

EXPLORATORY OPTIMAL STOPPING: A SINGULAR CONTROL FORMULATION

JODI DIANETTI, GIORGIO FERRARI, AND RENYUAN XU

ABSTRACT. This paper explores continuous-time and state-space optimal stopping problems from a reinforcement learning perspective. We begin by formulating the stopping problem using randomized stopping times, where the decision maker's control is represented by the probability of stopping within a given time—specifically, a bounded, non-decreasing, càdlàg control process. To encourage exploration and facilitate learning, we introduce a regularized version of the problem by penalizing it with the cumulative residual entropy of the randomized stopping time. The regularized problem takes the form of an $(n + 1)$ -dimensional degenerate singular stochastic control with finite-fuel. We address this through the dynamic programming principle, which enables us to identify the unique optimal exploratory strategy. For the specific case of a real option problem, we derive a semi-explicit solution to the regularized problem, allowing us to assess the impact of entropy regularization and analyze the vanishing entropy limit. Finally, we propose a reinforcement learning algorithm based on policy iteration. We show both policy improvement and policy convergence results for our proposed algorithm.

Keywords: Optimal stopping, singular stochastic control, free boundary problem, entropy regularization, reinforcement learning

AMS subject classification: 60G40, 93E20, 35R35, 68T01

CONTENTS

1. Introduction	2
1.1. Our work and contributions	3
1.2. Related literature	5
1.3. Outline of the paper	6
1.4. General notation	6
2. Exploratory formulation and entropy regularization of OS problem	7
2.1. Exploratory formulation via singular controls	8
2.2. Entropy regularization	9
2.3. Vanishing entropy limit	10
3. Solving the entropy regularized OS problem via dynamic programming	12
3.1. Preliminary estimates	14
3.2. Proof of Theorem 3.2: solution of the HJB	15
3.3. Proof of Theorem 3.2: uniqueness for the HJB	23
3.4. Proof of Theorem 3.3	24
4. A real option problem: analytical solutions and reinforcement learning algorithms	25
4.1. Analytical solution	25
4.2. Reinforcement learning algorithm	30
Appendix A. Proof of Lemma 2.3	40
Appendix B. Proof of the auxiliary estimates	41
References	46

1. INTRODUCTION

In optimal stopping (OS) problems, a decision maker chooses a time to take a given action based on an adaptively observed stochastic process in order to optimize an expected performance criterion. Problems of this type are found in the area of Statistics, where the action taken may be to test a hypothesis or to estimate a parameter; in the area of Operations Research, where the action may be to replace a machine, hire a secretary, or reorder stock; and in Mathematical Finance, in the context of pricing American contingent claims or in irreversible investment problems (we refer to the textbook [46] and to the references therein).

Despite the wide-ranging applications of optimal stopping across various fields, most of the existing literature relies on the *full knowledge* of the system, including both the underlying process and the reward function. Little is known about addressing these problems in a model-free context, where decision-makers interact with an unknown system and learn to improve their decisions over time. In this regard, machine learning techniques, particularly reinforcement learning (RL), offer a promising framework. However, the existing literature has largely focused on “frequently controlling the system”, leaving the RL framework for optimal stopping largely unexplored.

Generally speaking, RL aims to learn optimal strategies to control an unknown dynamical system (or random environment) by interacting with the system through exploration and exploitation [53]. In recent years, the availability of vast amounts of data and advances in computational resources have spurred extensive research in both theoretical developments and practical applications of RL. While much of the existing literature on RL has focused on discrete-time problems, continuous-time RL is rapidly growing and presents unique challenges and technical difficulties in algorithmic design, stability analysis, and efficiency guarantees. A key question is how to balance the trade-off between exploring unknown information in the continuous-time system and exploiting current information to take near-optimal actions. In their seminal work, [61] propose adding Shannon’s entropy to the objective function of a continuous-time linear-quadratic control problem to encourage exploration, leading to an optimal policy of Gaussian-type with the mean capturing the exploited policy and the variance quantifying the degree of exploration. Since then, there have been studies on continuous-time RL through the lens of partial differential equations and control theory under entropy regularization, primarily focusing on *regular controls*, such as controlling the drift and volatility coefficients of the underlying stochastic processes [3, 30, 31, 32, 40, 57, 60, 61].

One key challenge that distinguishes OS from regular control is that OS involves a non-smooth decision—whether to stop or continue—while regular controls gradually change the dynamics through drift and/or volatility terms. As a result, gradient-based RL algorithms, although popular for regular controls, cannot be directly applied to the “stop-or-continue” decisions in our setting [50, 52]. To overcome this difficulty, we replace these sharp rules based on the hitting times by stopping probabilities, leading to a *randomized stopping time* that follows the stopping probability. Unlike the heuristic of a fuzzy boundary considered in [50], where the agent stops with a probability proportional to the distance to the boundary, our approach to randomized stopping time is based on principled guidance—regularized objective function with cumulative residual entropy. Another advantage of randomized stopping time is its exploratory nature; it stops at different scenarios according to certain probability, thereby collecting more information from the unknown environment in the RL regime. Compared to classic RL settings, exploration is particularly essential for optimal stopping problems because the terminal reward can only be collected upon making the stopping decision, contributing to the *reward sparsity* challenge in RL [14, 29]. Therefore, randomized stopping time helps the decision-maker learn more about the unknown terminal reward.

Following the above idea, this paper proposes a framework to address continuous time and state-space OS problems via continuous-time RL. We first embed the perspective of exploration-exploitation into OS, and then use this construction to design a new type of RL algorithm for which we provide a convergence analysis.

1.1. Our work and contributions. In this section, we provide an informal discussion on the key ideas and contributions of this work. The precise assumptions and statements of the results can be found in Sections 2-4.

The classical OS problem (without regularization) is given, for any initial state $x \in \mathbb{R}^n$, as

$$V(x) := \sup_{\tau} \mathbb{E} \left[\int_0^{\tau} e^{-\rho s} \pi(X_s^x) ds + e^{-\rho \tau} G(X_{\tau}^x) \right],$$

subject to $dX_t^x = b(X_t^x)dt + \sigma(X_t^x)dW_t, \quad X_0^x = x,$

where the maximization is performed over stopping times τ of the (augmented) filtration generated by the Brownian motion W , $\rho > 0$, and the functions b, σ, π, G are the data of the problem (satisfying the assumptions in Section 2).

To model exploration, we follow the ideas in [59] and let the agent randomize the choice of the stopping time. The resulting strategy are characterized by nondecreasing processes $\xi = (\xi_t)_t$ with $0 \leq \xi_t \leq 1$, where ξ_t is the probability that the agent stops before time t according to $\xi_t = \mathbb{P}(\tau \leq t | \mathcal{F}_t^W)$. We refer to such processes ξ as *singular controls*. However, allowing for randomized stopping times *does not change* the optimal value function and, more importantly, it does not imply that optimal actions are necessarily randomized/exploratory (see Proposition 2.4). From the RL point of view, this non-exploratory behavior of the optimal control suggests that, without proper modification of the objective function, *optimizing* and *collecting information* do not naturally come together.

To overcome the above-mentioned non-exploratory phenomenon and to incentivize exploration, we reward intermediate values of ξ_t by perturbing the original OS problem via the cumulative residual entropy (CRE) [48] of the probability measures ξ

$$\text{CRE}(\xi) := - \int_0^{\infty} e^{-\rho t} (1 - \xi_t) \log(1 - \xi_t) dt.$$

The adoption of CRE in RL is *novel* in the literature, as Shannon's entropy and Kullback–Leibler (KL) divergence are more often used. It is worth noting that this CRE criterion is more suitable for encouraging randomized optimal stopping time by postponing the exercise time via a large cumulative probability.

The resulting *entropy regularized OS problem* takes the form of the singular control problem:

$$V^{\lambda}(x) := \sup_{\xi} \mathbb{E} \left[\int_0^{\infty} e^{-\rho t} \underbrace{(\pi(X_t^x)(1 - \xi_t)dt + G(X_t^x)d\xi_t)}_{\text{exploitation}} - \lambda \int_0^{\infty} e^{-\rho t} \underbrace{(1 - \xi_t) \log(1 - \xi_t) dt}_{\text{exploration}} \right],$$

where $\lambda > 0$ is a temperature parameter that balances the exploitation and exploration. For any $\lambda > 0$, such a problem admits a unique optimal control ξ^{λ} . The proposed regularization naturally approximates the original OS problem. Indeed, we show that (see formal statements in Propositions 2.9 and 2.10):

- $\sup_x |V^{\lambda}(x) - V(x)| \leq \lambda(\rho e)^{-1} \rightarrow 0$ as $\lambda \rightarrow 0$.
- For any sequence $(\lambda_k)_k \rightarrow 0$, the sequence $(\xi^{\lambda_k})_k$ converges (up to subsequences) to an optimal singular control ξ^* of the original OS with randomized stopping times.

These convergence results have also been discovered in [61] for the linear-quadratic control with Shannon's entropy, using explicit formulas for the value functions and optimal controls.

A solution to the entropy regularized OS problem is obtained by applying the dynamic programming principle (DPP). To this end, we introduce the additional controlled state variable

$Y_t^{y,\xi} = y - \xi_t$, $y \in [0, 1]$, and define, for $x \in \mathbb{R}^n$ and $\lambda > 0$, the extended problem

$$V^\lambda(x, y) := \sup_{\xi_t \leq y} \mathbb{E} \left[\int_0^\infty e^{-\rho t} (\pi(X_t^x) Y_t^{y,\xi} - \lambda Y_t^{y,\xi} \log(Y_t^{y,\xi})) dt + \int_0^\infty e^{-\rho t} G(X_t^x) d\xi_t \right].$$

This is an $(n+1)$ -dimensional *degenerate* singular stochastic control problem with finite-fuel. Notice that $V^\lambda(x, 1) = V^\lambda(x)$. In particular, we show that (see Theorems 3.2 and 3.3 for the formal statements):

- The value function V^λ is $W_{loc}^{2,2}(\mathbb{R}^n \times (0, 1))$ and it is the unique solution to the Hamilton-Jacobi-Bellman (HJB) variational inequality

$$\max \{ (\mathcal{L}_x - \rho) V^\lambda(x, y) + \pi(x)y - \lambda y \log y, -V_y^\lambda(x, y) + G(x) \} = 0,$$

with $V^\lambda(x, 0) = 0$, where $\mathcal{L}_x f(x, y) := b(x)D_x f(x, y) + \frac{1}{2} \text{tr}(\sigma \sigma^*(x) D_x^2 f(x, y))$.

- The optimal control ξ^λ is of reflecting type and is given by $\xi_t^\lambda := \sup_{s \leq t} (y - g_\lambda(X_s^x))^+$, with the reflecting free boundary defined as $g_\lambda(x) := \sup \{ y \in [0, 1] \mid -V_y^\lambda(x, y) + G(x) < 0 \}$.

These results illustrate that the optimal control of the entropy regularized OS problem is not related to any strict stopping time; that is, by introducing the entropy regularization, we obtain optimal strategies of exploratory type. From the RL point of view, the exploratory behavior of the optimal control is desirable, as it suggests that *optimizing* and *collecting information* from the environment can actually be achieved together.

From the technical point of view, showing the regularity of V^λ (required for the characterization of ξ^λ) involves several challenges. In particular, while the regularity in x is studied via semi-convexity estimates and PDE arguments, the regularity in y is obtained via a probabilistic connection with a different but related OS problem.

In a benchmark real option example, we can solve the entropy regularized OS problem semi-explicitly (see Theorem 4.2 for the details). For this particular example, we can characterize the free boundary g_λ and show its convergence, as $\lambda \rightarrow 0$, to the free boundary of the original OS problem.

Our theoretical analysis suggests a new approach to design RL algorithms for OS problems, which aims at learning the boundary of the λ -entropy regularized problem rather than the boundary of the original OS problem. Notice that the error introduced by the temperature parameter λ is of order $\mathcal{O}(\lambda)$. The proposed RL framework consists of two steps: a model-based step and a model-free step. In the first step, where all model parameters are known, we design a Policy Iteration algorithm to learn g_λ . In the second step, where the model parameters are unknown, we combine the Policy Iteration with a Policy Evaluation procedure. Notably, we do not directly estimate model parameters in the design, which enhances robustness against model misspecification and environmental shifts [1, 19, 28]. Our new approach also overcomes the instability issue often encountered when directly learning the boundary of the original OS problem [50].

We establish the theoretical foundation of the algorithm for the real option example. Specifically, in the k th iteration of the first step, we use Hessian information of $V_{g_k}^\lambda$ (i.e., the value function associated with the policy using reflecting boundary g_k) to update g_{k+1} , which lies in a region where $V_{g_k}^\lambda$ has a *better regularity*. Mathematically, define the new boundary $g_{k+1}(x)$ as:

$$\begin{cases} \max \left\{ y^* < g_k(x) \mid \partial_{xy} V_{g_k}^\lambda(x, y) = 0 \right\} & \text{if } \partial_{xy} V_{g_k}^\lambda(x, g_k(x)) < 0, \\ g_{k+1}(x) = g_k(x) & \text{o.w.,} \end{cases}$$

where we define the notation $\partial_{xy}^- f(x, y) := \lim_{h \rightarrow 0^-} \frac{\partial_x f(x, y+h) - \partial_x f(x, y)}{h}$ as the left y -derivative of $\partial_x f$ (if exists). With such an updating scheme, we have the following results (see Proposition 4.5 and Theorem 4.9 for formal statements):

- **Policy improvement:** $V_{g_{k+1}}^\lambda(x, y) \geq V_{g_k}^\lambda(x, y)$ for all $k \in \mathbb{N}$.
- **Policy convergence:** $\lim_{k \rightarrow \infty} V_{g_k}^\lambda(x, y) = V^\lambda(x, y)$.

Policy improvement is a result of Itô’s formula in the weak version. A similar result is shown in [61] for linear-quadratic controls. On the other hand, it is well-recognized in the literature that policy convergence is much harder to establish, with only a few exemptions for regular controls in the literature [3, 30, 40, 60]. Our result on policy convergence is established based on an induction argument that iteratively verifies the regularity, convexity, and a few estimation bounds held along the entire updating scheme.

In the second step, we combine the Policy Iteration with Policy Evaluation, where the value function of a given policy with reflecting barrier g_k is estimated using sample trajectories. Finally, we demonstrate the performance of the proposed algorithm on a few numerical examples.

1.2. Related literature. Our work is related to different streams of literature.

Randomized optimal stopping and singular control. First of all, we relate to the literature on randomized stopping times. These typically appear when proving the existence of equilibria in zero-sum or nonzero-sum optimal stopping games (see [12, 37, 51, 59], among others) and can be thought of as the “mixed strategies” in stopping times. Loosely speaking, according to a randomized stopping time, a decision maker chooses a probability of stopping and at any time decides or not to stop according to that distribution. As discussed in [51] and [59], different but equivalent definitions of randomized stopping times can be developed. In this paper, we follow the approach in [59], which identifies randomized stopping times with nondecreasing càdlàg processes starting from zero at the initial time and bounded above by one. This in turn leads to the fact that the exploratory formulation of the optimal stopping problem is mathematically equivalent to an $(n+1)$ -dimensional degenerate singular stochastic control with finite-fuel (see, e.g., [11, 34, 36]).

Continuous-time RL for regular control. Initial studies on continuous-time RL primarily focused on algorithm design and the potential for time discretization [17, 45, 56].

Following the seminal work of [61], which provided a mathematical perspective within the framework of linear-quadratic control, there has been a recent surge in employing stochastic control theory and partial differential equations (PDEs) to establish the theoretical foundation of continuous-time RL [3, 30, 32, 31, 40, 57, 60]. In particular, [57] generalized the framework in [61] and studied the exploratory Hamilton–Jacobi–Bellman (HJB) equation arising from a general entropy-regularized stochastic control problem; [32] developed q-Learning algorithm in continuous time; [31] established policy gradient and actor-critic learning frameworks in continuous time and space. Despite these developments, the overall convergence results of control-inspired or PDE-inspired algorithms, such as the Policy Improvement Algorithm (PIA), remain largely unexplored. Very recently, [30] proved a qualitative convergence result for the PIA with bounded coefficients when the diffusion term is not controlled. [60] generalized the result to unbounded coefficients or controlled diffusion terms by a uniform estimate for the value sequence generated by PIA. [40] provides a much simpler proof of the convergence of the PIA using Feynman-Kac-type probabilistic representation instead of sophisticated PDE estimates. Moreover, [3] showed the convergence of the PIA for an optimal dividend problem, taking advantage that the state process is one dimensional and taking non-negative values. In addition to the PIA method, another popular continuous-time RL algorithm is the policy gradient method. [49] demonstrated a linear convergence of this method for certain finite-horizon

non-linear control problems, while [24] showed its convergence for finite-horizon exploratory linear-quadratic control problems.

Another important aspect of quantifying the effectiveness of continuous-time RL algorithms is regret analysis. For theoretical developments in this area, see [6, 26, 54, 44, 55].

We note that all the works mentioned above focus on regular controls.

Machine learning for optimal stopping. RL for optimal stopping is closely related to the challenging RL scenario of sparse rewards [14, 29]. More specifically, the terminal reward $G(X_\tau)$ can only be collected upon making the stopping decision, contributing to the sparsity of the rewards. This sparsity significantly complicates the learning process compared to the more straightforward regular control problems in continuous-time settings or classical Markov decision processes in discrete time.

When model parameters are fully known to the decision maker, [50] and [52] developed deep-learning-based learning algorithms to learn the stopping boundaries. See also [2] for solving high-dimensional singular control problems using deep learning methods.

It is worth noting that [13] proposed a comprehensive framework for policy gradient methods tailored to continuous-time RL. This framework leverages the connection between stochastic control problems and randomized problems, allowing for applications beyond diffusion models, such as regular, impulse, and optimal stopping/switching problems. However, no theoretical convergence results have been established yet. The most relevant paper to our setting is [16], which considers a regularized American Put Option problem under Shannon's entropy in one dimension, which involves a single risky asset that follows a geometric Brownian motion. The author employs an intensity control formulation to encourage exploration and demonstrates the convergence of PIA for the regularized problem under a given temperature parameter. In comparison to the latter work, we highlight some novelties of our approach. Notably, our analysis of the regularized problem is based on regularity analysis that combines PDE arguments and a probabilistic connection with another OS problem. In contrast, the analysis in [16] relies primarily on explicit calculations, which are challenging to extend to higher dimensions. Moreover, within our framework, we can show the convergence of the optimal policies of the regularized problems to the original OS-optimal policy as the temperature parameter λ tends to zero. While our convergence holds regardless of the dimension of the problem, such a convergence is not clear in [16] even for the specific American Put Option problem.

1.3. Outline of the paper. Section 2 examines the properties of the original OS problem, introduces the entropy-regularized version, and explores its relationship with the original problem. The entropy-regularized problem is further explored in Section 3, including a regularity analysis of the value function and the construction of the optimal policy. Section 4 then shifts focus to a real option example, establishing its analytical properties and developing corresponding RL algorithms. Finally, Appendix A contains the proof of the uniqueness of the optimal stopping time, while the proofs of some auxiliary technical estimates are collected in Appendix B.

1.4. General notation. For $q \in [1, \infty]$, and a measure space (E, \mathcal{E}, m) , we define the set $\mathbb{L}^q = \mathbb{L}^q(E) = \mathbb{L}^q(E, m)$ of measurable functions $f : E \rightarrow \mathbb{R}$ s.t. $\int_E |f|^q dm < \infty$, if $q < \infty$, and $\text{ess sup}_E |f| < \infty$ for $q = \infty$. For $d \in \mathbb{N} \setminus \{0\}$, an open set $B \subset \mathbb{R}^d$, $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{N}^d$, and a function $f : B \rightarrow \mathbb{R}$, we denote by $D^\alpha f := D_1^{\alpha_1} \dots D_d^{\alpha_d} f$ the weak derivative of f , where $D_i f := f_{x_i} := \partial f / \partial x_i$, and we set $|\alpha| := \alpha_1 + \dots + \alpha_d$. Sometimes, we will denote by $D_x f$ and $D_x^2 f$ the gradient and the Hessian matrix of f (in the weak sense), respectively. Define the set $C(B)$ of continuous functions $f : B \rightarrow \mathbb{R}$, the set $C^\ell(B)$ of functions $f : B \rightarrow \mathbb{R}$ with continuous ℓ -order derivatives, and the Sobolev spaces $W^{\ell, q}(B)$ of functions $f \in \mathbb{L}^q(B)$ with $\sum_{|\alpha| \leq \ell} \int_E |D^\alpha f|^q dm < \infty$, if $q < \infty$, and $\sum_{|\alpha| \leq \ell} \text{ess sup}_E |D^\alpha f| < \infty$, for $q = \infty$. Finally,

define the Sobolev spaces $W_{loc}^{\ell,q}(B)$ of functions f s.t. $f \in W^{\ell,q}(D)$ for each bounded open set $D \subset B$.

2. EXPLORATORY FORMULATION AND ENTROPY REGULARIZATION OF OS PROBLEM

Consider a discount factor $\rho > 0$ and, for $n \in \mathbb{N} \setminus \{0\}$, continuous functions $b : \mathbb{R}^n \rightarrow \mathbb{R}$, $\sigma : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$, $\pi, G : \mathbb{R}^n \rightarrow \mathbb{R}$. Let $W = (W_t)_t$ be an m -dimensional Brownian motion on a complete probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ and let $\mathbb{F}^W = (\mathcal{F}_t^W)_t$ be the right-continuous extension of the filtration generated by W . Introduce the second-order differential operator

$$\mathcal{L}_x \phi(x) := b(x)D_x \phi(x) + \frac{1}{2} \text{tr}(\sigma^* \sigma(x) D_x^2 \phi(x)), \quad \phi \in W_{loc}^{2,2}(\mathbb{R}^n),$$

and denote by \mathcal{T} the set of \mathbb{F}^W -stopping times.

For an initial condition $x \in \mathbb{R}^n$ and a stopping time $\tau \in \mathcal{T}$, define the profit functional

$$J(x; \tau) := \mathbb{E} \left[\int_0^\tau e^{-\rho s} \pi(X_s^x) ds + e^{-\rho \tau} G(X_\tau^x) \right],$$

$$\text{subject to } dX_t^x = b(X_t^x)dt + \sigma(X_t^x)dW_t, \quad X_0^x = x,$$

and consider the optimal stopping problem

$$(2.1) \quad V(x) := \sup_{\tau \in \mathcal{T}} J(x; \tau).$$

A stopping time τ^* is said to be optimal if $J(x; \tau^*) = V(x)$ and the function V is referred to as the value function.

In order to have a well-defined problem, we enforce the some requirements on the data of the problem.

Assumption 2.1. *The following conditions hold true:*

(1) *There exists a constant $L > 0$ such that, for $\phi = b, \sigma$, we have*

$$|\phi(x)| \leq L(1 + |x|), \quad |\phi(\bar{x}) - \phi(x)| \leq L|\bar{x} - x|, \quad \text{for any } x, \bar{x} \in \mathbb{R}^n.$$

(2) *There exist $p \geq 2$ and $K > 0$ such that, for $\phi = G, \pi$, we have*

$$|\phi(x)| \leq K(1 + |x|^p), \quad \text{for any } x \in \mathbb{R}^n.$$

(3) *ρ is large enough: $\rho > p(\frac{3}{2}L + (2p-1)L^2)$.*

The following result is classical.

Proposition 2.2. *Under Assumption 2.1, there exists an optimal stopping time.*

Proof. We limit ourselves to directing the reader to a reference. More specifically, Assumption 2.1 implies (see estimate (3.5) below and the related (3.3)) that, for a suitable constant $C < \infty$, we have

$$\mathbb{E} \left[\int_0^\infty e^{-\rho t} |\pi(X_t^x)| dt + \sup_{t \geq 0} e^{-\rho t} |G(X_t^x)| \right] < C(1 + |x|^p).$$

Hence, the result follows by Corollary 2.9 at p. 46 in [46] (notice that, differently from [46], we do not require stopping times to be finite). \square

Some of the later results (see in particular Proposition 2.4 and Proposition 2.10 below) enjoy a sharper statement under uniqueness of the optimal stopping time, which is often satisfied in the one-dimensional case. For the sake of illustration, we discuss the following sufficient conditions.

Lemma 2.3. *Under Assumption 2.1, if*

- (2.2) *$n = 1$, $G \in C^2(\mathbb{R})$, $\sigma^2(x) \geq c > 0$ and $D_x G(x) \leq C(1 + |x|)$ for some $c, C \in (0, \infty)$, the function $\hat{\pi} := \pi + (\mathcal{L}_x - \rho)G$ is strictly increasing and there exists \bar{x} s.t. $\hat{\pi}(\bar{x}) = 0$, then the optimal stopping time is unique.*

The proof of this lemma relies on classical arguments in OS and it is provided in Appendix A.

2.1. Exploratory formulation via singular controls. Inspired by [59], we introduce a notion of randomized/exploratory stopping times. Define the set of *singular controls* as the set of processes

$$\mathcal{A}(1) := \left\{ \xi : \Omega \times [0, \infty) \rightarrow [0, 1], \mathbb{F}^W\text{-adapted, nondecreasing, càdlàg, with } \xi_{0-} = 0 \right\},$$

and assume the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ to be large enough to accommodate a W -independent uniformly distributed random variable $U : \Omega \rightarrow [0, 1]$.

Given any $\xi \in \mathcal{A}(1)$, we consider a random variable τ^ξ by setting $\tau^\xi := \inf \{t \geq 0 \mid \xi_t > U\}$, with the convention $\inf \emptyset = +\infty$, and we will refer to it as *randomized/exploratory stopping time*. Notice that τ^ξ is not necessarily an \mathbb{F}^W -stopping time. However, if $\tau \in \mathcal{T}$, the process $\xi^\tau \in \mathcal{A}(1)$ defined by $\xi^\tau := (\mathbb{1}_{\{t \geq \tau\}})_t$ is such that $\tau = \tau^{\xi^\tau}$, which gives a natural inclusion of \mathcal{T} into $\mathcal{A}(1)$. Moreover, we have

$$\mathbb{P}(\tau^\xi \leq t \mid \mathcal{F}_t^W) = \mathbb{P}(U \leq \xi_t \mid \mathcal{F}_t^W) = \int_0^{\xi_t} du = \xi_t,$$

so that ξ_t can be interpreted as the probability of stopping before time t . Notice also that $\xi_\infty := \lim_{t \rightarrow \infty} \xi_t$ is not necessarily equal to 1: indeed, the related randomized stopping time τ^ξ is not necessarily finite.

For $\xi \in \mathcal{A}(1)$, by evaluating the randomized stopping time τ^ξ , we obtain

$$\begin{aligned} J(x; \tau^\xi) &= \mathbb{E} \left[\int_0^{\tau^\xi} e^{-\rho t} \pi(X_t^x) dt + e^{-\rho \tau^\xi} G(X_{\tau^\xi}^x) \right] \\ &= \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) \mathbb{1}_{\{t \leq \tau^\xi\}} dt + \int_0^\infty e^{-\rho t} G(X_t^x) \mathbb{P}(\tau^\xi \in dt \mid \mathcal{F}^W) \right] \\ &= \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) \mathbb{1}_{\{\xi_t < U\}} dt + \int_0^\infty e^{-\rho t} G(X_t^x) d\xi_t \right] \\ &= \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) (1 - \xi_t) dt + \int_0^\infty e^{-\rho t} G(X_t^x) d\xi_t \right]. \end{aligned}$$

This suggests to define a profit functional in terms of the singular controls $\xi \in \mathcal{A}(1)$ by

$$J^0(x; \xi) := \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) (1 - \xi_t) dt + \int_0^\infty e^{-\rho t} G(X_t^x) d\xi_t \right], \quad \xi \in \mathcal{A}(1),$$

as a natural extension of J to randomized stopping times.

However, allowing for randomized stopping times does not change the optimal value and (crucial for our story) it does not imply that optimal actions are necessarily randomized, as discussed in the next proposition.

Proposition 2.4. *Under Assumption 2.1, for any $x \in \mathbb{R}^n$ we have*

$$V(x) = \sup_{\xi \in \mathcal{A}(1)} J^0(x; \xi) = J^0(x; \xi^*),$$

with $\xi^* := (\mathbb{1}_{\{t \geq \tau^*\}})_t$ and τ^* optimal for $J(x; \cdot)$. Moreover, if the optimal stopping time τ^* is unique (e.g., if (2.2) holds), then $\xi^* := (\mathbb{1}_{\{t \geq \tau^*\}})_t$ is the unique optimal control for $J^0(x; \cdot)$.

Proof. For generic $\xi \in \mathcal{A}(1)$ and $z \in [0, 1]$, we can define the stopping time $\tau^\xi(z)$ as

$$\tau^\xi(z) := \inf\{t \geq 0 \mid \xi_t \geq z\},$$

with the convention $\inf \emptyset := +\infty$. Indeed, notice that $\tau^\xi(z) = +\infty$ if $z > \xi_\infty$. By using Fubini's theorem and then the change of variable formula, we have

$$\begin{aligned} J^0(x; \xi) &= \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) \left(1 - \xi_\infty + \int_t^\infty d\xi_s \right) dt + \int_0^\infty e^{-\rho t} G(X_t^x) d\xi_t \right] \\ &= \mathbb{E} \left[(1 - \xi_\infty) \int_0^\infty e^{-\rho t} \pi(X_t^x) dt + \int_0^\infty \left(\int_0^t e^{-\rho s} \pi(X_s^x) ds + e^{-\rho t} G(X_t^x) \right) d\xi_t \right] \\ &= \mathbb{E} \left[(1 - \xi_\infty) \int_0^\infty e^{-\rho t} \pi(X_t^x) dt + \int_0^{\xi_\infty} \left(\int_0^{\tau^\xi(z)} e^{-\rho s} \pi(X_s^x) ds + e^{-\rho \tau^\xi(z)} G(X_{\tau^\xi(z)}^x) \right) dz \right] \\ &= \int_0^1 \mathbb{E} \left[\int_0^{\tau^\xi(z)} e^{-\rho s} \pi(X_s^x) ds + e^{-\rho \tau^\xi(z)} G(X_{\tau^\xi(z)}^x) \right] dz = \int_0^1 J(x; \tau^\xi(z)) dz. \end{aligned}$$

Thus, since $\tau^\xi(z)$ is a stopping time, by the optimality of τ^* we have

$$J^0(x; \xi) \leq \int_0^1 \sup_{\tau \in \mathcal{T}} J(x; \tau) dz = J(x; \tau^*) = V(x) = J^0(x; \xi^*),$$

where the last equality follows from the fact that ξ^* is such that $\tau^{\xi^*}(z) = \tau^*$ for any $z \in (0, 1]$. This proves the first part of the proposition.

Assume now that the optimal stopping time τ^* is unique. By repeating the previous argument with ξ being optimal, we obtain that $J(x; \tau^\xi(z)) = J(x; \tau^*)$ dz-a.e. in $(0, 1)$. By uniqueness of the optimal stopping time, we deduce that $\tau^\xi(z) = \tau^*$ dz-a.e. in $(0, 1)$, which in turn implies $\xi = \xi^*$. \square

Remark 2.5 (Non-exploratory behavior of the optimal controls). *We elaborate on Proposition 2.4 from the RL and exploration point of view. Under the uniqueness of the optimal stopping time τ^* , the optimal singular control ξ^* in fact corresponds to the strict stopping time τ^* . Thus, no randomization is needed for the optimal strategy. In particular, once an action becomes necessary, the agent stops the process, which prevents the gradual collection of information on the performance of other actions. From the RL point of view, this non-exploratory behavior of the optimal control ξ^* suggests that optimizing and collecting information do not naturally come together, and a modification is necessary in order to overcome this phenomenon (see the related Remark 3.5 below for a discussion on the consequences and benefits of our entropy regularization).*

2.2. Entropy regularization. In light of Proposition 2.4 and of the related Remark 2.5, we introduce a regularization term to incentivize exploration/randomization. This is achieved by incorporating an entropy term to regularize the problem, drawing motivation from the RL literature.

Since the sample paths of the processes $\xi \in \mathcal{A}(1)$ are not necessarily absolutely continuous with respect to the Lebesgue measure, we choose (a discounted version of) the *cumulative residual entropy* (see [48]) weighted by a parameter $\lambda > 0$; namely, we consider, for $\xi \in \mathcal{A}(1)$ the entropy

$$\begin{aligned} \Lambda^\lambda(\tau^\xi) &:= -\lambda \int_0^\infty e^{-\rho t} \mathbb{P}(\tau^\xi \geq t \mid \mathcal{F}_t^W) \log \left(\mathbb{P}(\tau^\xi \geq t \mid \mathcal{F}_t^W) \right) dt \\ (2.3) \quad &= -\lambda \int_0^\infty e^{-\rho t} (1 - \xi_t) \log(1 - \xi_t) dt =: \Lambda^\lambda(\xi). \end{aligned}$$

The entropy Λ^λ is nonnegative and achieves its highest values when the probability ξ_t is near the level e^{-1} , thus incentivising the use of randomized stopping times.

Building on this intuition, we define an entropy regularized OS problem

$$(2.4) \quad V^\lambda(x) := \sup_{\xi \in \mathcal{A}(1)} J^\lambda(x; \xi),$$

where the exploration–exploitation trade off is captured by the profit functional

$$J^\lambda(x; \xi) := \mathbb{E} \left[\int_0^\infty e^{-\rho t} \underbrace{(\pi(X_t^x)(1 - \xi_t)dt + G(X_t^x)d\xi_t)}_{\text{exploitation}} - \lambda \int_0^\infty e^{-\rho t} \underbrace{(1 - \xi_t) \log(1 - \xi_t) dt}_{\text{exploration}} \right].$$

A control $\xi^\lambda \in \mathcal{A}(1)$ is said to be optimal if $J^\lambda(x; \xi^\lambda) = V^\lambda(x)$.

Before studying the entropy regularized problem and its connections with the original stopping problem, we list here few remarks.

Remark 2.6. *For any $x \in \mathbb{R}^n$ and $\lambda > 0$, the existence of a unique optimal control ξ^λ can be established by exploiting the strict concavity of the functional $J^\lambda(x; \cdot)$ in the control variable ξ (see in particular Theorem 8 in [43]). Moreover, Theorem 3.3 below will characterize such an optimal control and illustrate that the randomized strategy maximizing J^λ does not correspond to a (classical) stopping time (see also Remark 3.5).*

Remark 2.7 (Connection between optimal stopping and singular control). *It is well-known in the stochastic control literature that optimal stopping problems and singular control problems enjoy intimate connections. On the one hand, the derivative (in the controlled state variable) of the value function of a singular control problem can be written as the value function of an optimal stopping problem [7, 8, 10, 27, 35]. On the other hand, randomized stopping strategies can be naturally seen as finite-fuel singular controls (see [59] or the discussion above), which has the advantage of compactifying the space of stopping strategies (thus explaining its heavy use in the Dynkin game literature).*

In this framework, our entropy regularization adds two crucial properties on the latter connection. First, while solely considering randomized strategies still leads to a degenerate singular control problem (where the optimal actions are pure jumps, as in Proposition 2.4), our cumulative residual entropy regularizes the problem by making it concave and allows for a more explicit characterization of the optimal control in terms of a reflecting strategy (see Theorem 3.3 below). Secondly, this entropy regularization very well approximates the original OS problem, both in terms of value and optimal strategy (see Propositions 2.9 and 2.10 in the next subsection). For these reasons, our entropy regularization seems to be a useful tool in order to approximate/characterize the equilibria also in the context of Dynkin games, which we leave for future research.

Remark 2.8. *Our choice to work with a Markovian formulation is motivated by RL. Since the typical optimization problem that RL aims to solve (in the discrete time case) is a Markov decision problem, starting our study with a Markovian optimal stopping problem seems to be a natural choice. However, it should be pointed out that our entropy regularization formulation (and its subsequent analysis) also allows us to treat optimal stopping problems in a non-Markovian framework. In this regard, we limit ourselves to mentioning that (while we will solve the Markovian entropy-regularized problem using the DDP approach in Section 3 below), in order to tackle the non-Markovian entropy-regularized OS problem, one can rely on representation theorems via backward stochastic differential equations (see [4, 5] among others).*

2.3. Vanishing entropy limit. With elementary arguments, one can show that the entropy regularized problem approximates the original OS problem as the entropy tends to zero.

Proposition 2.9. *Under Assumption 2.1, for any $\lambda, \bar{\lambda} \in [0, 1]$ we have*

$$\sup_{x \in \mathbb{R}^n} |V^\lambda(x) - V^{\bar{\lambda}}(x)| \leq |\lambda - \bar{\lambda}|(\rho e)^{-1}.$$

In particular, $V^\lambda \rightarrow V$ uniformly on \mathbb{R}^n , as $\lambda \rightarrow 0$.

Proof. Take $x \in \mathbb{R}$ and $\lambda, \bar{\lambda} \in [0, 1]$. If $\lambda \geq \bar{\lambda}$, for every $\xi \in \mathcal{A}(1)$ we have $J^\lambda(x; \xi) \geq J^{\bar{\lambda}}(x; \xi)$, so that

$$0 \leq V^\lambda(x) - V^{\bar{\lambda}}(x).$$

Moreover, upper bounding the right-hand side above, we find

$$\begin{aligned} 0 \leq V^\lambda(x) - V^{\bar{\lambda}}(x) &= \sup_{\xi \in \mathcal{A}(1)} J^\lambda(x; \xi) - \sup_{\bar{\xi} \in \mathcal{A}(1)} J^{\bar{\lambda}}(x; \bar{\xi}) \\ &\leq \sup_{\xi \in \mathcal{A}(1)} \left(J^\lambda(x; \xi) - J^{\bar{\lambda}}(x; \xi) \right) \\ &= (\lambda - \bar{\lambda}) \sup_{\xi \in \mathcal{A}(1)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} (-(1 - \xi_t) \log(1 - \xi_t)) dt \right] \\ &= (\lambda - \bar{\lambda}) \frac{1}{\rho} \sup_{z \in [0, 1]} (-z \log z) = (\lambda - \bar{\lambda})(\rho e)^{-1}, \end{aligned}$$

which is the claim bound. We conclude the proof by choosing $\bar{\lambda} = 0$ and taking limits as $\lambda \rightarrow 0$. \square

Proposition 2.10. *For any $x \in \mathbb{R}^n$ and $\lambda > 0$, let ξ^λ be the unique optimal control as in Remark 2.6. Under Assumption 2.1, the following statements hold true:*

- (1) *There exist a subsequence $(\lambda_k)_k$ with $\lambda_k \rightarrow 0$ and an optimal control ξ^* for $J^0(x; \cdot)$ such that $\xi^{\lambda_k} \rightarrow \xi^*$ weakly in $\mathbb{L}^2(\Omega \times [0, \infty))$ as $k \rightarrow \infty$.*
- (2) *If the optimal stopping time τ^* is unique (e.g., if (2.2) holds), then $\xi^\lambda \rightarrow (\mathbb{1}_{\{t \geq \tau^*\}})_t$ weakly in $\mathbb{L}^2(\Omega \times [0, \infty))$ as $\lambda \rightarrow 0$.*

Proof. Since the family $(\xi^\lambda)_\lambda$ is bounded, we can find a subsequence $(\lambda_k)_k$ with $\lambda_k \rightarrow 0$ and a limit point ξ^* such that $\xi^{\lambda_k} \rightarrow \xi^*$ weakly in $\mathbb{L}^2(\Omega \times [0, \infty))$ as $k \rightarrow \infty$. We need to show that ξ^* is optimal for $J^0(x; \cdot)$.

First of all, using Proposition 2.9, find

$$\begin{aligned} |V(x) - J(x; \xi^{\lambda_k})| &\leq |V(x) - J^{\lambda_k}(x; \xi^{\lambda_k})| + |J^{\lambda_k}(x; \xi^{\lambda_k}) - J(x; \xi^{\lambda_k})| \\ (2.5) \quad &\leq |V(x) - V^{\lambda_k}(x)| + \mathbb{E}[\Lambda^{\lambda_k}(\xi^{\lambda_k})] \\ &\leq C\lambda_k \rightarrow 0, \quad \text{as } k \rightarrow \infty. \end{aligned}$$

Moreover, since $\sup_{t \geq 0} \xi_t^{\lambda_k} \leq 1$, we can employ (after some minimal adjustment to take care of our infinite time horizon $[0, \infty)$) Lemma 3.5 in [33] in order to find a subsequence (not relabelled) $(\lambda_k)_k$ with $\lambda_k \rightarrow 0$ and a limit point $\xi \in \mathcal{A}(1)$ such that, setting $\zeta_t^m := \frac{1}{m} \sum_{k=1}^m \xi_t^{\lambda_k}$, we have

$$\int_0^\infty e^{-\rho t} f_t d\zeta_t^m \rightarrow \int_0^\infty e^{-\rho t} f_t d\xi_t, \quad \mathbb{P}\text{-a.s., as } m \rightarrow \infty, \text{ for any } f \in C_b([0, \infty)).$$

It is easy to show that $\xi = \xi^*$. Since $\zeta^m \in \mathcal{A}(1)$, by Proposition 2.4 and the concavity of $J^0(x; \cdot)$ we find

$$\begin{aligned} V(x) &\geq J^0(x; \zeta^m) \geq \frac{1}{m} \sum_{k=1}^m J^0(x; \xi^{\lambda_k}) \\ &= \frac{1}{m} \sum_{k=1}^m J^{\lambda_k}(x; \xi^{\lambda_k}) + \frac{1}{m} \sum_{k=1}^m (J^0(x; \xi^{\lambda_k}) - J^{\lambda_k}(x; \xi^{\lambda_k})) \\ &= \frac{1}{m} \sum_{k=1}^m J^{\lambda_k}(x; \xi^{\lambda_k}) - \frac{1}{m} \sum_{k=1}^m \mathbb{E}[\Lambda^{\lambda_k}(\xi^{\lambda_k})] \\ &\geq \frac{1}{m} \sum_{k=1}^m J^{\lambda_k}(x; \xi^{\lambda_k}) - (e\rho)^{-1} \frac{1}{m} \sum_{k=1}^m \lambda_k \rightarrow V(x), \quad \text{as } m \rightarrow \infty, \end{aligned}$$

where the limits in the right-hand side follows from (2.5). Thus, we obtain $\lim_m J^0(x; \zeta^m) = V(x)$. By using the convergence of ζ^m to ξ^* and the dominated convergence theorem, the latter limits implies that

$$J^0(x; \xi^*) = \lim_m J^0(x; \zeta^m) = V(x),$$

proving the optimality of ξ^* (cf. Proposition 2.4).

If the optimal stopping time τ^* is unique, then, by Proposition 2.4, the optimal randomized stopping time ξ^* is unique as well. Therefore, ξ^* is the unique limit point of any subsequence of $(\xi^\lambda)_\lambda$, so that the whole sequence converges to ξ^* . \square

Remark 2.11. *From the RL algorithmic design point of view, it is important to notice that Proposition 2.9 quantifies the error which is made when replacing the original OS problem (2.1) with the entropy regularized OS problem (2.4).*

Remark 2.12. *Regarding Propositions 2.9 and 2.10, it is natural to question whether, and in what sense, the reflecting boundary of the entropy-regularized OS problem (2.4) converges to the optimal stopping boundary of the original OS problem (2.1). Although establishing this result at a general level may be challenging, Section 4 (specifically the discussion in Subsection 4.1.3) offers a positive answer in a classical example of an OS problem arising from real option evaluation.*

3. SOLVING THE ENTROPY REGULARIZED OS PROBLEM VIA DYNAMIC PROGRAMMING

To approach problem (2.4) via DPP we introduce, for $y \in [0, 1]$, the additional controlled state process $Y_t^{y, \xi} = y - \xi_t$ and the set of admissible controls

$$(3.1) \quad \mathcal{A}(y) := \{\xi \in \mathcal{A}(1) \text{ with } \xi_t \in [0, y] \text{ a.s. for any } t \geq 0\}.$$

Thus, for $(x, y) \in \mathbb{R}^n \times [0, 1]$ and $\xi \in \mathcal{A}(y)$, define the profit functional

$$J^\lambda(x, y; \xi) := \mathbb{E} \left[\int_0^\infty e^{-\rho t} ((\pi(X_t^x) Y_t^{y, \xi} - \lambda Y_t^{y, \xi} \log(Y_t^{y, \xi})) dt + \int_0^\infty e^{-\rho t} G(X_t^x) d\xi_t \right],$$

$$\text{subject to } dX_t^x = b(X_t^x)dt + \sigma(X_t^x)dW_t, \quad X_0^x = x,$$

$$dY_t^{y, \xi} = -d\xi_t, \quad Y_{0-}^{y, \xi} = y.$$

We then consider the optimization problem

$$(3.2) \quad V^\lambda(x, y) := \sup_{\xi \in \mathcal{A}(y)} J^\lambda(x, y; \xi),$$

and we notice that $V^\lambda(x) = V^\lambda(x, 1)$.

In order to solve the entropy regularized problem, we introduce the following requirements, which imply Assumption 2.1.

Assumption 3.1. *The following conditions hold true:*

(1) *There exists a constant $L > 0$ such that, for $\phi = b, \sigma$, we have*

$$\begin{aligned} |\phi(x)| &\leq L(1 + |x|), \\ |\phi(\bar{x}) - \phi(x)| &\leq L|\bar{x} - x|, \\ |\phi(\delta\bar{x} + (1 - \delta)x) - \delta\phi(\bar{x}) - (1 - \delta)\phi(x)| &\leq L\delta(1 - \delta)|\bar{x} - x|^2, \end{aligned}$$

for any $x, \bar{x} \in \mathbb{R}^n$, $\delta \in [0, 1]$.

(2) *There exist $p \geq 2$ and $K > 0$ such that, for $\phi = G, \pi$, we have*

$$\begin{aligned} |\phi(x)| &\leq K(1 + |x|^p), \\ |\phi(\bar{x}) - \phi(x)| &\leq K(1 + |\bar{x}|^{p-1} + |x|^{p-1})|\bar{x} - x|, \\ |\phi(\delta\bar{x} + (1 - \delta)x) - \delta\phi(\bar{x}) - (1 - \delta)\phi(x)| &\leq K\delta(1 - \delta)(1 + |\bar{x}|^{p-2} + |x|^{p-2})|\bar{x} - x|^2, \end{aligned}$$

for any $x, \bar{x} \in \mathbb{R}^n$, $\delta \in [0, 1]$.

(3) *The matrix $a(x) := \sigma^* \sigma(x)$ is uniformly elliptic; that is, there exists a constant $\kappa_\sigma > 0$ such that*

$$\sum_{i,j=1}^n a_{i,j}(x) z_i z_j \geq \kappa_\sigma |z|^2, \quad \text{for any } z \in \mathbb{R}^n.$$

(4) *ρ is large enough: namely,*

$$\rho > \max \left\{ \frac{\hat{c}_0(4(p-1))}{2}, \hat{c}_1(2), \hat{c}_1(4), \frac{\hat{c}_0(2(p-1)) + \hat{c}_1(2)}{2}, \frac{\hat{c}_0(2p)}{2}, \hat{c}_0(p), \right. \\ \left. \frac{\hat{c}_0(2(p-1)) + \hat{c}_2(2)}{2}, \hat{c}_2(2), \frac{\hat{c}_2(4)}{2}, \frac{\hat{c}_1(8)}{2}, \frac{\hat{c}_0(4(p-2))}{2} \right\},$$

where, for generic q we define the constants

$$\begin{aligned} \hat{c}_0(q) &:= q \left(\frac{3}{2}L + (q-1)L^2 \right), \\ \hat{c}_1(q) &:= q \left(L + \frac{q-1}{2}L^2 \right), \\ \hat{c}_2(q) &:= 2q \left(L + \frac{2q-1}{2}L^2 \right) + (2q-1)L + 2(q-1)^2L^2, \end{aligned} \tag{3.3}$$

which are related to the growth behavior of the underlying diffusion (see the estimates (3.5)-(3.10) below).

The following theorem characterizes the value function of the entropy regularized OS problem. The proof of the existence of a solution to the HJB equation is given in Subsection 3.2, while the proof of its uniqueness is given in Subsection 3.3.

Theorem 3.2. *The value function V^λ is $C(\mathbb{R}^n \times [0, 1]) \cap W_{loc}^{2,2}(\mathbb{R}^n \times (0, 1))$, is concave in y , is such that $V(x, y) \leq C(1 + |x|^p)$ and it solves in the a.e. sense the HJB equation*

$$\max \{ (\mathcal{L}_x - \rho) V^\lambda(x, y) + \pi(x)y - \lambda y \log y, -V_y^\lambda(x, y) + G(x) \} = 0,$$

with boundary condition $V^\lambda(x, 0) = 0$.

Moreover, the HJB equation admits a unique solution in the class of $C(\mathbb{R}^n \times [0, 1]) \cap W_{loc}^{2,2}(\mathbb{R}^n \times (0, 1))$ functions which are concave in y and such that $V(x, y) \leq C(1 + |x|^p)$.

We next move to the characterization of the optimal control. Define the function $g_\lambda : \mathbb{R}^n \rightarrow [0, 1]$ by

$$(3.4) \quad g_\lambda(x) := \sup \{y \in [0, 1] \mid -V_y^\lambda(x, y) + G(x) < 0\},$$

where we set $g_\lambda(x) := 1$ if $\{y \in [0, 1] \mid -V_y^\lambda(x, y) + G(x) < 0\} = \emptyset$. Notice that the function g_λ is well defined by the concavity of V^λ . Intuitively, the function g_λ separates the “exploration region” $\mathcal{E}_\lambda := \{(x, y) \mid -V_y^\lambda(x, y) + G(x) < 0\}$ in which no action is required, from the “stopping region” $\mathcal{S}_\lambda := \{(x, y) \mid -V_y^\lambda(x, y) + G(x) = 0\}$, in which the agent has to act.

The following theorem characterizes the optimal control in terms of the function g_λ . Its proof can be essentially found in the Step 2 in Subsection 3.3 and it is briefly summarized in Subsection 3.4.

Theorem 3.3. *For any $(x, y) \in \mathbb{R}^n \times [0, 1]$, there exists a unique optimal control $\xi^\lambda \in \mathcal{A}(y)$, which is given by the reflection policy at the boundary g_λ ; that is,*

$$\xi_t^\lambda := \sup_{s \leq t} (y - g_\lambda(X_s^x))^+, \quad t \geq 0.$$

A few comments are in order.

Remark 3.4 (Optimal control and the related Skorokhod problem). *The optimal control ξ^λ is the control with minimal total variation such that*

$$(X_t^x, Y_t^{y, \xi^\lambda}) \in \overline{\mathcal{E}_\lambda}, \quad \text{for any } t \geq 0, \mathbb{P}\text{-a.s.},$$

where $\overline{\mathcal{E}_\lambda}$ is the closure of \mathcal{E}_λ . Such a process is also known as the solution to the Skorokhod reflection problem for the underlying state process in the waiting region \mathcal{E}_λ . Note that characterizing the optimal singular control in terms of the related Skorokhod problem represents a challenging open problem in singular control theory (we refer to the discussion in [15] for further details).

Remark 3.5 (Exploratory behavior of optimal controls). *Unlike the OS problem with exploratory strategies but without modification on the objective function (see Proposition 2.4 and the related Remark 2.5), Theorem 3.3 clarifies that the optimal control of the entropy regularized OS problem is no longer a strict stopping time. In other words, introducing entropy regularization leads to optimal strategies of an exploratory nature—more precisely, strategies that are exclusively exploratory.*

From an RL perspective, the exploratory behavior of the optimal control is crucial, as it allows for simultaneous optimization and information gathering from the environment. Specifically, once the state process reaches the boundary of the exploration region, not only is action required, but, given the regularity of the free boundary g_λ (as demonstrated in the example in Section 4), an infinite number of infinitesimal actions are necessary to prevent the state process from entering the stopping region. This enables the agent to continuously gather information on the optimality of selecting a particular reflecting boundary in a non-episodic manner.

3.1. Preliminary estimates. Before diving into the proof of Theorem 3.2, we discuss some estimates that will be used several times in the sequel.

First, for a generic $q \geq 1$ and $\hat{c}_0(q), \hat{c}_1(q), \hat{c}_2(q)$ as in (3.3), we have the estimates

$$(3.5) \quad \mathbb{E}[|X_t^x|^q] \leq C(1 + |x|^q)e^{\hat{c}_0(q)t}, \quad q \geq 1,$$

$$(3.6) \quad \mathbb{E}[|X_t^{\bar{x}} - X_t^x|^q] \leq |\bar{x} - x|^q e^{\hat{c}_1(q)t}, \quad q \geq 2,$$

$$(3.7) \quad \mathbb{E}[|\delta X_t^{\bar{x}} + (1 - \delta)X_t^x - X_t^{\delta\bar{x} + (1-\delta)x}|^q] \leq C\delta(1 - \delta)|\bar{x} - x|^{2q} e^{\hat{c}_2(q)t}, \quad q \geq 2,$$

for any $x, \bar{x} \in \mathbb{R}^n$, $\delta \in [0, 1]$. Even if these estimates are obtained through standard stochastic calculus techniques, it is difficult to provide a reference in which the exact exponential orders

of growth $\hat{c}_0, \hat{c}_1, \hat{c}_2$ appear (importantly, these constants are those appearing in our Condition 4 in Assumption 3.1). For this reason, we provide a proof in Appendix B.

When ρ is large enough (as in Condition 4 of Assumption 3.1), the estimates above together with the Burkholder-Davis-Gundy inequality imply that

$$(3.8) \quad \mathbb{E} \left[\sup_{t \leq \tau} e^{-\rho t} |X_t^x|^q \right] \leq C(1 + |x|^q), \quad q = 2(p-1), p, 2(p-2),$$

$$(3.9) \quad \mathbb{E} \left[\sup_{t \leq \tau} e^{-\rho t} |X_t^{\bar{x}} - X_t^x|^q \right] \leq C|\bar{x} - x|^q, \quad q = 2, 4$$

$$(3.10) \quad \mathbb{E} \left[\sup_{t \leq \tau} e^{-\rho t} |\delta X_t^{\bar{x}} + (1-\delta)X_t^x - X_t^{\delta\bar{x}+(1-\delta)x}|^2 \right] \leq C\delta(1-\delta)|\bar{x} - x|^4,$$

for any $x, \bar{x} \in \mathbb{R}^n$, $\delta \in [0, 1]$ and $\tau \in \mathcal{T}$. Again, we provide a proof in Appendix B.

Finally, using the growth of π and G in Condition 2 in Assumption 3.1 as well as the estimates (3.5) and (3.8), we find

$$(3.11) \quad \begin{aligned} V^\lambda(x, y) &\leq \sup_{\xi \in \mathcal{A}(y)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} |\Pi(X_t^x, Y_{t-}^{y, \xi})| dt + \int_0^\infty e^{-\rho t} |G(X_t^x)| d\xi_t \right] \\ &\leq \sup_{\xi \in \mathcal{A}(y)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} (1 + |X_t^x|^p) dt + \lim_{T \rightarrow \infty} \sup_{t \leq T} e^{-\rho t} (1 + |X_t^x|^p) \xi_T \right] \\ &\leq C(1 + |x|^p), \end{aligned}$$

where we have used the fact that $\xi_t \leq y \leq 1$ for any $\xi \in \mathcal{A}(y)$.

3.2. Proof of Theorem 3.2: solution of the HJB. The concavity of V^λ follows by straightforward arguments (see Remark 2.6 for a reference), while its growth rate was shown in (3.11). Thus, we will show in the sequel that V^λ has the claimed regularity and it solves the HJB equation in the almost everywhere sense.

The proof is divided into five steps. In the proof, we will often use the notation $\Pi(x, y) := y\pi(x) - \lambda y \log y$.

Step 1 (Penalization). Let $\beta : \mathbb{R} \rightarrow [0, \infty)$ be a C^∞ nondecreasing convex function such that

$$\beta(z) = \begin{cases} 0, & \text{if } z \leq 0, \\ 2z - 1, & \text{if } z \geq 1. \end{cases}$$

For any $\varepsilon > 0$, consider the control set $U_\varepsilon := \{(a^1, a^2) \in \mathbb{R}^2 \mid sa^1 - a^2 \leq \frac{\beta(s)}{\varepsilon} \ \forall s \in \mathbb{R} \text{ and } 0 \leq a^2 \leq \frac{1}{\varepsilon}\}$ and the functional

$$J^{\lambda, \varepsilon}(x, y; (u, \eta)) := \mathbb{E} \left[\int_0^\infty e^{-\rho t} (\Pi(X_t^x, Y_t^{y, u}) - G(X_t^x)u_t - \eta_t) dt \right], \quad (u, \eta) \in \mathcal{U}_\varepsilon(y)$$

$$\begin{aligned} \text{subject to } dX_t^x &= b(X_t^x)dt + \sigma(X_t^x)dW_t, \quad X_0^x = x, \\ dY_t^{y, u} &= -u_t dt, \quad Y_0^{y, u} = y, \end{aligned}$$

where $\mathcal{U}_\varepsilon(y) := \{(u, \eta) : \Omega \times [0, \infty) \rightarrow U_\varepsilon \mid Y_t^{y, u} \geq 0 \ \forall t \geq 0, \mathbb{P}\text{-a.s.}\}$. For any $\varepsilon > 0$, define the optimization problem

$$(3.12) \quad V^{\lambda, \varepsilon}(x, y) := \sup_{(u, \eta) \in \mathcal{U}_\varepsilon(y)} J^{\lambda, \varepsilon}(x, y; (u, \eta)).$$

Similarly to (3.11), we can find C such that, for y small enough, one has

$$\begin{aligned}
 (3.13) \quad & |V^{\lambda, \varepsilon}(x, y)| \\
 & \leq \sup_{(u, \eta) \in \mathcal{U}_\varepsilon(y)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} |\Pi(X_t^x, Y_t^{y, u}) - G(X_t^x)u_t - \eta_t| dt \right] \\
 & \leq \sup_{(u, \eta) \in \mathcal{U}_\varepsilon(y)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} y(|\pi(X_t^x)| - \lambda \log y) dt + \lim_{T \rightarrow \infty} \sup_{t \leq T} (e^{-\rho t} |G(X_t^x)|) \int_0^T u_t dt \right] \\
 & \leq Cy \left(\mathbb{E} \left[\int_0^\infty e^{-\rho t} (1 + |X_t^x|^p - \lambda \log y) dt + \lim_{T \rightarrow \infty} \sup_{t \leq T} e^{-\rho t} (1 + |X_t^x|^p) \right] \right) \\
 & \leq Cy(1 + |x|^p - \lambda \log y) \rightarrow 0, \quad \text{as } y \rightarrow 0,
 \end{aligned}$$

which implies the continuity of $V^{\lambda, \varepsilon}$ for $y = 0$.

Standard arguments (see e.g. Chapter IV.10 in [22]) give that $V^{\lambda, \varepsilon}$ is a generalized solution to the HJB equation:

$$(3.14) \quad \rho V^{\lambda, \varepsilon}(x, y) - \mathcal{L}_x V^{\lambda, \varepsilon}(x, y) - \Pi(x, y) - \frac{1}{\varepsilon} \beta \left(G(x) - V_y^{\lambda, \varepsilon}(x, y) \right) = 0,$$

with boundary condition $V^{\lambda, \varepsilon}(x, 0) = 0$.

Moreover, following the rationale of the proof of Theorem 3.2 in [62], one can show that $V(x, y) = \lim_{\varepsilon \rightarrow 0} V^{\lambda, \varepsilon}(x, y)$.

Step 2 (Uniform Lipschitz estimates for $V^{\lambda, \varepsilon}$ in x). Take $x, \bar{x} \in \mathbb{R}^n$, $y \in (0, 1]$, $\zeta > 0$ and $(u, \eta) \in \mathcal{U}_\varepsilon(y)$ such that $V^{\lambda, \varepsilon}(\bar{x}, y) - \zeta \leq J^{\lambda, \varepsilon}(\bar{x}, y; u, \eta)$. Since (u, η) is suboptimal for (x, y) , we have

$$\begin{aligned}
 (3.15) \quad & V^{\lambda, \varepsilon}(\bar{x}, y) - V^{\lambda, \varepsilon}(x, y) - \zeta \leq J^{\lambda, \varepsilon}(\bar{x}, y; u, \eta) - J^{\lambda, \varepsilon}(x, y; u, \eta) \\
 & = \mathbb{E} \left[\int_0^\infty e^{-\rho t} (Y_t^{y, u}(\pi(X_t^{\bar{x}}) - \pi(X_t^x)) + u_t(G(X_t^{\bar{x}}) - G(X_t^x))) dt \right].
 \end{aligned}$$

In order to continue this estimate, we notice that by monotone convergence theorem, we have

$$(3.16) \quad \mathbb{E} \left[\int_0^\infty e^{-\rho t} u_t (G(X_t^{\bar{x}}) - G(X_t^x)) dt \right] \leq \lim_{T \rightarrow \infty} \mathbb{E} \left[\int_0^T e^{-\rho t} u_t |G(X_t^{\bar{x}}) - G(X_t^x)| dt \right].$$

Furthermore, since $(u, \eta) \in \mathcal{U}_\varepsilon(y)$, we have $\int_0^t u_s ds \leq y \leq 1$. Thus, using Assumption 3.1 and Hölder inequality, we can estimate the right-hand side of (3.16) to obtain

$$\begin{aligned}
 & \mathbb{E} \left[\int_0^T e^{-\rho t} u_t |G(X_t^{\bar{x}}) - G(X_t^x)| dt \right] \\
 & \leq \mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} |G(X_t^{\bar{x}}) - G(X_t^x)|) \int_0^T u_t dt \right] \\
 & \leq C \mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} (1 + |X_t^{\bar{x}}|^{p-1} + |X_t^x|^{p-1}) |X_t^{\bar{x}} - X_t^x|) \right] \\
 & \leq C \left(\mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} (1 + |X_t^{\bar{x}}|^{2(p-1)} + |X_t^x|^{2(p-1)})) \right] \right)^{\frac{1}{2}} \left(\mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} |X_t^{\bar{x}} - X_t^x|^2) \right] \right)^{\frac{1}{2}}.
 \end{aligned}$$

Finally, using the estimates in (3.8) and (3.9), we get

$$\mathbb{E} \left[\int_0^T e^{-\rho t} u_t |G(X_t^{\bar{x}}) - G(X_t^x)| dt \right] \leq C(1 + |x|^{p-1} + |\bar{x}|^{p-1}) |\bar{x} - x|,$$

which allows to conclude thanks to (3.16):

$$(3.17) \quad \mathbb{E} \left[\int_0^\infty e^{-\rho t} u_t (G(X_t^{\bar{x}}) - G(X_t^x)) dt \right] \leq C(1 + |x|^{p-1} + |\bar{x}|^{p-1}) |\bar{x} - x|.$$

We next estimate the term in (3.15) which involves π . Since $Y_t^{y,u} \leq y \leq 1$, using Assumption 3.1 and Hölder inequality we find

$$\begin{aligned} & \mathbb{E} \left[\int_0^\infty e^{-\rho t} Y_t^{y,u} (\pi(X_t^{\bar{x}}) - \pi(X_t^x)) dt \right] \\ & \leq C \mathbb{E} \left[\int_0^T e^{-\rho t} (1 + |X_t^{\bar{x}}|^{p-1} + |X_t^x|^{p-1}) |X_t^{\bar{x}} - X_t^x| dt \right] \\ & \leq C \int_0^T e^{-\rho t} \left(1 + (\mathbb{E}[|X_t^{\bar{x}}|^{2(p-1)}])^{\frac{1}{2}} + (\mathbb{E}[|X_t^x|^{2(p-1)}])^{\frac{1}{2}} \right) (\mathbb{E}[|X_t^{\bar{x}} - X_t^x|^2])^{\frac{1}{2}} dt. \end{aligned}$$

Finally, using the estimates in (3.5) and (3.6), we get

$$\begin{aligned} & \mathbb{E} \left[\int_0^\infty e^{-\rho t} Y_t^{y,u} (\pi(X_t^{\bar{x}}) - \pi(X_t^x)) dt \right] \\ & \leq C(1 + |x|^{p-1} + |\bar{x}|^{p-1}) |\bar{x} - x| \int_0^T e^{(-\rho + \frac{\hat{c}_0(2(p-1))}{2} + \frac{\hat{c}_1(2)}{2})t} dt. \end{aligned}$$

Hence, thanks to Condition 4 in Assumption 3.1, we conclude that

$$(3.18) \quad \mathbb{E} \left[\int_0^\infty e^{-\rho t} Y_t^{y,u} (\pi(X_t^{\bar{x}}) - \pi(X_t^x)) dt \right] \leq C(1 + |x|^{p-1} + |\bar{x}|^{p-1}) |\bar{x} - x|.$$

Plugging (3.17) and (3.18) into (3.15), we obtain

$$V^{\lambda,\varepsilon}(\bar{x}, y) - V^{\lambda,\varepsilon}(x, y) - \zeta \leq C(1 + |x|^{p-1} + |\bar{x}|^{p-1}) |\bar{x} - x|,$$

which, by the arbitrariness of ζ , gives the local Lipschitz property of V^ε in x .

Step 3 (Uniform Lipschitz estimates for $V^{\lambda,\varepsilon}$ in y). Take $x \in \mathbb{R}^n$, $y, \bar{y} \in (0, 1]$, $\zeta > 0$ and $(\bar{u}, \bar{\eta}) \in \mathcal{U}_\varepsilon(\bar{y})$ such that $V^{\lambda,\varepsilon}(x, \bar{y}) - \zeta \leq J^{\lambda,\varepsilon}(x, \bar{y}; \bar{u}, \bar{\eta})$. We will distinguish the two cases $\bar{y} \leq y$ and $\bar{y} > y$.

If $\bar{y} \leq y$, then we have $Y_t^{y,\bar{u}} \geq 0$, so that $(\bar{u}, \bar{\eta}) \in \mathcal{U}_\varepsilon(y)$. By suboptimality of $(\bar{u}, \bar{\eta})$ for the initial condition (x, y) , we have

$$\begin{aligned} & V^{\lambda,\varepsilon}(x, \bar{y}) - V^{\lambda,\varepsilon}(x, y) - \zeta \leq J^{\lambda,\varepsilon}(x, \bar{y}; \bar{u}, \bar{\eta}) - J^{\lambda,\varepsilon}(x, y; \bar{u}, \bar{\eta}) \\ (3.19) \quad & = \mathbb{E} \left[\int_0^\infty e^{-\rho t} (\pi(X_t^x)(\bar{y} - y) - \lambda(Y_t^{\bar{y},\bar{u}} \log Y_t^{\bar{y},\bar{u}} - Y_t^{y,\bar{u}} \log Y_t^{y,\bar{u}})) dt \right] \\ & = C(1 + |x|^p) |\bar{y} - y| + \mathbb{E} \left[\int_0^\infty e^{-\rho t} (-\lambda(Y_t^{\bar{y},\bar{u}} \log Y_t^{\bar{y},\bar{u}} - Y_t^{y,\bar{u}} \log Y_t^{y,\bar{u}})) dt \right], \end{aligned}$$

where we have used Condition 4 in Assumption 3.1 together with (3.5). In order to continue the latter estimate, let \hat{z} be the maximum point of the function $z \mapsto -\lambda z \log z$ (i.e., $\hat{z} = e^{-1}$), and notice that $z \mapsto -\lambda z \log z$ is increasing in $[0, \hat{z}]$ and Lipschitz in $[\frac{\hat{z}}{2}, 1]$, with Lipschitz constant smaller than $C_\lambda := 1 + |\log(\hat{z}/2)|$. Suppose now that $y - \bar{y} \leq \frac{\hat{z}}{2}$. If $Y_t^{\bar{y},\bar{u}} \leq \frac{\hat{z}}{2}$, then $Y_t^{y,\bar{u}} \leq \hat{z}$, and since $Y_t^{\bar{y},\bar{u}} \leq Y_t^{y,\bar{u}}$, we obtain $-\lambda(Y_t^{\bar{y},\bar{u}} \log Y_t^{\bar{y},\bar{u}} - Y_t^{y,\bar{u}} \log Y_t^{y,\bar{u}}) \leq 0$ by monotonicity. On the other hand, if $Y_t^{\bar{y},\bar{u}} \geq \frac{\hat{z}}{2}$, then $Y_t^{y,\bar{u}} \geq \frac{\hat{z}}{2}$, and by Lipschitzianity we find $-\lambda(Y_t^{\bar{y},\bar{u}} \log Y_t^{\bar{y},\bar{u}} - Y_t^{y,\bar{u}} \log Y_t^{y,\bar{u}}) \leq C_\lambda(Y_t^{y,\bar{u}} - Y_t^{\bar{y},\bar{u}}) = C_\lambda(y - \bar{y})$. Going back to (3.19), this argument implies that

$$0 \leq y - \bar{y} \leq \frac{\hat{z}}{2} \implies V^{\lambda,\varepsilon}(x, \bar{y}) - V^{\lambda,\varepsilon}(x, y) - \zeta \leq C(1 + |x|^p + |\log(\hat{z}/2)|) |\bar{y} - y|,$$

which give the local Lipschitz property of V^ε in y in the case $\bar{y} \leq y$.

On the other hand, if $\bar{y} \geq y$, then the control $(\bar{u}, \bar{\eta})$ is not necessarily admissible for y , as $Y_t^{y, \bar{u}}$ could become smaller than 0. Define $u_t := \bar{u}_t \frac{y}{\bar{y}}$ and $\eta_t := \bar{\eta}_t \frac{y}{\bar{y}}$. Notice that, if $s\bar{u}_t - \bar{\eta}_t \leq 0$, then $su_t - \eta_t = \frac{y}{\bar{y}}(s\bar{u}_t - \bar{\eta}_t) \leq 0 \leq \frac{1}{\varepsilon}\beta(s)$, while $su_t - \eta_t = \frac{y}{\bar{y}}(s\bar{u}_t - \bar{\eta}_t) \leq s\bar{u}_t - \bar{\eta}_t \leq \frac{1}{\varepsilon}\beta(s)$ if $s\bar{u}_t - \bar{\eta}_t > 0$. Thus, the process (u, η) takes values in the set U_ε . Moreover, since

$$(3.20) \quad Y_t^{y, u} = y - \int_0^t u_s ds = \frac{y}{\bar{y}} \left(\bar{y} - \int_0^t \bar{u}_s ds \right) = \frac{y}{\bar{y}} Y_t^{\bar{y}, \bar{u}} \geq 0,$$

it follows that $(u, \eta) \in \mathcal{U}_\varepsilon(y)$. Hence, by the suboptimality of (u, η) for the initial condition (x, y) , using (3.20) we find

$$(3.21) \quad \begin{aligned} & V^{\lambda, \varepsilon}(x, \bar{y}) - V^{\lambda, \varepsilon}(x, y) - \zeta \\ & \leq J^{\lambda, \varepsilon}(x, \bar{y}; \bar{u}, \bar{\eta}) - J^{\lambda, \varepsilon}(x, y; u, \eta) \\ & = \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(-\lambda \int_0^1 \left(1 + \log(r Y_t^{\bar{y}, \bar{u}} + (1-r) Y_t^{y, u}) \right) (Y_t^{\bar{y}, \bar{u}} - Y_t^{y, u}) dr \right. \right. \\ & \quad \left. \left. + \pi(X_t^x) (\bar{y} - y) \frac{Y_t^{\bar{y}, \bar{u}}}{\bar{y}} + G(X_t^x) \frac{\bar{y} - y}{\bar{y}} \bar{u}_t - \frac{\bar{y} - y}{\bar{y}} \bar{\eta}_t \right) dt \right] \\ & \leq \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(-\lambda \int_0^1 \left(1 + \log(r \bar{y} + (1-r)y) + \log\left(\frac{Y_t^{y, \bar{u}}}{\bar{y}}\right) \right) \frac{Y_t^{y, \bar{u}}}{\bar{y}} (\bar{y} - y) dr \right. \right. \\ & \quad \left. \left. + |\pi(X_t^x)| (\bar{y} - y) + G(X_t^x) \frac{\bar{y} - y}{\bar{y}} \bar{u}_t \right) dt \right] \\ & \leq (\bar{y} - y) \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(\lambda \left(1 + |\log y| - \log\left(\frac{Y_t^{y, \bar{u}}}{\bar{y}}\right) \frac{Y_t^{y, \bar{u}}}{\bar{y}} \right) + |\pi(X_t^x)| + \frac{G(X_t^x)}{\bar{y}} \bar{u}_t \right) dt \right] \\ & \leq C(\bar{y} - y) \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(\lambda(1 + |\log y|) + |\pi(X_t^x)| + \frac{G(X_t^x)}{\bar{y}} \bar{u}_t \right) dt \right]. \end{aligned}$$

Similarly to the arguments that lead to (3.17), we have

$$\begin{aligned} \mathbb{E} \left[\int_0^\infty e^{-\rho t} \bar{u}_t G(X_t^x) dt \right] & \leq \lim_{T \rightarrow \infty} \mathbb{E} \left[\int_0^T e^{-\rho t} \bar{u}_t |G(X_t^x)| dt \right] \\ & \leq \lim_{T \rightarrow \infty} \mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} |G(X_t^x)|) \int_0^T \bar{u}_t dt \right] \\ & \leq \lim_{T \rightarrow \infty} C \mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} (1 + |X_t^x|^p)) \right] \\ & \leq C(1 + |x|^p), \end{aligned}$$

where we have used the estimate in (3.8). Plugging the last inequality into (3.21), we obtain

$$\begin{aligned} V^{\lambda, \varepsilon}(x, \bar{y}) - V^{\lambda, \varepsilon}(x, y) - \zeta & \leq C|\bar{y} - y| \left(\frac{1 + |x|^p}{\bar{y}} + |\log y| + \mathbb{E} \left[\int_0^\infty e^{-\rho t} (1 + |X_t^x|^p) dt \right] \right) \\ & \leq C|\bar{y} - y| \left(\frac{1 + |x|^p}{\bar{y}} + |\log y| \right) \int_0^\infty e^{(-\rho + \hat{c}_0(p))t} dt \\ & \leq C|\bar{y} - y| \left(\frac{1 + |x|^p}{\bar{y}} + |\log y| \right), \end{aligned}$$

where we have used the condition $\rho > \hat{c}_0(p)$ given in Assumption 3.1. This shows the local Lipschitz property of V^ε in y for $\bar{y} \geq y$.

Step 4 (Uniform estimates for $D_x^2 V^{\lambda, \varepsilon}$). We first show the following semiconvexity estimate: there exists a constant C such that

$$(3.22) \quad V^{\lambda, \varepsilon}(x^\delta, y) - \delta V^{\lambda, \varepsilon}(\bar{x}, y) - (1 - \delta)V^{\lambda, \varepsilon}(x, y) \leq C\delta(1 - \delta)(1 + |x|^{p-1} + |\bar{x}|^{p-1})|\bar{x} - x|^2,$$

for each $\delta \in [0, 1]$, $\bar{x}, x \in \mathbb{R}^n$, $y \in \mathbb{R}$ and $\varepsilon > 0$, and where $x^\delta := \delta\bar{x} + (1 - \delta)x$.

Take an arbitrary $\zeta > 0$ and $(u, \eta) \in \mathcal{U}_\varepsilon(y)$ such that $V^{\lambda, \varepsilon}(x^\delta, y) - \zeta \leq J^{\lambda, \varepsilon}(x^\delta, y; u, \eta)$. Notice that such a control is admissible also for the initial conditions (\bar{x}, y) and (x, y) . Since (u, η) is not necessarily optimal for (x, y) or (\bar{x}, y) , we have

$$(3.23) \quad \begin{aligned} & V^{\lambda, \varepsilon}(x^\delta, y) - \delta V^{\lambda, \varepsilon}(\bar{x}, y) - (1 - \delta)V^{\lambda, \varepsilon}(x, y) - \zeta \\ & \leq J^{\lambda, \varepsilon}(x^\delta, y; u, \eta) - \delta J^{\lambda, \varepsilon}(\bar{x}, y; u, \eta) - (1 - \delta)J^{\lambda, \varepsilon}(x, y; u, \eta) \\ & \leq \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(Y_t^{y, u} (\pi(X_t^\delta) - \delta\pi(X_t^{\bar{x}}) - (1 - \delta)\pi(X_t^x)) \right. \right. \\ & \quad \left. \left. + u_t (G(X_t^\delta) - \delta G(X_t^{\bar{x}}) - (1 - \delta)G(X_t^x)) \right) dt \right]. \end{aligned}$$

Next, set $\hat{X}_t := \delta X_t^{\bar{x}} + (1 - \delta)X_t^x$, $X_t^\delta := X_t^{x^\delta}$ and define, for a generic function $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$, the transformations

$$\begin{aligned} \Delta[\varphi]_t &:= |\varphi(X_t^\delta) - \varphi(\hat{X}_t)|, \\ \Gamma[\varphi]_t &:= |\varphi(\hat{X}_t) - \delta\varphi(X_t^{\bar{x}}) - (1 - \delta)\varphi(X_t^x)|. \end{aligned}$$

Then applying Fatou's lemma, we rewrite (3.23) as

$$(3.24) \quad \begin{aligned} & V^{\lambda, \varepsilon}(x^\delta, y) - \delta V^{\lambda, \varepsilon}(\bar{x}, y) - (1 - \delta)V^{\lambda, \varepsilon}(x, y) - \zeta \\ & \leq \lim_{T \rightarrow \infty} \mathbb{E} \left[\int_0^T e^{-\rho t} \left(Y_t^{y, u} \Delta[\pi]_t + u_t \Delta[G]_t + Y_t^{y, u} \Gamma[\pi]_t + u_t \Gamma[G]_t \right) dt \right]. \end{aligned}$$

We next estimate each term on the right-hand side of (3.24) separately.

First, using Condition 2 in Assumption 3.1, Hölder inequality, and the SDE estimates (3.5) and (3.7), we find

$$(3.25) \quad \begin{aligned} & \mathbb{E} \left[\int_0^T e^{-\rho t} Y_t^{y, u} \Delta[\pi]_t dt \right] \\ & \leq C \mathbb{E} \left[\int_0^T e^{-\rho t} (1 + |\hat{X}_t|^{p-1} + |X_t^\delta|^{p-1}) |\hat{X}_t - X_t^\delta| dt \right] \\ & \leq C \left(\int_0^T e^{-\rho t} \left(1 + (\mathbb{E}[|\hat{X}_t|^{2(p-1)}])^{\frac{1}{2}} + (\mathbb{E}[|X_t^\delta|^{2(p-1)}])^{\frac{1}{2}} \right) \left(\mathbb{E}[|\hat{X}_t - X_t^\delta|^2] \right)^{\frac{1}{2}} dt \right) \\ & \leq C\delta(1 - \delta)(1 + |x|^{p-1} + |\bar{x}|^{p-1})|\bar{x} - x|^2 \int_0^T e^{(-\rho + \frac{\hat{c}_0(2(p-1))}{2} + \frac{\hat{c}_1(2)}{2})t} dt \\ & \leq C\delta(1 - \delta)(1 + |x|^{p-1} + |\bar{x}|^{p-1})|\bar{x} - x|^2, \end{aligned}$$

where the last inequality follows by Condition 4 in Assumption 3.1.

Secondly, using Condition 2 in Assumption 3.1, Hölder inequality, and the SDE estimates (3.5) and (3.7), we find

$$\begin{aligned}
(3.26) \quad & \mathbb{E} \left[\int_0^T e^{-\rho t} u_t \Delta[G]_t dt \right] \\
& \leq \mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} \Delta[G]_t) \int_0^T u_t dt \right] \\
& \leq C \mathbb{E} \left[\sup_{t \leq T} \left(e^{-\rho t} (1 + |\hat{X}_t|^{p-1} + |X_t^\delta|^{p-1}) |\hat{X}_t - X_t^\delta| \right) \right] \\
& \leq C \left(\mathbb{E} \left[\sup_{t \leq T} \left(e^{-\rho t} (1 + |\hat{X}_t|^{2(p-1)} + |X_t^\delta|^{2(p-1)}) \right) \right] \right)^{\frac{1}{2}} \left(\mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} |\hat{X}_t - X_t^\delta|^2) \right] \right)^{\frac{1}{2}} \\
& \leq C \delta (1 - \delta) (1 + |x|^{p-1} + |\bar{x}|^{p-1}) |\bar{x} - x|.
\end{aligned}$$

where the last inequality follows from the estimates (3.8) and (3.10).

Thirdly, using Condition 2 in Assumption 3.1, Hölder inequality, and the SDE estimates (3.5) and (3.6), we find

$$\begin{aligned}
(3.27) \quad & \mathbb{E} \left[\int_0^T e^{-\rho t} Y_t^{y,u} \Gamma[\pi]_t dt \right] \\
& \leq C \delta (1 - \delta) \mathbb{E} \left[\int_0^T e^{-\rho t} \left(1 + |X_t^{\bar{x}}|^{p-2} + |X_t^x|^{p-2} \right) |X_t^{\bar{x}} - X_t^x|^2 dt \right] \\
& \leq C \delta (1 - \delta) \int_0^T e^{-\rho t} \left(1 + (\mathbb{E}[|X_t^{\bar{x}}|^{2(p-2)}])^{\frac{1}{2}} + (\mathbb{E}[|X_t^x|^{2(p-2)}])^{\frac{1}{2}} \right) \\
& \quad \times \left(\mathbb{E}[|X_t^{\bar{x}} - X_t^x|^4] \right)^{\frac{1}{2}} dt \\
& \leq C \delta (1 - \delta) (1 + |x|^{p-2} + |\bar{x}|^{p-2}) |\bar{x} - x|^2 \int_0^T e^{(-\rho + \frac{\hat{c}_0(2(p-2))}{2} + \frac{\hat{c}_1(4)}{2})t} dt \\
& \leq C \delta (1 - \delta) (1 + |x|^{p-2} + |\bar{x}|^{p-2}) |\bar{x} - x|^2,
\end{aligned}$$

where the finiteness of the integrals follows again from Condition 4 in Assumption 3.1.

Fourthly, using Condition 2 in Assumption 3.1, Hölder inequality, and the SDE estimates (3.5) and (3.6), we find

$$\begin{aligned}
(3.28) \quad & \mathbb{E} \left[\int_0^T e^{-\rho t} u_t \Gamma[G]_t dt \right] \\
& \leq \mathbb{E} \left[\sup_{t \leq T} (e^{-\rho t} \Gamma[G]_t) \int_0^T u_t dt \right] \\
& \leq C \delta (1 - \delta) \mathbb{E} \left[\sup_{t \leq T} \left(e^{-\rho t} (1 + |X_t^{\bar{x}}|^{p-2} + |X_t^x|^{p-2}) |X_t^{\bar{x}} - X_t^x|^2 \right) \right] \\
& \leq C \delta (1 - \delta) \left(\mathbb{E} \left[\sup_{t \leq T} \left(e^{-\rho t} (1 + |X_t^{\bar{x}}|^{2(p-2)} + |X_t^x|^{2(p-2)}) \right) \right] \right)^{\frac{1}{2}} \\
& \quad \times \left(\mathbb{E} \left[\sup_{t \leq T} \left(e^{-\rho t} |X_t^{\bar{x}} - X_t^x|^4 \right) \right] \right)^{\frac{1}{2}} \\
& \leq C \delta (1 - \delta) (1 + |x|^{p-2} + |\bar{x}|^{p-2}) |\bar{x} - x|^2 \int_0^T e^{(-\rho + \frac{\hat{c}_0(2(p-2))}{2} + \frac{\hat{c}_1(4)}{2})t} dt \\
& \leq C \delta (1 - \delta) (1 + |x|^{p-2} + |\bar{x}|^{p-2}) |\bar{x} - x|^2,
\end{aligned}$$

where we have used Condition 4 in Assumption 3.1.

Plugging (3.25), (3.26), (3.27) and (3.28) back into (3.24), we obtain (3.22).

We will now use (3.22) in the PDE (3.14) in order to obtain an estimate for the derivative $D_x^2 V^{\lambda, \varepsilon}$. Notice that (3.22) implies that for any $r > 0$ and $B_r(0) := \{x \in \mathbb{R}^n \mid |x| < r\}$ on can find a constant $C_r > 0$ such that

$$(3.29) \quad \partial_{zz}^2 V^{\lambda, \varepsilon}(x) \geq -C_r, \quad \text{for any } x \in B_r(0), z \in \mathbb{R}^n \text{ with } |z| = 1.$$

Choose an orthonormal basis z^1, \dots, z^n of \mathbb{R}^n such that

$$\text{tr}(\sigma \sigma^*(x) D_x^2 V^{\lambda, \varepsilon}(x, y)) = \sum_{i=1}^n c_\sigma^i(x) \partial_{z^i z^i}^2 V^{\lambda, \varepsilon}(x, y),$$

where $c_\sigma^1, \dots, c_\sigma^n$ are the eigenvalues of the symmetric, nonnegative definite matrix $\sigma \sigma^*(x)$. Since V^ε solves the PDE in (3.14), using the latter identity together with (3.29) and the estimates in the previous steps, unless to take a larger constant C_r (independently from ε) we obtain

$$\begin{aligned} -C_r &\leq \sum_{i=1}^n c_\sigma^i(x) \partial_{z^i z^i}^2 V^{\lambda, \varepsilon}(x, y) \\ &= \rho V^{\lambda, \varepsilon}(x, y) - b(x) D_x V^{\lambda, \varepsilon}(x, y) - \Pi(x, y) - \frac{1}{\varepsilon} \beta \left(G(x) - V_y^{\lambda, \varepsilon}(x, y) \right) \\ &\leq \rho V^{\lambda, \varepsilon}(x, y) - b(x) D_x V^{\lambda, \varepsilon}(x, y) - \Pi(x, y) \leq C_r, \quad \text{for any } (x, y) \in B_r(0) \times (0, 1). \end{aligned}$$

This in turn implies that

$$0 \leq \frac{1}{\varepsilon} \beta \left(G(x) - V_y^{\lambda, \varepsilon}(x, y) \right) \leq C_r, \quad \text{for any } (x, y) \in B_r(0) \times (0, 1),$$

so that, by using the uniform ellipticity of σ (cf. Condition 3 in Assumption 3.1), by Theorem 9.11 p. 235 in [25] we conclude that, for any $q > 0$, unless to enlarge the constant C_r (independently from ε), for any $y \in (0, 1)$ we have

$$(3.30) \quad \int_{B_{r/2}(0)} |D_x^2 V^{\lambda, \varepsilon}(x, y)|^q dx \leq C_r \left(1 + \int_{B_{r/2}(0)} \left| \frac{1}{\varepsilon} \beta \left(G(x) - V_y^{\lambda, \varepsilon}(x, y) \right) \right|^q dx \right) \leq C_r,$$

thus giving the desired estimate.

Step 5 (Lipschitzianity of V_y^λ via optimal stopping). Following the arguments in [35] (see in particular Theorem 3.4) we can show that, for any $(x, y) \in \mathbb{R}^n \times (0, 1]$, one has the representation

$$(3.31) \quad V_y^\lambda(x, y) = \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\int_0^\tau e^{-\rho t} \Pi_y(X_t^x, y) dt + e^{-\rho \tau} G(X_\tau^x) \right],$$

that we will now employ to prove the local Lipschitzianity of V_y^λ .

We first show the Lipschitzianity in y . Take $x \in \mathbb{R}^n$, $\bar{y}, y \in (0, 1]$ with $\bar{y} \leq y$. By the monotonicity of Π_y in y we have $0 \leq V_y^\lambda(x, \bar{y}) - V_y^\lambda(x, y)$. Moreover, for $\zeta > 0$ and $\tau \in \mathcal{T}$ such that $V_y^\lambda(x, \bar{y}) - \zeta \leq \mathbb{E} \left[\int_0^\tau e^{-\rho t} \Pi_y(X_t^x, \bar{y}) dt + e^{-\rho \tau} G(X_\tau^x) \right]$, we find

$$\begin{aligned} V_y^\lambda(x, \bar{y}) - V_y^\lambda(x, y) - \zeta &\leq \mathbb{E} \left[\int_0^\tau e^{-\rho t} (\Pi_y(X_t^x, \bar{y}) - \Pi_y(X_t^x, y)) dt \right] \\ &= \lambda \mathbb{E} \left[\int_0^\tau e^{-\rho t} (\log y - \log \bar{y}) dt \right] \\ &\leq \lambda \mathbb{E} \left[\int_0^\tau e^{-\rho t} dt \right] \frac{1}{\bar{y}} (y - \bar{y}) \leq \lambda C \frac{1}{\bar{y}} (y - \bar{y}). \end{aligned}$$

From the arbitrariness of ζ , we conclude that $0 \leq V_y^\lambda(x, \bar{y}) - V_y^\lambda(x, y) \leq C \frac{1}{\bar{y}}(y - \bar{y})$, showing the Lipschitzianity in y .

We next show the Lipschitzianity in x . For $\bar{x}, x \in \mathbb{R}^n$ and $y \in (0, 1]$, take $\zeta > 0$ and $\tau \in \mathcal{T}$ such that $V_y^\lambda(\bar{x}, y) - \zeta \leq \mathbb{E} \left[\int_0^\tau e^{-\rho t} \Pi_y(X_t^{\bar{x}}, y) dt + e^{-\rho \tau} G(X_\tau^{\bar{x}}) \right]$. We observe that

$$\begin{aligned}
 (3.32) \quad & V_y^\lambda(\bar{x}, y) - V_y^\lambda(x, y) - \zeta \\
 & \leq \mathbb{E} \left[\int_0^\tau e^{-\rho t} (\Pi_y(X_t^{\bar{x}}, y) - \Pi_y(X_t^x, y)) dt + e^{-\rho \tau} (G(X_\tau^{\bar{x}}) - G(X_\tau^x)) \right] \\
 & = \mathbb{E} \left[\int_0^\tau e^{-\rho t} (\pi(X_t^{\bar{x}}) - \pi(X_t^x)) dt + e^{-\rho \tau} (G(X_\tau^{\bar{x}}) - G(X_\tau^x)) \right].
 \end{aligned}$$

Thus, by repeating the arguments that lead to (3.17) and (3.18), we obtain

$$V_y^\lambda(\bar{x}, y) - V_y^\lambda(x, y) - \zeta \leq C(1 + |x|^{p-1} + |\bar{x}|^{p-1})|\bar{x} - x|.$$

Hence, the local Lipschitzianity in x follows by the arbitrariness of ζ and exchanging the role of x and \bar{x} .

Step 6 (Regularity of V^λ and variational inequality). We begin by noticing that the continuity of V^λ in $y = 0$ is shown as in (3.13).

By the estimates in Steps 2 and 3, for any $r, \delta > 0$ and any ball $B_r(0) := \{x \in \mathbb{R}^n \mid |x| < r\}$, we can find a subsequence $(\varepsilon_n)_n$ with $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$ such that

$$\begin{aligned}
 (3.33) \quad & V^{\lambda, \varepsilon_n} \rightarrow V^\lambda \text{ as } n \rightarrow \infty, \text{ uniformly in } B_r(0) \times (\delta, 1), \\
 & (D_x V^{\lambda, \varepsilon_n}, V_y^{\lambda, \varepsilon_n}) \rightarrow (D_x V^\lambda, V_y^\lambda) \text{ as } n \rightarrow \infty, \text{ weakly in } \mathbb{L}^2(B_r(0) \times (\delta, 1)).
 \end{aligned}$$

Moreover, thanks to the estimate in Step 4, for any fixed $y \in (0, 1]$ we can select a further subsequence $(\varepsilon_n^y)_n$ such that

$$\begin{aligned}
 (3.34) \quad & D_x V^{\lambda, \varepsilon_n^y}(\cdot, y) \rightarrow D_x V^\lambda(\cdot, y) \text{ as } n \rightarrow \infty, \text{ uniformly in } B_r(0), \\
 & D_x^2 V^{\lambda, \varepsilon_n^y}(\cdot, y) \rightarrow D_x^2 V^\lambda(\cdot, y) \text{ as } n \rightarrow \infty, \text{ weakly in } \mathbb{L}^2(B_r(0)).
 \end{aligned}$$

Therefore, taking limits in (3.14) we obtain that

$$\text{for any } y \in (0, 1) \quad (\rho - \mathcal{L}_x) V^\lambda(\cdot, y) - \Pi(\cdot, y) \geq 0 \quad \text{and} \quad V_y^\lambda(\cdot, y) - G \geq 0, \quad dx\text{-a.e.}$$

which shows the HJB equation with inequality; that is,

$$\max \{ (\mathcal{L}_x - \rho) V^\lambda(x, y) + \pi(x)y - \lambda y \log y, -V_y^\lambda(x, y) + G(x) \} \leq 0.$$

To complete the proof, if (\bar{x}, \bar{y}) is such that $V_y^\lambda(\bar{x}, \bar{y}) - G(\bar{x}) > 0$, by continuity of V_y^λ (cf. Step 5) we can find $r > 0$ such that $V_y^\lambda(x, \bar{y}) - G(x) > 0$ for any $x \in B_r(0)$. Thus, for a.a. $x \in B_r(0)$ there exists \bar{n}_x such that $V_y^{\lambda, \varepsilon_n^{\bar{y}}}(x, \bar{y}) - G(x) > 0$ for any $n \geq \bar{n}_x$, which gives $(\rho - \mathcal{L}_x) V^{\lambda, \varepsilon_n^{\bar{y}}}(x, \bar{y}) - \Pi(x, \bar{y}) = 0$ for any $n \geq \bar{n}_x$ so that

$$\lim_n (\rho - \mathcal{L}_x) V^{\lambda, \varepsilon_n^{\bar{y}}}(\cdot, \bar{y}) - \Pi(\cdot, \bar{y}) = 0, \quad \text{a.e. in } x \in B_r(0).$$

Since $(\rho - \mathcal{L}_x) V^\lambda(\cdot, \bar{y}) - \Pi(\cdot, \bar{y})$ is the weak limit in $\mathbb{L}^2(B_r(\bar{x}))$ of $((\rho - \mathcal{L}_x) V^{\lambda, \varepsilon_n^{\bar{y}}}(\cdot, \bar{y}) - \Pi(\cdot, \bar{y}))_n$, the latter limits in turn implies that $(\rho - \mathcal{L}_x) V^\lambda(\cdot, \bar{y}) - \Pi(\cdot, \bar{y}) \geq 0$ a.e. in $B_r(\bar{x})$, thus showing the HJB with equality and completing the proof.

3.3. Proof of Theorem 3.2: uniqueness for the HJB. Uniqueness of the solution follows from a verification argument, that we prove in two steps.

Step 1. Let \tilde{V} be a solution of the HJB equation as in the statement of the theorem. Fix $(x, y) \in \mathbb{R}^n \times [0, 1]$ and consider a generic $\xi \in \mathcal{A}(y)$. For $R > 0$, set $\tau_R := \inf\{t \geq 0 \mid |X_t^x| \geq R\}$. Thanks to the regularity of \tilde{V} , we can employ (a generalized) Itô's formula to the process $(e^{-\rho t} \tilde{V}(X_t^x, Y_t^{y, \xi}))_t$ (see [23, Chapter 8, Section VIII.4, Theorem 4.1] and [11, Theorem 4.2]), up to the stopping time $\tau_R \wedge T$, for some $T > 0$, to obtain

$$\begin{aligned} \mathbb{E} \left[e^{-\rho \tau_R \wedge T} \tilde{V}(X_{\tau_R \wedge T}^x, Y_{\tau_R \wedge T}^{y, \xi}) \right] &= \tilde{V}(x, y) + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} (\mathcal{L}_x - \rho) \tilde{V}(X_t^x, Y_t^{y, \xi}) dt \right] \\ &\quad + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \tilde{V}_y(X_t^x, Y_t^{y, \xi}) dY_t^{y, \lambda} \right] \\ &\quad + \mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} \left(\Delta \tilde{V}(X_t^x, Y_t^{y, \xi}) - \tilde{V}_y(X_t^x, Y_t^{y, \xi}) \Delta Y_t^{y, \xi} \right) \right], \end{aligned} \quad (3.35)$$

with the notations $\Delta \tilde{V}(X_t, Y_t^{y, \xi}) = \tilde{V}(X_t, Y_t^{y, \xi}) - \tilde{V}(X_t, Y_{t-}^{y, \xi})$ and $\Delta Y_t^{y, \xi} = Y_t^{y, \xi} - Y_{t-}^{y, \xi}$. The concavity of \tilde{V} in y now implies that

$$\sum_{0 \leq t \leq \tau_R \wedge T} \left(\Delta \tilde{V}(X_t^x, Y_t^{y, \xi}) - \tilde{V}_y(X_t^x, Y_t^{y, \xi}) \Delta Y_t^{y, \xi} \right) \leq 0.$$

Moreover, since \tilde{V} solves the HJB equation, we have

$$(\mathcal{L}_x - \rho) \tilde{V}(X_t^x, Y_{t-}^{y, \xi}) \leq -\Pi(X_t^x, Y_{t-}^{y, \xi}) \quad \text{and} \quad -\tilde{V}_y(X_t^x, Y_{t-}^{y, \xi}) \leq -G(X_t^x).$$

Plugging the latter two inequalities into (3.35) and then rearranging the terms, we obtain

$$\begin{aligned} \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \Pi(X_t^x, Y_{t-}^{y, \xi}) dt + \int_0^{\tau_R \wedge T} e^{-\rho t} G(X_t^x) d\xi_t \right] \\ \leq \tilde{V}(x, y) - \mathbb{E} \left[e^{-\rho \tau_R \wedge T} \tilde{V}(X_{\tau_R \wedge T}^x, Y_{\tau_R \wedge T}^{y, \xi}) \right]. \end{aligned}$$

Finally, thanks to the growth conditions of π and G , to the growth of \tilde{V} , to the estimates (3.5) and (3.8), we can use the dominated convergence theorem to take limits, first as $T \rightarrow \infty$ and then as $R \rightarrow \infty$, in order to obtain

$$J^\lambda(x, y; \xi) \leq \tilde{V}(x, y).$$

By arbitrariness of ξ , we deduce that $V^\lambda(x, y) \leq \tilde{V}(x, y)$.

Step 2. We adapt the arguments in [9] to construct a control ξ^λ which satisfies

$$J^\lambda(x, y; \xi^\lambda) = \tilde{V}(x, y).$$

This identity, together with the previous step, implies that $V^\lambda = \tilde{V}$ and that ξ^λ is optimal.

Define the function $\tilde{g}_\lambda : \mathbb{R}^n \rightarrow [0, 1]$ by

$$\tilde{g}_\lambda(x) := \sup \{ y \in [0, 1] \mid -\tilde{V}_y(x, y) + G(x) < 0 \},$$

where we set $\tilde{g}_\lambda(x) := 1$ if $\{y \in [0, 1] \mid -\tilde{V}_y(x, y) + G(x) < 0\} = \emptyset$. The function \tilde{g}_λ is well defined by concavity of \tilde{V} . Noticing that $\{(x, y) \in \mathbb{R}^n \times [0, 1] \mid y < \tilde{g}_\lambda(x)\} = \{(x, y) \mid -\tilde{V}_y(x, y) + G(x) < 0\}$, by continuity of \tilde{V}_y we deduce that the set $\{(x, y) \in \mathbb{R}^n \times [0, 1] \mid y < \tilde{g}_\lambda(x)\}$ is open, so that \tilde{g}_λ is lower semi-continuous.

Next, for fixed $(x, y) \in \mathbb{R}^n \times [0, 1]$, define the process ξ^λ by setting

$$\xi_t^\lambda := \sup_{s \leq t} (y - \tilde{g}_\lambda(X_s^x))^+, \quad t \geq 0.$$

Such a process is clearly monotone nondecreasing, so that it admits left limits. To show its right-continuity, use the upper semi-continuity of $x \mapsto (y - \tilde{g}_\lambda(x))^+$ to obtain

$$\begin{aligned} \lim_{s \downarrow t} \xi_s^\lambda &= \xi_t^\lambda \vee \lim_{s \downarrow t} \sup_{t \leq r \leq s} (y - \tilde{g}_\lambda(X_r^x))^+ \\ &= \xi_t^\lambda \vee \lim_{s \downarrow t} \sup (y - \tilde{g}_\lambda(X_s^x))^+ \leq \xi_t^\lambda \vee (y - \tilde{g}_\lambda(X_t^x))^+ = \xi_t^\lambda. \end{aligned}$$

By monotonicity of ξ^λ , this implies that $\lim_{s \downarrow t} \xi_s^\lambda = \xi_t^\lambda$. Thus, since this process is clearly adapted, it is progressively measurable, showing the admissibility of ξ^λ ; i.e., $\xi^\lambda \in \mathcal{A}(y)$.

We will now prove the optimality of ξ^λ by repeating, with some modifications, the passages in Step 1. For $Y^{y,\lambda} := Y^{y,\xi^\lambda}$, $\tau_R := \inf\{t \geq 0 \mid |X_t^x| \geq R\}$ and for $T > 0$ we find

$$\begin{aligned} \mathbb{E} \left[e^{-\rho \tau_R \wedge T} \tilde{V}(X_{\tau_R \wedge T}, Y_{\tau_R \wedge T}^{y,\lambda}) \right] &= \tilde{V}(x, y) + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} (\mathcal{L}_x - \rho) \tilde{V}(X_t^x, Y_{t-}^{y,\lambda}) dt \right] \\ &\quad + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \tilde{V}_y(X_t^x, Y_{t-}^{y,\lambda}) dY_t^{y,\lambda} \right] \\ &\quad + \mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} \left(\Delta \tilde{V}(X_t^x, Y_t^{y,\lambda}) - \tilde{V}_y(X_t^x, Y_{t-}^{y,\lambda}) \Delta Y_t^{y,\lambda} \right) \right], \end{aligned} \tag{3.36}$$

with the notations $\Delta \tilde{V}(X_t, Y_t^{y,\lambda}) = \tilde{V}(X_t, Y_t^{y,\lambda}) - \tilde{V}(X_t, Y_{t-}^{y,\lambda})$ and $\Delta Y_t^{y,\lambda} = Y_t^{y,\lambda} - Y_{t-}^{y,\lambda}$.

Notice that, by construction of ξ^λ , if the process $Y^{y,\lambda}$ jumps at time t , then we have $\tilde{g}_\lambda(X_t^x) < \tilde{g}_\lambda(X_{t-}^x)$ and

$$\tilde{V}_y(X_t^x, \zeta) = G(X_t^x), \quad \text{for any } \zeta \in [Y_t^{y,\lambda}, Y_{t-}^{y,\lambda}].$$

Therefore, we obtain

$$\Delta \tilde{V}(X_t^x, Y_t^{y,\lambda}) - \tilde{V}_y(X_t^x, Y_{t-}^{y,\lambda}) \Delta Y_t^{y,\lambda} = G(X_t^x)(Y_t^{y,\lambda} - Y_{t-}^{y,\lambda}) - G(X_t^x)(Y_t^{y,\lambda} - Y_{t-}^{y,\lambda}) = 0.$$

Using the latter equality in (3.36), together with the fact that the process ξ^λ increases only at times t in which $\tilde{V}_y(X_t^x, Y_{t-}^{y,\lambda}) = G(X_t^x)$ and the fact that $(\mathcal{L}_x - \rho) \tilde{V}(X_t^x, Y_{t-}^{y,\lambda}) = \Pi(X_t^x, Y_{t-}^{y,\lambda})$ $\mathbb{P} \otimes dt$ -a.e., we find

$$\begin{aligned} \tilde{V}(x, y) &= \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \Pi(X_t^x, Y_{t-}^{y,\lambda}) dt + \int_0^{\tau_R \wedge T} e^{-\rho t} G(X_t^x) d\xi_t^\lambda \right] \\ &\quad + \mathbb{E} \left[e^{-\rho \tau_R \wedge T} \tilde{V}(X_{\tau_R \wedge T}, Y_{\tau_R \wedge T}^{y,\lambda}) \right]. \end{aligned}$$

Taking limits (first as $T \rightarrow \infty$ and then as $M \rightarrow \infty$) we conclude that

$$\tilde{V}(x, y) = J^\lambda(x, y; \xi^\lambda),$$

thus proving the optimality of ξ^λ . This completes the proof.

3.4. Proof of Theorem 3.3. In order to show Theorem 3.3, it is enough to repeat the arguments in Step 2 in Subsection 3.3 with the value function V^λ and the free boundary g_λ defined in (3.4), and using the fact that V^λ is a solution of the HJB equation (cf. Theorem 3.2).

4. A REAL OPTION PROBLEM: ANALYTICAL SOLUTIONS AND REINFORCEMENT LEARNING ALGORITHMS

This section is devoted to a classical example of the optimal stopping problem arising from real option evaluation (see Chapter 5 in [47] and the references therein for further details). We provide analytical properties of the problem in Subsection 4.1 and develop the corresponding RL algorithms in Subsection 4.2.

We start with the introduction of the problem formulation. For parameters $\mu \in \mathbb{R}, \sigma, \kappa, \rho > 0$ and $\pi : \mathbb{R}_+ \rightarrow \mathbb{R}$, consider the optimal stopping problem

$$\begin{aligned} & \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\int_0^\tau e^{-\rho s} \pi(X_s^x) ds + \kappa e^{-\rho \tau} \right], \\ & \text{subject to } dX_t^x = \mu X_t^x dt + \sigma X_t^x dW_t, \quad X_0^x = x > 0. \end{aligned}$$

In order to have well posedness of the problem, we assume $\rho > \mu$. Noticing that,

$$\mathbb{E} \left[\int_0^\tau e^{-\rho s} \pi(X_s^x) ds + \kappa e^{-\rho \tau} \right] = \mathbb{E} \left[\int_0^\tau e^{-\rho s} (\pi(X_s^x) - \rho \kappa) ds \right] + \kappa,$$

we define the OS problem

$$(4.1) \quad V(x) := \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\int_0^\tau e^{-\rho s} (\pi(X_s^x) - \rho \kappa) ds \right].$$

Despite the latter OS problem can be explicitly solved, we are interested here in illustrating the relations between the classical solution approach and the entropy regularization approach (see in particular the discussion in Subsection 4.1.3 below).

Following the strategy outlined in Section 3, introduce the extra state variable $y \in [0, 1]$ and, for $\lambda > 0$, define the singular control problem

$$(4.2) \quad \begin{aligned} & V^\lambda(x, y) := \sup_{\xi \in \mathcal{A}(y)} \mathbb{E} \left[\int_0^\infty e^{-\rho t} (\pi(X_t^x) - \rho \kappa) Y_t^{y, \xi} - \lambda Y_t^{y, \xi} \log(Y_t^{y, \xi}) dt \right], \\ & \text{subject to } dX_t^x = \mu X_t^x dt + \sigma X_t^x dW_t, \quad X_0^x = x, \\ & dY_t^{y, \xi} = -d\xi_t, \quad Y_{0-}^{y, \xi} = y, \end{aligned}$$

with $\mathcal{A}(y)$ defined in (3.1). The HJB equation associated to V^λ reads as

$$\max \left\{ (\mathcal{L}_x - \rho)V^\lambda + (\pi(x) - \rho \kappa)y - \lambda y \log y, -V_y^\lambda \right\} = 0, \quad \text{a.e. in } (0, \infty) \times (0, 1),$$

with boundary condition $V(x, 0) = 0, x \in (0, \infty)$, where $\mathcal{L}_x \phi(x) = \mu x \phi_x(x) + \frac{1}{2} \sigma^2 x^2 \phi_{xx}(x)$.

The current problem differs from the general setting studied in Sections 2 and 3 due to the behavior of π at state 0. However, this model can be solved semi-explicitly with the –so called– *guess and verify* approach, which consists in guessing first a candidate value function and then proving that the candidate coincides with the true value function (see Theorem 4.2).

Assumption 4.1. $\pi \in C^2(\mathbb{R}_+)$ with $\pi(0) = 0$ satisfies the following conditions:

- π is non-decreasing and concave, π has a sublinear growth, and $x\pi'(x)$ is strictly increasing. More specifically, there exists some $\theta \in (0, 1)$ and $c > 0$ such that

$$|\pi(x)| \leq c(1 + |x|^\theta).$$

One example that satisfies Assumption 4.1 is $\pi(x) = c_0 x^{\theta_0}$ with $\theta_0 \in (0, \theta]$ and $c_0 \in (0, c]$.

4.1. Analytical solution. In this section, we analyze the solutions of the real option problem (4.2). Specifically, we construct a semi-explicit candidate solution in Subsection 4.1.1 and provide the verification theorem in Subsection 4.1.2.

4.1.1. *Construction of a candidate solution.* We will now search for a candidate solution $u : (0, \infty) \times [0, 1] \rightarrow \mathbb{R}$ and a nondecreasing function $g_\lambda : (0, \infty) \rightarrow [0, 1]$ such that

$$(4.3) \quad \begin{cases} (\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y = 0, & 0 < y \leq g_\lambda(x), \\ -u_y = 0, & y > g_\lambda(x), \end{cases}$$

and we will later verify that $u = V^\lambda$ (see Theorem 4.2 below).

If u and g_λ satisfy (4.3), for $0 < y \leq g_\lambda(x)$, we can integrate the equation to obtain

$$(4.4) \quad u(x, y) = A_2(y)x^{\alpha_-} + A_1(y)x^{\alpha_+} + \mathbb{E} \left[\int_0^\infty e^{-\rho t} ((\pi(X_t^x) - \rho\kappa)y - \lambda y \log y) dt \right],$$

with α_+, α_- solving $\frac{1}{2}\sigma^2\alpha(\alpha-1) + \mu\alpha - \rho = 0$, and such that $\alpha_+ > 1$ (as $\rho > \mu$) and $\alpha_- < 0$. Moreover, since $\rho > \mu$ and Assumption 4.1 hold, we obtain that

$$V^\lambda(x, y) \leq C + C\mathbb{E} \left[\int_0^\infty e^{-\rho t} (X_t^x)^\theta dt \right] = C + Cx^\theta \mathbb{E} \left[\int_0^\infty e^{(-\rho + \theta(\mu + \frac{\sigma^2}{2}(\theta-1)))t} dt \right] \leq C(1 + x^\theta).$$

Thus, we take $A_1 \equiv 0$, as otherwise u would explode as $x \uparrow \infty$, not matching the growth behavior of V^λ . Finally, when for $0 < y \leq g_\lambda(x)$, a direct integration leads to

$$u(x, y) = A_2(y)x^{\alpha_-} + H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y, \quad \text{with } H_\pi(x) = \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) dt \right].$$

To summarize, we have obtained the following candidate value function

$$u(x, y) = \begin{cases} A_2(y)x^{\alpha_-} + H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y, & 0 < y \leq g_\lambda(x), \\ A_2(g_\lambda(x))x^{\alpha_-} + \pi(x)y - \kappa g_\lambda(x) - \frac{\lambda}{\rho}g_\lambda(x) \log(g_\lambda(x)), & y > g_\lambda(x). \end{cases}$$

In order to determine A_2 and g_λ we impose

$$\lim_{z \uparrow g_\lambda(x)} u_y(x, z) = \lim_{z \uparrow g_\lambda(x)} u_{yx}(x, z) = 0.$$

From these conditions, we derive the system

$$\begin{cases} A_2'(g_\lambda(x))x^{\alpha_-} + H_\pi(x) - \kappa - \frac{\lambda}{\rho}(1 + \log g_\lambda(x)) = 0, \\ \alpha_- A_2'(g_\lambda(x))x^{\alpha_- - 1} + H_\pi'(x) = 0, \end{cases}$$

which can be solved as

$$(4.5) \quad \begin{cases} g_\lambda(x) = \exp \left(\frac{-\frac{H_\pi'(x)}{\alpha_-}x + H_\pi(x) - \kappa - \frac{\lambda}{\rho}}{\frac{\lambda}{\rho}} \right) \\ A_2(y) = \int_{g_\lambda(0)}^y \frac{\kappa + \frac{\lambda}{\rho} \log(u) + \frac{\lambda}{\rho} - H_\pi(g_\lambda^{-1}(u))}{(g_\lambda^{-1}(u))^{\alpha_-}} du, \end{cases}$$

Notice that

$$(4.6) \quad g_\lambda'(x) = g_\lambda(x) \left(\frac{-1}{\alpha_-} (H_\pi'(x)x)' + H_\pi'(x) \right) > 0,$$

as $(H_\pi'(x)x)' > 0$ and $H_\pi'(x) \geq 0$ according to Assumption 4.1. Hence g_λ is strictly increasing on $(0, \hat{x}_{g_\lambda}]$ with $\hat{x}_{g_\lambda} = \min\{x \in \mathbb{R}_+ : g_\lambda(x) = 1\}$.

Therefore, our construction of u and g_λ is valid only for $y \geq y^\lambda$ with

$$(4.7) \quad y^\lambda := g_\lambda(0) = e^{-1 - \frac{\kappa\rho}{\lambda}} \in (0, 1).$$

For $y < y^\lambda$ we notice that the function

$$u_0(x, y) := H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y,$$

is such that

$$(4.8) \quad \begin{cases} (\mathcal{L}_x - \rho)u_0 + (x^\theta - \rho\kappa)y - \lambda y \log y = 0, \\ -\partial_y u_0 = -H_\pi(x) + \left(\kappa + \frac{\lambda}{\rho}(\log y + 1)\right) < 0, \quad \text{for } y < y^\lambda. \end{cases}$$

Thus, $u_0(x, y)$ solves the HJB and for $y < y^\lambda$ and we expect $\xi^\lambda \equiv 0$ to be optimal.

Summarizing the previous heuristic, set $F(x, y) := A_2(y) x^{\alpha_-} + H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y$ and define the candidate value as

$$u(x, y) := \begin{cases} H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y, & y < y^\lambda, \\ F(x, y), & y \leq g_\lambda(x), \ y \geq y^\lambda, \\ F(x, b_\lambda^{-1}(x)), & y > g_\lambda(x), \ y \geq y^\lambda, \end{cases}$$

and we extend g_λ on $[0, \infty)$ as follows

$$g_\lambda(x) = \begin{cases} g_\lambda(x), & x \in [0, \hat{x}_{g_\lambda}] \\ 1, & x \geq \hat{x}_{g_\lambda}. \end{cases}$$

4.1.2. Verification theorem. The next result verifies that the function u constructed in the previous subsection is the value function.

Theorem 4.2. *The following claims hold true:*

- (1) *The candidate value function u coincides with the true value function V^λ ;*
- (2) *The reflection policy $\xi_t^\lambda := \sup_{0 \leq s \leq t} (y - g_\lambda(X_s^x))^+$ is optimal for the initial condition (x, y) , where g_λ is defined in (4.5).*

Proof. The proof is divided in three steps.

Step 1. In this step we show the C^2 -regularity of u . We split $(0, \infty) \times [0, 1]$ into the three sets $\mathcal{E}_0 := \{(x, y) \mid y < y^\lambda\}$, $\mathcal{E} := \{(x, y) \mid y \geq y^\lambda, \ y < g_\lambda(x)\}$ and $\mathcal{S} := \{(x, y) \mid y \geq g_\lambda(x), \ y \geq y^\lambda\}$. Clearly, u is regular in the interior of \mathcal{E}_0 , \mathcal{E} and \mathcal{S} . In order to show the C^2 -regularity across \mathcal{E}_0 and \mathcal{E} , it is sufficient to notice that $A_2(y^\lambda) = A_2'(y^\lambda) = A_2''(y^\lambda) = 0$.

We next study the behavior across \mathcal{S} and \mathcal{E} . To this end, recall the definition of F and notice that

$$\begin{cases} u(x, y) = F(x, y), & y < g_\lambda(x), \\ u(x, y) = F(x, g_\lambda(x)), & y \geq g_\lambda(x). \end{cases}$$

Clearly we have

$$\lim_{y \uparrow g_\lambda(x)} u(x, y) = F(x, g_\lambda(x)) = \lim_{y \downarrow g_\lambda(x)} u(x, y),$$

so that u is continuous across the boundary. Furthermore, after straightforward calculations we find

$$\begin{aligned} u_x(x, y) &= F_x(x, y), \quad y < g_\lambda(x), \\ u_x(x, y) &= F_x(x, g_\lambda(x)) + \underbrace{F_y(x, g_\lambda(x)) (g_\lambda)'(x)}_{=0}, \quad y > g_\lambda(x), \end{aligned}$$

so that

$$\lim_{y \uparrow g_\lambda(x)} u_x(x, y) = F_x(x, g_\lambda(x)) = \lim_{y \downarrow g_\lambda(x)} u_x(x, y),$$

which in turn implies the continuity of u_x . Finally,

$$\begin{aligned} u_{xx}(x, y) &= F_{xx}(x, y), \quad y < g_\lambda(x), \\ u_{xx}(x, y) &= F_{xx}(x, g_\lambda(x)) + \underbrace{F_{xy}(x, g_\lambda(x)) (g_\lambda)'(x)}_{=0}, \quad y > g_\lambda(x), \end{aligned}$$

from which we conclude that

$$(4.9) \quad \lim_{y \downarrow g_\lambda(x)} u_{xx}(x, y) = F_{xx}(x, g_\lambda(x)) = \lim_{y \uparrow g_\lambda(x)} u_{xx}(x, y),$$

which is the continuity of u_{xx} . For the derivatives involving y , by using the definitions of A_2 and g_λ in (4.5), a direct computation leads to

$$\begin{aligned} \lim_{x \downarrow g_\lambda^{-1}(y)} u_y(x, y) &= F_y(g_\lambda^{-1}(y), y) = 0 = \partial_y F(x, g_\lambda(x)), \\ \lim_{x \downarrow g_\lambda^{-1}(y)} u_{yy}(x, y) &= F_{yy}(g_\lambda^{-1}(y), y) = 0 = \partial_{yy}^2 F(x, g_\lambda(x)), \\ \lim_{x \downarrow g_\lambda^{-1}(y)} u_{yx}(x, y) &= F_{yx}(g_\lambda^{-1}(y), y) = 0 = \partial_{yx}^2 F(x, g_\lambda(x)), \end{aligned}$$

which gives us the desired regularity.

Step 2. In this step we show that u solves the HJB equation. Let \mathcal{E}_0 , \mathcal{E} and \mathcal{S} be as in the previous step. First, on the set \mathcal{E}_0 , u satisfies the HJB equation by (4.8). Secondly, on the set \mathcal{E} , by construction we have $(\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y = 0$, and we need to verify that $-u_y \leq 0$. To this end, use the definition of A_2 and b_λ to compute the derivative

$$\begin{aligned} -u_{yx}(x, y) &= -\alpha_- A'_2(y) x^{\alpha_- - 1} - H'_\pi(x) \\ &= -\alpha_- \frac{\kappa + \frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} - H_\pi(g_\lambda^{-1}(y))}{(g_\lambda^{-1}(y))^{\alpha_-}} x^{\alpha_- - 1} - H'_\pi(x) \\ &= -\alpha_- \frac{(\kappa + \frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} - H_\pi(g_\lambda^{-1}(y))) x^{\alpha_- - 1} + \frac{1}{\alpha_-} H'_\pi(x) (g_\lambda^{-1}(y))^{\alpha_-}}{(g_\lambda^{-1}(y))^{\alpha_-}} x^{\alpha_- - 1} \\ &= -\alpha_- \frac{\frac{H'_\pi(g_\lambda^{-1}(y))}{-\alpha_-} g_\lambda^{-1}(y) g_\lambda^{-1}(y) (x^{\alpha_- - 1} - (g_\lambda^{-1}(y))^{\alpha_- - 1})}{(g_\lambda^{-1}(y))^{\alpha_-}} x^{\alpha_- - 1} \leq 0 \end{aligned}$$

where the last inequality holds since $x^{\alpha_- - 1} - (g_\lambda^{-1}(y))^{\alpha_- - 1} < 0$, which is a direct consequence of $g_\lambda^{-1}(y) < x$. Thus, since $-u_y(x, g_\lambda(x)) = 0$, we have that $-u_y \leq 0$ on \mathcal{E} .

We next study the behavior on \mathcal{S} . By construction we have $-u_y = 0$ in \mathcal{S} . Moreover, using the expression for u , u_x and u_{xx} computed in the previous step, for any x such that $y > g_\lambda^{-1}(x)$ we find

$$\begin{aligned} &(\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y \\ &= \frac{1}{2} \sigma^2 x^2 \partial_{xx} (F(x, b_\lambda^{-1}(x))) + \mu x \partial_x (F(x, b_\lambda^{-1}(x))) - \rho F(x, b_\lambda^{-1}(x)) \\ &\quad + (\pi(x) - \rho\kappa)y - \lambda y \log y \\ &= \frac{1}{2} \sigma^2 x^2 F_{xx}(x, b_\lambda^{-1}(x)) + \mu x F_x(x, b_\lambda^{-1}(x)) - \rho F(x, b_\lambda^{-1}(x)) \\ &\quad + (\pi(x) - \rho\kappa)y - \lambda y \log y. \end{aligned}$$

Hence, using the ordinary differential equation for $y < g_\lambda(x)$, in the limit as $y \uparrow g_\lambda(x)$, we proceed with

$$\begin{aligned}
& (\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y \\
&= -(\pi(x) - \rho\kappa)g_\lambda(x) + \lambda g_\lambda(x) \log(g_\lambda(x)) + (\pi(x) - \rho\kappa)y - \lambda y \log y \\
&= \int_{g_\lambda(x)}^y \left(\pi(x) - \rho \left(k + \frac{\lambda}{\rho} (1 + \log z) \right) \right) dz \\
&< (y - g_\lambda(x)) \left(\pi(x) - \rho \left(k + \frac{\lambda}{\rho} (1 + \log(g_\lambda(x))) \right) \right) \\
&= \underbrace{(y - g_\lambda(x))}_{\geq 0} \left(\pi(x) - \rho \left(\frac{H'_\pi(x)}{-\alpha} x + H_\pi(x) \right) \right).
\end{aligned}$$

Note that H_π satisfies the following PDE:

$$(4.10) \quad \mu x H'_\pi(x) - \rho H_\pi(x) + \pi(x) = -\frac{1}{2} \sigma^2 x^2 H''_\pi(x).$$

Therefore

$$(4.11) \quad \pi(x) - \rho \left(\frac{H'_\pi(x)}{-\alpha} x + H_\pi(x) \right) = - \left(\mu + \frac{\rho}{-\alpha_-} \right) x H'_\pi(x) - \frac{1}{2} \sigma^2 x^2 H''_\pi(x) \leq 0,$$

where the last inequality holds since $\frac{1}{2} \sigma^2 \alpha_- (\alpha_- - 1) + \mu \alpha_- - \rho = 0$ and $\alpha_- < 0$. Hence $(\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y \leq 0$ on \mathcal{S} , so that u solves the HJB equation in the region \mathcal{S} .

Step 3. The last step consists of constructing a candidate optimal control ξ^λ as in the statement of the theorem and verifying its optimality by Itô's formula. The argument is analogous to that in Subsections 3.3 and 3.4, so that we omit the details to avoid repetitions. \square

4.1.3. *Vanishing exploration and free boundary limits.* In the case of no exploration (i.e., $\lambda = 0$), the standard OS problem (4.1) can be solved semi-explicitly. In particular, the optimal stopping time is given by

$$\tau^* = \inf \{t \geq 0 : X_t^x \geq b^*\}$$

with b^* solves the equation: $-\frac{H'_\pi(x)}{\alpha_-} x + H_\pi(x) = \kappa$. As a special case of $\pi(x) = x^\theta$ with $\theta \in (0, 1)$, we have $b^* = \left[-\frac{1}{P} \left(\frac{\alpha_-}{\theta - \alpha_-} \right) \kappa \right]^{1/\theta}$ with $P = \frac{1}{\rho + \frac{1}{2} \sigma^2 \theta (1 - \theta) - \theta \mu}$.

It is important to notice that, if $b_\lambda := g_\lambda^{-1} : [y^\lambda, 1] \rightarrow [0, \infty)$ denotes the free boundary of the entropy regularized OS problem with $\lambda > 0$ (determined in (4.5)) and with y^λ as in (4.7), we have the limits $\lim_{\lambda \downarrow 0} y^\lambda = 0$. Let $\lambda \downarrow 0$ in the following equation for any $y \in (0, 1]$

$$\frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} = \frac{H'_\pi(b_\lambda(y))}{-\alpha_-} b_\lambda(y) + H_\pi(b_\lambda(y)) - \kappa.$$

We observe that for a fixed $y \in (0, 1]$, $b_\lambda(y)$ is monotone increasing in λ , as we have $(H'_\pi(x)x)' > 0$ and $H'_\pi \geq 0$ according to Assumption 4.1. Hence $b_\lambda(y) \downarrow$ as $\lambda \downarrow 0$. Given that $b_\lambda(y) \geq 0$ for all $\lambda \geq 0$, there must exist a limit. By continuity argument, it is easy to show that the limit $\lim_{\lambda \downarrow 0} b_\lambda(y)$, which is independent of y , solves

$$0 = -\frac{H'_\pi(x)}{\alpha_-} x + H_\pi(x) - \kappa.$$

Hence we have $b^* \equiv \lim_{\lambda \downarrow 0} b_\lambda(y)$.

Thus, the boundary of the entropy regularized OS problem approximates the boundary of the (standard) OS problem (see also Propositions 2.9 and 2.10).

4.2. Reinforcement learning algorithm. In this section, we propose an RL framework to learn the reflection boundary g^λ and the corresponding value function V^λ , without prior knowledge on the form of g^λ . Our framework consists of two steps: a model-free step (see Subsection 4.2.1) and a model-based step (see Subsection 4.2.2). In the first step, where all model parameters are known, we design a Policy Iteration algorithm to learn g_λ . In the second step, where the model parameters are unknown, we combine the Policy Iteration with a Policy Evaluation procedure.

Note that the RL framework is currently formulated in terms of the real option example. We believe the results can be extended to a more general setting and in higher dimensions.

4.2.1. Model-based analysis. Motivated by the analysis in Subsection 4.1.1, here we focus on a subclass of control policies that can be fully characterized by some reflection boundary g that satisfies the following assumption.

Assumption 4.3. Assume $g : \mathbb{R}_+ \rightarrow (0, 1]$ is a non-decreasing function such that

$$g \in C^1([0, \hat{x}_g]), \quad \text{where} \quad \hat{x}_g = \inf_{x \in \mathbb{R}_+} \{g(x) = 1\}.$$

For a function g that satisfies Assumption 4.3, we define the associated policy ξ^g of the following form:

$$(4.12) \quad \xi_t^g = \sup_{s \leq t} (y - g(X_s^x))^+, t \geq 0.$$

Define the associated value function as:

$$(4.13) \quad \begin{aligned} V_g^\lambda(x, y) &:= J^\lambda(x, y; \xi^g) \\ &= \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left((\pi(X_t^x) - \rho\kappa) Y_t^{y, \xi^g} - \lambda Y_t^{y, \xi^g} \log(Y_t^{y, \xi^g}) \right) dt \right] \end{aligned}$$

subject to $Y_t^{y, \xi^g} = y - \xi_t^g$ with ξ_t^g defined in (4.12).

The policy in (4.12) defines two areas: the exploration area and the stopping area, both associated with g :

$$(4.14) \quad \mathcal{E}(g) = \{(x, y) \mid y \leq g(x)\},$$

$$(4.15) \quad \mathcal{S}(g) = \{(x, y) \mid y > g(x)\}.$$

Then we have the following results for the value function V_g^λ associated with policy ξ^g .

Theorem 4.4. Assume Assumptions 4.1-4.3 hold. Then the following HJB equation has a unique $C^1(\mathbb{R}_+ \times [0, 1]) \cap C^2(\overline{\mathcal{E}(g)})$ solution:

$$(4.16) \quad (\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y = 0 \quad \text{on} \quad \mathcal{E}(g),$$

$$(4.17) \quad -u_y = 0, \quad \text{on} \quad \mathcal{S}(g).$$

In addition, the solution satisfies $u_{xx} \in \mathbb{L}_{loc}^\infty(\mathbb{R}_+ \times [0, 1])$. Finally, we have the verification theorem holds, namely

$$(4.18) \quad V_g^\lambda(x, y) \equiv u(x, y).$$

Proof. Note that g^{-1} is well defined under Assumption 4.3 with $g^{-1}(x) > 0$ for all $x \in \mathbb{R}_+$. We seek for a semi-explicit solution to

$$(4.19) \quad (\mathcal{L}_x - \rho)u + (\pi(x) - \rho\kappa)y - \lambda y \log y = 0 \quad \text{on} \quad \mathcal{E}(g),$$

$$(4.20) \quad -u_y = 0, \quad \text{on} \quad \mathcal{S}(g).$$

Similar to the analysis in Section 4 and following the idea in [20, 21], take

$$(4.21) \quad u(x, y) = A(y)x^{\alpha_-} + H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y,$$

with α_- defined in (4.4) and $H_\pi(x)$ defined as

$$(4.22) \quad H_\pi(x) := \mathbb{E} \left[\int_0^\infty e^{-\rho t} \pi(X_t^x) dt \right].$$

Under Assumption 4.1, we have $H_\pi(x) \in C^2(\mathbb{R}_+)$ with sublinear growth [42, Eqn (2.11)].

Then,

$$(4.23) \quad u_y(x, y) = A'(y)x^{\alpha_-} + H_\pi(x) - \kappa - \frac{\lambda}{\rho} \log y - \frac{\lambda}{\rho}.$$

Setting $u_y(g^{-1}(y), y) = 0$, we have

$$(4.24) \quad A'(y)(g^{-1}(y))^{\alpha_-} + H_\pi(g^{-1}(y)) - \kappa - \frac{\lambda}{\rho} \log y - \frac{\lambda}{\rho} = 0.$$

Hence, $A'(y) = \frac{\kappa + \frac{\lambda}{\rho} \log y + \frac{\lambda}{\rho} - H_\pi(g^{-1}(y))}{(g^{-1}(y))^{\alpha_-}}$ and

$$(4.25) \quad A(y) = \int_{g(0)}^y \frac{\kappa + \frac{\lambda}{\rho} \log(u) + \frac{\lambda}{\rho} - H_\pi(g^{-1}(u))}{(g^{-1}(u))^{\alpha_-}} du.$$

It is easy to check that $u(x, y)$ is C^1 in y on $\mathbb{R}_+ \times [0, 1]$ and that, under condition $g \in C^1([0, \hat{x}_g])$, we have $u(x, y) \in C^2(\mathcal{E}(g))$.

Now, on the curve $y = g(x)$,

$$(4.26) \quad u_x^+(x, y) = A(g(x))\alpha_- x^{\alpha_- - 1} + R'_\pi(x)g(x).$$

On the other hand, we have $u(x, y) = u(x, g(x))$ for $(x, y) \in \mathcal{S}(g)$. Hence

$$\begin{aligned} & \frac{u(x, g(x)) - u(x - \delta, g(x))}{\delta} = \frac{u(x, g(x)) - u(x - \delta, g(x - \delta))}{\delta} \\ &= \frac{1}{\delta} \left(A(g(x))x^{\alpha_-} - A(g(x - \delta))(x - \delta)^{\alpha_-} \right) + \frac{1}{\delta} \left(H_\pi(x)g(x) - H_\pi(x - \delta)g(x - \delta) \right) \\ & \quad - \frac{\kappa}{\delta} \left(g(x) - g(x - \delta) \right) - \frac{\lambda}{\rho\delta} \left(g(x) \log(g(x)) - g(x - \delta) \log(g(x - \delta)) \right), \end{aligned}$$

and

$$\begin{aligned} u_x^-(x, y) &= \lim_{\delta \downarrow 0} \frac{u(x, g(x)) - u(x - \delta, g(x))}{\delta} \\ &= A(g(x))\alpha_- x^{\alpha_- - 1} + R'_\pi(x)g(x) + u_y(x, g(x)) \\ &= A(g(x))\alpha_- x^{\alpha_- - 1} + R'_\pi(x)g(x)g(x) = u_x^+(x, y), \end{aligned}$$

where the last equality holds since $u_y(x, g(x)) = 0$. Hence $u \in C^1(\mathbb{R}_+ \times [0, 1])$.

It should be noted that u_{xx} fails to be continuous across the boundary although it remains bounded on any compact subset of $\mathbb{R}_+ \times [0, 1]$.

Recall Y_t^{y, ξ^g} is the process under control ξ^g , defined in (4.12). Fix the initial condition $(X_0, Y_0^{y, \xi^g}) = (x, y)$ and take $R > 0$. Set $\tau_R := \inf\{t \geq 0 : X_t \geq R\}$. Applying Ito's formula in the weak version to u (see [23, Chapter 8, Section VIII.4, Theorem 4.1] and [11, Theorem 4.2]), up to the stopping time $\tau_R \wedge T$, for some $T > 0$, to obtain

$$\mathbb{E} \left[e^{-\rho(\tau_R \wedge T)} u(X_{\tau_R \wedge T}^x, Y_{\tau_R \wedge T}^{y, \xi^g}) \right] = u(x, y) + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\mathcal{L}u(X_t^x, Y_t^{y, \xi^g}) - \rho V_g^\lambda(X_t^x, Y_t^{y, \xi^g}) \right) dt \right]$$

$$\begin{aligned}
& + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \partial_y u(X_t^x, Y_{t-}^{y, \xi^g}) dY_{t-}^{y, \xi^g} \right] \\
& + \mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} \left(\Delta u(X_t^x, Y_t^{y, \xi^g}) - \partial_y u(X_t, Y_{t-}^{y, \xi^g}) \Delta Y_t^{y, \xi^g} \right) \right],
\end{aligned}$$

with the notations $\Delta u(X_t^x, Y_t^{y, \xi^g}) = u(X_t^x, Y_t^{y, \xi^g}) - u(X_t^x, Y_{t-}^{y, \xi^g})$ and $\Delta Y_t^{y, \xi^g} = Y_t^{y, \xi^g} - Y_{t-}^{y, \xi^g}$.

Recall that for any admissible control ξ^g , it can be decomposed into the sum of its continuous part and its pure jump part, i.e., $d\xi^g = d(\xi^g)^{cont} + \Delta \xi^g$. Hence we have the decomposition $dY^{y, \xi^g} = d(Y^{y, \xi^g})^{cont} + \Delta Y^{y, \xi^g}$, and therefore,

$$\begin{aligned}
\mathbb{E} \left[e^{-\rho(\tau_R \wedge T)} u(X_s^x, Y_s^{y, \xi^g}) \right] &= u(x, y) + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\mathcal{L}_x u(X_t^x, Y_{t-}^{y, \xi^g}) - \rho V_g^\lambda(X_t^x, Y_{t-}^{y, \xi^g}) \right) dt \right] \\
&+ \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \partial_y u(X_t^x, Y_{t-}^{y, \xi^g}) d(Y^{y, \xi^g})_t^{cont} \right] \\
(4.27) \quad &+ \mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} \Delta u(X_t^x, Y_t^{y, \xi^g}) \right]
\end{aligned}$$

$$(4.28) \quad = u(x, y) - \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\left(\pi(X_t^x) - \rho \kappa \right) Y_t^{y, \xi^g} - \lambda Y_t^{y, \xi^g} \log(Y_t^{y, \xi^g}) \right) dt \right].$$

The last equation holds due to the following facts:

(i) It holds that

$$\begin{aligned}
& \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\mathcal{L}_x u(X_t^x, Y_{t-}^{y, \xi^g}) - \rho V_g^\lambda(X_t^x, Y_{t-}^{y, \xi^g}) \right) dt \right] \\
&= -\mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\left(\pi(X_t^x) - \rho \kappa \right) Y_t^{y, \xi^g} - \lambda Y_t^{y, \xi^g} \log(Y_t^{y, \xi^g}) \right) dt \right],
\end{aligned}$$

as pure jump could only possibly happen at time 0, ξ^g keeps (X_t^x, Y_t^{y, ξ^g}) within $\overline{\mathcal{E}(g)}$ for $t > 0$ and the fact (4.19);

- (ii) $\mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} \Delta u(X_t^x, Y_t^{y, \xi^g}) \right] = 0$ as $-u_y = 0$ on $\overline{\mathcal{S}(g)}$ according to the facts that (4.20) holds and $u \in C^1(\mathbb{R}_+ \times [0, 1])$;
- (iii) $\mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \partial_y u(X_t^x, Y_{t-}^{y, \xi^g}) d(Y^{y, \xi^g})_t^{cont} \right] = 0$ as the infinitesimal control happens on the boundary $\overline{\mathcal{S}(g)} \cap \overline{\mathcal{E}(g)}$ where $-u_y = 0$.

When taking limits as $R \rightarrow \infty$ we have $\tau_R \wedge T \rightarrow T$, \mathbb{P} -a.s. By standard properties of Brownian motion, it is easy to prove that the integral terms in the last expression on the right-hand side of (4.28) are uniformly bounded in $L^2(\Omega, \mathbb{P})$, hence uniformly integrable. Moreover, u in (4.21) has sublinear growth by straightforward calculation. Then we also take limits as $T \rightarrow \infty$ and it follows that

$$u(x, y) = \mathbb{E} \left[\int_0^\infty e^{-\rho t} \left(\left(\pi(X_t^x) - \rho \kappa \right) Y_t^{y, \xi^g} - \lambda Y_t^{y, \xi^g} \log(Y_t^{y, \xi^g}) \right) dt \right].$$

Therefore $u \equiv V_g^\lambda$ and hence the verification theorem holds. \square

With the results in Theorem 4.4, we update the new boundary following:

$$(4.29) \quad \tilde{g}(x) = \begin{cases} \max \left\{ y^* < g(x) \mid \partial_{xy} V_g^\lambda(x, y^*) = 0 \right\} & \text{if } \partial_{xy}^- V_g^\lambda(x, g(x)) < 0, \text{ and} \\ \tilde{g}(x) = g(x) & \text{o.w.,} \end{cases}$$

where we define the notation $\partial_{xy}^- f(x, y) := \lim_{h \rightarrow 0^-} \frac{\partial_x f(x, y+h) - \partial_x f(x, y)}{h}$ as the left y -derivative of $\partial_x f$ (if exists).

We can show that the updating rule (4.29) always improves in terms of the value function.

Theorem 4.5 (Policy improvement). *Assume Assumptions 4.1-4.3 hold and \tilde{g} is updated according to (4.29), it holds that*

$$(4.30) \quad V_{\tilde{g}}^{\lambda}(x, y) \geq V_g^{\lambda}(x, y),$$

where $V_g^{\lambda}, V_{\tilde{g}}^{\lambda}$ are value functions associated with policies $\xi^g, \xi^{\tilde{g}}$, see (4.12).

Proof. First of all, $V_g^{\lambda} \in C^2(\overline{\mathcal{E}(g)})$ according to Theorem 4.4. Therefore, by the design of the algorithm, \tilde{g} satisfies the conditions in Assumption 4.3. Hence Theorem 4.4 also applies to \tilde{g} .

Denote $Y_t^{y, \xi^{\tilde{g}}}$ as the process under control $\xi^{\tilde{g}}$, with \tilde{g} defined in (4.29). Fix the initial condition (x, y) and take $R > 0$. Set $\tau_R := \inf\{t \geq 0 : X_t^x \geq R\}$. Similar to the proof of Theorem 4.21, apply Ito's formula in the weak version to V_g^{λ} (see [23, Chapter 8, Section VIII.4, Theorem 4.1] and [11, Theorem 4.2]) up to the stopping time $\tau_R \wedge T$ for some $T > 0$. We then obtain

$$\begin{aligned} & \mathbb{E} \left[e^{-\rho(\tau_R \wedge T)} V_g^{\lambda}(X_{\tau_R \wedge T}^x, Y_{\tau_R \wedge T}^{y, \xi^{\tilde{g}}}) \right] \\ &= V_g^{\lambda}(x, y) + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\mathcal{L}_x V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) - \rho V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) \right) dt \right] \\ & \quad + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \partial_y V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) dY_t^{\tilde{g}} \right] \\ & \quad + \mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} e^{-\rho t} \left(\Delta V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) - \partial_y V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) \Delta Y_t^{y, \xi^{\tilde{g}}} \right) \right] \\ &= V_g^{\lambda}(x, y) - \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\left(\pi(X_t^x) - \rho \kappa \right) Y_t^{y, \xi^{\tilde{g}}} - \lambda Y_t^{y, \xi^{\tilde{g}}} \log(Y_t^{y, \xi^{\tilde{g}}}) \right) dt \right] \\ & \quad + \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \partial_y V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) d(Y^{y, \xi^{\tilde{g}}})_t^{\text{cont}} \right] \\ & \quad + \mathbb{E} \left[\sum_{0 \leq t \leq \tau_R \wedge T} e^{-\rho t} \left(\Delta V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) \right) \right] \\ &\geq V_g^{\lambda}(x, y) - \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\left(\pi(X_t^x) - \rho \kappa \right) Y_t^{y, \xi^{\tilde{g}}} - \lambda Y_t^{y, \xi^{\tilde{g}}} \log(Y_t^{y, \xi^{\tilde{g}}}) \right) dt \right], \end{aligned}$$

with the notations $\Delta V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) = V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) - V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^g})$ and $\Delta Y_t^{y, \xi^{\tilde{g}}} = Y_t^{y, \xi^{\tilde{g}}} - Y_{t-}^{y, \xi^{\tilde{g}}}$. The second equation holds because

$$\begin{aligned} & \mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\mathcal{L}_x u(X_t^x, Y_{t-}^{y, \xi^g}) - \rho V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^g}) \right) dt \right] \\ &= -\mathbb{E} \left[\int_0^{\tau_R \wedge T} e^{-\rho t} \left(\left(\pi(X_t^x) - \rho \kappa \right) Y_t^{y, \xi^g} - \lambda Y_t^{y, \xi^g} \log(Y_t^{y, \xi^g}) \right) dt \right] \end{aligned}$$

as pure jump could only possibly happen at time 0, $\xi^{\tilde{g}}$ keeps $(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}})$ within $\overline{\mathcal{E}(g)}$ for $t > 0$ and the fact (4.19). The last inequality holds by the design of (4.29). In particular, for $(x, y) \in \mathcal{S}(\tilde{g}) \cap \mathcal{E}(g)$,

$$\partial_y V_g^{\lambda}(x, y) = \partial_y V_g^{\lambda}(g^{-1}(y), y) + \int_{g^{-1}(y)}^x \partial_{xy} V_g^{\lambda}(u, y) du \leq 0,$$

and

$$\Delta V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) = V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^{\tilde{g}}}) - V_g^{\lambda}(X_t^x, Y_{t-}^{y, \xi^g})$$

$$= \int_0^{\Delta Y_t^{y, \xi^{\bar{g}}}} \partial_y V(X_t^x, Y_{t-}^{y, \xi^{\bar{g}}} + u) du \geq 0,$$

as $\Delta Y_t^{y, \xi^{\bar{g}}} \leq 0$.

When taking limits as $R \rightarrow \infty$ we have $\tau_R \wedge T \rightarrow T$, \mathbb{P} -a.s. By standard properties of Brownian motion it is easy to prove that the integral terms in the last expression on the right-hand side of (4.28) are uniformly bounded in $\mathbb{L}^2(\Omega, \mathbb{P})$, hence uniformly integrable. Moreover, V_g^λ (taking the form as in (4.21)) has sublinear growth by straightforward calculation. Then we also take limits as $T \rightarrow \infty$ and it follows that

$$V_g^\lambda(x, y) \geq V_g^\lambda(x, y).$$

□

Repeat the updating procedure (4.29) iteratively, we have the general heuristic described in Algorithm 1. The idea is that we start with an initial $g_0(x)$ that has a sufficiently large exploration region $\mathcal{E}(g_0)$ (i.e., Assumption 4.6). Then in each iteration k , we update g_k function “downwards” into $\mathcal{E}(g_k)$ according to the Hessian information (cf. line 3 of Algorithm 1), in which u_k has a better regularity (cf. line 3 of Algorithm 1). See Figure 4.2.1 for a demonstration.

Algorithm 1 Policy Iteration for Optimal Stopping (PI-OS)

- 1: Initialize $g_0(x)$ for $x \in [0, \infty)$ according to Assumption 4.6.
- 2: **for** $k = 0, 1, \dots, K - 1$ **do**
- 3: Find $u_k(x, y)$ a $C^1(\mathbb{R}_+ \times [0, 1]) \cap C^2(\overline{\mathcal{E}(g_k)})$ solution to the following equations:

$$(4.31) \quad (\mathcal{L}_x - \rho)u + \left(\pi(x) - \rho\kappa\right)y - \lambda y \log y = 0 \quad \text{on } \mathcal{E}(g_k),$$

$$(4.32) \quad -u_y = 0, \quad \text{on } \mathcal{S}(g_k).$$

- 4: Update the strategy

$$g_{k+1}(x) = \begin{cases} \max \left\{ y^* < g_k(x) \mid \partial_{xy} u_k(x, y^*) = 0 \right\} & \text{if } \partial_{xy}^- u_k(x, g_k(x)) < 0, \\ g_{k+1}(x) = g_k(x) & \text{o.w.} \end{cases}$$

- 5: **end for**
-

This iterative scheme is inspired by [38] for a one-dimensional singular control problem. However, the convergence result in that work requires a second-order condition to hold throughout the entire iteration process, which is difficult to verify.

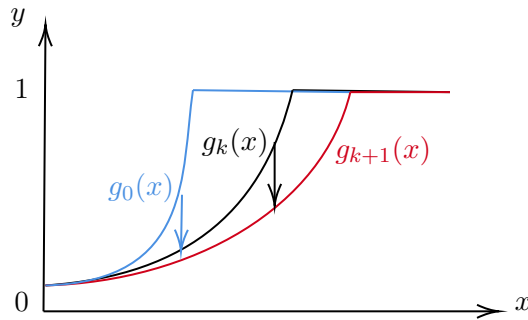


FIGURE 1. Demonstration of the Policy Iteration Algorithm.

In addition to the policy improvement result in Proposition 4.5, we provide the following policy convergence result under additional conditions on the initialization.

Assumption 4.6 (Initial policy). *Assume the initial policy g_0 satisfies the following conditions:*

- (a) $g_0 \in C^1([0, \hat{x}_{g_0}])$ is strictly increasing on $[0, \hat{x}_{g_0}]$,
- (b) $g_0(0) = e^{-(1+\frac{\kappa\rho}{\lambda})}$, and
- (c) $-\alpha_- \left(\kappa + \frac{\lambda}{\rho} \log(g_0(x)) + \frac{\lambda}{\rho} \right) + \alpha_- H_\pi(x) - H'_\pi(x) \cdot x \geq 0$ on $[0, \hat{x}_{g_0}]$.

Note that (c) implies that $\hat{x}_{g_0} \leq \hat{x}_{g_\lambda}$.

Remark 4.7 (Justification of Assumption 4.6). *Assumption 4.6-(a) is easy to satisfy and Assumption 4.6-(c) requires $\mathcal{E}(g_0)$ to be sufficiently large so that $\mathcal{E}(g_0) \supseteq \mathcal{E}(g_\lambda)$. Assumption 4.6-(b) seems to be the most restrictive one, as it requires the knowledge of $g_\lambda(0) = \exp(-(1+\frac{\kappa\rho}{\lambda}))$. In the model-free setting (such as in Section 4.2.2) and when $g_\lambda(0)$ is not available, the following algorithm can be applied to learn $g_\lambda(0)$ at a fast convergence rate.*

Algorithm 2 Learning initial value $g_0(0)$

- 1: Initialize $y_0 \in (0, 1)$, $c_0 > 0$ and $\{\eta_i\}_{i \geq 1}$.
- 2: **for** $i = 1, \dots$, **do**
- 3: Acquire the value functions $u_i^+ = J^\lambda(0, y_i + \varepsilon_i; \xi = 0)$ and $u_i^- = J^\lambda(0, y_i - \varepsilon_i; \xi = 0)$, with $\varepsilon_i = \min\{y_i, 1 - y_i, c_0/i\}$.
- 4: Update y_i using the following zeroth order gradient descent with a two-point estimator:

$$(4.33) \quad y_{i+1} = y_i - \eta_i \frac{u_i^+ - u_i^-}{2\varepsilon_i}.$$

5: **end for**

We know that $J^\lambda(0, y; \xi = 0) = -\kappa y - \frac{\lambda y}{\rho} \log y$ is convex in y with the minimizer taking value $y^* := e^{-(1+\frac{\kappa\rho}{\lambda})}$. When the functional form (hence the minimizer) is unknown to the decision-maker, Algorithm 2 converges linearly to the minimizer [18] when taking $\eta_i = \frac{\eta}{\sqrt{i}}$ with some $\eta > 0$. Mathematically,

$$(4.34) \quad |y_i - y^*|^2 \sim \mathcal{O}(\delta^i),$$

for some $\delta \in (0, 1)$ depending on model parameters.

Remark 4.8 (Examples that satisfy Assumption 4.6). *Assumption 4.6 is easy to satisfy and we provide two examples under the special case $\pi(x) = x^\theta$ with $0 < \theta < 1$: one is linear initialization and the other one is exponential initialization. For linear initialization, we can take*

$$(4.35) \quad g_0(x) = \min \left\{ e^{-(1+\frac{\kappa\rho}{\lambda})} + 2(1 - e^{-(1+\frac{\kappa\rho}{\lambda})}) \left(\frac{-\alpha_- (\kappa + \frac{\lambda}{\rho})}{C(\theta - \alpha_-)} \right)^\theta x, 1 \right\},$$

which interpolates (x_0, y_0) and (x_1, y_1) with $x_0 = 0$, $y_0 = e^{-(1+\frac{\kappa\rho}{\lambda})}$, $x_1 = \frac{1}{2} \left(\frac{-\alpha_- (\kappa + \frac{\lambda}{\rho})}{C(\theta - \alpha_-)} \right)^{-\theta}$ and $y_1 = 1$.

For exponential initialization, we can take

$$(4.36) \quad g_0(x) = \min \left\{ \exp \left(\frac{\rho C(\zeta - \alpha_-)(x)^\zeta}{-\alpha_- \lambda} - \frac{\rho}{\lambda} \left(\kappa + \frac{\lambda}{\rho} \right) \right), 1 \right\},$$

for some $\zeta \in (\theta, 1)$. This is inspired by the form of (4.5) that under $\pi(x) = x^\theta$

$$g_\lambda(x) = \min \left\{ \exp \left(\frac{\rho C(\theta - \alpha_-)(x)^\theta}{-\alpha_- \lambda} - \frac{\rho}{\lambda} \left(\kappa + \frac{\lambda}{\rho} \right) \right), 1 \right\}.$$

Theorem 4.9 (Policy convergence). *Assume Assumptions 4.1 and 4.6 hold. Then we have the following results*

$$(4.37) \quad \lim_{k \rightarrow \infty} g_k = g_\lambda,$$

$$(4.38) \quad \lim_{k \rightarrow \infty} V_{g_k}^\lambda = V_\lambda.$$

with g_λ defined in (4.5). In particular, we have g_k is strictly increasing on $[0, \hat{x}_{g_k}]$ and $g_k \in C^1([0, \hat{x}_{g_k}])$ for all $k \in \mathbb{N}_+$.

Note that the monotonicity and smoothness of the boundary g_k ensures the existence and uniqueness of the solution to (4.31)-(4.32), which coincides with the value function under policy ξ_{g_k} (see Theorem 4.4).

Remark 4.10. *It is worth noting that Theorem 4.9 presents the first policy improvement result for RL algorithms related to free boundary problems. Existing literature primarily focuses on the development of continuous-time regular controls. Compared to policy improvement, showing policy convergence remains significantly more challenging, even for regular controls. To the best of our knowledge, the only results on policy convergence for regular controls are found in [3, 30, 40, 60], and all of these results are technically intricate.*

Proof. Step 1. By the construction of the sequence in Algorithm 1, we have $0 \leq g_{k+1} \leq g_k \leq g_0$ for all $k \in \mathbb{N}_+$. Hence, by monotone convergence theorem, there exists a limit $\bar{g} \geq 0$ such that

$$(4.39) \quad \bar{g}(x) := \lim_{k \rightarrow \infty} g_k(x).$$

The main goal is to show that $\bar{g}(x) = g_\lambda(x)$ for $x \in \mathbb{R}_+$.

Step 2. Recall $V_{g_k}^\lambda(x, y)$ as the value function associated with strategy ξ^{g_k} . We use induction to prove the following iteratively:

- (a) g_k is strictly increasing on $[0, \hat{x}_{g_k}]$,
- (b) $g_k(0) = e^{-(1 + \frac{\kappa \rho}{\lambda})}$ for $k \in \mathbb{N}_+$,
- (c) $-\alpha_- \left(\kappa + \frac{\lambda}{\rho} \log(g_k(x)) + \frac{\lambda}{\rho} \right) + \alpha_- H_\pi(x) - H'_\pi(x) \cdot x \geq 0$ for $x \in [0, \hat{x}_{g_k}]$, namely, $g_k \geq g_\lambda$ on $[0, \hat{x}_{g_k}]$,
- (d) $\partial_{xy}^- u_k(x, g_k(x)) \leq 0$ for $x \in [0, \hat{x}_{g_k}]$.
- (e) $g_k \in C^1([0, \hat{x}_{g_k}])$,

The conditions (a)-(c), (e) hold for $k = 0$ by the design of the initial reflection boundary g_0 (see Assumption 4.6). Therefore, Theorems 4.4 and 4.5 apply to g_0 . Hence we have $u_0(x, y) = A_0(y)x^{\alpha_-} + H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y$, with $A_0(y) = \int_{g_0(0)}^y \frac{\kappa + \frac{\lambda}{\rho} \log(u) + \frac{\lambda}{\rho} - H_\pi(g_0^{-1}(u))}{(g_0^{-1}(u))^{\alpha_-}} du$. Combined with condition (c), we have for $x \in [0, \hat{x}_{g_0}]$,

$$\partial_{yx} u_0(x, g_0(x)) = \frac{\alpha_-}{x} \left(\kappa + \frac{\lambda}{\rho} \log(g_0(x)) + \frac{\lambda}{\rho} - H_\pi(x) \right) + H'_\pi(x) \leq 0.$$

Therefore condition (d) holds.

Assume the above property holds for k . Then by Theorem 4.4, the following HJB equation

$$(4.40) \quad (\mathcal{L}_x - \rho)u + \left(\pi(x) - \rho\kappa \right)y - \lambda y \log y = 0 \quad \text{on} \quad \mathcal{E}(g_k),$$

$$(4.41) \quad -u_y = 0, \quad \text{on} \quad \mathcal{S}(g_k).$$

has a unique $C^1(\mathbb{R}_+ \times [0, 1]) \cap C^2(\overline{\mathcal{E}(g_k)})$ solution, which we denoted as u_k . Also, $V_{g_k}^\lambda(x, y) \equiv u_k(x, y)$. Then for $y \leq g_k(x)$ the value function u_k satisfies the form:

$$(4.42) \quad u_k(x, y) = A_k(y)x^{\alpha_-} + H_\pi(x)y - \kappa y - \frac{\lambda}{\rho}y \log y,$$

with

$$(4.43) \quad A_k(y) = \int_{g_k(0)}^y \frac{\kappa + \frac{\lambda}{\rho} \log(u) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(u))}{(g_k^{-1}(u))^{\alpha_-}} du.$$

In this region, we also have,

$$\begin{aligned} \partial_y u_k(x, y) &= A'_k(y)x^{\alpha_-} + H_\pi(x) - k - \frac{\lambda}{\rho} - \frac{\lambda}{\rho} \log(y) \\ &= \frac{x^{\alpha_-}}{(g_k^{-1}(y))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log y + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(y)) \right) + H_\pi(x) - k - \frac{\lambda}{\rho} - \frac{\lambda}{\rho} \log(y) \\ &= \frac{x^{\alpha_-} - (g_k^{-1}(y))^{\alpha_-}}{(g_k^{-1}(y))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log y \right. \\ &\quad \left. + \frac{\lambda}{\rho} - H_\pi(x) \right) + \frac{x^{\alpha_-}}{(g_k^{-1}(y))^{\alpha_-}} \left(H_\pi(x) - H_\pi(g_k^{-1}(y)) \right), \end{aligned}$$

and

$$\begin{aligned} \partial_{xy} u_k(x, y) &= \alpha_- A'_k(y)x^{\alpha_- - 1} + H'_\pi(x) \\ &= \frac{\alpha_- x^{\alpha_- - 1}}{(g_k^{-1}(y))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log y + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(y)) \right) + H'_\pi(x). \end{aligned}$$

In the k th iteration there are two scenarios for each point $(x_0, g_k(x_0))$: $\partial_{xy}^- u_k(x_0, g_k(x_0)) = 0$ or $\partial_{xy}^- u_k(x_0, g_k(x_0)) < 0$. When $\partial_{xy}^- u_k(x_0, g_k(x_0)) = 0$, we have $g_{k+1}(x_0) = g_k(x_0)$. Therefore, for both scenarios, $\partial_{xy}^- u_k(x, y) = 0$ leads to $y = g_{k+1}(x)$, or equivalently, we solve $y = g_{k+1}(x) \leq g_k(x)$ from $\partial_{xy}^- u_k(x, g_{k+1}(x)) = 0$. Then we have

$$\begin{aligned} 0 &= -\frac{\alpha_- x^{\alpha_-}}{(g_k^{-1}(g_{k+1}(x)))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log(g_{k+1}(x)) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(g_{k+1}(x))) \right) - H'_\pi(x) x \\ (4.44) \quad &\leq -\alpha_- \left(\kappa + \frac{\lambda}{\rho} \log(g_{k+1}(x)) + \frac{\lambda}{\rho} - H_\pi(x) \right) - H'_\pi(x) x. \end{aligned}$$

The last inequality holds since for $y \geq g_\lambda(x)$,

$$(4.45) \quad \frac{k + \frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(g_{k+1}(x)))}{g_k^{-1}(g_{k+1}(x))^{\alpha_-}} \leq \frac{k + \frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} - H_\pi(x)}{x^{\alpha_-}}.$$

Equation (4.45) holds since $g_k^{-1}(g_{k+1}(x)) \leq x$ and

$$\begin{aligned} &\partial_x \left(\frac{k + \frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} - H_\pi(x)}{x^{\alpha_-}} \right) \\ (4.46) \quad &= \frac{-H'_\pi(x)x^{\alpha_-} - \alpha_- \left(k + \frac{\lambda}{\rho} \log(y) + \frac{\lambda}{\rho} - H_\pi(x) \right) x^{\alpha_- - 1}}{x^{2\alpha_-}} \geq 0. \end{aligned}$$

And finally (4.46) holds for $y \geq g_\lambda(x)$.

Hence (4.44) suggests that (c) is satisfied in the $k + 1$ iteration.

Next we show that $g_{k+1}(0) = e^{-(1+\frac{\kappa\rho}{\lambda})}$. For any points x_0 such that **(c)** holds with equality sign in the k th iteration, namely,

$$g_k(x_0) = \exp \left(\frac{\frac{-\alpha_- H_\pi(x_0) + H'_\pi(x_0)x_0}{-\alpha_-} - (\kappa + \frac{\lambda}{\rho})\rho}{\lambda} \right),$$

it holds that $\partial_{xy}u_k^-(x_0, g_k(x_0)) = 0$ by simple calculation. Hence it holds that $g_{k+1}(x_0) = g_k(x_0)$. Take $x_0 = 0$, we have $g_{k+1}(0) = g_k(0) = e^{-(1+\frac{\kappa\rho}{\lambda})}$.

We now show condition **(a)** by contradiction. Suppose there exists $0 < x_1 < x_2$ such that $\bar{y} = g_{k+1}(x_1) = g_{k+1}(x_2)$. Then we have

$$\begin{aligned} 0 &= x_1 \partial_{xy}u_k(x_1, \bar{y}) - x_2 \partial_{xy}u_k(x_2, \bar{y}) \\ &= \frac{\alpha_- x_1^{\alpha_-}}{(g_k^{-1}(\bar{y}))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log(\bar{y}) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(\bar{y})) \right) + H'_\pi(x_1)x_1 \\ &\quad - \frac{\alpha_- x_2^{\alpha_-}}{(g_k^{-1}(\bar{y}))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log(\bar{y}) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(\bar{y})) \right) - H'_\pi(x_2)x_2 \\ (4.47) \quad &= \frac{\alpha_- (x_1^{\alpha_-} - x_2^{\alpha_-})}{(g_k^{-1}(\bar{y}))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log(\bar{y}) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(\bar{y})) \right) + H'_\pi(x_1)x_1 - H'_\pi(x_2)x_2. \end{aligned}$$

Note that $x_1^{\alpha_-} > x_2^{\alpha_-}$ due to the facts that $x_1 < x_2$ and $\alpha_- < 0$. In addition, we have $x_1 H'_\pi(x_1) < x_2 H'_\pi(x_2)$ as $x\pi'(x)$ is strictly increasing according to Assumption 4.1. Hence we get a contradiction to (4.47). This suggests that g_{k+1} is either strictly increasing or decreasing. Given that $0 < g_\lambda \leq g_{k+1}$ on $[0, \hat{x}_{g_{k+1}}]$, then g_{k+1} must be strictly increasing.

(d) holds since, according to **(c)**

$$\partial_{yx}u_{k+1}(x, g_{k+1}(x)) = \frac{\alpha_-}{x} \left(\kappa + \frac{\lambda}{\rho} \log(g_{k+1}(x)) + \frac{\lambda}{\rho} - H_\pi(x) \right) + H'_\pi(x) \leq 0.$$

Finally, given that $u_k \in C^2(\overline{\mathcal{E}(g_k)})$, $\partial_{xy}u_k(x, g_{k+1}(x)) = 0$ for $x \in [0, \hat{x}_{g_{k+1}}]$ and $(x, g_{k+1}(x)) \in \overline{\mathcal{E}(g_k)}$, we have $g_{k+1} \in C^1([0, \hat{x}_{g_{k+1}}])$. Hence we have **(e)** holds.

Step 3. We study the property of the limiting function $\bar{g} \geq g_\lambda$ as $g_k \geq g_\lambda$ for all $k \in \mathbb{N}_+$.

According to Step 2, **(a)**-**(e)** holds iteratively. Therefore it holds for all $k \in \mathbb{N}_+$ that

$$(4.48) \quad \frac{-\alpha_- x^{\alpha_-}}{(g_k^{-1}(g_{k+1}(x)))^{\alpha_-}} \left(\kappa + \frac{\lambda}{\rho} \log(g_{k+1}(x)) + \frac{\lambda}{\rho} - H_\pi(g_k^{-1}(g_{k+1}(x))) \right) - H'_\pi(x)x = 0.$$

For each $x \in \mathbb{R}_+$ fixed, we have $\lim_{k \rightarrow \infty} g_k^{-1}(g_{k+1}(x)) = x$ hold. Combined with (4.48), we have

$$0 = -\alpha_- \left(\kappa + \frac{\lambda}{\rho} \log(\bar{g}(x)) + \frac{\lambda}{\rho} - H_\pi(x) \right) - H'_\pi(x)x.$$

Therefore, $\bar{g} = g_\lambda$ for $x \in [0, x_{g_\lambda}]$. Hence (4.37) holds.

In addition, according to Step 2, for $(x, y) \in \overline{\mathcal{E}(g_k)}$, we have u_k follows (4.42) with A_k defined in (4.42). As g_k converges to g_λ uniformly, we have A_k converges to A , which is defined in (4.25). Hence u_k converges to u^λ in $\mathcal{E}(g^\lambda)$. Hence $u_k = V_{g_k}^\lambda$ converges to $u^\lambda = V^\lambda$ on $\mathbb{R}_+ \times [0, 1]$. Finally, (4.38) holds. \square

4.2.2. Model-free implementation. When the model parameters are unknown, we are not able to calculate the solution u_k of (4.31)-(4.32). In this case, we need to approximate the value function by acquiring instantaneous rewards along multiple trajectories and take the average

(see lines 7-9 in Algorithm 3). This is referred to as the Policy Evaluation in the RL literature [53, 58].

Algorithm 3 Policy Iteration and Policy Evaluation for Optimal Stopping (PIPE-OS)

- 1: Initialize $g_0(x)$ according to Assumption 4.6. Specify a grid size δ_x for partitioning the x -axis and a grid size δ_y for partitioning the y -axis. Also, specify an upper bound $\bar{x} := N\delta$.
 - 2: **for** $k = 0, 1, \dots, K - 1$ **do**
 - 3: **for** $x \in \{0, \delta_x, 2\delta_x, \dots, N\delta_x\}$ and $y \in \{0, \delta_y, 2\delta_y, \dots, \lfloor 1/\delta_y \rfloor\}$ **do**
 - 4: **for** $m = 1, \dots, M$ **do**
 - 5: Simulate the m -th path $(X^{(m),x}, Y^{(m),y,\xi^{g_k}})$ under policy ξ^{g_k} (defined in (4.12))
 - 6: ▷ independent randomness across paths
 - 7: Calculate the instantaneous value function $u_k^{(m)}(x, y)$:

$$(4.49) u_k^{(m)}(x, y) = \int_0^\infty e^{-\rho t} \left(\left(\pi(X_t^{(m),x}) - \rho\kappa \right) Y_t^{(m),y,\xi^{g_k}} - \lambda Y_t^{(m),y,\xi^{g_k}} \log(Y_t^{(m),y,\xi^{g_k}}) \right) dt$$
 - 8: **end for**
 - 9: Calculate the approximated value function $\bar{u}_k(x, y) = \frac{\sum_{m=1}^M u_k^{(m)}(x, y)}{M}$
 - 10: Update the strategy

$$g_{k+1}(x) = \begin{cases} \max \left\{ y^* < g_k(x) \mid \partial_{xy} \bar{u}_k(x, y^*) = 0 \right\} & \text{if } \partial_{xy}^- \bar{u}_k(x, g_k(x)) < 0, \\ g_{k+1}(x) = g_k(x) & \text{o.w.} \end{cases}$$
 - 11: **end for**
 - 12: **end for**
-

Notably, we do not directly estimate model parameters in Algorithm 3, which enhances robustness against model misspecification and environmental shifts [1, 19, 28].

4.2.3. *Numerical performance.* We test the performance of Algorithm 3 on a few examples.

In the experiment, we set $\mu = 0.2$, $\sigma = 0.2$, $\rho = 0.5$, $k = 5$, and $\beta = 0.5$ (assumed to be unknown to the learner).

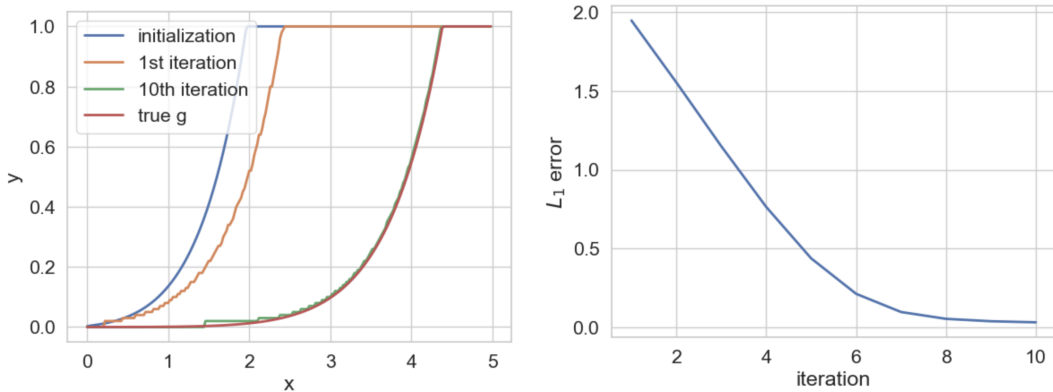


FIGURE 2. Exponential initialization. **Left:** Ground-truth and learned g function in selected iterations . **Right:** Convergence to the ground truth in L_1 norm

We evaluate the performance of the algorithm using two different initial policies: one with exponential initialization and the other with linear initialization (see Remark 4.8). As shown

in Figure 2, Algorithm 3 converges within 10 outer iterations when using exponential initialization, highlighting its effectiveness. In addition, it takes approximately 20 outer iterations for the algorithm to converge with linear initialization (see Figure 3). This suggests that learning the boundary is more challenging when x is smaller.

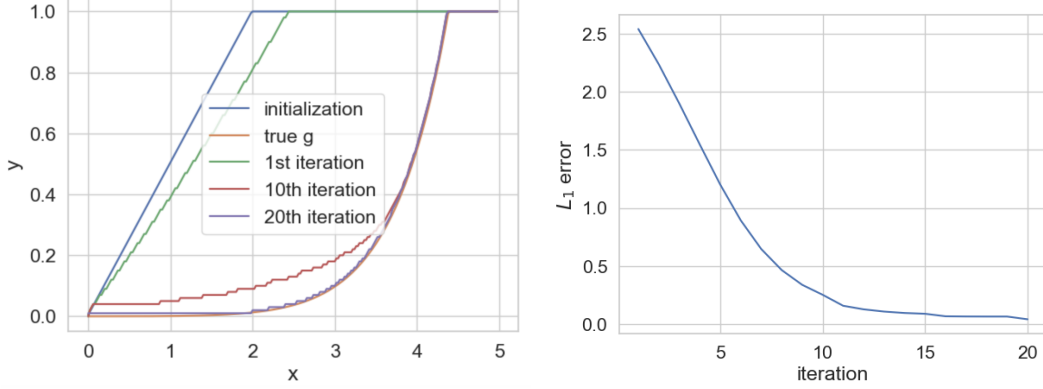


FIGURE 3. Linear initialization. **Left:** Ground-truth and learned g function in selected iterations . **Right:** Convergence to the ground truth in L_1 norm .

APPENDIX A. PROOF OF LEMMA 2.3

Setting $\hat{V} := V - G$, by Itô's formula, problem (2.1) reduces to the OS problem $\hat{V}(x) = \sup_{\tau \in \mathcal{T}} \mathbb{E} \left[\int_0^\tau e^{-\rho t} \hat{\pi}(X_t^x) dt \right]$, for which the value function \hat{V} is the unique $W_{loc}^{2,2}(\mathbb{R})$ -a.e. solution to the HJB variational inequality (see. e.g., Section 5.2 in [47])

$$(A.1) \quad \max \{ (\mathcal{L}_x - \rho) \hat{V} + \hat{\pi}, -\hat{V} \} = 0.$$

By Theorem 2.3.5 in [39], the minimal stopping time τ^* is the first hitting time to the region $\hat{\mathcal{S}} := \{\hat{V} = 0\}$. Also, since $\hat{\pi}$ is increasing, \hat{V} is nondecreasing so that $\hat{\mathcal{S}}$ is a closed interval (possibly empty). If $\hat{\mathcal{S}}$ is empty, then $\tau^* = \infty$ is the unique optimal stopping time, by minimality of τ^* . On the contrary, if $\hat{\mathcal{S}}$ is nonempty, then we can find $x^* \in \mathbb{R}$ such that $\hat{\mathcal{S}} := \{\hat{V} = 0\} = (-\infty, x^*]$. Notice moreover that necessarily we have $x^* \leq \bar{x}$, as otherwise we would have $(\mathcal{L}_x - \rho) \hat{V} + \hat{\pi} = \hat{\pi} > 0$ in the interval (\bar{x}, x^*) , thus contradicting the HJB equation. In other words,

$$(A.2) \quad (\mathcal{L}_x - \rho) \hat{V} + \hat{\pi} = \hat{\pi} < 0 \text{ in } (-\infty, x^*] \text{ and } (\mathcal{L}_x - \rho) \hat{V} + \hat{\pi} < 0, \text{ a.e. in } \mathbb{R}.$$

Next, assume by contradiction that there exists another optimal stopping time τ . By minimality of τ^* we have $\tau^* \leq \tau$. Moreover, by the martingale optimality principle (see e.g. Theorem 2.2 at p. 29 in [46]), the process $M = (M_t)_t$ defined by

$$M_t := e^{-\rho t \wedge \tau} \hat{V}(X_{t \wedge \tau}^x) + \int_0^{t \wedge \tau} e^{-\rho s} \hat{\pi}(X_s^x) ds,$$

is an \mathbb{F}^W -martingale and, using Itô's formula (thanks to the regularity of \hat{V}) we have

$$\mathbb{E}[M_t] := \hat{V}(x) + \mathbb{E} \left[\int_0^{t \wedge \tau} e^{-\rho s} ((\mathcal{L}_x - \rho) \hat{V} + \hat{\pi})(X_s^x) ds \right].$$

Notice that, since $\sigma^2(x) \geq c > 0$, after the time τ^* the Brownian motion pushes the state X_t^x in the region $(-\infty, x^*)$ with probability 1. Thus, by (A.2) and the characterization of τ^* we

have

$$\begin{aligned}\mathbb{E}[M_\tau] &= \hat{V}(x) + \mathbb{E}\left[\int_0^\tau e^{-\rho s}((\mathcal{L}_x - \rho)\hat{V} + \hat{\pi})(X_s^x)ds\right] \\ &= \hat{V}(x) + \mathbb{E}\left[\mathbb{1}_{\{\tau > \tau^*\}} \int_{\tau^*}^\tau e^{-\rho s}((\mathcal{L}_x - \rho)\hat{V} + \hat{\pi})(X_s^x)ds\right] < \hat{V}(x) = \mathbb{E}[M_0],\end{aligned}$$

where the last inequality follows from assuming that the event $\tau > \tau^*$ has positive probability. This contradicts the fact that M is a martingale, and thus show that the optimal stopping time τ^* is unique.

APPENDIX B. PROOF OF THE AUXILIARY ESTIMATES

We prove each of the estimate (3.5)–(3.10) separately.

Proof of (3.5). Noticing that $|x|(1 + |x|) \leq \frac{3}{2}(1 + |x|^2)$, we have

$$xb(x) + \frac{q-1}{2}|\sigma(x)|^2 \leq \left(\frac{3}{2}L + (q-1)L^2\right)(1 + |x|^2).$$

Hence, the estimate (3.5) follows from Theorem 4.1 at p. 60 in [41].

Proof of (3.6). We adopt the rationals in the proof of Theorem 4.1 at p. 60 in [41]. Set $\Delta_t := X_t^{\bar{x}} - X_t^x$. For any $\varepsilon > 0$, using Itô's formula together with elementary estimates, we find

$$\begin{aligned} & \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q}{2}} \\ &= \left(\varepsilon + |\bar{x} - x|^2\right)^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-2}{2}} \Delta_t^* (\sigma(X_t^{\bar{x}}) - \sigma(X_t^x)) dW_t \\ & \quad + q \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-2}{2}} \Delta_t^* (b(X_t^{\bar{x}}) - b(X_t^x)) dt \\ & \quad + q \frac{q-2}{2} \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-4}{2}} |\Delta_t^* (\sigma(X_t^{\bar{x}}) - \sigma(X_t^x))|^2 dt \\ & \quad + \frac{q}{2} \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-2}{2}} |\sigma(X_t^{\bar{x}}) - \sigma(X_t^x)|^2 dt \\ (B.1) \quad & \leq \left(\varepsilon + |\bar{x} - x|^2\right)^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-2}{2}} \Delta_t^* (\sigma(X_t^{\bar{x}}) - \sigma(X_t^x)) dW_t \\ & \quad + q \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-2}{2}} \left[\Delta_t^* (b(X_t^{\bar{x}}) - b(X_t^x)) + \frac{q-1}{2} |\sigma(X_t^{\bar{x}}) - \sigma(X_t^x)|^2 \right] dt \\ & \leq \left(\varepsilon + |\bar{x} - x|^2\right)^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q-2}{2}} \Delta_t^* (\sigma(X_t^{\bar{x}}) - \sigma(X_t^x)) dW_t \\ & \quad + q \left(L + \frac{q-1}{2} L^2 \right) \int_0^t \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q}{2}} dt. \end{aligned}$$

Taking expectations and using Grönwall inequality, we obtain

$$\mathbb{E} \left[\left(\varepsilon + |\Delta_t|^2\right)^{\frac{q}{2}} \right] \leq \left(\varepsilon + |\bar{x} - x|^2\right)^{\frac{q}{2}} e^{q(L + \frac{q-1}{2} L^2)t},$$

which completes the proof sending $\varepsilon \rightarrow 0$.

Proof of (3.7). Set

$$\begin{aligned} \hat{X}_t &:= \delta X_t^{\bar{x}} + (1 - \delta) X_t^x, \quad X_t^\delta := X_t^{\delta \bar{x} + (1 - \delta)x}, \quad \Gamma_t := \hat{X}_t - X_t^\delta, \\ \hat{b}_t &:= \delta b(X_t^{\bar{x}}) + (1 - \delta) b(X_t^x), \quad \hat{\sigma}_t := \delta \sigma(X_t^{\bar{x}}) + (1 - \delta) \sigma(X_t^x). \end{aligned}$$

By Itô's formula, with elementary estimates we obtain

$$\begin{aligned}
\left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} &= \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left[\left(\hat{b}_t - b\left(X_t^\delta\right)\right) dt + \left(\hat{\sigma}_t - \sigma\left(X_t^\delta\right)\right) dW_t \right] \\
&\quad + q \frac{q-2}{2} \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-4}{2}} \left| \Gamma_t^* \left(\hat{\sigma}_t - \sigma\left(X_t^\delta\right)\right) \right|^2 dt \\
&\quad + \frac{q}{2} \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left| \hat{\sigma}_t - \sigma\left(X_t^\delta\right) \right|^2 dt \\
&\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left[\Gamma_t^* \left(\hat{b}_t - b\left(X_t^\delta\right)\right) + \frac{q-1}{2} \left| \hat{\sigma}_t - \sigma\left(X_t^\delta\right) \right|^2 \right] dt \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma\left(X_t^\delta\right)\right) dW_t.
\end{aligned}$$

Using the properties of b and σ in Condition 1 in Assumption 3.1, we continue with

$$\begin{aligned}
&\left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} \\
&\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma\left(X_t^\delta\right)\right) dW_t \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left[\Gamma_t^* \left(b\left(\hat{X}_t\right) - b\left(X_t^\delta\right)\right) + (q-1) \left| \sigma\left(\hat{X}_t\right) - \sigma\left(X_t^\delta\right) \right|^2 \right] dt \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left[\Gamma_t^* \left(\hat{b}_t - b\left(\hat{X}_t\right)\right) + (q-1) \left| \hat{\sigma}_t - \sigma\left(\hat{X}_t\right) \right|^2 \right] dt \\
&\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma\left(X_t^\delta\right)\right) dW_t \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} (L + (q-1)L^2) |\Gamma_t|^2 dt \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left(|\Gamma_t^*| \delta(1-\delta)L \left| X_t^{\bar{x}} - X_t^x \right|^2 \right. \\
&\quad \quad \left. + (q-1)\delta^2(1-\delta)^2 L^2 \left| X_t^{\bar{x}} - X_t^x \right|^4 \right) dt.
\end{aligned}$$

Next, since $y\bar{y} \leq \frac{1}{2}(y^2 + \bar{y}^2)$, after rearranging some of the terms we arrive at the expression

$$\begin{aligned}
\left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} &\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma(X_t^\delta)\right) dW_t \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} (L + (q-1)L^2) |\Gamma_t|^2 dt \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left(\frac{L}{2} \left(|\Gamma_t|^2 + \delta^2(1-\delta)^2 |X_t^{\bar{x}} - X_t^x|^4 \right) \right. \\
&\quad \quad \left. + (q-1)L^2\delta^2(1-\delta)^2 |X_t^{\bar{x}} - X_t^x|^4 \right) dt \\
&\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma(X_t^\delta)\right) dW_t \\
&\quad + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \left[\left(\frac{3}{2}L + (q-1)L^2 \right) |\Gamma_t|^2 \right. \\
&\quad \quad \left. + \delta^2(1-\delta)^2 \left(\frac{L}{2} + (q-1)L^2 \right) |X_t^{\bar{x}} - X_t^x|^4 \right] dt.
\end{aligned}$$

Next, using Young inequality with exponent $\frac{q}{q-2}$ and its conjugate $\frac{q}{2}$, we obtain

$$\begin{aligned}
&\left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} \\
&\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma(X_t^\delta)\right) dW_t \\
&\quad + q \left(\frac{3}{2}L + (q-1)L^2 \right) \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} dt \\
&\quad + \delta^2(1-\delta)^2 q \left(\frac{L}{2} + (q-1)L^2 \right) \int_0^t \left(\frac{q-2}{q} \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} + \frac{2}{q} |X_t^{\bar{x}} - X_t^x|^{2q} \right) dt \\
&\leq \varepsilon^{\frac{q}{2}} + q \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma(X_t^\delta)\right) dW_t \\
&\quad + ((2q-1)L + 2(q-1)^2L^2) \int_0^t \left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} dt \\
&\quad + \delta^2(1-\delta)^2 (L + 2(q-1)L^2) \int_0^t |X_t^{\bar{x}} - X_t^x|^{2q} dt.
\end{aligned} \tag{B.2}$$

Taking expectations and using Grönwall, (3.6) implies that

$$\mathbb{E} \left[\left(\varepsilon + |\Gamma_t|^2\right)^{\frac{q}{2}} \right] \leq \left(\varepsilon^{\frac{q}{2}} + C\delta^2(1-\delta)^2 |\bar{x} - x|^{2q} e^{2q(L + \frac{2q-1}{2}L^2)t} \right) e^{((2q-1)L + 2(q-1)^2L^2)t},$$

which in turn gives (3.7) by sending $\varepsilon \rightarrow 0$.

Proof of (3.8). First, by Itô's formula we have

$$\begin{aligned}
& e^{-\rho t} \left(1 + |X_t^x|^2\right)^{\frac{q}{2}} \\
&= \left(1 + |x|^2\right)^{\frac{q}{2}} + q \int_0^t e^{-\rho s} \left(1 + |X_s^x|^2\right)^{\frac{q}{2}-1} (X_s^x)^* \sigma(X_s^x) dW_s \\
&\quad + \int_0^t e^{-\rho s} \left(1 + |X_s^x|^2\right)^{\frac{q}{2}} \left[-\rho + q \left(\frac{(X_s^x)^*}{1 + |X_s^x|^2} b(X_s^x) \right. \right. \\
&\quad \quad \quad \left. \left. + \left(\frac{q}{2} - 1\right) \frac{|(X_s^x)^* \sigma(X_s^x)|^2}{(1 + |X_s^x|^2)^2} + \frac{1}{2} \frac{|\sigma(X_s^x)|^2}{1 + |X_s^x|^2} \right) \right] ds \\
&\leq \left(1 + |x|^2\right)^{\frac{q}{2}} + q \int_0^t e^{-\rho s} \left(1 + |X_s^x|^2\right)^{\frac{q}{2}-1} (X_s^x)^* \sigma(X_s^x) dW_s \\
&\quad + C \int_0^t e^{-\rho s} \left(1 + |X_s^x|^2\right)^{\frac{q}{2}} ds.
\end{aligned}$$

Next, for $\tau \in \mathcal{T}$ and $T > 0$, Burkholder-Davis-Gundy inequality gives

$$\begin{aligned}
\mathbb{E} \left[\sup_{t \leq \tau \wedge T} e^{-\rho t} \left(1 + |X_t^x|^2\right)^{\frac{q}{2}} \right] &\leq \left(1 + |x|^2\right)^{\frac{q}{2}} + C \int_0^T e^{-\rho t} \mathbb{E} \left[\left(1 + |X_t^x|^2\right)^{\frac{q}{2}} \right] dt \\
&\quad + C \mathbb{E} \left[\left(\int_0^T e^{-2\rho t} \left(1 + |X_t^x|^2\right)^{q-2} |X_t^x|^2 |\sigma(X_t^x)|^2 dt \right)^{\frac{1}{2}} \right].
\end{aligned}$$

Thus, using (3.5) and Jensen inequality we find

$$\begin{aligned}
\mathbb{E} \left[\sup_{t \leq \tau \wedge T} e^{-\rho t} \left(1 + |X_t^x|^2\right)^{\frac{q}{2}} \right] &\leq C(1 + |x|^q) + C(1 + |x|^q) \int_0^\infty e^{(-\rho + \hat{c}_0(q))t} dt \\
&\quad + C \left(\int_0^\infty e^{-2\rho t} \mathbb{E} \left[\left(1 + |X_t^x|^2\right)^q \right] dt \right)^{\frac{1}{2}} \\
&\leq C(1 + |x|^q) + C(1 + |x|^q) \left(\int_0^\infty e^{(-2\rho + \hat{c}_0(2q))t} dt \right)^{\frac{1}{2}} \\
&\leq C(1 + |x|^q),
\end{aligned}$$

where the finiteness of the integrals follows from the conditions $\rho > \hat{c}_0(q)$ and $2\rho > \hat{c}_0(2q)$ and the particular choice of q in (3.8) (which are part of Condition 4 in Assumption 3.1). The estimate (3.8) follows taking limits as $T \rightarrow \infty$ and using the dominated convergence theorem.

Proof of (3.9). With the notation $\Delta_t := X_t^{\bar{x}} - X_t^x$, similarly to (B.1) we find

$$\begin{aligned}
& \left(\varepsilon + |\Delta_t|^2\right)^{\frac{q}{2}} \varepsilon (\varepsilon + |\bar{x} - x|^2)^{\frac{q}{2}} + \left(-\rho + q \left(L + \frac{q-1}{2} L^2\right)\right) \int_0^t e^{-\rho s} \left(\varepsilon + |\Delta_s|^2\right)^{\frac{q}{2}} ds \\
&\quad + q \int_0^t e^{-\rho s} \left(\varepsilon + |\Delta_s|^2\right)^{\frac{q-2}{2}} \Delta_s^* (\sigma(X_s^{\bar{x}}) - \sigma(X_s^x)) dW_s.
\end{aligned}$$

For $\tau \in \mathcal{T}$ and $T > 0$, by using Burkholder-Davis-Gundy inequality and Jensen inequality, thanks to (3.6) we obtain

$$\begin{aligned} \mathbb{E} \left[\sup_{t \leq \tau \wedge T} \left(\varepsilon + |\Delta_t|^2 \right)^{\frac{q}{2}} \right] &\leq C(\varepsilon^{\frac{q}{2}} + |\bar{x} - x|^q) + C \int_0^T e^{-\rho t} \mathbb{E} [|\Delta_t|^q] dt \\ &\quad + C \left(\int_0^T e^{-2\rho t} \mathbb{E} [|\Delta_t|^{2q}] dt \right)^{\frac{1}{2}} \\ &\leq C \left(\varepsilon^{\frac{q}{2}} + |\bar{x} - x|^q \left(1 + \int_0^\infty \left(e^{(-\rho + \hat{c}_1(q))t} + e^{(-2\rho + \hat{c}_1(2q))t} \right) dt \right) \right) \\ &\leq C \left(\varepsilon^{\frac{q}{2}} + |\bar{x} - x|^q \right), \end{aligned}$$

where the finiteness of the integrals follows from the conditions $\rho > \hat{c}_1(q)$ and $2\rho > \hat{c}_1(2q)$ and the choice of q in (3.9) (which are part of Condition 4 in Assumption 3.1).

From the latter estimate, take limits as $\varepsilon \rightarrow 0$ and $T \rightarrow \infty$ and use monotone convergence theorem to derive (3.9).

Proof of (3.10). With the same notations as in the proof of (3.7) above, by repeating the steps leading to (B.2), we find

$$\begin{aligned} e^{-\rho t} \left(\varepsilon + |\Gamma_t|^2 \right)^{\frac{q}{2}} &\leq \varepsilon^{\frac{q}{2}} + C \int_0^t e^{-\rho t} \left(\varepsilon + |\Gamma_t|^2 \right)^{\frac{q}{2}} dt + C\delta^2(1-\delta)^2 \int_0^t e^{-\rho t} |X_t^{\bar{x}} - X_t^x|^{2q} dt \\ &\quad + q \int_0^t e^{-\rho t} \left(\varepsilon + |\Gamma_t|^2 \right)^{\frac{q-2}{2}} \Gamma_t^* \left(\hat{\sigma}_t - \sigma(X_t^\delta) \right) dW_t \end{aligned}$$

Hence, for $\tau \in \mathcal{T}$ and $T > 0$, by using Burkholder-Davis-Gundy inequality and Jensen inequality, thanks to (3.6) and (3.7) we get

$$\begin{aligned} \mathbb{E} \left[\sup_{t \leq \tau \wedge T} e^{-\rho t} \left(\varepsilon + |\Gamma_t|^2 \right)^{\frac{q}{2}} \right] &\leq C\varepsilon^{\frac{q}{2}} + C \int_0^T e^{-\rho t} \mathbb{E} [|\Gamma_t|^q] dt \\ &\quad + C\delta^2(1-\delta)^2 \int_0^T e^{-\rho t} \mathbb{E} [|X_t^{\bar{x}} - X_t^x|^{2q}] dt \\ &\quad + C \left(\int_0^T e^{-2\rho t} \mathbb{E} \left[\left(\varepsilon + |\Gamma_t|^2 \right)^{q-2} |\Gamma_t|^2 \left| \hat{\sigma}_t - \sigma(X_t^\delta) \right|^2 \right] dt \right)^{\frac{1}{2}}. \end{aligned} \tag{B.3}$$

To estimate the last term in the right-hand side, we first use the regularity of σ to obtain

$$\begin{aligned} &\int_0^T e^{-2\rho t} \mathbb{E} \left[\left(\varepsilon + |\Gamma_t|^2 \right)^{q-2} |\Gamma_t|^2 \left| \hat{\sigma}_t - \sigma(X_t^\delta) \right|^2 \right] dt \\ &\leq C \int_0^T e^{-2\rho t} \mathbb{E} \left[\left(\varepsilon + |\Gamma_t|^2 \right)^{q-1} \left(\left| \hat{\sigma}_t - \sigma(\hat{X}_t) \right|^2 + \left| \sigma(\hat{X}_t) - \sigma(X_t^\delta) \right|^2 \right) \right] dt \\ &\leq C \int_0^T e^{-2\rho t} \mathbb{E} \left[\left(\varepsilon + |\Gamma_t|^2 \right)^{q-1} \left(\delta^2(1-\delta)^2 \left| \hat{X}_t - X_t^\delta \right|^4 + |\Gamma_t|^2 \right) \right] dt, \end{aligned}$$

and then employ Young inequality with exponent $\frac{q}{q-1}$ and its conjugate q to conclude that

$$\begin{aligned} & \int_0^T e^{-2\rho t} \mathbb{E} \left[\left(\varepsilon + |\Gamma_t|^2 \right)^{q-2} |\Gamma_t|^2 \left| \hat{\sigma}_t - \sigma \left(X_t^\delta \right) \right|^2 \right] dt \\ & \leq C \left(\varepsilon^{\frac{q}{2}} + \int_0^T e^{-2\rho t} \mathbb{E} \left[|\Gamma_t|^{2q} \right] dt + \delta^2 (1 - \delta)^2 \int_0^T e^{-2\rho t} \mathbb{E} \left[\left| \hat{X}_t - X_t^\delta \right|^{4q} \right] dt \right). \end{aligned}$$

By plugging the latter inequality into (B.3), and then using (3.6) and (3.7), we find

$$\begin{aligned} & \mathbb{E} \left[\sup_{t \leq \tau \wedge T} e^{-\rho t} \left(\varepsilon + |\Gamma_t|^2 \right)^{\frac{q}{2}} \right] \\ & \leq C \left(\varepsilon^{\frac{q}{2}} + \int_0^T e^{-\rho t} \mathbb{E} \left[|\Gamma_t|^q \right] dt + \delta^2 (1 - \delta)^2 \int_0^T e^{-\rho t} \mathbb{E} \left[\left| X_t^{\bar{x}} - X_t^x \right|^{2q} \right] dt \right. \\ & \quad \left. + \left(\varepsilon^{\frac{q}{2}} + \int_0^T e^{-2\rho t} \mathbb{E} \left[|\Gamma_t|^{2q} \right] dt + \delta^2 (1 - \delta)^2 \int_0^T e^{-2\rho t} \mathbb{E} \left[\left| \hat{X}_t - X_t^\delta \right|^{4q} \right] dt \right)^{\frac{1}{2}} \right) \\ & \leq C \left(\varepsilon^{\frac{q}{2}} + \delta^2 (1 - \delta)^2 |\bar{x} - x|^{2q} \int_0^\infty \left(e^{(-\rho + \hat{c}_2(q))t} + e^{(-\rho + \hat{c}_1(2q))t} \right) dt \right. \\ & \quad \left. + \left(\varepsilon^{\frac{q}{2}} + \delta^2 (1 - \delta)^2 |\bar{x} - x|^{4q} \int_0^\infty \left(e^{(-2\rho + \hat{c}_2(2q))t} + e^{(-2\rho + \hat{c}_1(4q))t} \right) dt \right)^{\frac{1}{2}} \right). \end{aligned}$$

Finally, by sending $\varepsilon \rightarrow 0$, we conclude that

$$\begin{aligned} & \mathbb{E} \left[\sup_{t \leq \tau \wedge T} e^{-\rho t} |\Gamma_t|^q \right] \leq C \delta (1 - \delta) |\bar{x} - x|^{2q} \int_0^\infty \left(e^{(-\rho + \hat{c}_2(q))t} + e^{(-\rho + \hat{c}_1(2q))t} \right) dt \\ & \quad + C \delta (1 - \delta) |\bar{x} - x|^{2q} \left(\int_0^\infty \left(e^{(-2\rho + \hat{c}_2(2q))t} + e^{(-2\rho + \hat{c}_1(4q))t} \right) dt \right)^{\frac{1}{2}} \\ & \leq C \delta (1 - \delta) |\bar{x} - x|^{2q}, \end{aligned}$$

where the finiteness of the integrals follows from taking $q = 2$ and from the conditions $\rho > \hat{c}_1(4q)$ and $2\rho > \hat{c}_2(2q)$ (which are part of Condition 4 in Assumption 3.1). The estimate (3.9) follows taking limits as $T \rightarrow \infty$ and using the dominated convergence theorem.

Acknowledgements. Supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - Project-ID 317210226 - SFB 1283 and the NSF CAREER award DMS-2339240.

REFERENCES

- [1] A. AGARWAL, S. M. KAKADE, J. D. LEE, AND G. MAHAJAN, *On the theory of policy gradient methods: Optimality, approximation, and distribution shift*, J. Mach. Learn. Res., 22 (2021), pp. 1–76.
- [2] B. ATA, J. M. HARRISON, AND N. SI, *Singular control of (reflected) Brownian motion: A computational method suitable for queueing applications*, arXiv preprint arXiv:2312.11823, (2023).
- [3] L. BAI, T. GAMAGE, J. MA, AND P. XIE, *Reinforcement learning for optimal dividend problem under diffusion model*, arXiv preprint arXiv:2309.10242, (2023).
- [4] P. BANK AND N. EL KAROUI, *A stochastic representation theorem with applications to optimization and obstacle problems*, Ann. Probab., 32 (2004), pp. 1030–1067.
- [5] P. BANK AND F. RIEDEL, *Optimal consumption choice with intertemporal substitution*, Ann. Appl. Probab., 11 (2001), pp. 750–788.
- [6] M. BASEI, X. GUO, A. HU, AND Y. ZHANG, *Logarithmic regret for episodic continuous-time linear-quadratic reinforcement learning over a finite-time horizon*, J. Mach. Learn. Res., 23 (2022), pp. 1–34.
- [7] J. BATHER, *A continuous time inventory model*, J. Appl. Probab., 3 (1966), pp. 538–549.

- [8] F. BOETIUS AND M. KOHLMANN, *Connections between optimal stopping and singular stochastic control*, Stochastic Process. Appl., 77 (1998), pp. 253–281.
- [9] T. DE ANGELIS, S. FEDERICO, AND G. FERRARI, *Optimal boundary surface for irreversible investment with stochastic costs*, Math. Oper. Res., 42 (2017), pp. 1135–1161.
- [10] T. DE ANGELIS AND G. FERRARI, *Stochastic nonzero-sum games: a new connection between singular control and optimal stopping*, Adv. in Appl. Probab., 50 (2018), pp. 347–372.
- [11] T. DE ANGELIS, G. FERRARI, AND J. MORIARTY, *A solvable two-dimensional degenerate singular stochastic control problem with nonconvex costs*, Math. Oper. Res., 44 (2019), pp. 512–531.
- [12] T. DE ANGELIS, N. MERKULOV, AND J. PALCZEWSKI, *On the value of non-Markovian Dynkin games with partial and asymmetric information*, Ann. Appl. Probab., 32 (2022), pp. 1774–1813.
- [13] H. DENKERT, ROBERT PHAM AND X. WARIN, *Control randomisation approach for policy gradient and application to reinforcement learning in optimal switching*, arXiv preprint arXiv:2404.17939, (2024).
- [14] R. DEVIDZE, P. KAMALARUBAN, AND A. SINGLA, *Exploration-guided reward shaping for reinforcement learning under sparse rewards*, Adv. Neural Inf. Process. Syst., 35 (2022), pp. 5829–5842.
- [15] J. DIANETTI AND G. FERRARI, *Multidimensional singular control and related Skorokhod problem: sufficient conditions for the characterization of optimal controls*, Stochastic Process. Appl., 162 (2023), pp. 547–592.
- [16] Y. DONG, *Randomized optimal stopping problem in continuous time and reinforcement learning algorithm*, SIAM J. Control Optim., 62 (2024), pp. 1590–1614.
- [17] K. DOYA, *Reinforcement learning in continuous time and space*, Neural Comput., 12 (2000), pp. 219–245.
- [18] J. C. DUCHI, M. I. JORDAN, M. J. WAINWRIGHT, AND A. WIBISONO, *Optimal rates for zero-order convex optimization: The power of two function evaluations*, IEEE Trans. Inf. Theory, 61 (2015), pp. 2788–2806.
- [19] M. FAZEL, R. GE, S. KAKADE, AND M. MESBAHI, *Global convergence of policy gradient methods for the linear quadratic regulator*, in International conference on machine learning, PMLR, 2018, pp. 1467–1476.
- [20] G. FERRARI, *On an integral equation for the free-boundary of stochastic, irreversible investment problems*, Ann. Appl. Probab., 25 (2015), pp. 150 – 176.
- [21] G. FERRARI AND P. SALMINEN, *Irreversible investment under lévy uncertainty: An equation for the optimal boundary*, Adv. in Appl. Probab., 48 (2016), pp. 298–314.
- [22] W. H. FLEMING AND H. M. SONER, *Controlled Markov Processes and Viscosity Solutions*, Springer Science & Business Media, 2006.
- [23] ———, *Controlled Markov processes and viscosity solutions*, vol. 25, Springer Science & Business Media, 2006.
- [24] M. GIEGRICH, C. REISINGER, AND Y. ZHANG, *Convergence of policy gradient methods for finite-horizon exploratory linear-quadratic control problems*, SIAM J. Control Optim., 62 (2024), pp. 1060–1092.
- [25] D. GILBARG AND N. S. TRUDINGER, *Elliptic Partial Differential Equations of Second Order*, vol. 224, Springer Science & Business Media, 2001.
- [26] X. GUO, A. HU, AND Y. ZHANG, *Reinforcement learning for linear-convex models with jumps via stability analysis of feedback controls*, SIAM J. Control Optim., 61 (2023), pp. 755–787.
- [27] X. GUO AND P. TOMECEK, *A class of singular control problems and the smooth fit principle*, SIAM J. Control Optim., 47 (2009), pp. 3076–3099.
- [28] B. HAMBLY, R. XU, AND H. YANG, *Policy gradient methods for the noisy linear quadratic regulator over a finite horizon*, SIAM J. Control Optim., 59 (2021), pp. 3359–3391.
- [29] J. HARE, *Dealing with sparse rewards in reinforcement learning*, arXiv preprint arXiv:1910.09281, (2019).
- [30] Y.-J. HUANG, Z. WANG, AND Z. ZHOU, *Convergence of policy improvement for entropy-regularized stochastic control problems*, arXiv preprint arXiv:2209.07059, (2022).
- [31] Y. JIA AND X. Y. ZHOU, *Policy gradient and actor-critic learning in continuous time and space: Theory and algorithms*, J. Mach. Learn. Res., 23 (2022), pp. 1–50.
- [32] ———, *q-learning in continuous time*, J. Mach. Learn. Res., 24 (2023), pp. 1–61.
- [33] Y. M. KABANOV, *Hedging and liquidation under transaction costs in currency markets*, Finance Stoch., 3 (1999), pp. 237–248.
- [34] I. KARATZAS, D. OCONE, H. WANG, AND M. ZERVOS, *Finite-fuel singular control with discretionary stopping*, Stochastics, 71 (2000), pp. 1–50.
- [35] I. KARATZAS AND S. E. SHREVE, *Connections between optimal stopping and singular stochastic control I. Monotone follower problems*, SIAM J. Control Optim., 22 (1984), pp. 856–877.
- [36] N. E. KAROUI AND I. KARATZAS, *Probabilistic aspects of finite-fuel, reflected follower problems*, Acta Appl. Math., 11 (1988), pp. 223–258.
- [37] Y. KIFER, *Dynkin’s games and Israeli options*, International Scholarly Research Notices, 2013 (2013), p. 856458.
- [38] S. KUMAR AND K. MUTHURAMAN, *A numerical method for solving singular stochastic control problems*, Oper. Res., 52 (2004), pp. 563–582.

- [39] D. LAMBERTON ET AL., *Optimal stopping and American options*, Ljubljana Summer School on Financial Mathematics, 134 (2009).
- [40] J. MA, G. WANG, AND J. ZHANG, *On convergence and rate of convergence of policy improvement algorithms*, arXiv preprint arXiv:2406.10959, (2024).
- [41] X. MAO, *Stochastic Differential Equations and Applications*, Elsevier, 2007.
- [42] J.-L. MENALDI AND M. ROBIN, *On some cheap control problems for diffusion processes*, Trans. Amer. Math. Soc., 278 (1983), pp. 771–802.
- [43] J. L. MENALDI AND M. I. TAKSAR, *Optimal correction problem of a multidimensional stochastic system*, Automatica J. IFAC, 25 (1989), pp. 223–232.
- [44] E. NEUMAN, W. STOCKINGER, AND Y. ZHANG, *An offline learning approach to propagator models*, arXiv preprint arXiv:2309.02994, (2023).
- [45] M. PALANISAMY, H. MODARES, F. L. LEWIS, AND M. AURANGZEB, *Continuous-time q-learning for infinite-horizon discounted cost linear quadratic regulator problems*, IEEE Trans. Cybern., 45 (2014), pp. 165–176.
- [46] G. PESKIR AND A. SHIRYAEV, *Optimal stopping and free-boundary problems*, Springer, 2006.
- [47] H. PHAM, *Continuous-time Stochastic Control and Optimization with Financial Applications*, vol. 61, Springer Science & Business Media, 2009.
- [48] M. RAO, Y. CHEN, B. C. VEMURI, AND F. WANG, *Cumulative residual entropy: a new measure of information*, IEEE Trans. Inform. Theory, 50 (2004), pp. 1220–1228.
- [49] C. REISINGER, W. STOCKINGER, AND Y. ZHANG, *Linear convergence of a policy gradient method for some finite horizon continuous time control problems*, SIAM J. Control Optim., 61 (2023), pp. 3526–3558.
- [50] A. M. REPPEN, H. M. SONER, AND V. TISSOT-DAGUETTE, *Neural optimal stopping boundary*, arXiv preprint arXiv:2205.04595, (2022).
- [51] E. SOLAN, B. TSIRELSON, AND N. VIEILLE, *Random stopping times in stopping problems and stopping games*, arXiv preprint arXiv:1211.5802, (2012).
- [52] H. M. SONER AND V. TISSOT-DAGUETTE, *Stopping times of boundaries: Relaxation and continuity*, arXiv preprint arXiv:2305.09766, (2023).
- [53] R. S. SUTTON AND A. G. BARTO, *Reinforcement learning: An introduction*, MIT press, 2018.
- [54] L. SZPRUCH, T. TREETANTHPILOET, AND Y. ZHANG, *Exploration-exploitation trade-off for continuous-time episodic reinforcement learning with linear-convex models*, arXiv preprint arXiv:2112.10264, (2021).
- [55] ———, *Optimal scheduling of entropy regularizer for continuous-time linear-quadratic reinforcement learning*, SIAM J. Control Optim., 62 (2024), pp. 135–166.
- [56] C. TALLEC, L. BLIER, AND Y. OLLIVIER, *Making deep q-learning methods robust to time discretization*, in International Conference on Machine Learning, PMLR, 2019, pp. 6096–6104.
- [57] W. TANG, Y. P. ZHANG, AND X. Y. ZHOU, *Exploratory HJB equations and their convergence*, SIAM J. Control Optim., 60 (2022), pp. 3191–3216.
- [58] P. THOMAS AND E. BRUNSKILL, *Data-efficient off-policy policy evaluation for reinforcement learning*, in International Conference on Machine Learning, PMLR, 2016, pp. 2139–2148.
- [59] N. TOUZI AND N. VIEILLE, *Continuous-time Dynkin games with mixed strategies*, SIAM J. Control Optim., 41 (2002), pp. 1073–1088.
- [60] H. V. TRAN, Z. WANG, AND Y. P. ZHANG, *Policy iteration for exploratory Hamilton–Jacobi–Bellman equations*, arXiv preprint arXiv:2406.00612, (2024).
- [61] H. WANG, T. ZARIPHPOULOU, AND X. Y. ZHOU, *Reinforcement learning in continuous time and space: A stochastic control approach*, J. Mach. Learn. Res., 21 (2020), pp. 1–34.
- [62] S. A. WILLIAMS, P.-L. CHOW, AND J.-L. MENALDI, *Regularity of the free boundary in singular stochastic control*, J. Differential Equations, 111 (1994), pp. 175–201.