider is the following:

Find $(\hat{k}, \hat{\pi}) \in \mathcal{T} \times \Pi$ and $J^{\hat{k}, \hat{\pi}}[s]$ s.th.

$$\max_{\pi \in \Pi} \left(\min_{k \in \mathcal{T}} J^{k,\pi}[s] \right) = J^{\hat{k}, \hat{\pi}}[s], \quad (2)$$

where the minimisation is taken pointwise and $\mathcal{T} \subseteq \{0, 1, 2, \dots\}$ is a set of (\mathcal{F} -measurable) stopping times.

Hereon, we employ the following shorthand $R(s, a) \equiv R_s^a$ for any $s \in \mathcal{S}, a \in \mathcal{A}$.

The dual objective problem in (2) consists of finding both a stopping time that minimises J and an optimal policy that maximises J . By considering the tasks as being delegated to two individual *players*, the problem becomes an SG. The SG occurs between a ‘controller’ that seeks to maximise J by manipulating state visitations through its actions and an adversary or ‘stopper’ that chooses a time to stop the process to minimise J (i.e. at the worst possible time). The structure of the game combines an OSP and a Markov decision process (MDP). We consider a setting in which neither player has up-front knowledge of the transition model or objective function but each only observes their realised rewards.

2.2 Robust Optimal Stopping

The second problem we consider is robust optimal stopping. In OSPs, the goal is to determine a criterion for stopping the system at a time that maximises some state-dependent payoff. OSPs are ubiquitous in finance e.g. for options pricing [23] and in economics for characterising optimal market entry/exit strategies [18]. OSPs are closely related to multi-armed bandits and clinical trials [15].

OSP in worst-case scenarios regularly arise in economic decision-making when an agent seeks to determine an optimal time to exit the financial market [31] or terminate some costly industrial process [33] under worst-case scenarios. Examples of *worst-case* OSPs are agents that seek to determine when to arrest a costly industrial process or experiment (e.g. clinical trials) and, within finance, investors that seek to determine market entry/exit decisions; each under worst-case scenarios.

We later prove an equivalence of SGs of control and stopping and robust OSPs, the latter of which we now introduce:

The problem involves an agent that seeks to find an optimal stopping time $\hat{\tau}$ under the adverse *non-linear expectation* $\mathcal{E}_P := \min_{\pi \in \Pi} \mathbb{E}_{P,\pi}$ such that:

$$\hat{\tau} \in \arg \max_{\tau \in \mathcal{T}} \mathcal{E}_P [Y_\tau] = \arg \max_{\tau \in \mathcal{T}} \left(\min_{\pi \in \Pi} \mathbb{E}_{P,\pi} [Y_\tau] \right), \quad (3)$$

where $Y_k := \sum_{t=0}^{k \wedge T} \gamma^t R(s_t, a_t) + \gamma^{k \wedge T} G(s_{k \wedge T})$.

The problem describes an agent that seeks to find an optimal stopping time under a worst-case scenario.

3 EXAMPLES

To elucidate the ideas, in this section we provide concrete applications of the problems we are concerned with.

As the following example illustrates, the framework applies to actuator failure within RL applications.

3.1 Example: Control with random actuator failure

Consider an adaptive learner, for example a robot that uses a set of actuators to perform actions. Given full operability of its set of actuators, the agent’s actions are determined by a policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ which maps from the state space \mathcal{S} and the set of actions \mathcal{A} to a probability. In many systems, there exists some risk of actuator failure at which point the agent thereafter can affect the state transitions by operating only a *subset* of its actuators. In this instance, the agent’s policy determines actions using only a subset of its action space $\hat{\mathcal{A}} \subset \mathcal{A}$. In this scenario, the agent is now restricted to policies $\pi_{\text{partial}} : \mathcal{S} \times \hat{\mathcal{A}} \rightarrow [0, 1]$ which map from from a subset of operative actuators — thereafter its expected return is given by the value function $V^{\pi_{\text{partial}}}$. In order to perform robustly against actuator failure, it is therefore necessary to consider a set of stopping times $\mathcal{T} \subseteq \{0, 1, 2, \dots\}$ after which, the robot can no longer select actions that require functionality of the full set of actuators. In particular, in order to construct a robust policy against catastrophic outcomes, it is useful to consider actuator failure in worst-case scenarios.

The problem involves finding a pair $(\hat{\tau}, \hat{\pi}) \in \mathcal{T} \times \Pi$ which consists of a stopping time and control policy s.th.

$$\min_{k' \in \mathcal{T}} \left(\max_{\pi' \in \Pi} \mathbb{E} \left[H^{\pi', k'}(s) \right] \right) = \mathbb{E} \left[H^{\hat{\pi}, \hat{\tau}}(s) \right],$$

where $a_t \sim \pi'$ and $H^{\pi, k}(s) := \sum_{t=0}^{k \wedge \infty} \gamma^t R(s_t, a_t) + \gamma^{k \wedge \infty} V^{\pi_{\text{partial}}}(s_{k \wedge \infty})$. The resulting policy $\hat{\pi}$ is robust against actuator failure in worst-case scenarios.

3.2 Example: Optimal selling in an adversarial market

An investor (**I**) seeks to exit the market (sell all market holdings) at an optimal stopping time $\tau \in \mathcal{T}$. It is assumed that the market acts in such a way to minimise risk-free profit opportunities for the investor.² When **I** exits the market, **I** receives a return of $\lambda^\tau X_\tau$ where $X_t \equiv X(t, \omega) \in [0, \infty[\times \Omega$ is a Markov process that determines the asset price at time t and $\lambda \in]0, 1]$ is **I**’s discount factor. Classically, the exit time is computed as the solution to the following problem:

$$\max_{k \in \mathcal{T}} \mathbb{E}_P \left[\gamma^k X_k \right]. \quad (4)$$

In (4), the expectation is taken with respect to a risk-neutral measure P . However, the above formulation does not include the adversarial effect of the market. To accommodate this, we modify the objective

²This is the *no arbitrage principle* [7].

to the following:

$$\max_{k \in \mathcal{T}} \left(\min_{\pi \in \Pi} \mathbb{E}_{P, \pi} \left[\gamma^k X_k \right] \right). \quad (5)$$

In (5), the worst-case dynamics are induced by choice of adversarial probability measure π that alters the neutral measure P over which the objective expectation is defined. This captures the observed effect that financial markets adversarially eliminate investment opportunities. Now the goal of the agent is to find an optimal time to exit a financial market *under an adversarial market scenario*.

4 STOCHASTIC GAMES

Embedded within problem (2) is an interdependence between the actions of the players — that is, the solution to the problem is jointly determined by the actions of both players and their responses to each other. The appropriate framework to tackle this problem is therefore an SG [25]. An SG is an augmented MDP which proceeds by two players tacking actions that *jointly* manipulate the transitions of a system over K rounds which may be infinite. At each round, the players receive some immediate reward or cost which is a function of the players’ joint actions. The framework is zero-sum so that a reward for player 1 simultaneously represents a cost for player 2.

Formally, a two-player zero-sum **stochastic game** is a 6-tuple $\langle \mathcal{S}, \mathcal{A}_{i \in \{1,2\}}, P, R, \gamma \rangle$ where $\mathcal{S} = \{s_1, s_2, \dots, s_n\}$ is a set of states, \mathcal{A}_i is an action set for each player $i \in \{1, 2\}$. The map $P : \mathcal{S} \times \mathcal{A}_1 \times \mathcal{A}_2 \times \mathcal{S} \rightarrow [0, 1]$ is a Markov transition probability matrix i.e. $P(s'; s, a_1, a_2)$ is the probability of the state s' being the next state given the system is in state s and actions $a_1 \in \mathcal{A}_1$ and $a_2 \in \mathcal{A}_2$ are applied by player 1 and player 2 (resp.). The function $R : \mathcal{S} \times \mathcal{A}_1 \times \mathcal{A}_2$ is the one-step reward for player 1 and represents one-step cost for player 2 when player 1 takes action $a_1 \in \mathcal{A}_1$ and player 2 takes action $a_2 \in \mathcal{A}_2$ and $\gamma \in [0, 1[$ is a discount factor. The goal of each player is to maximise its expected cumulative return — since the game is antagonistic, the total expected reward received by player 1 which we denote by J , represents a total expected cost for player 2.

Denote by $H_{t \leq K}$ the set of all finite histories and by $\mathcal{H} \equiv \cup_{t \leq K} H_t$ so that for each $h_{j \leq K} = ((s_0, (a_1, a_2)), (s_1, (a_1, a_2)), \dots, (s_j, (a_1, a_2))) \in H_j$ which is a sequence of state and joint action pairs. For each player $i \in \{1, 2\}$, a **pure strategy** is a map $\pi_i : (\mathcal{H}) \times \mathcal{A}_i \rightarrow [0, 1]$ that assigns to every finite history $h_{j \leq K} \in \mathcal{H}$ an action $\pi(h)$ in \mathcal{A}_i . Similarly, for each player $i \in \{1, 2\}$, a **behavioural strategy** is a map $\pi_i : \mathcal{H} \times \mathcal{A}_i \rightarrow [0, 1]$ that assigns to every finite history $h \in \mathcal{H}$ a probability distribution $\pi(h)$ in \mathcal{A}_i . We denote the space of strategies for each player $i \in \{1, 2\}$ by Π_i . Note that pure strategies are a degenerate class of behavioural strategies which assign to any history $h \in \mathcal{H}$ the Dirac measure with its probability mass concentrated at a single point $\pi(h)$.

For SGs with Markovian transition dynamics, we can safely dispense with path dependencies in the space of strategies.³ Consequently, w.log. we restrict ourselves to the class of behavioural strategies that depend only on the current state and round, namely **Markov strategies**, hence for each player i , the strategy space Π_i consists of strategies of the form $\pi_i : \mathcal{S} \times \mathcal{A}_i \rightarrow [0, 1]$. It is well-known that for SGs, an equilibrium exists in Markov strategies even when the opponent can draw from non-Markovian strategies [14].

³There are some exceptions for games with payoff structures not considered here for example, limiting average (Ergodic) payoffs [6].

In SGs, it is usual to consider the case $\mathcal{A}_1 = \mathcal{A}_2$ so that the players’ actions are drawn from the same set. We depart from this model and consider a game in which player 2 can choose a time to stop the process so that the action set for player 2 is the set $\mathcal{T} \subseteq \{0, 1, 2, \dots\}$ which consists of (\mathcal{F} -measurable) stopping times. In this setting, player 1 can manipulate the system dynamics by taking actions drawn from \mathcal{A}_1 (we hereon use \mathcal{A}) and at each point, player 2 can decide to intervene to stop the game.

Let us define by $\text{val}^+[J] := \min_{k \in \mathcal{T}} \max_{\pi \in \Pi} J^{k, \pi}$ the *upper value function* and by $\text{val}^-[J] := \max_{\pi \in \Pi} \min_{k \in \mathcal{T}} J^{k, \pi}$, the *lower value function*. The upper (lower) value function represents the minimum payoff that player I (player II) can guarantee itself irrespective of the actions of the opponent.

The **value** of the game exists if we can commute the max and min operators:

$$\text{val}^-[J] = \max_{\pi \in \Pi} \min_{k \in \mathcal{T}} J^{k, \pi}[\cdot] = \min_{k \in \mathcal{T}} \max_{\pi \in \Pi} J^{k, \pi}[\cdot] = \text{val}^+[J]. \quad (6)$$

We denote the value by $J^* := \text{val}^+[J] = \text{val}^-[J]$ and denote by $(\hat{k}, \hat{\pi}) \in \mathcal{T} \times \Pi$ the pair that satisfies $J^{\hat{k}, \hat{\pi}} \equiv J^*$. The value, should it exist, is the minimum payoff each player can guarantee itself under the equilibrium strategy. In general, the functions $\text{val}^+[J]$ and $\text{val}^-[J]$ may not coincide.

Should the value J^* exist, it constitutes a **saddle point equilibrium** of the game in which neither player can improve their payoff by playing some other control — an analogous concept to a Nash equilibrium for the case of two-player zero-sum games [26]. Thus the central task to establish an equilibrium involves unambiguously assigning a value to the game, that is proving the existence of J^* .

5 MAIN ANALYSIS

In this section, we present the key results and perform the main analysis of the paper. In the main analysis, our first task is to establish the existence of a value of the game then secondly, we perform analyses that enables us to construct an approximate dynamic programming method. In particular, we construct a Bellman operator for the game and show that the operator is a contraction mapping. We show that the value is unique and that the value coincides with a fixed point of the Bellman operator. Using these results, we construct an equivalence between robust OSPs and games of control and stopping. We defer some of the proofs to the appendix.

5.1 Summary of Main Results

Our results develop the theory of risk within RL to cover instances in which the agent has concern for stopping the process at an optimal time. We develop the theory of SGs to cover games of control and stopping when neither player has up-front environment knowledge. In particular, we establish the existence of a value of the game in a discrete-time setting and show that the value can be obtained using a value-iterative method. This, as we show in Sec. 9, underpins a simulation-based scheme that is suitable for settings in which the transition model and reward function is a priori unknown.

5.2 Theoretical Analysis

The purpose of this section is to twofold: our first task is to establish the existence of a value of the game. Secondly, we perform analyses

that enables us to construct an approximate dynamic programming method. In particular, we construct a Bellman operator for the game and show that the operator is a contraction mapping. We show that the value is unique and that the value coincides with a fixed point of the Bellman operator. Using these results, we construct an equivalence between robust OSPs and games of control and stopping. We defer some of the proofs to the appendix.

Our results are built under the following assumptions:

Assumption A.1. Stationarity: the expectations \mathbb{E} are taken w.r.t. a stationary distribution so that for any measurable function f we have $\mathbb{E}[f(s_0)] = \mathbb{E}[f(s_k)]$ for any $k \geq 0$.

Assumption A.2. Ergodicity: i) Any invariant random variable of the state process is P -almost surely (P -a.s.) a constant.

Assumption A.3. Markovian transition dynamics: the transition probability function P satisfies the following equality: $P(s_{k+1} \in A | \mathcal{F}_k) = P(s_{k+1}, A)$ for any $A \in \mathcal{B}(\mathbb{R}^P)$.

Assumption A.4. The constituent functions $\{R, G\}$ in (1) are square integrable i.e. $R, G \in \mathbb{L}_2(\mu)$.

We begin the analysis with some preliminary lemmata which will become useful for proving the main results:

LEMMA 5.1. *Define $\text{val}^+[f] := \min_{b \in \mathbb{B}} \max_{a \in \mathbb{A}} f(a, b)$ and define $\text{val}^-[f] := \max_{a \in \mathbb{A}} \min_{b \in \mathbb{B}} f(a, b)$, then for any $b \in \mathbb{B}$ we have that for any $f, g \in \mathbb{L}$ and for any $c \in \mathbb{R}_{>0}$:*

$$\begin{aligned} & \left| \max_{a \in \mathbb{A}} f(a, b) - \max_{a \in \mathbb{A}} g(a, b) \right| \leq c \\ \implies & |\text{val}^+[f] - \text{val}^+[g]| \leq c. \end{aligned} \quad (7)$$

The proof is straightforward.

LEMMA 5.2. *For any $f, g, h \in \mathbb{L}$ and for any $c \in \mathbb{R}_{>0}$ we have that:*

$$\|f - g\| \leq c \implies \|\min\{f, h\} - \min\{g, h\}\| \leq c$$

LEMMA 5.3. *Let the functions $f, g, h \in \mathbb{L}$ then*

$$\|\max\{f, h\} - \max\{g, h\}\| \leq \|f - g\|. \quad (8)$$

We now introduce some concepts that relate to estimates on the operators of the game. These concepts will be useful for proving the existence of a fixed point.

Definition 5.4. An operator T is said to be a **contraction** w.r.t a norm $\|\cdot\|$ if there exists a constant $c \in [0, 1[$ s.th for any $V_1, V_2 \in \mathcal{V}$ we have that:

$$\|TV_1 - TV_2\| \leq c\|V_1 - V_2\|. \quad (9)$$

A central task of the paper is to prove that the Bellman operator for the game is a contraction mapping. Thereafter, convergence to the unique value of the game using T is proved

Definition 5.5. An operator $T : \mathcal{V} \rightarrow \mathcal{V}$ is **non-expansive** if $\forall V_1, V_2 \in \mathcal{V}$ we have that:

$$\|TV_1 - TV_2\| \leq \|V_1 - V_2\|. \quad (10)$$

Consider a Borel measurable function which is absolutely integrable w.r.t. the probability transition kernel P then:

$$\mathbb{E}[J[s'] | \mathcal{F}_t] = \int_{\mathcal{S}} J[s'] P_{ss'}^a. \quad (11)$$

In this paper, we denote by $(PJ)(s) := \int_{\mathcal{S}} J[s'] P_{sd s'}^a$.

We now introduce the operator of the game which is of central importance:

$$TJ := \min_{\tau \in \mathcal{T}} \left\{ \max_{a \in \mathbb{A}} R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a J^{\tau, \pi}[s'], G(S_\tau) \right\}. \quad (12)$$

The operator T enables the game to be broken down into a sequence of sub minimax problems. It will later play a crucial role in establishing a value iterative method for computing the value of the game.

LEMMA 5.6. *The operator T satisfies the following:*

- (1) (*Monotonicity*) For any $J_1, J_2 \in \mathbb{L}_2$ s.th. $J_1(s) \leq J_2(s), \forall s \in \mathcal{S}$ then $TJ_1(s) \leq TJ_2(s), \forall s \in \mathcal{S}$.
- (2) (*Constant shift*) Let $I(s) \equiv I$ be the unit function, then for any $J \in \mathbb{L}_2$ and for any scalar $\alpha \in \mathbb{R}$, T satisfies $T(J + \alpha I)(s) = TJ(s) + \alpha I(s)$.

Before performing the analysis on T however, we firstly prove the existence of the value of the game:

THEOREM 5.7.

$$\text{val}^+[J] = \text{val}^-[J] \equiv J^*. \quad (13)$$

Theorem 5.7 establishes the existence of the game which permits commuting the max and min operators of the objective (2). Crucially, the theorem secures the existence of an optimal pair of controls $(\hat{\tau}, \hat{\pi}) \in \mathcal{T} \times \Pi$ — the computation of which, is the subject of the next section.

PROOF. We begin by noting the following inequality holds:

$$\text{val}^+[J] = \min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} \mathbb{E}[J^{\tau, \pi}[s]] \geq \max_{\pi \in \Pi} \min_{\tau \in \mathcal{T}} \mathbb{E}[J^{\tau, \pi}[s]] = \text{val}^-[J]. \quad (14)$$

The inequality follows by noticing $J^{k, \pi} \leq \max_{\pi \in \Pi} J^{k, \pi}$ and thereafter applying the $\min_{k \in \mathcal{T}}$ and $\max_{\pi \in \Pi}$ operators.

The proof is now settled by reversing the inequality in (14). To begin, choose a sequence of open intervals $\{D_m\}_{m=1}^\infty$ s.th. for each $m = 1, 2, \dots$ \bar{D}_m is compact and $\bar{D}_m \supset \bar{D}_{m+1}$ and $[0, T] = \bigcap_{m=1}^\infty \bar{D}_m$ and define $\tau_D(m) := \inf_{k \in D_m} \mathbb{E}[J^{k, \pi}[s]]$.

We now observe that:

$$\begin{aligned}
& \mathbb{E}[J^{\hat{\tau}, \hat{\pi}}[s]] \\
&= \max_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^{\tau_D(m)} \gamma^t (R(s_t, a_t) + G(s_{\tau_D(m)})) \right] \\
&\quad - \mathbb{E} \left[\sum_{t=\hat{\tau}}^{\tau_D(m)} \gamma^t (R(s_t, a_t) + G(s_{\tau_D(m)})) \right] \\
&\geq \mathbb{E} [J^{\tau_D(m), \pi}[s]] - \left| \mathbb{E} \left[\sum_{t=\hat{\tau}}^{\tau_D(m)} \gamma^t (R(s_t, a_t) + G(s_{\tau_D(m)})) \right] \right| \\
&\geq \mathbb{E} [J^{\tau_D(m), \pi}[s]] - \sum_{t=\hat{\tau}}^{\tau_D(m)} \gamma^t (\mathbb{E}[|R(s_t, a_t)|] + \mathbb{E}[|G(s_{\tau_D(m)})|]) \\
&\geq \mathbb{E} [J^{\tau_D(m), \pi}[s]] - \sum_{t=\hat{\tau}}^{\tau_D(m)} \gamma^t (\mathbb{E}[|R(s_0, \cdot)|] + \mathbb{E}[|G(s_0)|]) \\
&= \mathbb{E} [J^{\tau_D(m), \pi}[s]] + \gamma^{\tau_D(m)+1} \frac{1 - \gamma^{\hat{\tau} - \tau_D(m)}}{1 - \gamma} c \\
&= \lim_{m \rightarrow \infty} \inf \mathbb{E}[J^{\tau_D(m), \pi}[s]] + \lim_{m \rightarrow \infty} \left[\gamma^{\tau_D(m)+1} \frac{1 - \gamma^{\hat{\tau} - \tau_D(m)}}{1 - \gamma} \right] c \\
&\geq \mathbb{E}[J^{\hat{\tau}, \pi}[s]], \tag{16}
\end{aligned}$$

where we have used the stationarity property and, in the limit $m \rightarrow \infty$ and the Fatou lemma to generate the inequality after taking the limit. The constant c is given by $c := (\mathbb{E}[|R(s_0, \cdot)|] + \mathbb{E}[|G(s_0)|]) \in \mathbb{L}$.

Hence, we now find that

$$\mathbb{E}[J^{\hat{\tau}, \hat{\pi}}[s]] \geq \mathbb{E}[J^{\hat{\tau}, \pi}[s]]. \tag{17}$$

Now since (17) holds $\forall \pi \in \Pi$ we find that:

$$\mathbb{E}[J^{\hat{\tau}, \hat{\pi}}[s]] \geq \max_{\pi \in \Pi} \mathbb{E}[J^{\hat{\tau}, \pi}[s]]. \tag{18}$$

Lastly, applying min operator we observe that:

$$\mathbb{E}[J^{\hat{\tau}, \hat{\pi}}[s]] \geq \min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} \mathbb{E}[J^{\tau, \pi}[s]] = \text{val}^+[J]. \tag{19}$$

It now remains to show the reverse inequality holds:

$$\mathbb{E}[J^{\hat{\tau}, \hat{\pi}}[s]] \leq \max_{\pi \in \Pi} \min_{\tau \in \mathcal{T}} \mathbb{E}[J^{\tau, \pi}[s]] = \text{val}^-[J]. \tag{20}$$

Indeed, we observe that

$$\mathbb{E} [J^{\hat{\tau}, \hat{\pi}}[s]] \tag{21}$$

$$\leq \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau \wedge m, \hat{\pi}}[s]] + \mathbb{E} \left[\sum_{t=m}^{\infty} \gamma^t (|R(s_t, a_t)| + |G(s_t)|) \right] \tag{22}$$

$$\leq \lim_{m \rightarrow \infty} \left[\min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau \wedge m, \hat{\pi}}[s]] + c(m) \right] \tag{23}$$

$$= \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau, \hat{\pi}}[s]] \leq \max_{\pi \in \Pi} \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau, \pi}[s]], \tag{24}$$

since $\gamma \in [0, 1[$, where $c(m) := \frac{\gamma^m}{1-\gamma} (\mathbb{E}[|R(s_0, \cdot)|] + \mathbb{E}[|G(s_0)|])$ (using the stationarity of the state process) and where we have used Lebesgue's Dominated Convergence Theorem in the penultimate step.

Hence, by (24) we have that:

$$\mathbb{E} [J^{\hat{\tau}, \hat{\pi}}[s]] \leq \max_{\pi \in \Pi} \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau, \pi}[s]] = \text{val}^-[J]. \tag{25}$$

Hence putting (19) and (25) together gives:

$$\text{val}^-[J] = \max_{\pi \in \Pi} \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau, \pi}[s]] \tag{26}$$

$$\geq \mathbb{E}[J^{\hat{\tau}, \hat{\pi}}[s]] \geq \min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} \mathbb{E}[J^{\tau, \pi}[s]] \tag{27}$$

$$= \text{val}^+[J]. \tag{28}$$

After combining (28) with (14) we deduce the thesis. \square

Theorem 5.7 proves the existence of a value and that we can commute the max and min operators. Crucially, Theorem 5.7 establishes the existence of an equilibrium, the value of which is given by J^* . In turn, we can now establish the optimal strategies for each player. To this end, we now define best-response strategies which shall be useful for further characterising the equilibrium:

Definition 5.8. The set of **best-response (BR) strategies** for player I against the stopping time $\tau \in \mathcal{T}$ (BR strategies for player II against the control policy $\pi \in \Pi$) is defined by $\hat{\pi} \in \arg \max_{\pi' \in \Pi} \mathbb{E}[J^{\tau, \pi'}[\cdot]]$ (resp., $\hat{\tau} \in \arg \min_{\tau' \in \mathcal{T}} \mathbb{E}[J^{\tau', \pi}[\cdot]]$).

The question of computing the value of the game remains. To this end, we now prove the contraction mapping property of the operator T . We then show that repeatedly applying T produces a sequence that converges to the value.

Definition 5.9. The **residual** of a vector $V \in \mathcal{V}$ w.r.t the operator $T : \mathcal{V} \rightarrow \mathcal{V}$ is:

$$\epsilon_T(V) := \|TV - V\|. \tag{29}$$

The following lemma is a required result for proving the contraction mapping property of the operator T .

LEMMA 5.10. *The map P is non-expansive, that is:*

$$\|PV_1 - PV_2\| \leq \|V_1 - V_2\|. \tag{30}$$

PROPOSITION 5.11. *The operator T in (12) is a contraction.*

PROOF. We wish to prove that:

$$\|TJ - T\tilde{J}\|_{\pi} \leq \gamma \|J - \tilde{J}\|. \tag{31}$$

Firstly, we observe that:

$$\begin{aligned}
& \left\| \max_{a \in A} \left\{ R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a J^{\tau, \pi}[s'], G(s_k) \right\} \right. \\
& \quad \left. - \left(\max_{a \in A} \left\{ R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a \tilde{J}^{\pi}[s'], \tilde{G}(s_k) \right\} \right) \right\| \\
& \leq \gamma \max_{a \in A} \left\| \sum_{s' \in \mathcal{S}} P_{ss'}^a (J_{s-1}^{\tau, \pi}[s'] - \tilde{J}_{s-1}^{\pi}[s']) \right\| \\
& \leq \gamma \|J_{s-1}^{\tau, \pi} - \tilde{J}_{s-1}^{\pi}\|, \tag{32}
\end{aligned}$$

using Cauchy-Schwartz (and that $\gamma \in [0, 1[$) and (59). The result follows after applying Lemma 5.2 and Lemma 5.3. \square

Prop. 5.11 underscores a *fixed point property* which is stated in the following:

THEOREM 5.12. *1. The sequence $(T^n J)_{n=0}^\infty$ converges (in \mathbb{L}_2).
2. There exists a unique function $J^\star \in \mathbb{L}_2$ s.th.:*

$$J^\star = TJ^\star \text{ and } \lim_{n \rightarrow \infty} T^n J = J^\star. \quad (33)$$

PROOF. *Part 1:* We note that the contraction property of T (c.f. Prop. 5.11) allows us to demonstrate that the game has a unique fixed point to which a sequence $(T^n J)_{n=0}^\infty$ converges (in \mathbb{L}_2). In particular, by Prop. 5.11 we have that $\|T^2 J - TJ\| \leq \gamma \|TJ - J\|$ which proves that the sequence $(T^n J)_{n=0}^\infty$ converges to a fixed point.

Part 2: We observe that the fixed point is unique since if $\exists J, M \in \mathbb{L}_2$ s.th. $TJ = J$ and $TM = M$ we find that $\|M - J\| = \|TM - TJ\| = \gamma \|M - J\|$, so that $M = J$ (since $\gamma \in [0, 1[$) which gives the desired result.

Adopting notions in dynamic programming, denote by:

$$T^n J[s] = \min_{\tau \in \mathcal{T}} \max_{\pi_0, \pi_1, \dots, \pi_{n-1}} \mathbb{E} \left[\sum_{t=0}^{\{n-1 \wedge \tau\}} \gamma^t R(s_t, a_t) + \gamma^n J(s_{n \wedge \tau}) \right].$$

We begin the proof by invoking similar reasoning as (21) - (23) to deduce that:

$$\mathbb{E} [J^{\hat{\tau}, \hat{\pi}}[s]] \leq \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau \wedge n, \hat{\pi}}[s]] + \frac{\gamma^n}{1-\gamma} c,$$

where $c := (\mathbb{E}[|R(s_0, \cdot)|] + \mathbb{E}[|G(s_0)|])$.

Hence,

$$T^n J[s] \leq \max_{\pi \in \Pi} \min_{\tau \in \mathcal{T}} \mathbb{E} [J^{\tau, \pi}[s]] + \frac{\gamma^n}{1-\gamma} c \quad (34)$$

$$= J^\star[s] + \frac{\gamma^n}{1-\gamma} c. \quad (35)$$

By analogous reasoning we can deduce that:

$$T^n J[s] \geq \min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} \mathbb{E} [J^{\tau, \pi}[s]] - \frac{\gamma^n}{1-\gamma} c \quad (36)$$

$$= J^\star[s] - \frac{\gamma^n}{1-\gamma} c. \quad (37)$$

Putting (35) and (37) together implies:

$$J^\star[s] - \frac{\gamma^n}{1-\gamma} c \leq T^n J[s] \leq J^\star[s] + \frac{\gamma^n}{1-\gamma} c. \quad (38)$$

By Lemma 5.6, i.e. invoking the monotonicity and constant shift properties of T , we apply T to (38) and preserve the inequalities:

$$TJ^\star[s] - \frac{\gamma^n}{1-\gamma} c \leq T^{n+1} J[s] \leq TJ^\star[s] + \frac{\gamma^n}{1-\gamma} c. \quad (39)$$

After taking the limit in (39) and, using the sandwich theorem of calculus, we deduce the result. \square

Theorem 5.12 establishes the existence of a fixed point of T . Crucially, it underpins a value iterative method which we formally develop in Sec. 6.

Definition 5.13. The pair $(\hat{\tau}, \hat{\pi}) \in \mathcal{T} \times \Pi$ is a saddle point equilibrium iff $\forall s \in \mathcal{S}$:

$$J^{\hat{\tau}, \hat{\pi}}[s] = \max_{\tau \in \mathcal{T}} J^{\tau, \hat{\pi}}[s] = \min_{\pi \in \Pi} J^{\hat{\tau}, \pi}[s]. \quad (40)$$

A saddle point equilibrium therefore defines a strategic configuration in which both players play their BR strategies.

PROPOSITION 5.14. *The pair $(\hat{\tau}, \hat{\pi}) \in \mathcal{T} \times \Pi$ consists of BR strategies and constitutes a saddle point equilibrium.*

By Prop. 5.14, when the pair $(\hat{\tau}, \hat{\pi})$ is played, each player executes its BR strategy in response to their opponent. In the context of the problem in Sec. 2, the strategic response induces risk minimising behaviour by the controller.

PROOF. Prop. 5.14 follows from the fact that if either player plays a Markov strategy then their opponent's best-response is also Markovian. Moreover, by Theorem 5.12, $\hat{\tau}$ is a BR strategy for player 2 (recall Definition 5.8). By Theorem 5.7 (commuting the max and min operators) we observe that $\hat{\pi}$ is a BR strategy for player 1. \square

We now turn to proving the existence and characterising the optimal stopping time for player 2. The following result establishes its existence.

THEOREM 5.15. *There exists an \mathcal{F} -measurable stopping time:*

$$\hat{\tau} = \min \left\{ k \in \mathcal{T} \mid G(s_k) \leq \min_{k \in \mathcal{T}} \max_{\pi \in \Pi} J^{k, \pi}[s_k] \right\}, a.s.$$

Theorem 5.15 establishes the existence and characterises the player 2 optimal stopping time. The stopping time $\hat{\tau}$ is a best-response for player 2 against the equilibrium policy played by player 1. The theorem plays a vital role in the robust optimal stopping problem of Sec. 4.

PROOF. For any $m \in \mathbb{N}$ we have that:

$$\max_{\pi \in \Pi} J^{\tau, \pi}[s] \geq \max_{\pi \in \Pi} J^{\tau \wedge m, \pi}[s] - \sum_{t=m}^{\infty} \gamma^t \max_{\pi \in \Pi} (|R(s_t, a_t)| + |G(s_t)|). \quad (41)$$

We now apply the min operator to both sides of (41) which gives:

$$\begin{aligned} & \min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau, \pi}[s] \\ & \geq \min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau \wedge m, \pi}[s] - \sum_{t=m}^{\infty} \gamma^t \max_{\pi \in \Pi} (|R(s_t, a_t)| + |G(s_t)|). \end{aligned} \quad (42)$$

After taking expectations, we find that:

$$\begin{aligned} & \mathbb{E} \left[\min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau, \pi}[s] \right] \\ & \geq \mathbb{E} \left[\min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau \wedge m, \pi}[s] \right] - \sum_{t=m}^{\infty} \gamma^t \mathbb{E} \left[\max_{\pi \in \Pi} (|R(s_t, a_t)| + |G(s_t)|) \right]. \end{aligned} \quad (43)$$

Now by Jensen's inequality and, using the stationarity of the state process (recall the expectation is taken under π) we have that:

$$\begin{aligned} & \mathbb{E} \left[\max_{\pi \in \Pi} (|R(s_t, a_t)| + |G(s_t)|) \right] \\ & \geq \max_{\pi \in \Pi} \mathbb{E} [(|R(s_t, a_t)| + |G(s_t)|)] \end{aligned} \quad (44)$$

$$= \mathbb{E}[|R(s_0, \cdot)|] + \mathbb{E}[|G(s_0)|]. \quad (45)$$

By standard arguments of dynamic programming, the value of the game with horizon n can be obtained from n iterations of the dynamic recursion; in particular, we have that:

$$\min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau \wedge m, \pi}[s] = T^m G(s). \quad (46)$$

Inserting (45) and (46) into (43) gives:

$$\begin{aligned} \mathbb{E} \left[\min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau, \pi}[s] \right] &\geq \mathbb{E} [T^m G(s)] - c(m) \\ &= \lim_{m \rightarrow \infty} [\mathbb{E} [T^m G(s)] - c(m)] = \mathbb{E} [J^{\hat{\tau}, \hat{\pi}}[s]], \end{aligned} \quad (47)$$

where $c(m) := \frac{\gamma^m}{1-\gamma} (\mathbb{E}[|R(s_0, \cdot)|] + \mathbb{E}[|G(s_0)|])$ so that $\lim_{m \rightarrow \infty} c(m) = 0$. Hence, we find that:

$$\mathbb{E} [J^{\hat{\tau}, \hat{\pi}}[s]] \leq \mathbb{E} \left[\min_{\tau \in \mathcal{T}} \max_{\pi \in \Pi} J^{\tau, \pi}[s] \right]. \quad (48)$$

The result follows since $G(s_\tau) = J^{\tau, \cdot}[s_\tau]$ by definition of G . \square

Having proved the existence of the optimal stopping time τ^* , by Theorem 5.15 and Theorem 5.7, we find the following:

THEOREM 5.16. *Let $\hat{\tau}$ be the player 2 optimal stopping time defined in (5.15) and let τ^* be the optimal stopping time for the robust OSP (c.f. (3)) then $\tau^* = \hat{\tau}$.*

Theorem 5.16 establishes an equivalence between the robust OSP and the SG of control and stopping. Computing the optimal stopping time of the SG yields a solution to the robust OSP.

6 SIMULATION-BASED VALUE ITERATION

We now develop a *simulation-based* value-iterative scheme. We show that the method produces an iterative sequence that converges to the value of the game from which the optimal controls can be extracted. The method is suitable for environments in which the transition model and reward functions are not known to either player. Our approach is related to approximated dynamic programming methods [5]. However, our problem requires generalisation to an SG involving a controller and stopper which alters the proofs throughout.

The fixed point property of the game established in Theorem 5.12 immediately suggests a solution method for finding the value. In particular, we may seek to solve the fixed point equation (FPE) $J^* = T J^*$. Direct approaches at solving the FPE are not generally fruitful as closed solutions are typically unavailable.

To compute the value function, we develop an iterative method that tunes weights of a set of basis functions $\{\phi_k : \mathbb{R}^P \rightarrow \mathbb{R} | k \in 1, 2, \dots, D\}$ to approximate J^* through simulated system trajectories and associated costs. Algorithms of this type were first introduced by Watkins [29] as an approximate dynamic programming method and have since been augmented to cover various settings. Therefore the following can be considered as a generalised Q-learning algorithm for zero-sum controller stopper games.

Let us denote by $\Phi r := \sum_{j=1}^D r(j) \phi_j$ an operator representation of the basis expansion. The algorithm is initialised with weight vector $r_0 = (r_0(1), \dots, r_0(P))' \in \mathbb{R}^d$. Then as the trajectory $\{s_t | t = 0, 1, 2, \dots\}$ is simulated, the algorithm produces an updated series of vectors $\{r_t | t = 0, 1, 2, \dots\}$ by the update:

$$r_{t+1} = r_t + \gamma \phi(s_t) \left(\max_{a \in \mathcal{A}} R_{s_t}^a + \gamma \min \{(\phi r_t)(s_{t+1}), G(s_{t+1})\} - (\phi r_t)(s_t) \right).$$

Theorem 6.1 demonstrates that the method converges to an approximation of J^* . We provide a bound for the approximation error in terms of the basis choice.

We define the function Q^* which the algorithm approximates by:

$$Q^*(s) = \max_{a \in \mathcal{A}} R_s^a + \gamma P J^*, \quad \forall s \in \mathcal{S} \quad (49)$$

We later show that Q^* serves to approximate the value J^* . In particular, we show that the algorithm generates a sequence of weights r_n that converge to a vector r^* and that Φr^* , in turn approximates Q^* . To complete the connection, we then provide a bound between the outcome of the game when the players use controls generated by the algorithm.

First, we introduce our player 2 stopping criterion which now takes the form:

$$\hat{t} = \min\{t | G(s_t) \leq Q^*(s_t)\} \quad (50)$$

We define a orthogonal projection Π and the function F by:

$$\Pi Q := \arg \min_{\bar{Q} \in \{\Phi r | r \in \mathbb{R}^P\}} \|\bar{Q} - Q\|, \quad (51)$$

$$FQ := \max_{a \in \mathcal{A}} R_s^a + \gamma P \min\{G, Q\}. \quad (52)$$

We now state the main results of the section:

THEOREM 6.1. *Under (49), r_n converges to r^* where r^* is the unique solution: $\Pi F(\Phi r^*) = \Phi r^*$, a.e.*

THEOREM 6.2. *Let $\hat{t} = \min\{k \in \mathcal{T} | G(s_k) \leq (\Phi r^*)(s_k)\}$, then the following hold:*

1. $\|\Phi r^* - Q^*\| \leq (\sqrt{1-\gamma^2})^{-1} \|\Pi Q^* - Q^*\|$,
2. $\mathbb{E} [J^*[s] - J^{\hat{\tau}, \hat{\pi}}[s]] \leq \frac{2}{(1-\gamma)\sqrt{1-\gamma^2}} \|\Pi Q^* - Q^*\|$.

Theorem 6.2 says the error bound in algorithm approximation of the value is determined by the goodness of the projection.

CONCLUSION

In this paper, we tackled the problem of risk within an RL setting in which the controller seeks to obtain a fault-tolerant control that is robust to catastrophic failures. To formally analyse the problem and characterise the optimal behaviour, we performed an in-depth analysis of a stochastic game (SG) of control and stopping. We established the existence of an equilibrium value then, using a contraction mapping argument, showed that the game can be solved by iterative application of a Bellman operator. We proved that the method leads to an approximate dynamic programming algorithm so that the game can be solved by simulation. By proving an equivalence between the SG and robust optimal stopping problems, we showed that the method developed in the paper serves to compute solutions to optimal stopping problems in worst-case scenarios.

REFERENCES

- [1] Pieter Abbeel, Adam Coates, and Andrew Y Ng. 2010. Autonomous helicopter aerobatics through apprenticeship learning. *The International Journal of Robotics Research* 29, 13 (2010), 1608–1639.
- [2] Itamar Arel, Cong Liu, T Urbanik, and AG Kohls. 2010. Reinforcement learning-based multi-agent system for network traffic signal control. *IET Intelligent Transport Systems* 4, 2 (2010), 128–135.

- [3] Fouzia Baghery, Sven Haadem, Bernt Øksendal, and Isabelle Turpin. 2013. Optimal stopping and stochastic control differential games for jump diffusions. *Stochastics An International Journal of Probability and Stochastic Processes* 85, 1 (2013), 85–97.
- [4] Erhan Bayraktar, Xueying Hu, and Virginia R Young. 2011. Minimizing the probability of lifetime ruin under stochastic volatility. *Insurance: Mathematics and Economics* 49, 2 (2011), 194–206.
- [5] Dimitri P Bertsekas. 2008. Approximate dynamic programming. (2008).
- [6] David Blackwell and Tom S Ferguson. 1968. The big match. *The Annals of Mathematical Statistics* 39, 1 (1968), 159–163.
- [7] Peter Carr and Dilip B Madan. 2005. A note on sufficient conditions for no arbitrage. *Finance Research Letters* 2, 3 (2005), 125–130.
- [8] Jean-Philippe Chancelier, Bernt Øksendal, and Agnès Sulem. 2002. Combined stochastic control and optimal stopping, and application to numerical approximation of combined stochastic and impulse control. 237, 0 (2002), 149–172.
- [9] Marc Peter Deisenroth, Gerhard Neumann, Jan Peters, et al. 2013. A survey on policy search for robotics. *Foundations and Trends® in Robotics* 2, 1–2 (2013), 1–142.
- [10] EB Dynkin. 1967. Game variant of a problem on optimal stopping. In *Soviet Math. Dokl.*, Vol. 10. 270–274.
- [11] Jerzy A Filar, Todd A Schultz, Frank Thuijsman, and OJ Vrieze. 1991. Nonlinear programming and stationary equilibria in stochastic games. *Mathematical Programming* 50, 1-3 (1991), 227–237.
- [12] Javier Garcia and Fernando Fernández. 2015. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* 16, 1 (2015), 1437–1480.
- [13] Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. 2019. Guidelines for reinforcement learning in healthcare. *Nature medicine* 25, 1 (2019), 16–18.
- [14] Theodore Preston Hill. 1979. On the existence of good Markov strategies. *Trans. Amer. Math. Soc.* 247 (1979), 157–176.
- [15] Christopher Jennison and Bruce W Turnbull. 2013. Interim monitoring of clinical trials: Decision theory, dynamic programming and optimal stopping. *Kuwait Journal of Science* 40, 2 (2013).
- [16] Ying Jiao and Huyên Pham. 2011. Optimal investment with counterparty risk: a default-density model approach. *Finance and Stochastics* 15, 4 (2011), 725–753.
- [17] Ioannis Karatzas and William Sudderth. 2006. Stochastic games of control and stopping for a linear diffusion. In *Random Walk, Sequential Analysis And Related Topics: A Festschrift in Honor of Yuan-Shih Chow*. World Scientific, 100–117.
- [18] Thomas Kruse and Philipp Strack. 2015. Optimal stopping with private information. *Journal of Economic Theory* 159 (2015), 702–727.
- [19] A Maitra and T Parthasarathy. 1970. On stochastic games. *Journal of Optimization Theory and Applications* 5, 4 (1970), 289–300.
- [20] David Mguni. 2018. A Viscosity Approach to Stochastic Differential Games of Control and Stopping Involving Impulsive Control. *arXiv preprint arXiv:1803.11432* (2018).
- [21] Jun Morimoto and Kenji Doya. 2001. Robust reinforcement learning. In *Advances in Neural Information Processing Systems*. 1061–1067.
- [22] Goran Peskir and Albert Shiryaev. 2006. *Optimal stopping and free-boundary problems*. Springer.
- [23] Huyên Pham. 1997. Optimal stopping, free boundary, and American option in a jump-diffusion model. *Applied Mathematics and Optimization* 35, 2 (1997), 145–164.
- [24] Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua. 2016. Safe, multi-agent, reinforcement learning for autonomous driving. *arXiv preprint arXiv:1610.03295* (2016).
- [25] Lloyd S Shapley. 1953. Stochastic games. *Proceedings of the national academy of sciences* 39, 10 (1953), 1095–1100.
- [26] Alan Wilbor Starr and Yu-Chi Ho. 1969. Nonzero-sum differential games. *Journal of Optimization Theory and Applications* 3, 3 (1969), 184–206.
- [27] Aviv Tamar, Yinlam Chow, Mohammad Ghavamzadeh, and Shie Mannor. 2015. Policy gradient for coherent risk measures. In *Advances in Neural Information Processing Systems*. 1468–1476.
- [28] John N Tsitsiklis and Benjamin Van Roy. 1999. Optimal stopping of Markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. *IEEE Trans. Automat. Control* 44, 10 (1999), 1840–1851.
- [29] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8, 3-4 (1992), 279–292.
- [30] Toshiyuki Yasuda, Kazuhiro Ohkura, and Kanji Ueda. 2006. A homogeneous mobile robot team that is fault-tolerant. *Advanced Engineering Informatics* 20, 3 (2006), 301–311.
- [31] Virginia R Young. 2004. Optimal investment strategy to minimize the probability of lifetime ruin. *North American Actuarial Journal* 8, 4 (2004), 106–126.
- [32] Dapeng Zhang and Zhiwei Gao. 2018. Reinforcement learning-based fault-tolerant control with application to flux cored wire system. *Measurement and Control* 51, 7-8 (2018), 349–359.
- [33] Guozhen Zhao and Wen Chen. 2009. Enhancing R&D in science-based industry: An optimal stopping model for drug discovery. *International Journal of Project Management* 27, 8 (2009), 754–764.

SUPPLEMENTARY MATERIAL

APPENDIX

ADDITIONAL LEMMATA

The following estimates provide bounds on the value J^\star which we use later in the development of the iterative algorithm. We defer the proof of the results to the appendix.

LEMMA B.4. *Let $T : \mathcal{V} \rightarrow \mathcal{V}$ be a contraction mapping in $\|\cdot\|$ and let V^\star be a fixed point so that $TJ^\star = J^\star$ then there exists a constant $c \in [0, 1[$ s.th:*

$$\|J^\star - J\| \leq (1 - c)^{-1} \epsilon_T(J). \quad (53)$$

LEMMA B.5. *Let $T_1 : \mathcal{V} \rightarrow \mathcal{V}, T_2 : \mathcal{V} \rightarrow \mathcal{V}$ be contraction mappings and suppose there exists vectors J_1^\star, J_2^\star s.th $T_1 J_1^\star = J_1^\star$ and $T_2 J_2^\star = J_2^\star$ (i.e. J_1^\star, J_2^\star are fixed points w.r.t T_1 and T_2 respectively) then $\exists c_1, c_2 \in [0, 1[$ s.th:*

$$\|J_1^\star - J_2^\star\| \leq (1 - \{c_1 \wedge c_2\})^{-1} (\epsilon_{T_1}(J) - \epsilon_{T_2}(J)).$$

PROOF OF RESULTS

Proof of Lemma 5.1.

PROOF. We begin by noting the following inequality for any $f : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}, g : \mathcal{V} \times \mathcal{V} \rightarrow \mathbb{R}$ s.th. $f, g \in \mathbb{L}$ we have that for all $b \in \mathcal{V}$:

$$\left| \max_{a \in \mathcal{V}} f(a, b) - \max_{a \in \mathcal{V}} g(a, b) \right| \leq \max_{a \in \mathcal{V}} |f(a, b) - g(a, b)| \quad (54)$$

From (54) we can straightforwardly derive the fact that for any $b \in \mathcal{V}$:

$$\left| \min_{a \in \mathcal{V}} f(a, b) - \min_{a \in \mathcal{V}} g(a, b) \right| \leq \max_{a \in \mathcal{V}} |f(a, b) - g(a, b)| \quad (55)$$

(this can be seen by negating each of the functions in (54) and using the properties of the max operator).

Assume that for any $b \in \mathcal{V}$ the following inequality holds:

$$\max_{a \in \mathcal{V}} |f(a, b) - g(a, b)| \leq c \quad (56)$$

Since (55) holds for any $b \in \mathcal{V}$ and, by (54), we have in particular that

$$\begin{aligned} & \left| \max_{b \in \mathcal{V}} \min_{a \in \mathcal{V}} f(a, b) - \max_{b \in \mathcal{V}} \min_{a \in \mathcal{V}} g(a, b) \right| \\ & \leq \max_{b \in \mathcal{V}} \left| \min_{a \in \mathcal{V}} f(a, b) - \min_{a \in \mathcal{V}} g(a, b) \right| \\ & \leq \max_{b \in \mathcal{V}} \max_{a \in \mathcal{V}} |f(a, b) - g(a, b)| \leq c \end{aligned} \quad (57)$$

whenever (56) holds which gives the required result. \square

Proof of Lemma 5.6

PROOF. The proof uses an application of Lemma 5.2.

Part 2 immediately follows from the properties of the max and min operators. It remains only to prove part 1.

We seek to prove that for any $s \in \mathcal{S}$, if $J \leq \bar{J}$ then

$$\begin{aligned} & \min_{\tau \in \mathcal{T}} \left\{ \max_{a \in A} R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a J^{\tau, \pi}[s'], G(S_\tau) \right\} \\ & - \min_{\tau \in \mathcal{T}} \left\{ \max_{a \in A} R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a \bar{J}^\pi[s'], G(S_\tau) \right\} \leq 0. \end{aligned}$$

We begin by firstly making the following observations:

1. For any $x, y, h \in \mathcal{V}$

$$x \leq y \implies \min\{x, h\} \leq \min\{y, h\}. \quad (58)$$

2. For any $f, g, h \in \mathbb{L}_2$

$$\left| \max_{x \in \mathcal{V}} f(x) - \max_{x \in \mathcal{V}} g(x) \right| \leq \max_{x \in \mathcal{V}} |f(x) - g(x)| \quad (59)$$

Assume that $J \leq \bar{J}$, then we observe that

$$\begin{aligned} & \max_{a \in \mathcal{A}} \left\{ R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a J^{\tau, \pi}[s'] \right\} - \max_{a \in \mathcal{A}} \left\{ R_s^a + \gamma \sum_{s' \in \mathcal{S}} P_{ss'}^a \bar{J}^\pi[s'] \right\} \\ & \leq \gamma \max_{a \in A} \left\{ \sum_{s' \in \mathcal{S}} P_{ss'}^a (J^{\tau, \pi}[s'] - \bar{J}^\pi[s']) \right\} \\ & = \gamma ((PJ) - (P\bar{J})) \leq J - \bar{J} \leq 0, \end{aligned}$$

where we have used (59) in the penultimate line. The result immediately follows after applying (58). \square

Proof of Lemma 5.10.

PROOF. The proof is standard, we give the details for the sake of completion. Indeed, using the Tonelli-Fubini theorem and the iterated law of expectations, we have that:

$$\begin{aligned} \|PJ\|^2 &= \mathbb{E} [(PJ)^2[s]] \\ &= \mathbb{E} \left[\mathbb{E} [J[s_1]|s]^2 \right] \leq \mathbb{E} \left[\mathbb{E} [J^2[s_1]|s] \right] = \mathbb{E} [J^2[s_1]] = \|J\|^2, \end{aligned}$$

where we have used Jensen's inequality to generate the inequality. This completes the proof. \square

Proof of Lemma B.4.

PROOF. The proof follows almost immediately from the triangle inequality, indeed for any $J \in \mathbb{L}_2$:

$$\|J^\star - J\| = \|TJ^\star - J\| \leq \gamma \|J^\star - J\| + \|TJ - J\| \quad (60)$$

where we have added and subtracted TJ to produce the inequality. The result then follows after inserting the definition of $\epsilon_T(J)$. \square

Proof of Lemma B.5.

PROOF. The proof follows directly from Lemma B.4. Indeed, we observe that for any $J \in \mathbb{L}_2$ we have

$$\|J_1^\star - J_2^\star\| \leq \|J_1^\star - J\| + \|J_2^\star - J\| \quad (61)$$

where we have added and subtracted J to produce the inequality. The result then follows from Lemma B.4. \square

The proofs of the results in Sec. 9 are constructed in a similar fashion that in (Bertsekas, 2008) (approximate dynamic programming). However, the analysis incorporates some important departures due to the need to accommodate the actions of two players that operate antagonistically.

We now prove the first of the two results of Sec. 9.

Proof of Theorem 6.1.

PROOF. We firstly notice the construction of \hat{t} given by

$$\hat{t} = \min\{t|G(s_t) \leq Q^*\}, \quad (62)$$

is sensible since we observe that

$$\begin{aligned} & \min\{t|G(s_t) \leq J^*\} \\ &= \min\{t|G(s_t) \leq \min\{G(s_t), Q^*(s_t)\}\} \\ &= \min\{t|G(s_t) \leq Q^*\}. \end{aligned}$$

Result 1

Step 1 Our first step is to prove the following bound:

$$\|FQ - F\tilde{Q}\| \leq \gamma \|Q - \tilde{Q}\|. \quad (63)$$

PROOF.

$$\begin{aligned} & \left\| \max_{a \in \mathcal{A}} R_s^a + \gamma P \min\{G, Q\} - \left(\max_{a \in \mathcal{A}} R_s^a + \gamma P \min\{G, \tilde{Q}\} \right) \right\| \\ &= \gamma \|P \min\{G, Q\} - P \min\{G, \tilde{Q}\}\| \\ &\leq \gamma \|\min\{G, Q\} - \min\{G, \tilde{Q}\}\| \\ &\leq \gamma \|Q - \tilde{Q}\|. \end{aligned}$$

which is the required result. \square

Step 2

Our next task is to prove that the quantity Q^* is a fixed point of F and hence we can apply the operator F to achieve the approximation of the value.

PROOF. Using the definition of T (c.f. (13)) we find that:

$$\begin{aligned} J^* &= TJ^* \iff \max_{a \in \mathcal{A}} R_s^a + \gamma PJ^* \\ &= \max_{a \in \mathcal{A}} R_s^a + \gamma P \min \left\{ \max_{a \in \mathcal{A}} R_s^a + \gamma PJ, G \right\} \\ &\iff \\ Q^* &= \max_{a \in \mathcal{A}} R_s^a + \gamma P \min \{Q^*, G\} \\ &\iff \\ Q^* &= FQ^*. \end{aligned}$$

\square

Step 3

We now prove that the operator ΠF is a contraction on Q , that is the following inequality holds:

$$\|\Pi FQ - \Pi F\tilde{Q}\| \leq \gamma \|Q - \tilde{Q}\|.$$

PROOF. The proof follows straightforwardly by the properties of a projection mapping:

$$\|\Pi FQ - \Pi F\tilde{Q}\| \leq \|FQ - F\tilde{Q}\| \leq \gamma \|Q - \tilde{Q}\|. \quad \square$$

Step 4

$$\|\Phi r^* - Q^*\| \leq \frac{1}{\sqrt{1-\gamma^2}} \|\Pi Q^* - Q^*\|. \quad (64)$$

The result is proven using the orthogonality of the (orthogonal) projection and by the Pythagorean theorem. Indeed, we have that:

PROOF.

$$\begin{aligned} \|\Phi r^* - Q^*\|^2 &= \|\Phi r^* - \Pi Q^*\|^2 + \|\Pi Q^* - Q^*\|^2 \\ &= \|\Pi F \Phi r^* - \Pi Q^*\|^2 + \|\Pi Q^* - Q^*\|^2 \\ &= \|\Pi F \Phi r^* - \Pi Q^*\|^2 + \|\Pi Q^* - Q^*\|^2 \\ &\leq \gamma^2 \|\Phi r^* - Q^*\|^2 + \|\Pi Q^* - Q^*\|^2. \end{aligned}$$

Hence, we find that

$$\|\Phi r^* - Q^*\| \leq \frac{1}{\sqrt{1-\gamma^2}} \|\Pi Q^* - Q^*\|,$$

which is the required result. \square

Result 2

$$\mathbb{E}[J^*[s]] - \mathbb{E}[J^{\tilde{t}, \tilde{\pi}}[s]] \leq \frac{2}{(1-\gamma)\sqrt{1-\gamma^2}} \|\Pi Q^* - Q^*\|. \quad (65)$$

PROOF. The proof by Jensen's inequality, stationarity and the non-expansive property of P . In particular, we have

$$\begin{aligned} & \mathbb{E}[J^*[s]] - \mathbb{E}[J^{\tilde{t}, \tilde{\pi}}[s]] \\ &= \mathbb{E}[PJ^*[s]] - \mathbb{E}[PJ^{\tilde{t}, \tilde{\pi}}[s]] \\ &\leq \left| \mathbb{E}[PJ^*[s]] - \mathbb{E}[PJ^{\tilde{t}, \tilde{\pi}}[s]] \right| \\ &\leq \|PJ - PJ^{\tilde{t}, \tilde{\pi}}\|. \end{aligned} \quad (66)$$

Inserting the definitions of Q^* and \tilde{Q} into (66) then gives:

$$\mathbb{E}[J^*[s]] - \mathbb{E}[J^{\tilde{t}, \tilde{\pi}}[s]] \leq \frac{1}{\gamma} \|Q^* - \tilde{Q}\|. \quad (67)$$

It remains therefore to place a bound on the term $\|Q^* - \tilde{Q}\|$. We observe that by the triangle inequality and the fixed point properties of F on Q and \tilde{F} on \tilde{Q} we have

$$\|Q^* - \tilde{Q}\| \leq \|Q^* - F(\Phi r^*)\| + \|\tilde{Q} - F(\Phi r^*)\| \quad (68)$$

$$\leq \gamma \{\|Q^* - \Phi r^*\| + \|\tilde{Q} - \Phi r^*\|\} \quad (69)$$

$$\leq \gamma \{2\|Q^* - \Phi r^*\| + \|Q^* - \tilde{Q}\|\}. \quad (70)$$

So that:

$$\|Q^* - \tilde{Q}\| \leq \frac{2\gamma}{1-\gamma} \|Q^* - \Phi r^*\|. \quad (71)$$

The result then follows after substituting the result of step 4 (64). \square

Let us now define the following quantity:

$$HQ(s) := \begin{cases} G(s) & \text{if } G(s) \leq (\Phi r^*)(s) \\ Q(s) & \text{otherwise,} \end{cases} \quad (72)$$

and

$$\tilde{F}Q := \max_{a \in \mathcal{I}} R_s^a + \gamma PHQ. \quad (73)$$

Step 5

$$\|\tilde{F}Q - \tilde{F}\bar{Q}\| \leq \gamma \|Q - \bar{Q}\| \quad (74)$$

PROOF.

$$\begin{aligned} \|\tilde{F}Q - \tilde{F}\bar{Q}\| &= \left\| \max_{a \in \mathcal{I}} R_s^a + \gamma PHQ - \left(\max_{a \in \mathcal{I}} R_s^a + \gamma PH\bar{Q} \right) \right\| \\ &= \gamma \|PHQ - PH\bar{Q}\| \\ &\leq \gamma \|HQ - H\bar{Q}\| \\ &= \gamma \|\min\{G, Q\} - \min\{G, \bar{Q}\}\| \\ &\leq \gamma \|Q - \bar{Q}\|. \end{aligned}$$

We now prove that $\tilde{Q} = \max_{a \in \mathcal{I}} R_s^a + \gamma PJ^{\pi, \tilde{\tau}}$ is a fixed point.

$$\begin{aligned} H\tilde{Q} &= H \left(\max_{a \in \mathcal{I}} R_s^a + \gamma PJ^{\pi, \tilde{\tau}} \right) \\ &= \begin{cases} G(s) & \text{if } G(s) \leq (\Phi r^\star)(s) \\ \max_{a \in \mathcal{I}} R_s^a + \gamma PJ^{\pi, \tilde{\tau}} & \text{otherwise} \end{cases} \\ &= J^{\pi, \tilde{\tau}} \end{aligned}$$

□

Let us now define the following quantity:

$$s(z, r) := \phi(s) \left(\max_{a \in \mathcal{I}} R_s^a + \gamma \min \{(\Phi r)(y), G(y)\} - (\Phi r)(s) \right).$$

Additionally, we define \bar{s} by the following:

$$\bar{s}(z, r) := \mathbb{E}[s(z_0, r)].$$

The components of $s(z, r)$ are then given by:

$$s_k \equiv \mathbb{E} \left[\phi_k(s) \left(\max_{a \in \mathcal{I}} R_s^a + \gamma \min \{(\phi r)(s), G(s)\} - (\phi r)(s) \right) \right].$$

We now observe that s_k can be described in terms of an inner product. Indeed, using the iterated law of expectations we have that

$$\begin{aligned} s_k &\equiv \mathbb{E} \left[\Phi_k(s) \left(\max_{a \in \mathcal{I}} R_s^a + \gamma \min \{(\Phi r)(s), G(s)\} - (\Phi r)(s) \right) \right] \\ &= \mathbb{E} \left[\Phi_k(s) \left(\max_{a \in \mathcal{I}} R_s^a + \gamma \mathbb{E}[\min \{(\Phi r)(s), G(s)\} | s] - (\Phi r)(s) \right) \right] \\ &= \mathbb{E} \left[\Phi_k(s) \left(\max_{a \in \mathcal{I}} R_s^a + \gamma P \min \{(\Phi r)(s), G(s)\} - (\Phi r)(s) \right) \right] \\ &= \langle \Phi_k, F(\Phi r) - F(\Phi r) \rangle. \end{aligned}$$

□

Proof of Theorem 6.2

Step 5 enables us to use classic arguments for approximate dynamic programming. In particular, following step 5, Theorem 6 follows directly from Theorem 2 in (Tsitsiklis & Van Roy, 1999) with only a minor adjustment in substituting the max operator with min.