Singular Control, Ambiguity and Real Options

Arnon Archankul

PhD

UNIVERSITY OF YORK
MATHEMATICS

September 2025

Abstract

Uncertainty in decision making is a core challenge in economics, finance, and operations research. While classical models address risk with known probabilities, real-world environments often involve ambiguity, where probabilities are unknown or misspecified. Ambiguity arises in corporate cash holdings, climate-driven water management, and inventory control, each structured as an inventory-type problem in which resources accumulate or deplete and interventions occur once thresholds are reached. This thesis develops singular stochastic control models that explicitly incorporate ambiguity, extending the scope of real options theory beyond risk-based formulations.

The thesis progresses from restrictive to more flexible ambiguity frameworks. The first study analyzes cash management under maxmin utility with κ-ignorance, showing via Dynkin games that extreme ambiguity aversion narrows inaction regions and raises expected costs, providing a tractable benchmark. The second study considers reservoir management under smooth ambiguity with a finite-state hidden variable, capturing flood-drought dynamics under climate uncertainty. Using forward-backward stochastic differential equations (FBSDEs), Hamilton-Jacobi-Bellman variational inequalities, and an efficient Markov chain approximation scheme, it shows how ambiguity aversion accelerates interventions while learning gradually mitigates this conservatism. The third study advances to smooth ambiguity with Gaussian hidden variables in a finite-horizon inventory setting, where the problem is formulated through an FBSDE with quadratic growth. This analysis uncovers novel patterns: while higher risk typically delays action, under deep ambiguity, it can instead prompt earlier and more cautious interventions.

Together, these contributions enrich the theory of singular control by embedding ambiguity preferences within nonlinear expected utility, advance analytical and numerical techniques of independent interest, and offer practical insights for managing cash, water, and inventory under deep uncertainty. Conceptually, all models act as forms of ambiguity insurance, quantifying the additional strategic costs required to safeguard decisions in environments where probabilities are themselves uncertain.

Contents

List of Figures

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Jacco Thijssen, for his unwavering support and guidance throughout my PhD journey. His time, patience, and dedication in meeting with me weekly, always ready to listen and share ideas, not only on research but also on life and the future, have been invaluable. His profound insights have broadened my perspective on both academia and the wider world, inspiring me to embrace and enjoy research rather than view it as a burden. I am especially grateful for his encouragement of my ambitious ideas and the freedom he has given me to think and explore, even when the challenges carried risks for my PhD. I am proud of the researcher I have become, thanks in large part to his mentorship. Above all, I am sincerely thankful to him for accepting me as one of his students.

I would also like to extend my sincere thanks to my thesis advisory panel members, Dr. Alet Roux and Dr. Nick Huberts, for their support, valuable insights, and thoughtful guidance throughout my time as a PhD student in York.

I would also like to thank Prof. Giorgio Ferrari and Dr. Tobias Hellmann for their mathematical insights and constructive suggestions, which contributed significantly to the development of my mathematical writing and to the completion of my first research paper at York.

My PhD studies would not have been possible without the generous support of the Development and Promotion of Science and Technology Talents Project (DPST), the Thai government scholarship. I am deeply grateful to DPST for providing me with this invaluable opportunity to pursue my development as a researcher and to contribute to the advancement of science and technology in Thailand.

None of my achievements would have been possible without the everlasting support of my beloved family–my mother, Boonjiranon; my brother and sister, Kiattikul and Tibanan; and my grandparents, Wirot and Ranong. Words cannot adequately express my gratitude for their unconditional love, and for their support and understanding of the challenges I faced while being thousands of miles away from home.

Last but not least, I would like to thank all of my friends in York for their incredible support and for making my time there both memorable and enjoyable, as well as everyone who contributed, directly or indirectly, to the completion of my PhD.

Author’s declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for an award at this, or any other, University. All sources are acknowledged as References.

Chapter 3 has been published as

> Archankul, A., Ferrari, G., Hellmann, T., & Thijssen, J. J. (2025). Singular control in a cash management model with ambiguity. *European Journal of Operational Research*, *327*(2), 500–514.

Chapters 4 and 5 have not yet been published but are prepared with the intention of being submitted to journals in environmental economics and operations research, respectively.

— 1 —

Introduction

Decision-making under uncertainty is a central problem in economics, finance, and operations research. In practice, managers frequently face environments where outcomes are influenced not only by risk (known probabilities), but also by ambiguity (unknown or misspecified probabilities). Ambiguity typically arises when the decision maker (DM) lacks reliable information about the underlying uncertainty, or does not trust that the available information accurately reflects its future evolution, and therefore cannot reduce uncertainty to a single, known probability. The distinction between risk and ambiguity, which goes back to Knight (1921), has profound implications for both theoretical modeling and managerial practice. Within this context, real options theory, as formalized in the seminal work of Dixit and Pindyck (1994), has established itself as a strong foundation for valuing flexibility: investment and operational decisions are viewed as options that can be exercised contingent on the evolution of uncertain variables. Real options models, however, are almost universally developed under risk, assuming known probabilistic laws for the underlying processes. In many important applications, this assumption is too restrictive: firms must manage cash under uncertain investors' preference (Breuer et al., 2017), water authorities must operate reservoirs in the face of climate-driven precipitation uncertainty (Brugnach et al., 2025), and supply chain managers must manage inventories under incomplete knowledge of demand fluctuations (Kocabıyıkoğlu et al., 2024; Ma & Aloysius, 2022). Each of these settings shares a common structure: they are inventory-type problems, where resources such as cash, water, or stock are accumulated or depleted over time, and adjustments are made when thresholds are reached. This thesis advances the development of mathematical frameworks for such inventory-type problems by explicitly incorporating ambiguity, thereby extending the scope of real options theory to settings where probabilities themselves are uncertain.

Singular stochastic control is particularly well suited to modeling inventory-type decisions. Unlike regular control, where adjustments are made through instantaneous rates, singular control operates through cumulative interventions that may be applied continuously once critical thresholds are reached. This structure naturally reflects the logic of resource management: firms inject or withdraw cash in lumps, reservoirs release or store water in response to inflows, and managers adjust inventories when stocks cross target levels. The mathematical structure of such problems has been studied extensively in the literature, beginning with early contributions on Brownian inventory models by Harrison and Taksar (1983) and dividend/cash management models by Jeanblanc-Picqué and Shiryaev (1995). More generally, the theory is closely linked to free-boundary problems and Dynkin games (see, e.g., Taksar, 1985), where optimal policies are characterized by inaction regions bounded by thresholds that act as endogenous barriers. These works establish the classical foundation on which this thesis builds by incorporating ambiguity into the singular control framework.

This thesis incorporates different frameworks of ambiguity into singular control models. The literature on decision-making under ambiguity spans a wide range of approaches. The behavior of decision makers under ambiguity was first highlighted by Ellsberg (1961) in the well-known *Ellsberg urn experiment*. In this experiment, subjects choose between bets on two urns: one with a known distribution of colored balls and another with an unknown distribution. The majority prefer the former, revealing an aversion to ambiguity, that is, a preference for risk over ambiguity.

A first axiomatic treatment of ambiguity within utility theory is provided by Gilboa and Schmeidler (1989). They postulate that ambiguity-averse DMs behave as if nature consistently selects the worst-case distribution from a given set of priors. Translated into the urn experiment, this corresponds to assuming that the urn is effectively replaced before each draw, with the decision maker evaluating bets under the worst possible prior each time. This framework is known as *maxmin utility*. Building on this, Chen and Epstein (2002) incorporate maxmin utility into stochastic differential utility in the spirit of Duffie and Epstein (1992). Their work provides a conservative benchmark for continuous-time decision-making under ambiguity. The first part of this thesis extends their framework by embedding it into singular control to address an optimization problem in cash management.

While analytically tractable, maxmin utility is often regarded as excessively pessimistic, since it excludes belief updating or learning from observed signals. A more flexible alternative is offered by the *smooth ambiguity preferences* of Klibanoff et al. (2005), which separate beliefs from attitudes toward ambiguity and, in their dynamic extension (Klibanoff et al., 2009), allow for Bayesian learning. Returning to the urn analogy, a decision maker with smooth ambiguity preferences treats the urns composition as fixed but unknown, and gradually learns it through repeated draws. Hansen and Sargent (2011) further connect smooth ambiguity with stochastic differential utility, thereby providing a richer and more realistic account of managerial behavior under ambiguity. This thesis adopts smooth ambiguity as a natural development of maxmin utility and applies it to singular control problems under different specifications of belief over ambiguity. In doing so, the second part of the thesis studies reservoir management under smooth ambiguity with finite-state beliefs, and the third part considers inventory management under smooth ambiguity with continuous-state beliefs.

1.1 ORGANIZATION OF THESIS

The central aim of this thesis is to trace a progression from restrictive to more flexible ambiguity frameworks, applying them to singular control problems motivated by cash, water, and inventory management. The thesis adopts a journal-style format and is organized into six chapters, three of which correspond to individual research papers. Chapter 2 provides an essential mathematical background for the analysis of each paper. Chapter 3, representing the first paper, introduces the simplest setting, analyzing cash management under maxmin utility. Chapter 4, representing the second paper, turns to reservoir management under smooth ambiguity, where ambiguity is represented by a finite-state hidden variable. Finally, Chapter 5, representing the second paper, extends the analysis to inventory management under smooth ambiguity with a continuous hidden variable, offering a more general and flexible representation of uncertainty. The thesis concludes with Chapter 6, which summarizes the main findings and outlines directions for future research.

1.2 OVERVIEW OF THESIS

Chapter 3 considers the singular control of cash management under maxmin utility with κ-ignorance. Cash inflows are modeled as a diffusion process, either arithmetic Brownian motion or Ornstein-Uhlenbeck. The problem is formulated as a Dynkin game, which permits analytical characterization of the value function and comparative statics of the optimal control policy. The paper shows that maxmin utility, while analytically convenient, is restrictive in several important respects. It embodies an extreme form of ambiguity aversion: the worst-case prior is held fixed throughout the decision horizon, precluding any form of ambiguity learning. The main contribution of the first paper is twofold: it provides a tractable benchmark model linking singular control, Dynkin games, and ambiguity, and it motivates the shift to more flexible ambiguity frameworks. This chapter has been published as Archankul et al. (2025).

Chapter 4 addresses limitations of Chapter 3 by turning to smooth ambiguity preferences. The application is reservoir management under climate-driven uncertainty. Reservoir inflows are modeled as an Ornstein-Uhlenbeck process, while ambiguity is represented by a Bernoulli-distributed hidden variable corresponding to two extreme climate regimes: flood and drought. This formulation captures the challenge of managing water resources under climate change, where managers cannot commit to a single probabilistic model of inflows. The optimal control problem leads to a Hamilton-Jacobi-Bellman (HJB) equation in the viscosity sense, which has no closed-form solution. To compute the value function and optimal policy, the paper develops an efficient numerical scheme based on Markov chain approximation combined with a coordinate transformation that improves stability and tractability. The analysis reveals that while the Bernoulli model provides a useful first step in incorporating smooth ambiguity into singular control, it is limited in scope:

it forces the manager to view ambiguity in binary terms. Extending the model to richer ambiguity structures, such as multi-state Wonham filters, would be computationally demanding. This limitation motivates the third paper.

Chapter 5 generalizes the analysis by considering inventory management under smooth ambiguity. Net inventory flows follow an arithmetic Brownian motion, while ambiguity is represented by a Gaussian hidden variable. Unlike the first two papers, the problem is set on a finite horizon, which is often more realistic for inventory and production planning. The continuous representation of ambiguity removes the need to commit to simplistic assumptions such as "flood vs. drought" or "good vs. bad state." Instead, it allows managers to model their uncertainty in a flexible way, accommodating a continuum of possible scenarios. The associated HJB equation is solved using an efficient Markov chain approximation scheme, building on the computational techniques of the Chapter 4.

Taken together, the three papers provide a coherent exploration of singular stochastic control under ambiguity. The thesis contributes to the literature at three levels. First, it enriches the theory of singular control by embedding ambiguity preferences, typically represented through multiple priors and nonlinear expectations, thereby broadening its scope beyond the classical risk-based model. Second, it develops novel analytical and numerical techniques, drawing on Dynkin games, viscosity solutions, and Markov chain approximation, that are of independent methodological interest. Third, it delivers managerial insights across multiple disciplines, demonstrating how different models of ambiguity affect optimal policies in cash, water, and inventory management. By starting with maxmin utility and advancing toward smooth ambiguity with continuous hidden variables, the thesis highlights both the limitations of extreme conservatism and the benefits of flexible modeling in ambiguous environments. In essence, the models developed here function as forms of *ambiguity insurance*: they quantify and address the additional strategic costs required to safeguard against uncertainty that cannot be captured by risk alone.

— 2 —

Mathematical Background

2.1 LIST OF NOTATIONS

2.1.1 General notations

1) $\triangleq$ refers to "is defined to be".
2) $\mathbb{N} \triangleq \{1, 2, 3, \ldots\}$.
3) $\mathbb{R}^d$ refers to the d-dimensional Euclidean space, i.e., the set of d-tuples real numbers. Denote $\mathbb{R} = \mathbb{R}^1$.
4) $\mathbb{R}^{n\times d}$ refers to the set of real-valued $n \times d$ matrices.

- Let $a, b \in \mathbb{R}$, $x \triangleq (x_i)_{1\le i\le n},\ y \triangleq (y_i)_{1\le i\le n} \in \mathbb{R}^{n\times 1}$ and $A \triangleq (a_{ij})_{1\le i\le n, 1\le j\le d},\ B \triangleq (b_{ij})_{1\le i\le n, 1\le j\le d} \in \mathbb{R}^{n\times d}$.

5) $a \wedge b \triangleq \min\{a, b\}$.
6) $a \vee b \triangleq \max\{a, b\}$.
7) $a- \triangleq \lim_{x\downarrow a} x$.
8) $a+ \triangleq \lim_{x\uparrow a} x$.
9) $1_A(x) \triangleq \begin{cases} 1 & \text{if } x \in A \\ 0 & \text{otherwise.} \end{cases}$
10) $\text{sign}(a) \triangleq \begin{cases} 1 & \text{if } a > 0 \\ 0 & \text{if } a = 0 \\ -1 & \text{if } a < 0. \end{cases}$
11) $A^\top = (a_{ji})_{1\le j\le d, 1\le i\le n}$ refers the *transpose* of A.
12) If $n = d$, then $\text{Tr}(A) \triangleq \sum_{i=1}^n a_{ii}$ refers to the *trace* of A.
13) $|\cdot|$ refers to the *Euclidean norm*, i.e., $|x| = \sqrt{x.x}$ where $x.y = \sum_{i=1}^n x_i y_i$.

2.1.2 FUNCTION SPACES

Let $T \in [0, \infty)$.

1) $C^k(\mathbb{R}^n)$ is the space of continuous functions with real values in $\mathbb{R}^n$ with derivatives that exist and are continuous up to k order.
2) $C([0,T] \times \mathbb{R}^n)$ is the space of real-valued continuous functions on $C([0,T] \times \mathbb{R}^n)$.
3) $C^{1,2}([0,T] \times \mathbb{R}^n)$ is the space of continuous functions f with real values in $[0,T] \times \mathbb{R}^n$ with partial derivatives $\frac{\partial}{\partial t} f$, $\frac{\partial}{\partial x_i} f$, $\frac{\partial^2}{\partial x_i \partial x_j} f$, $1 \le i, j \le n$ which exist and are continuous on $[0,T] \times \mathbb{R}^n$.

- Let $f \in C^{1,2}([0,T] \times \mathbb{R}^n)$. Then

4) The *gradient* of f with respect to the state variable $x = (x_1, \ldots, x_n)$ is the column vector

$$\nabla f(t,x) \triangleq \begin{pmatrix} \frac{\partial}{\partial x_1} f(t,x) \\ \vdots \\ \frac{\partial}{\partial x_n} f(t,x) \end{pmatrix} \in \mathbb{R}^n.$$

5) The *Hessian matrix* of f with respect to x is

$$D^2 f(t,x) \triangleq \left(\frac{\partial^2}{\partial x_i \partial x_j} f(t,x) \right)_{1 \le i,j \le n} \in \mathbb{R}^{n \times n}.$$

- Let $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t \in [0,T]}, \mathsf{P})$ be a filtered probability space. Then

6) $\mathbb{L}^\infty(\mathscr{F}, \mathsf{P})$ is the set of the almost surely bounded $\mathscr{F}$-measurable random variable ξ, i.e., $|\xi| < \infty$, $\mathsf{P}-$a.s.
7) $\mathbb{L}^2(\mathscr{F}, \mathsf{P})$ is the set of $\mathscr{F}$-measurable random variable ξ such that $\mathsf{E}^{\mathsf{P}}\left[|\xi|^2\right] < \infty$.
8) $\mathbb{L}^\infty_T(\mathbf{F}, \mathsf{P})$ is the set of the almost surely bounded $\mathbf{F}$-adapted process Y, i.e., $|Y_t| < \infty$, P-a.s., for any $0 \le t \le T$,
9) $\mathbb{L}^2_T(\mathbf{F}, \mathsf{P})$ is the set of $\mathbf{F}$-adapted process Y such that $\mathsf{E}^{\mathsf{P}}\left[\sup_{0 \le s \le T} |Y_s|^2\right] < \infty$.
10) $\mathbb{H}^2_T(\mathbf{F}, \mathsf{P})$ is the set of $\mathbf{F}$-adapted process Y such that $\mathsf{E}^{\mathsf{P}}\left[\int_0^T |Y_s|^2 \mathrm{d}s\right] < \infty$.

2.2 STOCHASTIC ANALYSIS FOR SINGULAR CONTROL

This section assembles the stochastic analysis used in the thesis that covers the following attributes.

1) In this thesis, all of the state dynamics are modeled as *Itō diffusions with singular control*, i.e., semimartingales driven by Brownian motion with an additional finite-variation control process.
2) Ambiguity is represented by families of probability measures connected through Girsanov transformations.
3) Optimal policies are characterized via Hamilton-Jacobi-Bellman (HJB) variational inequalities, or in a special case through *Dynkin games* (Chapter 3).
4) Problems in Chapters 4 and 5 require numerical methods based on the *Markov Chain Approximation (MCA)* scheme to solve the associated HJB equations. Therefore, it is compulsory that the controlled process is a Markov process and its discrete generator is *locally consistent* with the continuous counterpart.

To prepare for these developments, we introduce filtered probability spaces and Brownian motion, stopping times, martingales, and stochastic differential equations; establish Markov properties and transition kernels; present semimartingales and Itō's formula; and conclude with a change-of-measure result used to model ambiguity. Together, these elements provide the analytical foundation for the control problems and guarantee the consistency and convergence of the MCA scheme.

Unless otherwise stated, all definitions and notations in this section are based on the adaptation from Protter (2010) and Øksendal (2010).

2.2.1 STOCHASTIC PROCESSES

Let $T < \infty$. Throughout this thesis, we assume that a defined filtered probability space $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t\in[0,T]}, \mathsf{P})$ satisfies the *usual conditions*. That is, $\mathbf{F}$ right continuous, i.e., $\mathscr{F}_{t+} \triangleq \cap_{s>t}\mathscr{F}_s$, for any $t \in [0, T]$, and every subset of a P-null set lies in $\mathbf{F}$. We write $\mathbf{F} = (\mathscr{F}_t)_{t\geq 0}$ if $T = \infty$. The associated expectation operator is denoted by E.

Remark 1 (On the usual conditions)**.** *The* usual conditions *are essential for constructing a mathematically coherent probabilistic framework.*

1) Right-continuity. *The requirement that the filtration be right-continuous ensures that no "unexpected jump" of information occurs at an observation time. This property is crucial for the well-definedness of conditional expectations.*
2) Completeness. *The inclusion of all* P*-null sets in the filtration guarantees consistency of the measurable structure* $(\Omega, \mathscr{F}, \mathsf{P})$*. In particular, if* $A \in \mathscr{F}$ *satisfies* $\mathsf{P}(A) = 1$*, then every*

subset of the negligible complement A^c is also in $\mathscr{F}$ and has probability zero. This ensures that events holding almost surely (a.s.) are properly measurable and compatible with the filtration.

Definition 1 (Random variable)**.** Let $(\Omega, \mathscr{F}, \mathsf{P})$ be a probability space. A function $\xi : \Omega \to \mathcal{B}(E)$ is said to be a *random variable* if it is $\mathscr{F}$-measurable, i.e.,

$$\{\omega \in \Omega : \xi(\omega) \in A\} \in \mathscr{F}, \text{ for all } A \in \mathcal{B}(E).$$

Definition 2 (Normal distribution)**.** If the random variable X has *normal distribution* (or is *Gaussian*) with mean μ and variance σ^2, we write $X \sim \mathsf{N}(\mu, \sigma^2)$.

Definition 3 (Stopping time)**.** A random time $\tau : \Omega \to [0, \infty]$ is an $\mathbf{F}$-stopping time if $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$. The filtration generated by a stopping time τ is $\mathcal{F}_\tau := \{A \in \mathcal{F} : \ A \cap \{\tau \leq t\} \in \mathcal{F}_t, \ \forall t \geq 0\}$.

Definition 4 (Stochastic process)**.** Let $(\Omega, \mathscr{F}, \mathsf{P})$ be a probability space, and $(E, \mathcal{B}(E))$, $E \subseteq \mathbb{R}$, a measurable state space. A *stochastic process* (or just *process*) with state space E is a family $X = (X_t)_{t \geq 0}$ of random variables

$$X_t : (\Omega, \mathscr{F}) \to (E, \mathcal{B}(E)).$$

Definition 5 (Adaptedness)**.** Given a filtration $\mathbf{F}$, the process X is said to be $\mathbf{F}$-*adapted* if, for every $t > 0$, the map X_t is $\mathscr{F}_t$-measurable.

We simply say that X is an adapted process when there is no confusion about $\mathbf{F}$.

Definition 6. For each $\omega \in \Omega$, the function $t \mapsto X_t(\omega)$ mapping from $[0, \infty)$ to $\mathcal{B}(E)$ is called a *sample path* of the process X.

Definition 7 (Brownian motion)**.** Let $(\Omega, \mathscr{F}, \mathbf{F}, \mathsf{P})$ be a filtered probability space satisfying the usual conditions. An $\mathbb{R}$-valued, $\mathbf{F}$-adapted process $B = (B_t)_{t \geq 0}$ is a *(standard) Brownian motion* under P if

1) $B_0 = 0$, P-a.s.
2) (Independent increments) For $0 \leq t_0 < t_1 < \cdots < t_n$, the increments $B_{t_1} - B_{t_0}, \ldots, B_{t_n} - B_{t_{n-1}}$ are independent;
3) (Stationary Gaussian increments) For $s < t$, $B_t - B_s \sim \mathsf{N}(0,\, t - s)$;
4) (Continuity) $t \mapsto B_t(\omega)$ is continuous for all $\omega \in \Omega$.

Definition 8 (d-dimensional Brownian motion)**.** An $\mathbb{R}^d$-valued $\mathbf{F}$-adapted process $W = (W^{(1)}, \ldots, W^{(d)})^\top$ is an *d-dimensional (standard) Brownian motion* if its components are independent standard one-dimensional Brownian motions.

Definition 9 (Natural filtration generated by a process)**.** Let $X = (X_t)_{t\geq 0}$ be a stochastic process with state space E on $(\Omega, \mathscr{F}, \mathsf{P})$. The *natural filtration* generated by X is the family of σ-algebras $\mathbf{F}^X \triangleq \left(\mathscr{F}_t^X\right)_{t\geq 0}$, where

$$\mathscr{F}_t^X \triangleq \sigma(X_s : 0 \leq s \leq t)\,, \quad t \geq 0,$$

that is, the smallest σ-algebra making $\{X_s : 0 \leq s \leq t\}$ measurable. Moreover, $\mathbf{F}^X$ satisfies the usual conditions if it is right continuous and contains every subset of P-null set.

Definition 10. (càdlàg process) The process X is said to be càdlàg if its sample paths are right continuous a.s. with left limits.

Definition 11 (Hitting time)**.** Let X be a stochastic process. Suppose that $A \in \mathcal{B}(E)$. Then the term defined by

$$\tau(\omega) \triangleq \inf\{t > 0 : X_t(\omega) \in A\}$$

is called a *hitting time* of A for X.

Theorem 1 (Protter, 2010, Theorem I.3)**.** *Let X be an $\mathbf{F}$-adapted càdlàg process and suppose that $A \in \mathcal{B}(E)$ is an open set. Then the hitting time of A for X is a stopping time.*

Theorem 2 (Protter, 2010, Theorem I.5)**.** *Let τ and η be stopping times. Then $\tau \wedge \eta$, $\tau \vee \eta$, $\tau + \eta$ and $a\tau$ where $a > 1$, are also stopping times.*

2.2.2 MARTINGALES

Definition 12 (Martingale)**.** Let $(\Omega, \mathscr{F}, (\mathscr{F}_t)_{t\geq 0}, \mathsf{P})$ be a filtered probability space. An adapted process $M = (M_t)_{t\geq 0}$ with $\mathsf{E}[|M_t|] < \infty$ for all t is called

1) a *martingale* if
$$\mathsf{E}[M_t \mid \mathscr{F}_s] = M_s \quad \text{a.s. for all } 0 \leq s < t;$$

2) a *submartingale* if
$$\mathsf{E}[M_t \mid \mathscr{F}_s] \geq M_s \quad \text{a.s. for all } 0 \leq s < t;$$

3) a *supermartingale* if
$$\mathsf{E}[M_t \mid \mathscr{F}_s] \leq M_s \quad \text{a.s. for all } 0 \leq s < t.$$

Definition 13. A process M is called a *local martingale* if there exists an increasing sequence of stopping times $(\tau_n)_{n\in\mathbb{N}}$ with $\tau_n \uparrow \infty$ a.s. such that the stopped processes $M^{\tau_n} = (M_{t\wedge\tau_n})_{t\geq 0}$ are martingales for all n.

Definition 14. A martingale M is called a *uniformly integrable (UI) martingale* if

$$\sup_{t\geq 0} \mathsf{E}[|M_t|] < \infty \quad \text{and} \quad \lim_{K\to\infty} \sup_{t\geq 0} \mathsf{E}\left[|M_t| \mathbf{1}_{\{|M_t|>K\}}\right] = 0.$$

2.2.3 STOCHASTIC DIFFERENTIAL EQUATIONS

Definition 15 (Stochastic differential equation)**.** Let $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t\geq 0}, \mathsf{P})$ be a filtered probability space satisfying the usual conditions, and let $B = (B_t)_{t\geq 0}$ be a Brownian motion. Suppose that $\alpha : [0,\infty) \times E \to \mathbb{R}$ and $\sigma : [0,\infty) \times E \to \mathbb{R}$ be measurable functions. An $\mathbb{R}$-valued adapted process $X = (X_t)_{t\geq 0}$ is a solution of *stochastic differential equation* (SDE) if

$$X_t = X_0 + \int_0^t \alpha(s, X_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\,\mathrm{d}B_s, \qquad t \geq 0,$$

or in a concise differential form:

$$\mathrm{d}X_t = \alpha(t, X_t)\mathrm{d}t + \sigma(t, X_t)\mathrm{d}B_t, \quad X_0 = x \in E.$$

We distinguish two common categories:

1) *Time-homogeneous SDE.* If the coefficients depend only on the current state $x \in E$, i.e.

$$\alpha(s, x) = \alpha(x), \quad \text{and} \quad \sigma(s, x) = \sigma(x),$$

then X is called a *time-homogeneous SDE*.

2) *Time-inhomogeneous SDE.* If the coefficients depend explicitly on time as well as the state,

$$\alpha(s, x) \neq \alpha(x) \quad \text{or} \quad \sigma(s, x) \neq \sigma(x),$$

then X is called a *time-inhomogeneous SDE*.

In both cases, α is referred to as the *drift* and σ as the *diffusion coefficient*.

Theorem 3 (Existence and uniqueness of SDE: Øksendal, 2010, Theorem 5.2.1)**.** *Let* $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t\geq 0}, \mathsf{P})$ *be a filtered probability space, and let* $\alpha : [0,\infty) \times E \to \mathbb{R}$ *and* $\sigma : [0,\infty) \times E \to \mathbb{R}$ *be measurable functions. Suppose that for every* $T > 0$ *there exists a constant* $C > 0$ *such that for all* $t \in [0, T]$ *and all* $x, y \in E$*:*

1) Lipschitz condition:

$$|\alpha(t, x) - \alpha(t, y)| + |\sigma(t, x) - \sigma(t, y)| \leq C|x - y|;$$

2) Linear growth condition:

$$|\alpha(t, x)| + |\sigma(t, x)| \leq C\big(1 + |x|\big).$$

Then for any initial value $X_0 = x \in E$ *there exists a unique (up to indistinguishability) continuous adapted process* $X = (X_t)_{t\geq 0}$ *satisfying the SDE*

$$X_t = x + \int_0^t \alpha(s, X_s)\mathrm{d}s + \int_0^t \sigma(s, X_s)\,\mathrm{d}B_s, \quad t \in [0, T]$$

in $\mathbb{H}^2_T(\mathbf{F}, \mathsf{P})$*. This process is called the* unique strong solution *of the SDE.*

Definition 16 (Itō diffusion)**.** A process X is called an *Itō diffusion* (or simply a *diffusion*) if it solves an SDE which admits the unique strong solution.

Example 1 (Examples of Itō diffusion models)**.** *We list four fundamental examples of Itō diffusions which play central roles in stochastic control and mathematical economics.*

1) Arithmetic Brownian motion (ABM). *Let $\alpha \in \mathbb{R}$ and $\sigma > 0$. The process $X = (X_t)_{t\geq 0}$ satisfies*

$$\mathrm{d}X_t = \alpha\,\mathrm{d}t + \sigma\,\mathrm{d}B_t, \quad X_0 = x.$$

The explicit solution is

$$X_t = x + \alpha t + \sigma B_t.$$

2) Ornstein-Uhlenbeck (OU) process. *Let $\beta, \sigma > 0$ and $\tilde{x} \in \mathbb{R}$. The process $X = (X_t)_{t\geq 0}$ satisfies*

$$\mathrm{d}X_t = \beta(\tilde{x} - X_t)\,\mathrm{d}t + \sigma\,\mathrm{d}B_t, \quad X_0 = x.$$

Its explicit solution is

$$X_t = xe^{-\beta t} + \tilde{x}\big(1 - e^{-\beta t}\big) + \sigma \int_0^t e^{-\beta(t-s)}\,\mathrm{d}B_s.$$

3) Geometric Brownian motion (GBM). *Let $\mu \in \mathbb{R}$ and $\sigma > 0$. The process $X = (X_t)_{t\geq 0}$ satisfies*

$$\mathrm{d}X_t = \mu X_t\,\mathrm{d}t + \sigma X_t\,\mathrm{d}B_t, \quad X_0 = x > 0.$$

Its explicit solution is

$$X_t = x\exp\Big(\Big(\mu - \tfrac{1}{2}\sigma^2\Big)t + \sigma B_t\Big).$$

4) Kalman-Bucy filter (conditional mean of a linear Gaussian system). *Consider the hidden signal $\mathrm{d}\theta_t = a\theta_t\,\mathrm{d}t + \sigma_\theta\,\mathrm{d}B_t^\theta$ observed through $\mathrm{d}X_t = c\theta_t\,\mathrm{d}t + \sigma_X\,\mathrm{d}B_t^X$, where B^θ and B^X are independent Brownian motions. The conditional mean $M_t = \mathsf{E}[\theta_t \mid \mathscr{F}_t^X]$ satisfies the Kalman-Bucy SDE*

$$\mathrm{d}M_t = aM_t\,\mathrm{d}t + K_t\mathrm{d}\overline{B}_t, \quad \pi_0 = \mathsf{E}[\theta_0],$$

where the process $\overline{B}$ solves the SDE

$$\mathrm{d}\overline{B}_t = \mathrm{d}X_t - cM_t\,\mathrm{d}t,\ \overline{B}_0 = 0.$$

$\overline{B}$ is called the innovation process*, which is an $\mathscr{F}_t^X$-measurable Brownian motion. $K_t = \frac{cS_t}{\sigma_X^2}$ is called the Kalman gain and S_t is the conditional variance evolving by a Riccati equation. This is an example of time-inhomogeneous Itō diffusion with a deterministic diffusion coefficient depending on t through S_t. See, for example, (Øksendal, 2010, Chapter 6) or (Liptser & Shiryaev, 2013b, Chapter 12) for the detailed derivation of this problem.*

2.2.4 MARKOV PROPERTIES

Definition 17 (Markov property and transition kernel)**.** Let X be a stochastic process. We say that X is a *Markov process* if for all $0 \le s < t$ and bounded measurable $f : E \to \mathbb{R}$,

$$\mathsf{E}\big[f(X_t) \,\big|\, \mathscr{F}_s\big] = \mathsf{E}\big[f(X_t) \,\big|\, X_s\big] \quad \mathsf{P}\text{-a.s.}$$

The conditional distribution of X_t given $X_s = x$ is described by the *transition kernel*

$$P_{s,t}(x, A) \triangleq \mathsf{P}\big(X_t \in A \,\big|\, X_s = x\big), \qquad A \in \mathcal{B}(E).$$

We distinguish two types:

1) *Time-homogeneous Markov process.* There exists a one-parameter family of transition kernels $\{P_t\}_{t\ge 0}$ such that

$$P_{s,t}(x, A) = P_{t-s}(x, A), \qquad 0 \le s < t,$$

and hence

$$\mathsf{E}[f(X_t) \mid X_s] = (P_{t-s}f)(X_s), \quad (P_t f)(x) := \int_E f(y)\, P_t(x, \mathrm{d}y).$$

2) *Time-inhomogeneous Markov process.* There exists a two-parameter family $\{P_{s,t}\}_{0\le s<t}$ such that

$$\mathsf{E}[f(X_t) \mid X_s] = (P_{s,t}f)(X_s), \quad (P_{s,t} f)(x) := \int_E f(y)\, P_{s,t}(x, \mathrm{d}y).$$

Remark 2 (Transition kernels of Example 1 diffusions)**.** *Suppose that*

$$\varphi(y \mid \mu, \sigma^2) \triangleq \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(y-\mu)^2}{2\sigma^2}} \quad \textit{and} \quad \phi(y \mid \mu, \sigma^2) \triangleq \frac{1}{y\sqrt{2\pi\sigma^2}} e^{-\frac{(\log y-\mu)^2}{2\sigma^2}}$$

are probability density functions of normal and log-normal distributions, respectively. The transition kernels of the processes in Example 1 can be written explicitly by the followings.

1) Arithmetic Brownian motion (ABM). *For* $\mathrm{d}X_t = \alpha \mathrm{d}t + \sigma \mathrm{d}B_t$*, the kernel is*

$$P_t(x, dy) = \varphi(y \mid x + \mu t,\ \sigma^2 t)\mathrm{d}y.$$

2) Ornstein-Uhlenbeck (OU) process. *For* $\mathrm{d}X_t = \beta(\tilde{x} - X_t)\mathrm{d}t + \sigma \mathrm{d}B_t$*, the kernel is*

$$P_t(x, dy) = \varphi\left(y \mid xe^{-\beta t} + \tilde{x}(1 - e^{-\beta t}),\ \frac{\sigma^2}{2\beta}(1 - e^{-2\beta t})\right) \mathrm{d}y.$$

3) Geometric Brownian motion (GBM). *For* $\mathrm{d}X_t = \mu X_t \mathrm{d}t + \sigma X_t \mathrm{d}B_t$*,* X_t *is lognormal, i.e.,*

$$P_t(x, dy) = \phi\left(x \mid \log x + \left(\mu - \tfrac{1}{2}\sigma^2\right)t,\ \ \sigma^2 t\right) \mathrm{d}y.$$

4) Kalman-Bucy filter. *For* $\mathrm{d}M_t = aM_t\,\mathrm{d}t + K_t\mathrm{d}\overline{B}_t,\ K_t = \frac{S_t c}{\sigma_X^2}$,
the transition kernel is

$$P_{s,t}(x, dy) = \varphi(y \mid m_{s,t}(x),\, v_{s,t})\,\mathrm{d}y,$$

with mean $m_{s,t}(x)$ *and variance* $v_{s,t}$ *determined explicitly by* K_u *and* S_u *for* $u \in [s,t]$. *This is a time-inhomogeneous Gaussian Markov process.*

The first three processes are time-homogeneous diffusions, hence $P_{s,t}(x,\cdot) = P_{t-s}(x,\cdot)$, *while the Kalman-Bucy filter is time-inhomogeneous, with* $P_{s,t}$ *depending on both* s *and* t.

2.2.5 SEMIMARTINGALES

Definition 18 (Finite-variation process)**.** A càdlàg adapted process A has finite variation if

$$V_t(A) := \sup_{\pi\in\mathcal{P}} \sum_{t_i\in\pi} |A_{t_i} - A_{t_{i-1}}| < \infty, \tag{2.2.1}$$

for all t, where the supremum is over finite partitions $\mathcal{P}$ of $[0,t]$.

Definition 19 (Semimartingale)**.** A process X is a (continuous) semimartingale if $X = M + A$ with M a continuous local martingale and A a finite-variation process.

Definition 20 (Itō diffusions with singular control)**.** An *Itō diffusion with singular control* has the form

$$X_t = X_0 + \int_0^t \alpha(s, X_s)\,\mathrm{d}s + \int_0^t \sigma(s, X_s)\,\mathrm{d}B_s + A_t,$$

or in a differential form

$$\mathrm{d}X_t = \alpha(t, X_t)\mathrm{d}t + \sigma(t, X_t)\mathrm{d}B_t + \mathrm{d}A_t,\ X_0 = x \in E,$$

where α and σ are measurable functions satisfying the Lipschitz and linear growth conditions, and A is an $\mathbf{F}$-adapted finite-variation process (representing the singular control).

Remark 3 (Semimartingale property)**.** *Any Itō diffusion with singular control is a continuous semimartingale. Indeed, writing*

$$X_t = X_0 + \underbrace{\int_0^t \sigma(s, X_s)\,\mathrm{d}B_s}_{M_t} + \underbrace{\int_0^t \alpha(s, X_s)\,\mathrm{d}s + A_t}_{A'_t},$$

one can see that M *is a continuous local martingale and* A' *is an adapted finite-variation process. Thus* $X = X_0 + M + A'$ *is a semimartingale.*

Remark 4 (SDEs driven by semimartingales). *Let $Z = (Z_t)_{t\geq 0}$ be a semimartingale, and consider the SDE*

$$Y_t = Y_0 + \int_0^t \bar{\alpha}(s, Y_s)\, \mathrm{d}s + \int_0^t \bar{\beta}(s, Y_s)\, \mathrm{d}Z_s,$$

where $\bar{\alpha}, \bar{\beta}$ are measurable functions satisfying the usual Lipschitz and linear growth conditions. Then, by the general theory of stochastic integration with respect to semimartingales (see Protter, 2010, Theorem V.5), there exists a unique strong solution Y which is also a semimartingale.

In particular, since Itō diffusions with singular control are semimartingales (Remark 3), one can safely use them as state dynamics in stochastic control problems without leaving the class of well-posed SDEs.

Definition 21 (Quadratic (co)variation: Protter, 2010, Section II.6). Let $X = (X_t)_{t\geq 0}$ and $Y = (Y_t)_{t\geq 0}$ be càdlàg semimartingales on a filtered probability space $(\Omega, \mathscr{F}, \mathbf{F}, \mathsf{P})$.

1) *Quadratic variation.* The *quadratic variation* (or bracket process) of X, denoted $[X] = ([X]_t)_{t\geq 0}$, is the unique càdlàg, adapted, increasing process of finite variation such that

$$X_t^2 - [X]_t$$

is a local martingale. Equivalently, for a refining sequence of partitions $\pi^n = \{0 = t_0^n < t_1^n < \cdots < t_{k_n}^n = t\}$ of $[0, t]$ with mesh $|\pi^n| \to 0$,

$$[X]_t = \lim_{n\to\infty} \sum_{i=0}^{k_n - 1} (X_{t_{i+1}^n} - X_{t_i^n})^2,$$

where the limit is in probability.

2) *Quadratic covariation.* The *quadratic covariation* (or cross-variation) of X and Y, denoted $[X, Y] = ([X, Y]_t)_{t\geq 0}$, is the unique càdlàg, adapted, finite-variation process such that

$$X_t Y_t - [X, Y]_t$$

is a local martingale. Equivalently,

$$[X, Y]_t = \frac{1}{4}\Big([X + Y]_t - [X - Y]_t\Big),$$

or in Riemann-sum form,

$$[X, Y]_t = \lim_{n\to\infty} \sum_{i=0}^{k_n - 1} (X_{t_{i+1}^n} - X_{t_i^n})(Y_{t_{i+1}^n} - Y_{t_i^n}),$$

with convergence in probability.

Remark 5.

1) For a standard Brownian motion B, one has $[B]_t = t$.

2) Implying from (2.2.1), if A is a process of finite variation, then $[A] = 0$ and $[A, M] = 0$ for every local martingale M.

Theorem 4 (Itō's formula for continuous semimartingales: Protter, 2010, Theorem II.33)**.** *Let $X = (X_t)_{t \geq 0}$ be an $\mathbb{R}^m$-valued continuous semimartingale on $(\Omega, \mathscr{F}, \mathbf{F}, \mathsf{P})$ with the canonical decomposition*

$$X_t = X_0 + M_t + A_t,$$

where M is an $\mathbb{R}^m$-valued continuous local martingale and A is an $\mathbb{R}^m$-valued finite-variation adapted process. Let $f \in C^{1,2}([0,T] \times \mathbb{R}^m)$. Then for all $t \in [0,T]$,

$$\begin{aligned} f(t, X_t) = f(0, X_0) &+ \int_0^t \frac{\partial}{\partial s} f(s, X_s)\,\mathrm{d}s + \sum_{i=1}^m \int_0^t \frac{\partial}{\partial x_i} f(s, X_s)\,\mathrm{d}X_s^i \\ &+ \tfrac{1}{2} \sum_{i,j=1}^m \int_0^t \frac{\partial^2}{\partial x_i \partial x_j} f(s, X_s)\,\mathrm{d}[X^i, X^j]_s, \end{aligned}$$

where $[X^i, X^j]$ denotes the quadratic covariation of the i-th and j-th components of X. Equivalently, the differential form is

$$\mathrm{d}f(t, X_t) = \frac{\partial}{\partial t} f(t, X_t)\,\mathrm{d}t + \nabla f(t, X_t)^\top \mathrm{d}A_t + \nabla f(t, X_t)^\top \mathrm{d}M_t + \tfrac{1}{2}\,\mathrm{Tr}\left(D^2 f(t, X_t)\,\mathrm{d}[X, X]_t\right).$$

Remark 6 (Two-dimensional Itō diffusion with singular control)**.** *Let $X = (X^{(1)}, X^{(2)})^\top$ solve*

$$\mathrm{d}X_t = \alpha(t, X_t)\,\mathrm{d}t + \sigma(t, X_t)\,\mathrm{d}B_t + \mathrm{d}A_t,$$

where $\alpha : [0,T] \times \mathbb{R}^2 \to \mathbb{R}^2$, $\sigma : [0,T] \times \mathbb{R}^2 \to \mathbb{R}^{2 \times m}$, B is an m-dimensional Brownian motion, and A is an $\mathbb{R}^2$-valued process of finite variation.

Write the local martingale part $M_t := \int_0^t \sigma(s, X_s)\,\mathrm{d}B_s$ and the finite-variation part $A'_t := \int_0^t \alpha(s, X_s)\,\mathrm{d}s + A_t$. Then $X = M + A'$ with $[A'] = 0$ and $[A', M] = 0$. Hence the quadratic covariation matrix of X is

$$\mathrm{d}[X, X]_t = \mathrm{d}[M, M]_t = \sigma(t, X_t)\,\sigma(t, X_t)^\top\,\mathrm{d}t.$$

For $f \in C^{1,2}([0,T] \times \mathbb{R}^2)$, Itōs formula reads

$$\begin{aligned} \mathrm{d}f(t, X_t) &= \partial_t f(t, X_t)\,\mathrm{d}t + \nabla f(t, X_t)^\top \big(\alpha(t, X_t)\,\mathrm{d}t + \mathrm{d}A_t\big) \\ &\quad + \tfrac{1}{2}\,\mathrm{Tr}\big(\sigma\sigma^\top(t, X_t)\,D^2 f(t, X_t)\big)\,\mathrm{d}t + \nabla f(t, X_t)^\top \sigma(t, X_t)\,\mathrm{d}B_t. \end{aligned}$$

2.2.6 GIRSANOV THEOREM

Theorem 5 (Girsanov Theorem: Karatzas and Shreve, 1991, Theorem 5.1 in Chapter 3.5)**.** *Let $\theta \in \mathcal{S}^2$. Suppose that the process $\xi = (\xi_t)_{t\geq 0}$ solves the following SDE,*

$$\mathrm{d}\xi_t = -\theta\xi_t \mathrm{d}B_t,\ \xi_0 = 1.$$

If ξ is a martingale, then the measure Q defined via the Radon-Nikodym derivative $\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} |_{\mathscr{F}_T} = \xi_T$ makes the process

$$B_t^\theta \triangleq B_t + \int_0^t \theta_s \,\mathrm{d}s$$

a Brownian motion under Q.

2.3 FILTERING THEORY

In Chapters 4 and 5 we model ambiguity as arising from a *hidden variable* whose distribution is updated dynamically from noisy observations.

1) In Chapter 4, the ambiguity concerns a *static Bernoulli parameter* (flood vs. drought drift), updated through a Wonham filter. This corresponds to regime ambiguity, where the prior distribution is Bernoulli and the posterior evolves via Bayesian updating in continuous time.
2) In Chapter 5, the ambiguity is represented by a *Gaussian latent factor*, updated continuously through a Kalman-Bucy filter. This corresponds to distributional ambiguity about drift estimated from Gaussian signals.

Both settings can be understood through the lens of Bayesian filtering: the filter is the conditional law of the hidden state given the observation history. In discrete time, the classical recursive Bayes' rule (Bayes, 1763 and Laplace, 1812). In continuous time, Bayesian filtering appears as stochastic differential equations for the conditional law: in the linear-Gaussian case this is the Kalman-Bucy filter (Kalman & Bucy, 1961); for finite-state Markov chains observed in Gaussian noise, the Wonham filter (Wonham, 1965). Thus, the filtering problem provides the natural connection between Bayesian learning and stochastic control under ambiguity. In this thesis, we use Liptser and Shiryaev (2013a, 2013b) as standard references for the associated filtering theory.

We now recall the fundamental results for each case, presented in theorem form for clarity.

2.3.1 LINEAR-GAUSSIAN CASE: KALMAN-BUCY FILTER

Theorem 6 (Kalman-Bucy filter: Liptser and Shiryaev, 2013b, Theorem 12.1)**.** *Let (X, θ) be Itō diffusions in a filtered probability space $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t\geq 0}, \mathsf{P})$, satisfying*

$$\mathrm{d}X_t = (A_0(t, X_t) + A_1(t, X_t)\theta_t)\,\mathrm{d}t + C(t, X_t)\,\mathrm{d}B_t, \qquad \text{(observable process)}$$

$$\mathrm{d}\theta_t = (a_0(t, X_t) + a_1(t, X_t)\theta_t)\,\mathrm{d}t + c(t, X_t)\,\mathrm{d}W_t, \qquad \text{(unobservable process)}$$

where B, W are independent Brownian motions, and $C, c > 0$, P-a.s., for any $t \geq 0$. Suppose that $\mathbf{F}^X \triangleq (\mathscr{F}_t^X)_{t\geq 0}$ is a filtration generated by the observable process X, and $\theta_0 \sim \mathsf{N}(m, s)$ where $m \in \mathbb{R}$ and $s > 0$. Then the conditional law $\Pi_t = \mathsf{P}(\theta_t \mid \mathscr{F}_t^X)$ is Gaussian with mean M_t and variance S_t satisfying

$$\begin{aligned}
\mathrm{d}M_t &= (a_0(t, X_t) + a_1(t, X_t)M_t)\mathrm{d}t + K_t\mathrm{d}\overline{B}_t, && M_0 = 0,\\
K_t &\triangleq \frac{c(t, X_t)C(t, X_t) + S_tA_1(t, X_t)}{C^2(t, X_t)},\\
\mathrm{d}S_t &= \Big(2a_1(t, X_t)S_t + c^2(t, X_t) - (C(t, X_t)K_t)^2\Big)\mathrm{d}t, && S_0 = s.
\end{aligned}$$

where the process $\overline{B}$ is such that

$$\overline{B}_t = \int_0^t \frac{1}{C(s, X_s)}\Big(\mathrm{d}X_s - (A_0(s, X_s) + A_1(s, X_s)M_s)\Big)\,\mathrm{d}s, \tag{2.3.1}$$

which is an $\mathscr{F}_t^X$-measurable Brownian motion, and is called the innovation process *(of the Kalman-Bucy filter)*

Remark 7. *Notice from (2.3.1) that X solves*

$$\mathrm{d}X_t = (A_0(t, X_t) + A_1(t, X_t)M_t)\,\mathrm{d}t + C(t, X_t)\,\mathrm{d}\overline{B}_t,$$

and thus the pair (X, M) is now time-inhomogeneous Markov processes under $\mathbf{F}^X \triangleq \left(\mathscr{F}_t^X\right)_{t\geq 0}$.

2.3.2 FINITE-STATE CASE: WONHAM FILTER

Theorem 7 (Wonham filter: Liptser and Shiryaev, 2013a, Theorem 9.1)**.** *Let (X, θ) be stochastic processes in a filtered probability space $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t\geq 0}, \mathsf{P})$. The unobservable state $\theta_t \in \{1, \ldots, N\}$ is assumed to be a continuous-time Markov chain with generator $Q = (q_{ij})_{i,j}$. Suppose we observe an Itō diffusion X such that*

$$\mathrm{d}X_t = A(\theta_t, X_t)\,\mathrm{d}t + C(X_t)\,\mathrm{d}B_t,$$

where B is a Brownian motion, $A : \{1, \ldots, N\} \times E \to \mathbb{R}$ is the observation map, and $C > 0$, P-a.s., for any $t \geq 0$. Then the posterior probabilities $\Pi_t^i = \mathsf{P}(\theta_t = i \mid \mathscr{F}_t^X)$ satisfy the SDE

$$\mathrm{d}\Pi_t^i = \sum_{j=1}^N q_{ji}\Pi_t^j\,\mathrm{d}t + \Pi_t^i\frac{A(i, X_t) - \overline{A}(X_t)}{C^2(X_t)}\,\mathrm{d}\overline{B}_t, \qquad \Pi_0^i \in (0, 1),$$

where $\overline{A}(X_t) = \sum_{i=1}^N \Pi_t^i A(i, X_t)$, and the process $\overline{B}$ is such that

$$\overline{B}_t = \int_0^t \frac{1}{C(X_s)}\Big(\mathrm{d}X_s - \overline{A}(X_s)\Big)\,\mathrm{d}s, \tag{2.3.2}$$

which is an $\mathscr{F}_t^X$-measurable Brownian motion, and is called the innovation process *(of the Wonham filter).*

Remark 8. *One can see that (2.3.2) implies that X solves*

$$\mathrm{d}X_t = \overline{A}(X_t)\,\mathrm{d}t + C(t, X_t)\,\mathrm{d}\overline{B}_t,$$

and thus the tuple (X, Π^i), $i = 1, \ldots, N$, are now time-homogeneous Markov processes under $\mathbf{F}^X \triangleq \left(\mathscr{F}_t^X\right)_{t\geq 0}$.

Remark 9.

In our framework:

1) *Chapter 4 employs the two-state Wonham filter with no state switching (Bernoulli prior with $Q \equiv 0$) to represent ambiguity about drift regimes. The resulting posterior Π drives the ambiguity aggregator and control.*
2) *Chapter 5 employs the Kalman-Bucy filter to update a hidden constant ($\mathrm{d}\theta_t = 0$, P-a.s.) in models with continuous ambiguity. Here the belief process is fully characterized by (M, S).*

2.4 FORWARD-BACKWARD STOCHASTIC DIFFERENTIAL EQUATIONS

Forward-backward stochastic differential equations (FBSDEs) provide a powerful probabilistic representation of nonlinear PDEs and variational inequalities, and play a central role in stochastic control under ambiguity. In particular, the value function of a singular control problem can often be characterized both by a stochastic formulation in terms of BSDEs and by a deterministic formulation in terms of HJB equation. This connection is made precise by the nonlinear Feynman-Kac formula of Pardoux and Peng (1992), which links BSDEs with solutions of nonlinear PDEs. The BSDE perspective is especially useful when ambiguity is modeled via a change of measure, while the HJB perspective provides the foundation for the free-boundary and viscosity solution methods used later in the thesis. We recall here the well-posedness results under standard Lipschitz conditions (to be applied in Chapter 4) and, in the quadratic growth condition (to be applied in Chapter 5), thereby connecting the stochastic and deterministic representations of the control problem. The associated monograph used for this thesis belongs to Zhang (2017).

2.4.1 GENERAL FORM OF AN FBSDE

Fix $T > 0$. Let $(\Omega, \mathscr{F}, \mathbf{F} \triangleq (\mathscr{F}_t)_{t\in[0,T]}, \mathsf{P})$ be a filtered probability space and let $B = (B_t)_{t\in[0,T]}$ be a d-dimensional Brownian motion. Consider the system

$$\mathrm{d}X_t = \alpha(t, X_t)\,\mathrm{d}t + \sigma(t, X_t)\,\mathrm{d}B_t, \qquad X_0 = x \tag{2.4.1}$$

$$-\mathrm{d}Y_t = f(t, X_t, Y_t, Z_t)\,\mathrm{d}t - Z_t\,\mathrm{d}B_t, \qquad\qquad Y_T = \xi, \tag{2.4.2}$$

for $t \in [0, T]$, where:

- (Y, Z) is the pair of adapted processes, with $Y = (Y_t)_{t\in[0,T]}$ real-valued (or vector-valued) and $Z = (Z_t)_{t\in[0,T]}$ $\mathbb{R}^d$-valued, solving the backward SDE with driver $f : [0, T] \times \mathbb{R}^m \times \mathbb{R} \times \mathbb{R}^d \to \mathbb{R}$;
- ξ is the terminal condition, $\mathscr{F}_T$-measurable, typically depending on X_T.

2.4.2 Existence and uniqueness under Lipschitz conditions

The classical result is the following.

Theorem 8 (Lipschitz BSDEs: Pardoux and Peng, 1990)**.** *Assume:*

1) *α, σ are measurable functions satisfying the usual Lipschitz and linear growth conditions;*
2) *the driver f is Lipschitz continuous in (y, z) uniformly in (t, x), i.e., there exists $L > 0$ such that*
$$|f(t, x, y, z) - f(t, x, y', z')| \;\le\; L\big(|y - y'| + |z - z'|\big);$$
3) *$\xi \in \mathbb{L}^2(\mathscr{F}_T, \mathsf{P})$.*

Then the FBSDE system (2.4.1)-(2.4.2) admits a unique adapted solution (X, Y, Z) with $X \in \mathbb{L}^2_T(\mathbf{F}, \mathsf{P})$, $Y \in \mathbb{L}^2_T(\mathbf{F}, \mathsf{P})$, $Z \in \mathbb{H}^2_T(\mathbf{F}, \mathsf{P})$.

Remark 10.

1) *For the extension of Theorem 8 with singular control, see, e.g., El Karoui, Peng, and Quenez (1997).*
2) *In the maxmin utility framework, Chen and Epstein (2002) show that when the density generator takes the form of κ-ignorance, the associated BSDE driver is Lipschitz, namely $f(z) = -\kappa|z|$. This framework is adopted to solve problems in Chapter 3.*
3) *Similarly, as we demonstrate in Chapter 4, a smooth ambiguity problem with a Bernoulli prior also yields a value function characterized by a BSDE with a Lipschitz-type driver.*

2.4.3 Existence and uniqueness under quadratic growth

In many applications, such as robust control (Hansen & Sargent, 2001), utility maximisation in incomplete markets (Hu et al., 2005; Mania & Schweizer, 2005), or settings involving smooth ambiguity with a Gaussian prior in Hansen and Sargent (2011) and Chapter 5, the driver f exhibits quadratic growth in z. In such cases, standard Lipschitz theory no longer applies. However, existence and uniqueness can still be established under suitable structural conditions.

Theorem 9 (Quadratic BSDEs: Kobylański, 2000). *Suppose that:*

1) α, σ *are measurable functions satisfying the usual Lipschitz and linear growth conditions;*

2) f *is continuous in* (y, z) *and convex in* z*;*

3) f *satisfies the quadratic growth condition in* z*, i.e., there exists* $C > 0$ *such that*

$$
\begin{aligned}
|f(t,x,y,z)| &\le C(1+|y|+|z|^2) \\
|f(t,x,y,z) - f(t,x,y',z')| &\le C\Big(|y-y'| \\
&\qquad + (1+|y|+|y'|+|z|+|z'|)|z-z'|\Big).
\end{aligned}
$$

4) $\xi \in \mathbb{L}^\infty(\mathscr{F}_T, \mathsf{P})$.

Then the backward equation (2.4.2) *admits a unique solution* (Y, Z) *with* $Y \in \mathbb{L}^\infty_T(\mathbf{F}, \mathsf{P})$ *and* $Z \in \mathbb{H}^2_T(\mathbf{F}, \mathsf{P})$*. In particular,* (X, Y, Z) *form a well-defined FBSDE system* (2.4.1)-(2.4.2).

Remark 11. *See, Kobylański et al. (2002), for the treatment of FBSDEs with quadratic growth and singular control.*

— 3 —

Optimal Cash Management under Ambiguity: A Singular Control Model with Maxmin Utility

Abstract

We consider a singular control model of cash reserve management, driven by a diffusion under ambiguity. The manager is assumed to have maxmin preferences over a set of priors characterized by κ-ignorance. A verification theorem is established to determine the firm's cost function and the optimal cash policy; the latter taking the form of a control barrier policy. In a model driven by arithmetic Brownian motion, we use Dynkin games to show that an increase in ambiguity leads to higher expected costs under the worst-case prior and a narrower inaction region. The latter effect can be used to provide an ambiguity-driven explanation for observed cash management behavior. Our findings can be applied to broader applications of singular control in managing inventories under ambiguity.
Keywords: Singular control, Ambiguity, Inventory models

3.1 INTRODUCTION

An important question in corporate finance is that of optimal cash management. On the one hand, firms require cash to finance the firm as a going concern. On the other hand, shareholders require dividend payouts as a reward for providing capital. The seminal contribution by Jeanblanc-Picqué and Shiryaev (1995) uses a stochastic storage models *à la* Harrison and Taksar (1983) to find the optimal size of a firm's cash hoard in the face of stochastically evolving net cash flows. In this paper, we are interested in optimal cash management under ambiguity, i.e., a situation where the manager is not able to reduce the uncertainty over future net cash flows into a single probability measure. We are interested in the interplay between traditional concerns over risk (as measured by, e.g., confidence intervals provided by a given probability measure) and ambiguity (as measured by the "size" of the set of probability measures considered by the manager) under the assumption that the manager is ambiguity averse. Furthermore, we demonstrate that our results apply more broadly to singular control of inventories under ambiguity.

Our motivation for incorporating ambiguity into the singular control framework of cash holding stems from empirical evidence that ambiguity exerts a first-order effect on corporate cash management. A particularly relevant study is that of Breuer et al. (2017), which investigates how firms adjust their cash policies in response to investors' ambiguity aversion. They find that ambiguity averse investors tend to undervalue uncertain future investment opportunities, thereby diminishing the perceived value of holding cash for such purposes. In turn, firm managers respond by actively reducing cash holdings, often through dividend payouts, to align with investor preferences and preserve firm value. This adjustment is not a passive by-product of other financial decisions but, rather, a direct policy response to ambiguity.

While a broader empirical literature links ambiguity to firms financial decisions, much of it centers on investment behavior, with cash policy treated as a secondary outcome. For example, studies such as Goodell et al. (2021), Luo and Tian (2022), Neamtiu et al. (2014), and Zhao et al. (2023) document how macroeconomic ambiguity, arising from such as inflation, GDP volatility, unemployment, or economic policy uncertainty, affects managerial incentives and capital expenditures, with subsequent implications for cash holdings. In these cases, ambiguity affects cash indirectly through its impact on investment timing, implying a second-order effect.

From a theoretical perspective, such dynamics are well captured by models of investment under ambiguity, notably the real options framework developed by Nishimura and Ozaki (2007). However, existing models largely abstract from the possibility that cash itself may be the principal adjustment mechanism to ambiguity. To fill this gap, our paper develops a stochastic control model of cash management under ambiguity, where cash holding responds directly and optimally to ambiguity. This allows us to characterize the first-order role of ambiguity in shaping firms' cash policies, and to explore the implications for corporate financial behavior in uncertain environments.

In this paper, we develop a framework in which a firms decision maker (DM) can dynamically adjust cash holdings, where holding cash is costly and cash adjustments may incur additional costs. We assume that the DM has a *reference prior*, possibly elicited from available data or based on their industry experience, but is *ambiguous* about the true probability measure. The DM then dynamically uses the worst-case prior to determine the optimal cash policy.

The distinction between uncertainty resulting from randomness governed by a distribution ("known unknowns") and uncertainty over the correct distribution ("unknown unknowns") goes back to Knight (1921). In his seminal work he refers to the former as *risk* and the latter as *uncertainty* or *ambiguity*. The effect of ambiguity on decision making has been studied extensively, most famously by Ellsberg (1961). The overwhelming conclusion of the experimental literature is that DM are *ambiguity averse*. In the classical Ellsberg experiment, a DM has to place bets on one of two urns, both with 100 red or blue balls. For the first urn it is known that half the balls are red. For the second urn no such information is available. Since most people are observed to choose bets on the first urn over bets on the second urn, Savage's "sure thing principle" is being violated. When compared to previous empirical findings, this experiment can be interpreted as if the firms' manager are confronted with the second Ellsberg urn. This, as a result, prompts them to assign their subjective priors.

Note that the Ellsberg paradox is not really a paradox, because it does not result from a cognitive bias or irrationality. Rather, observed behavior is driven by a lack of information. It is perfectly possible for DMs to make consistent decisions under ambiguity. This has been shown by Gilboa and Schmeidler (1989), who incorporate an ambiguity aversion axiom into the subjective expected utility framework. They then show that a rational DM acts *as if* she maximizes expected utility over the worst–case prior within a (subjectively chosen) set of priors. This approach has been successfully used in many applications in economics, finance, and operation research.[1]

Our contribution is to apply the maxmin multiple prior model to a singular control model of optimal storage inventory, with an application to a firm's cash management. On a regular basis, firms are faced with operational costs (e.g. rent, capital stock, labor's wage, etc.) that have to be settled promptly with reserved cash. The fact that this cash generates no (or low) return means holding it results in an opportunity loss, which can be interpreted as a holding cost as it could potentially be used for income-generating activities, such as investments or paying out dividends. Therefore, excessive cash holding is undesirable. On the other hand, having a shortage of cash reserves results in a delay of cost settlement, which often incurs a penalty fee or credit loss. Therefore, the firm has an incentive to inject some amount of cash into the system. This could, for example, be done by selling some assets or issuing bonds. These two circumstances create a trade-off that suggests the existence of target level of cash. In a model where cash adjustments are costly, we show that there exists an optimal *control band policy*, where the firm keeps its cash hoard between an upper and lower bound. The cash reserve problem was first addressed in the literature by Baumol (1952) and Tobin (1956), who studied the cash balance problem under the assumption that demand is deterministic, which is far from realistic. The stochastic treatment was later established under a discrete-time (Markov chain) framework by e.g. Eppen and Fama (1969). A more general approach for storage systems in continuous time, in particular, with demand driven by Brownian motion, has been developed over the past decades. Bather (1966), Constantinides

[1] See, for example of applications in investment decisions, Cheng and Riedel (2013), Hellmann and Thijssen (2018), Nishimura and Ozaki (2007), Thijssen (2011), and Trojanowska and Kort (2010) delve into timing game under ambiguous environment, while Asano and Osaki (2021) and Driouchi et al. (2020) incorporate ambiguity into the model by means of technology shocks, and cultural biases, respectively. The works of Fouque et al. (2016), Jin and Yu Zhou (2015), and Lin and Riedel (2021) apply ambiguity to portfolio management. For the broader theory of ambiguity in volatility and interest rate in asset pricing, we refer to Epstein and Ji (2013) and Lin and Riedel (2021), respectively. For the literature related to decision making under *smooth ambiguity*, which is another approach to model of ambiguity introduced by Klibanoff et al. (2005), we refer, for example, to the work of Balter and Pelsser (2020), Balter et al. (2021), Borgonovo and Marinacci (2015), Hansen and Miao (2018, 2022), and Suzuki (2018), among other notable authors.

(1976), Dai and Yao (2013a, 2013b), Harrison (1978), Harrison and Taksar (1983), and Vial (1972) and many others are among the notable authors. To get an overview of the related papers, we refer the reader to Harrison (2013).

While this is no different from a standard model under risk, ambiguity does bring some new aspects to the comparative statics of the optimal policy. For example, as in the standard model without ambiguity, the higher the risk, the higher the long-term discounted cost of cash. Ambiguity amplifies this effect, even though an increase in the degree of ambiguity leads the manager to exert control *earlier*. This is in contrast to the risk-only model, where an increase in risk leads the manager to exert control *later*.

The reason for this result is that a more ambiguous DM expects the cash level to increase (when positive) or decrease (when negative) more rapidly (in expectation) than a less ambiguous DM. Since holding costs are increasing in the absolute value of the cash hoard, a more ambiguous DM will, thus, exert control sooner. In our model, this behavior is not due to irrationality, but an aspect of the uncertain environment that the manager faces.

One of the first papers to axiomatize ambiguity is Gilboa and Schmeidler (1989). They model ambiguity as a set of priors, among which the DM (subjectively) selects the one that maximizes the DM's expected utility. Under an axiom of ambiguity aversion, the prior that is chosen is called the *worse-case prior*, which captures the intuition that an ambiguity-averse DM is cautious about their beliefs and heavily weighs the possibility of undesirable consequences of their decision. The Gilboa–Schmeidler criterion has become known as *maxmin utility*. However, the Gilboa–Schmeidler framework is a static one and is, thus, insufficient for dealing with situations where the worst-case prior might change over time. An inter-temporal version was proposed by Epstein and Wang (1994) in discrete time and by Chen and Epstein (2002) in continuous time. In these models, the worst-case prior is updated in a Bellman principle-like one-step-ahead procedure. In order to make this work, attention is restricted to sets of priors that are called *strongly rectangular*. We use the Chen and Epstein (2002) approach to modeling multiple priors.

In fact, we use a stronger assumption, also introduced in Chen and Epstein (2002), and assume that ambiguity takes the form of κ-*ignorance*. That is to say, the DM has a reference probability, which is distorted through a density generator. The density generator is assumed to take values in an interval $[-\kappa, \kappa]$, so that the reference prior together with the parameter κ determines the set of priors that is considered by the DM. While restrictive, an advantage of this approach is that the degree of ambiguity can be seen to be measured by κ.

Importantly, in our model the worst-case prior is not constant but varies over time, depending on the evolution of the actual amount of cash currently held. This unusual feature has been observed by Cheng and Riedel (2013) in the context of pricing a straddle option and Hellmann and Thijssen (2018). The latter paper models a timing game between two firms contemplating an investment opportunity under ambiguity and show that ambiguity aversion has two effects: ambiguity over future demand (fear of the market), as in the standard literature, but also ambiguity over the other firm's investment decision (fear of the competitor). These have opposite effects on what constitutes the worst-case prior. It turns out there is a threshold to distinguish which type of ambiguity dominates through time. Our model has a similar feature in that control costs are incurred whether at the upper or lower barrier. The worst-case priors at each of these barriers are opposite and this leads, in turn, to the existence of a threshold somewhere in the inaction region (endogenously determined) that separates two regions where different measures constitute the worst-case prior.

The closest contribution to our work is Chakraborty et al. (2023) in which a one-side singular control of a firm's dividend payout policy is considered under ambiguity. They assume, in addition to the classical singular control, that there is a penalty cost associated with a change of measure, which is determined by the Kullback-Leibler divergence. The use of Kullback-Leibler divergence as a model for multiple priors is well established in the literature on robust control; see, e.g., Anderson et al. (2003), Ferrari et al. (2022), Hansen and Sargent (2010), Hansen and Miao (2018, 2022), Hansen and Sargent (2011), Hansen et al. (2006), and Maenhout (2004) and references therein. The more behavioral approach that motivates κ-ignorance is, in fact, closely related to the robust control approach. In both cases, the solution to the control problem takes the form of a control band policy. However, the robust control of approach of Chakraborty et al. (2023) gives rise to a nonlinear Bellman equation, which poses significant challenges for analytically deriving comparative statics results. In contrast, our model admits analytical results by reformulating the classical singular control problem under κ-ignorance into a Dynkin game, as detailed in Section 3.4. To the best of our knowledge, this is the first analytical exploration of comparative statics for singular control under ambiguity.

It is important to recognize that while our model is motivated by ambiguity in cash management, it can be easily adapted to models of singular control for other applications. Essentially, our model allows the inventory process to take on various diffusion types, as outlined in Section 3.4. For instance, while we focus on cash management with a simple arithmetic Brownian motion diffusion, one could apply the same concept to managing, for example, a firm's stock of outstanding shares, where the market maker continuously adjusts shares to maintain a desired price trajectory. This type of inventory process exhibits mean-reversion characteristics (cf. Cadenillas et al., 2010). Furthermore, our model can extend to non-monetary domains such as dam reservoir control, as demonstrated by Jiang et al. (2022), where reservoir levels exhibit seasonal mean reversion. Here, ambiguity can be addressed by incorporating seasonal pattern uncertainty, potentially influenced by climate change. Thus, our model represents a general contribution to singular control with maxmin preference ambiguity.

The structure of this paper is as follows: In Section 3.2 we construct a general formulation for singular control of the Brownian cash reserve under ambiguity. We provide a verification theorem for the optimal control band policy and the existence of the ambiguity trigger in Section 3.3. In Section 3.4 we provide a simplification of the verification theorem and the associated Dynkin game for the case where the present value of the (uncontrolled) expected holding costs is affine in the current value of the cash holdings. This includes, e.g., the case where the uncontrolled cash process follows an arithmetic Brownian motion, or a mean-reverting Ornstein-Uhlenbeck process. A theoretical comparative static analysis for the arithmetic Brownian motion case is given in Section 3.5.

3.2 SIMPLE CASH MANAGEMENT MODEL WITH DRIFT AMBIGUITY

Let $E \subseteq \mathbb{R}$ be a connected state space endowed with the Euclidean topology and such that $0 \in E$. Given $(\Omega, \mathscr{F}, \mathsf{P})$ a complete probability space. On $(\Omega, \mathscr{F}, \mathsf{P})$, we assume that $\alpha : E \to \mathbb{R}$ and $\sigma : E \to \mathbb{R}$ are continuously differentiable functions such that for all $x \in E$

$$\begin{aligned} &\sigma'(x) \text{ is locally Lipschitz continuous} \\ &|\alpha(x)| + |\sigma(x)| \leq C(1 + |x|) \end{aligned} \tag{3.2.1}$$

for some $C > 0$. Then a time-homogeneous diffusion, $X \triangleq (X_t)_{t \geq 0}$, taking values in E, is the unique strong solution to the stochastic differential equation (SDE),

$$\mathrm{d}X_t = \alpha(X_t)\mathrm{d}t + \sigma(X_t)\mathrm{d}B_t, \quad \mathsf{P}(X_0 = x) = 1,$$

where $B \triangleq (B_t)_{t \geq 0}$ is a standard Brownian motion. Dynamic revelation of information is modeled by the natural filtration $\mathbf{F} = (\mathscr{F}_t)_{t \geq 0}$ generated by X, assumed to satisfy the usual conditions. We assume that the end points of E are P-a.s. unattainable, given $\mathsf{P}(X_0 = x) = 1$. For brevity of notation we write $\mathsf{P}_x(\cdot) \triangleq \mathsf{P}(\cdot | X_0 = x)$, associated with an expectation operator E_x.

A *control policy* is a pair of processes (L, U), where L and U are adapted, non-decreasing, and non-negative. These processes are associated with increases and decreases, respectively, of X at times at which control is exerted. With the policy (L, U) we associate the *controlled process* $X^{L,U}$ and we say that a control policy (L, U) is *feasible* if for all $x \in E$, there exists a unique $X^{L,U}$ that strongly solves

$$\mathrm{d}X_t^{L,U} = \alpha(X_t^{L,U})\mathrm{d}t + \sigma(X_t^{L,U})\mathrm{d}B_t + \mathrm{d}L_t - \mathrm{d}U_t, \;\; X_0 = x, \;\; \mathsf{P}_x\text{-a.s.},$$

and if there exist $A > 0$ and $B < 0$, such that

$$\mathsf{P}_x\left(\sup_{t \geq 0} X_t^{L,U} < A, \inf_{t \geq 0} X_t^{L,U} > B\right) = 1 \tag{3.2.2}$$

The set of feasible control policies is denoted by $\mathscr{D}$, while we denote by X^0, the uncontrolled process; that is, $X^0 \triangleq X^{0,0}$.

The instantaneous holding costs are given by an almost everywhere differentiable function $c : \mathbb{R} \to \mathbb{R}_+$. For simplicity we will assume that

$$c(x) = \begin{cases} \hat{c}|x| & \text{if } x \geq 0 \\ \check{c}|x| & \text{if } x < 0. \end{cases}$$

for some $\hat{c}, \check{c} > 0$. The instantaneous and proportional costs of lower and upper control are denoted by $\ell > 0$ and $u > 0$, respectively. Our results can easily be extended to more general convex holding costs with $c(0) = 0$, albeit at the cost of more cumbersome notation. In a cash management setting, one could think of 0 as the *target level* of cash. When $x > 0$, the firm has excess cash while if $x < 0$ the firm needs to access cash on the markets. When the cash reserves get too low the firm may need to issue new equity, which incurs costs ℓ, whereas when x gets too large, the firm may wish to pay out dividends, which incurs a cost u.

The DM discounts costs at the constant rate $\rho > 0$. We, furthermore, assume that

$$\mathsf{E}_x \left[\int_0^\infty e^{-\rho t} |X_t^0| \mathrm{d}t \right] < \infty, \quad x \in E.$$

A typical process that satisfies all the assumptions made so far is the arithmetic Brownian motion (ABM), defined on the state space $E = \mathbb{R}$, being the strong solution of the SDE

$$\mathrm{d}X_t^0 = \alpha \mathrm{d}t + \sigma \mathrm{d}B_t,$$

with constant drift $\alpha \in \mathbb{R}$ and standard deviation $\sigma > 0$. For this specification the uncontrolled cash process is

$$X_t^0 = x + \alpha t + \sigma B_t,$$

whereas for any feasible control policy $(L, U) \in \mathscr{D}$, the controlled cash process satisfies

$$X_t^{L,U} = x + \alpha t + \sigma B_t + L_t - U_t.$$

Another process that can be used is the mean-reverting Ornstein-Uhlenbeck (OU) process

$$\mathrm{d}X_t^0 = -\beta X_t^0 \mathrm{d}t + \sigma \mathrm{d}B_t,$$

where $\beta > 0$ is the speed of mean-reversion. In this case

$$X_t^0 = x e^{-\beta t} + \sigma \int_0^t e^{-\beta(t-s)} \mathrm{d}B_s.$$

It is assumed that the DM is *ambiguous* about the measure P_x and, consequently, considers a set of priors $\mathscr{P}^\Theta$. Each of these priors is constructed from the reference measure P_x by means of a *density generator* $\theta \in \Theta$. A process $\theta = (\theta_t)_{t \geq 0}$ is a density generator if the process $\left(M_t^\theta\right)_{t \geq 0}$, with

$$\mathrm{d}M_t^\theta = -\theta_t M_t^\theta \mathrm{d}B_t, \quad M_0^\theta = 1,$$

is a P_x-martingale. Such a process θ generates a new measure P_x^θ on $(\Omega, \mathscr{F}^B)$ via the Radon–Nikodym derivative $\frac{\mathrm{d}\mathsf{P}_x^\theta}{\mathrm{d}\mathsf{P}_x} \mid \mathscr{F}_T^B = M_T^\theta$ for any $T > 0$. Here, $\mathscr{F}^B \triangleq \mathscr{F}_\infty^B$, where $\mathbf{F}^B \triangleq \left(\mathscr{F}_t^B\right)_{t\geq 0}$ is the (uncompleted) filtration generated by B. Indeed, if $\theta \in \Theta$, then it follows from Girsanov's theorem (see, Corollary 5.2 in Chapter 3.5 of Karatzas and Shreve, 1991) that under the measure P_x^θ the process $B^\theta \triangleq \left(B_t^\theta\right)_{t\geq 0}$, defined by

$$B_t^\theta \triangleq B_t + \int_0^t \theta_s \mathrm{d}s, \tag{3.2.3}$$

is a Brownian motion on $(\Omega, \mathscr{F}^B, \mathbf{F}^B, \mathsf{P}_x^\theta)$ and that, under P_x^θ, the process $X^{L,U,\theta}$ is the unique strong solution to the SDE

$$\begin{aligned}\mathrm{d}X_t^{L,U,\theta} &= \alpha(X_t^{L,U,\theta})\mathrm{d}t + \sigma(X_t^{L,U,\theta})\mathrm{d}B_t^\theta + \mathrm{d}L_t - \mathrm{d}U_t \\ &= \left(\alpha(X_t^{L,U,\theta}) + \sigma(X_t^{L,U,\theta})\theta_t\right)\mathrm{d}t + \sigma(X_t^{L,U,\theta})\mathrm{d}B_t + \mathrm{d}L_t - \mathrm{d}U_t, \quad \mathsf{P}_x\text{-a.s.}\end{aligned}$$

In the remainder we restrict attention to so-called *κ-ignorance*, i.e., we only use density generators θ for which $\theta_t \in [-\kappa, +\kappa]$ for all $t \geq 0$ and some $\kappa \geq 0$. Note that $\mathscr{P}^\Theta = \{\mathsf{P}_x\}$ if $\kappa = 0$.

To model *ambiguity aversion*, it is assumed that the DM uses maxmin utility *à la* Gilboa and Schmeidler (1989). That is, the *worst-case cost function* associated with the feasible policy $(L, U) \in \mathscr{D}$ is given by $J^{L,U} : E \to \mathbb{R}$, where

$$J^{L,U}(x) \triangleq \sup_{\theta\in\Theta} \mathsf{E}_x^\theta \left[\int_0^\infty e^{-\rho t} \left(c(X_t^{L,U,\theta})\mathrm{d}t + \ell \mathrm{d}L_t + u\mathrm{d}U_t\right)\right].$$

The DM's objective is to find the feasible policy that minimizes the worst-case expected costs over the set of priors $\mathscr{P}^\Theta$. The firm's *minimal cost function* is

$$V(x) \triangleq \inf_{(L,U)\in\mathscr{D}} J^{L,U}(x).$$

From Chen and Epstein (2002, Theorem 2.1) it follows that there exists an *upper-rim generator* $\theta^* \in \Theta$ so that

$$J^{L,U}(x) = \mathsf{E}_x^{\theta^*} \left[\int_0^\infty e^{-\rho t} \left\{c(X_t^{L,U,\theta^*})\mathrm{d}t + \ell \mathrm{d}L_t + u\mathrm{d}U_t\right\}\right].$$

Furthermore, from Chen and Epstein (2002, Section 3.3) it follows that under κ-ignorance it holds that $\theta_t^* \in \{-\kappa, \kappa\}$ for all $t \geq 0$.

Finally, in many cases the optimal policy consists of exerting control only when the process X exits an interval $(\underline{x}, \overline{x})$. Therefore, with each pair $(\underline{x}, \overline{x}) \in E \times E$, $\underline{x} < \overline{x}$, we associate the *control band policy* $(L, U) \in \mathscr{D}$ for which $\underline{x}$ is an (upward) reflecting barrier for L and $\overline{x}$ is a (downward) reflecting barrier for U. For such policies it holds that

1) $X_t^{L,U,\theta^*} \in [\underline{x}, \overline{x}]$, P_x-a.s. for all $t \geq 0$, and
2) $\int_0^\infty 1_{(\underline{x},\overline{x})}(X_t^{L,U,\theta^*})\mathrm{d}(L_t + U_t) = 0$, P_x-a.s.

Following Tanaka (1979), our assumptions on X are sufficient to guarantee the existence of control band policies.

3.3 A GENERAL VERIFICATION THEOREM

Let $\mathscr{L}$ denote the characteristic operator on $C^2(E)$ of the killed process $\big(e^{-\rho t}X_t\big)_{t\geq 0}$ under P_x, i.e.

$$\mathscr{L}\varphi(x) = \frac{1}{2}\sigma^2(x)\varphi''(x) + \alpha(x)\varphi'(x) - \rho\varphi(x).$$

On $C^1(E)$ we also define the density generator,

$$\theta^\varphi \triangleq \theta^\varphi(x) \triangleq \begin{cases} +\kappa & \text{if } \varphi'(x) \geq 0 \\ -\kappa & \text{if } \varphi'(x) < 0. \end{cases}$$

We get the following verification theorem.

Theorem 10. *Suppose there exists a pair $(\underline{x}, \overline{x}) \in E \times E$, $\underline{x} < 0 < \overline{x}$, and a non-negative, convex, and C^2-function φ on $(\underline{x}, \overline{x})$ such that*

1) $\mathscr{L}\varphi(x) + \theta^\varphi(x)\sigma(x)\varphi'(x) + c(x) = 0$ *on* $(\underline{x}, \overline{x})$,
2) $\varphi'(\underline{x}+) = -\ell$, $\varphi'(\overline{x}-) = u$,
3) $\varphi''(\underline{x}+) = \varphi''(\overline{x}-) = 0$,
4) $\check{c}(\underline{x} - x) \geq \ell\big[\rho(\underline{x} - x) - \big(\alpha(\underline{x}) - \alpha(x)\big) - \kappa\big(\sigma(\underline{x}) - \sigma(x)\big)\big]$, *for all* $x < \underline{x}$,
5) $\hat{c}(x - \overline{x}) \geq u\big[\rho(x - \overline{x}) - \big(\alpha(x) - \alpha(\overline{x})\big) + \kappa\big(\sigma(x) - \sigma(\overline{x})\big)\big]$, *for all* $x > \overline{x}$, *and*
6) $\lim_{T\to\infty} e^{-\rho T}\mathsf{E}_x^{\theta^\varphi}\Big[\varphi\big(X_T^{L,U}\big)\Big] = 0$, *for all* $(L, U) \in \mathscr{D}$.

Then the optimal policy (L^*, U^*) is the control band policy associated with $(\underline{x}, \overline{x})$ and the minimal cost function is

$$V(x) = \begin{cases} \ell|\underline{x} - x| + \varphi(\underline{x}+) & \text{if } x \leq \underline{x} \\ \varphi(x) & \text{if } \underline{x} < x < \overline{x} \\ u|x - \overline{x}| + \varphi(\overline{x}-) & \text{if } x \geq \overline{x} \end{cases}.$$

Remark 12. *The optimal control policy (L^*, U^*) is such that no control takes place as long as the cash stock is between $\underline{x}$ and $\overline{x}$. This implies that the processes L^* and U^* remain constant. Whenever the lower bound $\underline{x}$ is hit, the processes L^* increases to keep the cash stock at $\underline{x}$. Similarly, whenever the lower bound $\overline{x}$ is hit, the processes U^* increases to keep the cash stock at $\overline{x}$.*

Remark 13. *Conditions 4 and 5 guarantee existence of a feasible policy under ambiguity. For the case of an uncontrolled ABM,*

$$\mathrm{d}X_t^0 = \alpha\mathrm{d}t + \sigma\mathrm{d}B_t,$$

these conditions reduce to

$$\check{c} \geq \rho\ell, \quad \textit{and} \quad \hat{c} \geq \rho u.$$

That is, the discounted perpetual holding costs of positive (negative) cash balances should exceed the control costs of reducing (increasing) the cash balance. As another example, for the case of an uncontrolled mean-reverting OU process,

$$\mathrm{d}X_t^0 = -\beta X_t^0 \mathrm{d}t + \sigma \mathrm{d}B_t,$$

In this case, conditions 4 and 5 reduce to

$$\check{c} \geq (\rho+\beta)\ell, \quad \textit{and} \quad \hat{c} \geq (\rho+\beta)u.$$

Proof of Theorem 10. Let φ and $\underline{x} < 0 < \overline{x}$ satisfy conditions 1–6. Extend φ to E, in a twice-continuously differentiable way, as follows:

$$\varphi(x) = \begin{cases} \ell|\underline{x}-x| + \varphi(\underline{x}+) & \text{if } x \leq \underline{x} \\ \varphi(x) & \text{if } \underline{x} < x < \overline{x} \\ u|x-\overline{x}| + \varphi(\overline{x}-) & \text{if } x \geq \overline{x} \end{cases}. \tag{3.3.1}$$

Let (L^*, U^*) be the control band policy associated with $(\underline{x}, \overline{x})$. The proof proceeds in two steps. First we prove that $J^{L^*,U^*} = \varphi$. Then we show that for any other feasible policy (L, U) it holds that $J^{L,U} \geq J^{L^*,U^*}$, so that $J^* = J^{L^*,U^*}$. Note that

$$\theta^{\varphi}(x) = \arg\max_{\theta\in[-\kappa,+\kappa]}\big(\theta\sigma(x)\mathrm{sign}(\varphi'(x))\big), \tag{3.3.2}$$

so that the worst-case prior is generated by

$$\theta_t^*(\omega) = \theta^{\varphi}\big(X_t(\omega)\big). \tag{3.3.3}$$

1. Fix $T > 0$, $x \in E$, $\theta \in \Theta$ and set $\theta_t^{\varphi} \triangleq \theta^{\varphi}(X_t^{L^*,U^*})$. From Itō's lemma it then follows that

$$\begin{aligned}
&\mathsf{E}_x^{\theta}\left[e^{-\rho T}\varphi\big(X_T^{L^*,U^*,\theta}\big)\right] \\
&= \varphi(x) + \mathsf{E}_x^{\theta}\Bigg[\int_0^T e^{-\rho s}\varphi'\big(X_s^{L^*,U^*,\theta}\big)\mathrm{d}\left(L_s^* + U_s^*\right) \\
&\qquad + \int_0^T e^{-\rho s}\Big\{\mathscr{L}\varphi\big(X_s^{L^*,U^*,\theta}\big) + \theta_s\sigma\big(X_s^{L^*,U^*,\theta}\big)\varphi'\big(X_s^{L^*,U^*,\theta}\big)\Big\}\mathrm{d}s\Bigg] \\
&\leq \varphi(x) + \mathsf{E}_x^{\theta}\Bigg[\int_0^T e^{-\rho s}\varphi'\big(X_s^{L^*,U^*,\theta}\big)\mathrm{d}\left(L_s^* + U_s^*\right) \\
&\qquad + \int_0^T e^{-\rho s}\big\{\mathscr{L}\varphi\big(X_s^{L^*,U^*,\theta}\big) + \theta_s^{\varphi}\sigma\big(X_s^{L^*,U^*,\theta}\big)\varphi'\big(X_s^{L^*,U^*,\theta}\big)\big\}\mathrm{d}s\Bigg] \\
&= \varphi(x) - \mathsf{E}_x^{\theta}\left[\int_0^T e^{-\rho s}c\left(X_s^{L^*,U^*,\theta}\right)\mathrm{d}s + \int_0^T e^{-\rho s}\left(u\mathrm{d}U_s^* + \ell\mathrm{d}L_s^*\right)\right],
\end{aligned}$$

where the inequality follows from (3.3.2) and (3.3.3), and the final equality follows from conditions 1 and 2.

Sending $T \to \infty$ and exploiting the non-negativity of φ, c, u, and ℓ, we find that

$$\varphi(x) \geq \mathsf{E}_x^{\theta}\left[\int_0^{\infty} e^{-\rho s} c\left(X_s^{L^*,U^*,\theta}\right) \mathrm{d}s + \int_0^{\infty} e^{-\rho s}\left(u \mathrm{d}U_s^* + \ell \mathrm{d}L_s^*\right)\right].$$

Since $\theta \in \Theta$ was chosen arbitrarily, this implies that

$$\varphi(x) \geq \inf_{(L,U)\in\mathscr{D}} \sup_{\theta\in\Theta} \mathsf{E}_x^{\theta}\left[\int_0^{\infty} e^{-\rho s} c\left(X_s^{L,U,\theta}\right) \mathrm{d}s + \int_0^{\infty} e^{-\rho s}\left(u \mathrm{d}U_s + \ell \mathrm{d}L_s\right)\right].$$

2. Next, note that Conditions 4 and 5 ensure that

$$\mathscr{L}\varphi(x) + \theta^{\varphi}(x)\sigma(x)\varphi'(x) + c(x) \geq 0, \quad \text{on } E.$$

On $(\underline{x}, \overline{x})$ this holds by construction. To see that it holds for $x \leq \underline{x}$, note that condition 4 and (3.3.1) implies that

$$\begin{aligned}
c(x) = -\check{c}x &\geq \rho\ell(\underline{x} - x) - \check{c}\underline{x} - \ell(\alpha(\underline{x}) - \alpha(x) - \kappa(\sigma(\underline{x}) - \sigma(x)) \\
&= \rho(\varphi(x) - \varphi(\underline{x})) - \check{c}\underline{x} + \varphi'(\underline{x})(\alpha(\underline{x}) - \kappa\sigma(\underline{x})) - \varphi'(x)(\alpha(x) - \kappa\sigma(x)) \\
&= \theta^{\varphi}(\underline{x})\sigma(x)\varphi'(x) - \mathscr{L}\varphi(x)
\end{aligned}$$

The equalities hold since $\varphi'(x) = -\ell, \varphi''(x) = 0$ for any $x \leq \underline{x}$ and that $\theta^{\varphi}(\underline{x}) = \kappa$. Similarly, Condition 5 ensures that the results holds for $x \geq \overline{x}$. Then, from convexity it follows that

$$-\ell \leq \varphi'(x) \leq u, \quad \text{on } E. \tag{3.3.4}$$

3. Let $(\bar{L}, \bar{U})$ be a feasible control policy. Fix $T > 0$. An application of Itō's lemma now gives that

$$\begin{aligned}
&\mathsf{E}_x^{\theta^{\varphi}}\left[\int_0^T e^{-\rho s}\Big\{c(X_s^{\bar{L},\bar{U},\theta^{\varphi}})\mathrm{d}s + u\mathrm{d}\bar{U}_s + \ell\mathrm{d}\bar{L}_s\Big\}\right] \\
&\underset{(3.3.4)}{\geq} -\mathsf{E}_x^{\theta^{\varphi}}\left[\int_0^T e^{-\rho s}\Big\{\theta^{\varphi}\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big)\sigma\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big)\varphi'\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big) + \mathscr{L}\varphi\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big)\Big\}\mathrm{d}t \right. \\
&\qquad\qquad \left. + \int_0^T e^{-\rho s}\varphi'\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big)\big(\mathrm{d}\bar{L}_s - \mathrm{d}\bar{U}_s\big)\right] \\
&= \varphi(x) - \mathsf{E}_x^{\theta^{\varphi}}\left[e^{-\rho T}\varphi\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big)\right].
\end{aligned}$$

Therefore,

$$\varphi(x) \leq \mathsf{E}_x^{\theta^{\varphi}}\left[\int_0^T e^{-\rho s}\Big\{c(X_s^{\bar{L},\bar{U},\theta^{\varphi}})\mathrm{d}s + u\mathrm{d}\bar{U}_s + \ell\mathrm{d}\bar{L}_s\Big\} + e^{-\rho T}\varphi\big(X_s^{\bar{L},\bar{U},\theta^{\varphi}}\big)\right].$$

Sending $T \to \infty$ and by exploiting Condition 6, the monotone convergence theorem gives

$$\varphi(x) \leq \mathsf{E}_x^{\theta^\varphi}\left[\int_0^\infty e^{-\rho s}\Big\{c(X_s^{\bar{L},\bar{U},\theta^\varphi})\mathrm{d}s + u\mathrm{d}\bar{U}_s + \ell\mathrm{d}\bar{L}_s\Big\}\right]$$

By arbitrariness of $(\bar{L},\bar{U})$, it then follows that

$$\begin{aligned}\varphi(x) &\leq \sup_{\theta\in\Theta}\inf_{(L,U)\in\mathscr{D}}\mathsf{E}_x^{\theta}\left[\int_0^\infty e^{-\rho s}\Big\{c(X_s^{L,U,\theta})\mathrm{d}s + u\mathrm{d}U_s + \ell\mathrm{d}L_s\Big\}\right] \\ &= \sup_{\theta\in\Theta}\mathsf{E}_x^{\theta}\left[\int_0^\infty e^{-\rho s}\Big\{c(X_s^{L^*,U^*,\theta})\mathrm{d}s + u\mathrm{d}U_s^* + \ell\mathrm{d}L_s^*\Big\}\right].\end{aligned}$$

4. Combining the results from Steps 1 and 3 gives that

$$\begin{aligned}\varphi(x) &= \inf_{L,U\mathscr{D}}\sup_{\theta\in\Theta}\mathsf{E}_x^{\theta}\left[\int_0^\infty e^{-\rho s}\Big\{c(X_s^{L,U,\theta})\mathrm{d}s + u\mathrm{d}U_s + \ell\mathrm{d}L_s\Big\}\right] \\ &= \sup_{\theta\in\Theta}\inf_{L,U\in\mathscr{D}}\mathsf{E}_x^{\theta}\left[\int_0^\infty e^{-\rho s}\left\{c(X_s^{L,U,\theta})\mathrm{d}s + u\mathrm{d}U_s + \ell\mathrm{d}L_s\right\}\right]\end{aligned}$$

and that $(\theta^\varphi,(L^*,U^*))$ realise a saddle-point. ■

3.4 AFFINE PERPETUAL HOLDING COSTS

Under some additional assumptions, it is often possible to write down sufficient conditions that are easier to check. In order to pursue this program, we first derive an expression for the perpetual holding costs of the *uncontrolled* process. First we let $\hat{\varphi}_{\pm\kappa}$ and $\check{\varphi}_{\pm\kappa}$ denote the increasing and decreasing *fundamental* solutions to the ordinary differential equation (ODE)

$$\mathscr{L}\varphi(x) + \theta^{\pm\kappa}(x)\sigma(x)\varphi'(x) = 0, \quad \text{on } E,$$

respectively. Here $\theta^{\pm\kappa}$ is the density generator $\theta^{\pm\kappa}(x) = \pm\kappa$, all $x \in E$. The measure generated by $\theta^{\pm\kappa}$ is denoted by $\mathsf{P}_x^{\pm\kappa}$. We normalize $\hat{\varphi}_{\pm\kappa}(0) = \check{\varphi}_{\pm\kappa}(0) = 1$, and denote

$$f_{\pm\kappa}(x) \triangleq \mathsf{E}_x^{\theta^{\pm\kappa}}\left[\int_0^\infty e^{-\rho t}X_t^0\mathrm{d}t\right],$$

where we assume that $f_{\pm\kappa}$ is affine in x. We summarize our assumptions on X^0 for future reference.

Assumption 1. The process X^0 is such that

1) the present value of its expected evolution is affine in its current state, i.e.

$$f_{\pm\kappa}(x) = \mathsf{E}_x^{\theta^{\pm\kappa}}\left[\int_0^\infty e^{-\rho t}X_t^0\mathrm{d}t\right] = ax + b, \quad (a \neq 0),$$

2) the increasing and decreasing solutions, $\hat{\varphi}$ and $\check{\varphi}$, to $\mathscr{L}\varphi(x) - \theta^{\pm\kappa}(x)\sigma(x)\varphi'(x) = 0$ are convex (see Alvarez, 2003 for sufficient conditions) and such that

$$\hat{\varphi}(0) = \check{\varphi}(0) = 1,$$

and

3) the holding costs $\check{c}$ and $\hat{c}$ are such that

$$\ell \leq \check{c} f'_{-\kappa}(x), \quad \text{and} \quad u \leq \hat{c} f'_{+\kappa}(x). \tag{3.4.1}$$

It can be verified that both ABM and OU process satisfy this assumption. For example, if X^0 follows an ABM, then

$$f_{\pm\kappa}(x) = \frac{x}{\rho} + \frac{\alpha \pm \kappa\sigma}{\rho^2}.$$

Moreover,

$$\hat{\varphi}_{\pm\kappa}(x) = e^{\beta_{\pm\kappa} x}, \quad \text{and} \quad \check{\varphi}_{\pm\kappa}(x) = e^{\gamma_{\pm\kappa} x},$$

where $\beta_{\pm\kappa} > 0$ and $\gamma_{\pm\kappa} < 0$ are the roots of the quadratic equation

$$\mathscr{Q}_{\pm\kappa}(\chi) \equiv \frac{1}{2}\sigma^2\chi^2 + (\alpha \pm \kappa\sigma)\chi - \rho = 0.$$

Condition (3.4.1) now reduces to

$$\ell \leq \check{c}/\rho, \quad \text{and} \quad u \leq \hat{c}/\rho.$$

If, on the other hand, X^0 follows the mean-reverting process

$$\mathrm{d}X_t^0 = -\beta X_t^0 \mathrm{d}t + \sigma \mathrm{d}B_t, \quad (\eta > 0),$$

under P_x, then under $\mathsf{P}_x^{\pm\kappa}$ it holds that,

$$\mathrm{d}X_t^0 = (-\beta X_t^0 \pm \kappa\sigma)\mathrm{d}t + \sigma \mathrm{d}B_t^{\pm\kappa},$$

where $B^{\pm\kappa}$ is a $\mathsf{P}_x^{\pm\kappa}$-Brownian motion. This process can be seen as an Ornstein-Uhlenbeck (OU) process with long-run mean $\tilde{x}_{\pm\kappa}$, i.e.

$$\mathrm{d}X_t^0 = \beta(\tilde{x}_{\pm\kappa} - X_t^0)\mathrm{d}t + \sigma \mathrm{d}B_t^{\pm\kappa}, \quad \text{where } \tilde{x}_{\pm\kappa} = \pm\kappa\sigma/\beta.$$

Therefore,

$$f_{\pm\kappa}(x) = \frac{x - \tilde{x}_{\pm\kappa}}{\rho + \beta} + \frac{\tilde{x}_{\pm\kappa}}{\rho}.$$

Here, the fundamental solutions of which X^0 follows the OU process are

$$\hat{\varphi}_{\pm\kappa}(x) = e^{\frac{\beta(x - \tilde{x}_{\pm\kappa})^2}{2\sigma^2}} D_{-\frac{\rho}{\beta}}\Big(\frac{x - \tilde{x}_{\pm\kappa}}{\sigma}\sqrt{2\beta}\Big), \quad \text{and}$$

$$\check{\varphi}_{\pm\kappa}(x) = e^{\frac{\beta(x-\tilde{x}_{\pm\kappa})^2}{2\sigma^2}} D_{-\frac{\rho}{\beta}}\Big(-\frac{x-\tilde{x}_{\pm\kappa}}{\sigma}\sqrt{2\beta}\Big),$$

where D_z is the parabolic cylinder function with index z (see, for example, Jeanblanc et al., 2009, chapter 5).

Condition (3.4.1) now reduces to

$$\ell \leq \check{c}/(\rho+\beta), \quad \text{and} \quad u \leq \hat{c}/(\rho+\beta).$$

The perpetual holding costs of the uncontrolled process can be found using the Feynman-Kac formula in the standard way:

$$R_{\pm\kappa}(x) \triangleq \mathsf{E}_x^{\theta^{\pm\kappa}}\left[\int_0^\infty e^{-\rho t}c(X_t^0)\mathrm{d}t\right] = \begin{cases} -\check{c}f_{\pm\kappa}(x) + \hat{E}_{\pm\kappa}\hat{\varphi}_{\pm\kappa}(x) & \text{if } x<0 \\ +\hat{c}f_{\pm\kappa}(x) + \check{E}_{\pm\kappa}\check{\varphi}_{\pm\kappa}(x) & \text{if } x \geq 0. \end{cases}$$

Here, $\hat{E}_{\pm\kappa}$ and $\check{E}_{\pm\kappa}$ are constants that are determined by "value-matching" and "smooth-pasting" conditions at 0, i.e.,

$$R_{\pm\kappa}(0-) = R_{\pm\kappa}(0+), \text{ and } R'_{\pm\kappa}(0-) = R'_{\pm\kappa}(0+),$$

respectively. This gives

$$\begin{aligned} \hat{E}_{\pm\kappa} &= (\hat{c}+\check{c})\frac{f'_{\pm\kappa}(0) - f_{\pm\kappa}(0)\check{\varphi}'_{\pm\kappa}(0)}{\hat{\varphi}'_{\pm\kappa}(0) - \check{\varphi}'_{\pm\kappa}(0)}, \quad \text{and} \\ \check{E}_{\pm\kappa} &= (\hat{c}+\check{c})\frac{f'_{\pm\kappa}(0) - f_{\pm\kappa}(0)\hat{\varphi}'_{\pm\kappa}(0)}{\hat{\varphi}'_{\pm\kappa}(0) - \check{\varphi}'_{\pm\kappa}(0)}, \end{aligned}$$

so that

$$R_{\pm\kappa}(x) = \begin{cases} -\check{c}f_{\pm\kappa}(x) + (\hat{c}+\check{c})\frac{f'_{\pm\kappa}(0)-f_{\pm\kappa}(0)\check{\varphi}'_{\pm\kappa}(0)}{\hat{\varphi}'_{\pm\kappa}(0)-\check{\varphi}'_{\pm\kappa}(0)}\hat{\varphi}_{\pm\kappa}(x) & \text{if } x<0 \\ +\hat{c}f_{\pm\kappa}(x) + (\hat{c}+\check{c})\frac{f'_{\pm\kappa}(0)-f_{\pm\kappa}(0)\hat{\varphi}'_{\pm\kappa}(0)}{\hat{\varphi}'_{\pm\kappa}(0)-\check{\varphi}'_{\pm\kappa}(0)}\check{\varphi}_{\pm\kappa}(x) & \text{if } x\geq 0. \end{cases}$$

Without ambiguity ($\kappa = 0$), in order to construct the function φ of Theorem 10, one would now find constants A and B, and control barriers $\underline{x}$ and $\overline{x}$ such that the following value-matching and smooth-pasting conditions hold:

$$\begin{aligned} R'_0(\underline{x}) + A\hat{\varphi}'_0(\underline{x}) + B\check{\varphi}'_0(\underline{x}) &= -\ell \\ R'_0(\overline{x}) + A\hat{\varphi}'_0(\overline{x}) + B\check{\varphi}'_0(\overline{x}) &= u \\ R''_0(\underline{x}) + A\hat{\varphi}''_0(\underline{x}) + B\check{\varphi}''_0(\underline{x}) &= 0 \\ R''_0(\overline{x}) + A\hat{\varphi}''_0(\overline{x}) + B\check{\varphi}''_0(\overline{x}) &= 0. \end{aligned}$$

One then proceeds by showing that the resulting function,

$$\varphi(x) = R_0(x) + A\varphi_0(x) + B\varphi_0(x), \quad \text{on } (\underline{x}, \overline{x}),$$

and the constants $\underline{x} < 0 < \overline{x}$ satisfy the conditions of the verification Theorem 10.

Under ambiguity ($\kappa > 0$) matters are a bit more complicated. Intuitively speaking, the main issue is that the "worst–case drift" is different at $\underline{x}$ and $\overline{x}$. In particular, at the lower control bound the worst case drift is $\alpha(\underline{x}) - \kappa\sigma(\underline{x})$, because the worst that can happen is that the cash hoard depletes even more and, thus, increases the control costs. Similarly, at the upper control bound the worst case drift is $\alpha(\overline{x}) + \kappa\sigma(\overline{x})$, because the worst that can happen is that the cash hoard increases even more and, thus, increases the control costs.

So, at $\underline{x}$ and $\overline{x}$ we need to work with functions R, $\hat{\varphi}$, and $\check{\varphi}$ under $-\kappa$ and $+\kappa$, respectively. That is, we will look for constants A, B, C, and D, as well as control bounds $\underline{x}$ and $\overline{x}$ such that the following value-matching and smooth-pasting conditions hold:

$$
\begin{aligned}
R'_{-\kappa}(\underline{x}) + A\hat{\varphi}'_{-\kappa}(\underline{x}) + B\check{\varphi}'_{-\kappa}(\underline{x}) &= -\ell \\
R''_{-\kappa}(\underline{x}) + A\hat{\varphi}''_{-\kappa}(\underline{x}) + B\check{\varphi}''_{-\kappa}(\underline{x}) &= 0 \\
R'_{+\kappa}(\overline{x}) + C\hat{\varphi}'_{+\kappa}(\overline{x}) + D\check{\varphi}'_{+\kappa}(\overline{x}) &= u \\
R''_{+\kappa}(\overline{x}) + C\hat{\varphi}''_{+\kappa}(\overline{x}) + D\check{\varphi}''_{+\kappa}(\overline{x}) &= 0.
\end{aligned}
\tag{3.4.2}
$$

Now, of course, we have too few equations to determine all the constants. The "missing" constraints come from the fact that there is a point x^* where the worst-case drift changes. This is the point where the firm's cost function changes from being decreasing to increasing. At this point we also impose a value-matching and smooth-pasting condition, i.e., we find $x^* \in (\underline{x}, \overline{x})$ such that

$$
\begin{aligned}
R'_{-\kappa}(x^*-) + A\hat{\varphi}'_{-\kappa}(x^*-) + B\check{\varphi}'_{-\kappa}(x^*-) &= 0 \\
R'_{+\kappa}(x^*+) + C\hat{\varphi}'_{+\kappa}(x^*+) + D\check{\varphi}'_{+\kappa}(x^*+) &= 0
\end{aligned}
\tag{3.4.3}
$$

$$
\begin{aligned}
&R''_{-\kappa}(x^*-) + A\hat{\varphi}''_{-\kappa}(x^*-) + B\check{\varphi}''_{-\kappa}(x^*-) \\
&\quad = R''_{+\kappa}(x^*+) + C\hat{\varphi}''_{+\kappa}(x^*+) + D\check{\varphi}''_{+\kappa}(x^*+).
\end{aligned}
\tag{3.4.4}
$$

We show below that if this system of 7 equations in 7 unknowns has a solution, then a function φ can be constructed on $(\underline{x}, \overline{x})$ so that the conditions of verification Theorem 10 are satisfied. A similar approach has also been used by Cheng and Riedel (2013) to price a straddle option under ambiguity and by Hellmann and Thijssen (2018) to analyse preemptive investment behavior in a duopoly under ambiguity.

Theorem 11. *Suppose that the system of equations (3.4.2)–(3.4.4) admits a solution $(A, B, C, D, \underline{x}, \overline{x}, x^*)$ with $\underline{x} < x^* < \overline{x}$. Then the optimal policy (L^*, U^*) is the control band policy associated with $(\underline{x}, \overline{x})$ and the firm's cost function is*

$$
V(x) = \begin{cases}
\ell(\underline{x} - x) + R_{-\kappa}(\underline{x}+) + A\hat{V}_{-\kappa}(\underline{x}+) + B\check{V}_{-\kappa}(\underline{x}+) & \text{if } \; x \le \underline{x} \\
R_{-\kappa}(x) + A\hat{V}_{-\kappa}(x) + B\check{V}_{-\kappa}(x) & \text{if } \; \underline{x} < x < x^* \\
R_{+\kappa}(x) + C\hat{V}_{+\kappa}(x) + D\check{V}_{+\kappa}(x) & \text{if } \; x^* \le x < \overline{x} \\
u(x - \overline{x}) + R_{+\kappa}(\overline{x}-) + C\hat{V}_{+\kappa}(\overline{x}-) + D\check{V}_{+\kappa}(\overline{x}-) & \text{if } \; x \ge \overline{x}.
\end{cases}
\tag{3.4.5}
$$

Proof of Theorem 11. See Appendix A. ■

Now we propose a probabilistic representation of the first derivative of the value function, formulated as an optimal stopping game for an uncontrolled process. This approach will be highly beneficial for the comparative statics analysis in the following section. Such a problem is known as a zero-sum Dynkin game for singular control, initially introduced by Taksar (1985). For further applications in operations research, see also Federico et al. (2023), Ferrari and Vargiolu (2020), and Guo and Tomecek (2008).

Theorem 12. *Let $V(x)$ be the value function of (3.4.5) and assume furthermore that*

$$\rho > \sup_{x\in E} |\alpha'(x)| \pm \kappa|\sigma'(x)|.$$

Then for any $x \in E$, we have

$$V'(x) \triangleq v(x) = \inf_{\eta}\sup_{\tau} w(x,\tau,\eta) = \sup_{\tau}\inf_{\eta} w(x,\tau,\eta) \tag{3.4.6}$$

where τ and η are $\mathscr{F}_t$-stopping times under $\mathsf{P}_x^{\theta^V}$,

$$\begin{aligned} w(x,\tau,\eta) \triangleq \mathsf{E}_x^{\theta^V}\Bigg[\int_0^{\tau\wedge\eta} & e^{-\int_0^t \hat\rho(\hat X_s^{\theta^V})\mathrm{d}s}\zeta(\hat X_t^{\theta^V})\mathrm{d}t \\ & - 1_{(0,\eta)}(\tau)e^{-\int_0^\tau \hat\rho(\hat X_s^{\theta^V})\mathrm{d}s}\ell + 1_{(0,\tau)}(\eta)e^{-\int_0^\eta \hat\rho(\hat X_s^{\theta^V})\mathrm{d}s}u\Bigg], \end{aligned}$$

$$\begin{aligned} \theta^V(x) &\triangleq -\kappa 1_{(-\infty,x^*)}(x) + \kappa 1_{[x^*,\infty)}(x), \\ \zeta(x) &\triangleq \hat c \cdot 1_{(0,\infty)}(x) - \check c \cdot 1_{(-\infty,0)}(x) \ \ \textit{and}, \\ \hat\rho(x) &\triangleq \rho - (\alpha'(x) + \theta^V(x)\sigma'(x)). \end{aligned}$$

Here $\hat X^{\theta^V}$ solves

$$\begin{aligned} \mathrm{d}\hat X_t^{\theta^V} &= \left(\alpha(\hat X_t^{\theta^V}) + \sigma'(\hat X_t^{\theta^V})\sigma(\hat X_t^{\theta^V})\right)\mathrm{d}t + \sigma(\hat X_t^{\theta^V})\mathrm{d}B_t^{\theta^V}, \\ &= \left(\alpha(\hat X_t^{\theta^V}) + (\sigma'(\hat X_t^{\theta^V}) + \theta^V(\hat X_t^{\theta^V}))\sigma(\hat X_t^{\theta^V})\right)\mathrm{d}t + \sigma(\hat X_t^{\theta^V})\mathrm{d}B_t, \quad \mathsf{P}_x^{\theta^V}\text{-a.s.} \end{aligned} \tag{3.4.7}$$

where B^{θ^V} is a Brownian motion in the sense of (3.2.3) when $\theta = \theta^V$. Moreover, the saddle-point stopping times (τ^*, η^*) are given by

$$\tau^* \triangleq \inf\{t > 0 : \hat X_t^{\theta^V} \le \underline{x}\}, \quad \textit{and} \quad \eta^* \triangleq \inf\{t > 0 : \hat X_t^{\theta^V} \ge \overline{x}\}, \quad \mathsf{P}_x^{\theta^V}\text{-a.s.} \tag{3.4.8}$$

That is,

$$w(x,\tau,\eta^*) \le v(x) \le w(x,\tau^*,\eta) \ \ \textit{for any} \ \ x \in E$$

Proof of Theorem 12. The idea of what follows is adopted from Ferrari and Vargiolu (2020). One can see from Theorem 11 that V'' is bounded in E, non-decreasing in $E \setminus (x^*, \infty)$ and non-increasing in $E \setminus (-\infty, x^*]$. According to Proposition 7.7 of Federico et al. (2023), this implies that v'' is locally bounded on E. Therefore, we can deduce from (3.4.5) that

$$\widehat{\mathscr{L}}v(x) + \zeta(x) = 0 \ \text{ and } \ -\ell \le v(x) \le u \qquad \text{if } \ \underline{x} < x < \overline{x}, \tag{3.4.9}$$

$$\widehat{\mathscr{L}}v(x) + \zeta(x) \le 0 \ \text{ and } \ v(x) = -\ell \qquad \text{if } \ x \le \underline{x}, \tag{3.4.10}$$

$$\widehat{\mathscr{L}}v(x) + \zeta(x) \ge 0 \ \text{ and } \ v(x) = u \qquad \text{if } \ x \ge \overline{x}, \tag{3.4.11}$$

where

$$\widehat{\mathscr{L}}v(x) \triangleq \frac{1}{2}\sigma^2(x)v''(x) + (\alpha(x) + (\sigma'(x) + \theta^V(x))\sigma(x))v'(x) + \hat{\rho}(x)v(x), \quad \text{on } C^2(E).$$

The first equation in (3.4.9) comes from taking the derivative of V with respect to x over the intervals $(\underline{x}, x^*)$. The first inequality in (3.4.10) and (3.4.11) follow from Condition 4 and 5 in Theorem 10, respectively.

Notice that $\widehat{\mathscr{L}}$ is the characteristic operator on $C^2(E)$ of the killed process

$$\left(e^{-\int_0^t \hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}\hat{X}_t\right)_{t\ge 0},$$

which is well-defined since it is implied by Condition (3.2.1) that $\hat{X}$ also admits a strong unique solution. Therefore, by the variational inequalities of optimal stopping times (see, for example, Øksendal, 2010, Theorem 10.4.1, or Øksendal and Sulem, 2019, Theorem 6.1), we can infer from (3.4.9) and (3.4.10) that on E, $v(x) \le w(x, \tau^*, \eta)$ for any stopping time η. This means that $v(x) \le \sup_\tau \inf_\eta w(x, \tau, \eta)$. On the other hand, (3.4.9) and (3.4.10) induces $v(x) \ge w(x, \tau, \eta^*)$, for any stopping time τ, implying that $v(x) \ge \inf_\eta \sup_\tau w(x, \tau, \eta)$. Since $\inf_\eta \sup_\tau w(x, \tau, \eta) \ge \sup_\tau \inf_\eta w(x, \tau, \eta)$, we finally obtain (3.4.6), which concludes the proof. ■

The following result is the refinement of Theorem 12, which later used to examine the comparative static of ambiguity as well as giving an insight toward managerial concept of the Dynkin game in singular control.

Corollary 1. *Suppose that v solves (3.4.6) and (τ^*, η^*) is the saddle point (3.4.8). Then*

$$v(x) = \begin{cases} \underline{w}(x, \tau^*, \phi^*) & \text{if } \ x < x^* \\ \overline{w}(x, \eta^*, \phi^*) & \text{if } \ x \ge x^* \end{cases} \tag{3.4.12}$$

where

$$\underline{w}(x, \tau^*, \phi^*) \triangleq \mathsf{E}_x^{\theta^V = +\kappa}\left[\int_0^{\tau^* \wedge \phi^*} e^{-\int_0^t \hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t - 1_{(0,\phi^*)}(\tau^*)e^{-\int_0^{\tau^*} \hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}\ell\right],$$

$$\overline{w}(x,\eta^*,\phi^*) \triangleq \mathsf{E}_x^{\theta^V=-\kappa}\left[\int_0^{\eta^*\wedge\phi^*} e^{-\int_0^t \hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t + 1_{(0,\phi^*)}(\eta^*)e^{-\int_0^{\eta^*}\hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}u\right]$$

and $\phi^* \triangleq \inf\{t>0 : \hat{X}_t^{\theta^V} = x^*\}$.

Proof of Corollary 1. We provide details only for $\overline{w}$ since the proof of the representation of $\underline{w}$ is analogous. Given that $x^* < \overline{x}$, by (3.4.9), (3.4.11), and (3.4.3)-(3.4.4) we have that $(v, x^*, \overline{x})$ satisfy

$$\begin{aligned} \widehat{\mathscr{L}}v(x) + \zeta(x) &= 0 \quad \text{if } x^* < x < \overline{x}, \\ v(x) &= u \quad \text{if } x \geq \overline{x}, \\ v(x^*) &= 0 \end{aligned}$$

Since $v \in C^1([x^*,\infty]) \cap C^2([x^*,\infty]\setminus\{\overline{x}\})$, an application of Dynkin's formula yields that, for any $x \geq x^*$,

$$v(x) = \mathsf{E}_x^{\theta^V}\left[\int_0^{\eta^*\wedge\phi^*} e^{-\int_0^t \hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t + 1_{(0,\phi^*)}(\eta^*)e^{-\int_0^{\eta^*}\hat{\rho}(\hat{X}_s^{\theta^V})\mathrm{d}s}u\right].$$

The claim representation of $\overline{w}$ follows by noticing that $\theta^V(x) = -\kappa$ for any $x \geq x^*$. ■

Remark 14. *In managerial terms, Corollary 1 can be interpreted as two timing games, one on* $(-\infty, x^*)$ *and the other on* $[x^*,\infty)$*, involving two players: the DM and a hypothetical player (nature). On the interval* $(-\infty, x^*)$*, the game requires the DM and nature selecting stopping times* τ *and* ϕ*, respectively. The game ends when time reaches* $\tau\wedge\phi$*, at which point the DM receives a payoff of* $\underline{w}(x,\tau,\phi)$ *from nature. The DM's objective is to choose* τ *to maximize* $\underline{w}(x,\tau,\phi)$*, while nature select* ϕ *to minimize its payment.*

On the other hand, if the inventory level lies within $[x^*,\infty)$*, the DM and nature must choose stopping times* η *and* ϕ*, respectively. When the game concludes at* $\eta\wedge\phi$*, the DM pays* $\overline{w}(x,\eta^*,\phi^*)$ *to nature. Thus, the DM's goal is to minimize the payment by selecting* η*, while nature seeks to maximize its payoff by determining* ϕ*.*

From (3.4.12), it follows that the DM's ideal payoff occurs when $\underline{w}(x^*,\tau^*,\phi^*) = \overline{w}(x^*,\eta^*,\phi^*) = 0$*. In essence, the DM's optimal strategy is to keep the running marginal cost of holding inventory as close to zero as possible under timing games against the uncertain nature.*

3.5 COMPARATIVE STATICS WITH ARITHMETIC BROWNIAN MOTION

Suppose that the uncontrolled cash inventory X^0 follows, under P_x, the ABM

$$X_t^0 = x + \alpha t + \sigma B_t, \tag{3.5.1}$$

so that

$$\hat{V}_{\pm\kappa}(x) = e^{\beta_{\pm\kappa} x}, \quad \text{and} \quad \check{V}_{\pm\kappa}(x) = e^{\gamma_{\pm\kappa} x}, \tag{3.5.2}$$

where $\beta_{\pm\kappa} > 0$ and $\gamma_{\pm\kappa} < 0$ are the positive and negative roots, respectively, of the quadratic equation

$$\mathscr{Q}_{\pm\kappa}(\chi) \equiv \frac{1}{2}\sigma^2\chi^2 + (\alpha \pm \kappa\sigma)\,\chi - \rho = 0.$$

Recall that the holding costs of cash are given by

$$c(x) = \hat{c}x \cdot 1_{[0,\infty)}(x) - \check{c}x \cdot 1_{(-\infty,0)}(x),$$

for some $\hat{c}, \check{c} > 0$. As mentioned before, if $\underline{x} < 0 < \overline{x}$, then Conditions 4 and 5 in Theorem 10 reduce to

$$\check{c} \geq \rho\ell, \quad \text{and} \quad \hat{c} \geq \rho u,$$

i.e. the control costs must not exceed the expected discounted uncontrolled holding costs; otherwise, it would never be optimal to exercise control.

The expected discounted uncontrolled holding costs in this case are given by

$$R_{\pm\kappa}(x) = \begin{cases} -\frac{\check{c}}{\rho}\left[x + \frac{\alpha\pm\kappa\sigma}{\rho}\right] + \hat{E}_{\pm\kappa}e^{\beta_{\pm\kappa}x} & \text{if } x < 0 \\ +\frac{\hat{c}}{\rho}\left[x + \frac{\alpha\pm\kappa\sigma}{\rho}\right] + \check{E}_{\pm\kappa}e^{\gamma_{\pm\kappa}x} & \text{if } x \geq 0 \end{cases},$$

where

$$\begin{aligned} \hat{E}_{\pm\kappa} &= \frac{\hat{c}+\check{c}}{\rho^2}\frac{\rho - \gamma_{\pm\kappa}(\alpha \pm \sigma\kappa)}{\beta_{\pm\kappa} - \gamma_{\pm\kappa}}, \quad \text{and} \\ \check{E}_{\pm\kappa} &= \frac{\hat{c}+\check{c}}{\rho^2}\frac{\rho - \beta_{\pm\kappa}(\alpha \pm \sigma\kappa)}{\beta_{\pm\kappa} - \gamma_{\pm\kappa}}. \end{aligned}$$

Since

$$\mathscr{Q}_{\pm\kappa}(\beta_{\pm\kappa}) = \mathscr{Q}_{\pm\kappa}(\gamma_{\pm\kappa}) = 0,$$

the constants $\hat{E}$ and $\check{E}$ can be written as

$$\begin{aligned} \hat{E} &= \frac{\hat{c}+\check{c}}{2\rho^2}\frac{\sigma^2\gamma_{\pm\kappa}^2}{\beta_{\pm\kappa} - \gamma_{\pm\kappa}} > 0, \quad \text{and} \\ \check{E} &= \frac{\hat{c}+\check{c}}{2\rho^2}\frac{\sigma^2\beta_{\pm\kappa}^2}{\beta_{\pm\kappa} - \gamma_{\pm\kappa}} > 0, \end{aligned}$$

respectively. That is, the expected discounted holding costs of uncontrolled cash inventory equals

$$R_{\pm\kappa}(x) = \begin{cases} \frac{-\check{c}}{\rho}\left[x + \frac{\alpha\pm\kappa\sigma}{\rho}\right] + \frac{(\hat{c}+\check{c})\sigma^2\gamma_{\pm\kappa}^2}{2\rho^2(\beta_{\pm\kappa}-\gamma_{\pm\kappa})}e^{\beta_{\pm\kappa}x} & \text{if } x < 0 \\ \frac{+\hat{c}}{\rho}\left[x + \frac{\alpha\pm\kappa\sigma}{\rho}\right] + \frac{(\hat{c}+\check{c})\sigma^2\beta_{\pm\kappa}^2}{2\rho^2(\beta_{\pm\kappa}-\gamma_{\pm\kappa})}e^{\gamma_{\pm\kappa}x} & \text{if } x \geq 0 \end{cases}. \tag{3.5.3}$$

We recall that the general solution of the value function takes the form of (3.4.5). Therefore, in the case an arithmetic Brownian motion, substituting (3.5.2) and (3.5.3) into (3.4.5) reads

$$V(x) = \begin{cases} \ell|\underline{x} - x| - \overbrace{\frac{1}{\rho}\left[\ell(\alpha - \kappa\sigma) + \check{c}\underline{x}\right]}^{\triangleq \Gamma^-(\underline{x})} & \text{if } x \leq \underline{x} \\ R_{-\kappa}(x) + A\hat{V}_{-\kappa}(x) + B\check{V}_{-\kappa}(x) & \text{if } \underline{x} < x < x^* \\ R_{+\kappa}(x) + C\hat{V}_{+\kappa}(x) + D\check{V}_{+\kappa}(x) & \text{if } x^* \leq x < \overline{x} \\ u|x - \overline{x}| + \underbrace{\frac{1}{\rho}\left[u(\alpha + \kappa\sigma) + \hat{c}\overline{x}\right]}_{\triangleq \Gamma^+(\overline{x})} & \text{if } x \geq \overline{x}. \end{cases} \tag{3.5.4}$$

for some constant A, B, C and D.

In the following, we perform a sensitivity analysis of the optimal control boundaries $\underline{x}$ and $\overline{x}$ with respect to same model parameters.

To keep the argument concise, we assume from this point onward that Assumption (1) holds, that X^0 satisfies (3.5.1), and that $(\underline{x}, \overline{x})$ is a solution to (3.5.4).

Proposition 1 (Comparative statics of risk)**.** *The control barriers $(\underline{x}, \overline{x})$ expand as σ increases. That is, $\sigma \mapsto \underline{x}(\sigma)$ is non-increasing and $\sigma \mapsto \overline{x}(\sigma)$ is non-decreasing.*

Proof of Proposition 1. The idea of the proof is adopted from Matomäki (2012), Theorem 6.1 (also shown in Ferrari and Vargiolu, 2020.) Suppose that σ_1, where $\sigma \leq \sigma_1$, is a diffusion term of the uncontrolled arithmetic Brownian motion $X_t^{0,\sigma_1} = x + \alpha t + \sigma_1 B_t$. Let $\mathscr{L}_{\sigma_1}$ be the characteristic operator on $C^2(E)$ of the killed process $\left(e^{-\rho t} X_t^{0,\sigma_1}\right)_{t \geq 0}$, that is, $\mathscr{L}_{\sigma_1}\varphi(x) \triangleq \frac{1}{2}\sigma_1^2\varphi''(x) + \alpha\varphi'(x) - \rho\varphi(x)$, and $V_{\sigma_1}(x)$ is the associated value function. Then a straightforward calculation gives

$$\mathscr{L}_{\sigma_1}V(x) - \theta^*\sigma_1 V'(x) + c(x) = \begin{cases} \Gamma^-(\underline{x}) - \Gamma^-(x) & \text{if} \quad x \leq \underline{x} \\ \frac{1}{2}(\sigma_1^2 - \sigma^2)V''(x) - \kappa(\sigma_1 - \sigma)V'(x) & \text{if} \quad \underline{x} < x < x^* \\ \frac{1}{2}(\sigma_1^2 - \sigma^2)V''(x) + \kappa(\sigma_1 - \sigma)V'(x) & \text{if} \quad x^* \leq x < \overline{x} \\ \Gamma^+(x) - \Gamma^+(\overline{x}) & \text{if} \quad x \geq \overline{x}. \end{cases}$$

One can see that $\mathscr{L}_{\sigma_1}V(x) + \theta^*\sigma_1 V'(x) + c(x) \geq 0$ for any $x \in E$. This is because $\Gamma^-(\underline{x}) - \Gamma^-(x) \geq 0$ if $x \leq \underline{x}$ and also $\Gamma^+(x) - \Gamma^+(\overline{x}) \geq 0$ if $x \geq \overline{x}$. The claim holds on $(\underline{x}, \overline{x})$ because V is convex is x, thanks to Theorem 11, and the fact that x^* situates the point where worst-case changes, i.e. $V'(x) \leq 0$ if $x < x^*$ and $V'(x) \geq 0$, otherwise. Since $V_{\sigma_1}(x)$ is the value function when $\sigma = \sigma_1$, Theorem 10 asserts that

$$\mathscr{L}_{\sigma_1}V(x) + \theta^*\sigma_1 V'(x) + c(x) \geq \mathscr{L}_{\sigma_1}V_{\sigma_1}(x) + \theta^*\sigma_1 V'_{\sigma_1}(x) + c(x) = 0$$

Since this inequality holds for any $\sigma_1 \geq \sigma$, it is implied by Itō's lemma that $V_{\sigma_1}(x) \geq V(x)$ for all $x \in E$. This completes the first part of the proof.

Now, we demonstrate that $(\underline{x}, \overline{x})$ expands as σ increases. To do this, we assume that $\sigma_1 > \sigma \triangleq \sigma_0$ and denote $(\underline{x}_1, \overline{x}_1)$ as the optimal control barriers of $V_{\sigma_1}(x)$. If we assume in contrary that $\underline{x}_1 > \underline{x} \triangleq \underline{x}_0$ and $\overline{x}_1 < \overline{x} \triangleq \overline{x}_0$, then Condition 4 in Theorem 10 gives

$$\check{c}(\underline{x}_i - x) \geq \ell\big[\rho(\underline{x}_i - x) - \big(\big(\alpha(\underline{x}_i) - \alpha(x)\big) - \kappa\big(\sigma_i(\underline{x}_i) - \sigma_i(x)\big)\big],$$

for $x \leq \underline{x}_0 < \underline{x}_1,\ i = 0, 1$, implying that

$$V_{\sigma_0}(\underline{x}_0) = -\frac{1}{\rho}(\ell(\alpha - \kappa\sigma_0) + \check{c}\underline{x}_0) > \ell(\underline{x}_1 - \underline{x}_0) - \frac{1}{\rho}(\ell(\alpha - \kappa\sigma_1) + \check{c}\underline{x}_1) = V_{\sigma_1}(\underline{x}_0),$$

which is a contradiction. Therefore, we conclude that $\underline{x}_1 \leq \underline{x}_0$. A similar argument also gives that $\overline{x}_1 \geq \overline{x}_0$. Hence, these result in an expansion of the control barriers as σ increases, concluding the proof. ■

Not surprisingly, this aligns with the familiar result that greater risk increases the firms cost function. The same applies to ambiguity, as becomes evident by extending the first part of the proof of Proposition 1 to account for a higher level of κ. These can be depicted by Figure 3.1a and 3.1b. Specifically, a manager with maxmin utility assigns a higher expected cost to inventory management under such conditions. Furthermore, the expansion of the control barriers suggests that increased risk causes delays in taking action. This occurs because a rise in σ makes extreme scenarios more probable, meaning it becomes, on average, just as likely for the inventory to reach the target level as it is to trigger a control action. Consequently, the optimal singular control policy indicates that acting too frequently becomes increasingly costly. It is evident that Conditions 4 and 5 in Theorem 10 are essential for the previous result. However, these conditions alone are not enough to offer a comparative static analysis of ambiguity (or of the other model parameters). This is where the Dynkin game of Theorem 12 becomes relevant and is therefore utilized to support the following arguments.

It is important to note that, when ambiguity arises, the comparative statics may exhibit non-monotonicity, depending on the *symmetry* of the model parameters. This phenomenon does not appears in the standard comparative statics of two-sided singular control, as discussed in Matomäki (2012) and Ferrari and Vargiolu (2020), for instance. To explore this, we begin by introducing the notion of symmetry as defined below.

Definition 22. Suppose that $\alpha, \rho, \kappa, \sigma, \hat{c}, \check{c}, \ell$ and u satisfy Condition 4 and 5 in Theorem 10. Then the model parameters are said to be

- (symmetric) if $\alpha = 0, \check{c} = \hat{c}$ and $\ell = u$,
- (asymmetric) if $\alpha \neq 0$, $\check{c} \neq \hat{c}$ or $\ell \neq u$.

The following lemma will be useful for our comparative statics arguments.

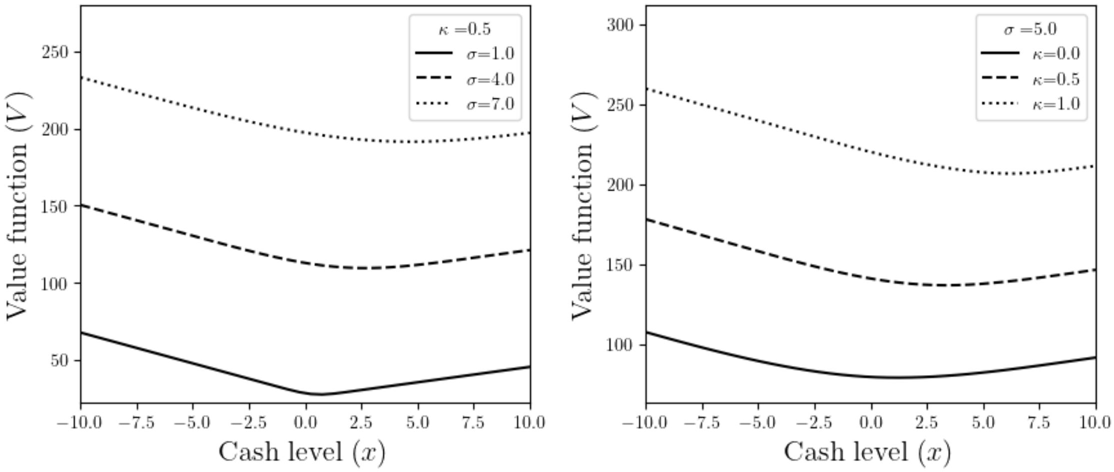

(a) Value functions given $\kappa = 0.5$ for $\sigma \in \{1, 4, 7\}$.

(b) Value functions given $\sigma = 5$ for $\kappa \in \{0, 0.5, 1\}$.

Figure 3.1: Value functions with different levels of (3.1a) risk and (3.1b) ambiguity, given sample parameters $\alpha = 0$, $\rho = 0.1$, $u = 3$, $\ell = 5$, and $\check{c} = \hat{c} = 1$.

Lemma 1. *Suppose that the model parameters are symmetric, i.e. $\alpha = 0$, $u = \ell$, and $\hat{c} = \check{c}$. Then $x^* = 0$, $|\underline{x}| = |\overline{x}|$, and the marginal value function is an odd function: $v(x) = -v(-x)$ for all $x \in E$.*

Prove of Lemma 1. Suppose, ad absurdum, that $x^* > 0$ and $|\underline{x}| = |\overline{x}|$. Adapting the argument in Harrison (2013, Section 3.3), it holds that,

$$\mathsf{P}^{\theta^V}(\phi^* < \tau^* | \hat{X}_0^{\theta^V} = x < x^*) = \begin{cases} \frac{x - \underline{x}}{x^* - \underline{x}} & \text{if } \kappa = 0 \\ \frac{e^{\frac{2\kappa}{\sigma}x} - e^{\frac{2\kappa}{\sigma}\underline{x}}}{e^{\frac{2\kappa}{\sigma}x^*} - e^{\frac{2\kappa}{\sigma}\underline{x}}}, & \text{if } \kappa > 0 \end{cases}$$

$$\mathsf{P}^{\theta^V}(\phi^* < \eta^* | \hat{X}_0^{\theta^V} = x > x^*) = \begin{cases} \frac{x - \overline{x}}{x^* - \overline{x}} & \text{if } \kappa = 0 \\ \frac{e^{-\frac{2\kappa}{\sigma}x} - e^{-\frac{2\kappa}{\sigma}\overline{x}}}{e^{-\frac{2\kappa}{\sigma}x^*} - e^{-\frac{2\kappa}{\sigma}\overline{x}}}, & \text{if } \kappa > 0 \end{cases}$$

respectively. Now suppose that $\underline{\phi} \triangleq \inf\{t > 0 : \hat{X}_t^{\theta^V} = -x^*\}$. If $x^* > 0$, then for any $\delta > 0$ such that $\underline{x} + \delta \leq x^*$ and $\overline{x} - \delta \geq x^*$, we have

$$\begin{aligned} \frac{\overline{x} - \delta - \overline{x}}{x^* - \overline{x}} = \frac{-\delta}{x^* + \underline{x}} > \frac{-\delta}{-x^* + \underline{x}} = \frac{\underline{x} + \delta - \underline{x}}{x^* - \underline{x}} \qquad & \text{if } \kappa = 0, \text{ and} \\ \frac{e^{-\frac{2\kappa}{\sigma}(\overline{x} - \delta)} - e^{-\frac{2\kappa}{\sigma}\overline{x}}}{e^{-\frac{2\kappa}{\sigma}x^*} - e^{-\frac{2\kappa}{\sigma}\overline{x}}} > \frac{e^{\frac{2\kappa}{\sigma}(\underline{x} + \delta)} - e^{\frac{2\kappa}{\sigma}\underline{x}}}{e^{\frac{2\kappa}{\sigma}x^*} - e^{\frac{2\kappa}{\sigma}\underline{x}}} \qquad & \text{if } \kappa > 0, \end{aligned}$$

which implies that

$$\mathsf{P}_{\overline{x}-\delta}^{\theta^V}(\phi^* < \eta^*) > \mathsf{P}_{\underline{x}+\delta}^{\theta^V}(\phi^* < \tau^*). \tag{3.5.5}$$

Suppose now that $\hat{X}_0^{-\kappa} = x$ and $\hat{X}_0^{+\kappa} = -x$, with $x > x^* \, (> 0)$. Then it is more likely that $\hat{X}_0^{\theta^V}$ starting from above zero arrives at x^* earlier. Consequently, it follows from Corollary 1 that

$$\begin{aligned}
&\underline{w}(\underline{x}+\delta) = \mathsf{E}_{\underline{x}+\delta}^{\theta^V}\left[\int_0^{\tau^*\wedge\phi^*} e^{-\rho t}(\hat{c}\cdot 1_{(0,\infty)}(\hat{X}_t^{\theta^V}) - \check{c}\cdot 1_{(-\infty,0)}(\hat{X}_t^{\theta^V}))\mathrm{d}t - 1_{(0,\phi^*)}(\tau^*)e^{-\rho\tau^*}\ell\right] \\
&= -\mathsf{E}_{\underline{x}+\delta}^{\theta^V}\left[\int_0^{\tau^*\wedge\phi^*} e^{-\rho t}(\hat{c}\cdot 1_{(-\infty,0)}(\hat{X}_t^{\theta^V}) - \check{c}\cdot 1_{(0,\infty)}(\hat{X}_t^{\theta^V}))\mathrm{d}t + 1_{(0,\phi^*)}(\tau^*)e^{-\rho\tau^*}u\right] \\
&\underset{(3.5.5)}{<} -\mathsf{E}_{\overline{x}-\delta}^{\theta^V}\left[\int_0^{\eta^*\wedge\phi^*} e^{-\rho t}(\hat{c}\cdot 1_{(0,\infty)}(\hat{X}_t^{\theta^V}) - \check{c}\cdot 1_{(-\infty,0)}(\hat{X}_t^{\theta^V}))\mathrm{d}t + 1_{(0,\phi^*)}(\eta^*)e^{-\rho\eta^*}u\right] \\
&= -\overline{w}(\overline{x}-\delta).
\end{aligned}$$

Note that the inequality holds by (3.5.5) alone. This is due to symmetry in the marginal holding cost around zero, i.e.

$$\mathsf{P}_{\underline{x}-\delta}^{\theta^V}(\hat{X}_t^{\theta^V} < 0) = \Phi\left(\frac{-(\underline{x}-\delta)-(-\kappa\sigma)t}{\sigma\sqrt{t}}\right) = \Phi\left(\frac{(\overline{x}+\delta)+\kappa\sigma t}{\sigma\sqrt{t}}\right) = \mathsf{P}_{\overline{x}+\delta}^{\theta^V}(\hat{X}_t^{\theta^V} > 0),$$

where Φ is cumulative distribution function of the standard normal distribution, and that $\check{c} = \hat{c}$. Therefore, we can imply that $\ell = -\underline{w}(\underline{x}) > \overline{w}(\overline{x}) = u$, which is a contradiction. A similar technique can be easily adopted to examine a contradiction when $x^* = 0$ and $|\underline{x}| \neq |\overline{x}|$, so omitted for brevity. This completes the proof. ■

Using Lemma 1 we can derive several comparative statics results, starting from a symmetric base case. First we show that the continuation region $(\underline{x}, \overline{x})$ "shrinks" as κ increases.

Proposition 2 (Comparative static of ambiguity)**.** *Suppose that the model parameters are symmetric. Then, the control barriers $(\underline{x}, \overline{x})$ shrink as κ increases. That is, $\kappa \mapsto \underline{x}(\kappa)$ is non-decreasing and $\kappa \mapsto \overline{x}(\kappa)$ is non-increasing.*

Prove of Proposition 2. Suppose that κ_1, where $\kappa_0 \triangleq \kappa \leq \kappa_1$, and that denoted by $v_{\kappa_i}(x)$ is the first derivative with respect to x of the value function $V(x;\kappa_i)$, and $\theta_i^V \triangleq \kappa_i 1_{(-\infty,x^*)}(x) - \kappa_i 1_{[x^*,\infty)}(x)$ for each $\kappa_i, i = 0, 1$. According to the comparison theorem for Itō's processes (Karatzas and Shreve, 1991, Chapter 5.2, or Protter, 2010, Theorem 52), together with (3.4.7), one can see that

$$\hat{X}_t^{\theta_0^V} \geq \hat{X}_t^{\theta_1^V}, \text{ on } (-\infty, x^*), \text{ and } \hat{X}_t^{\theta_1^V} \geq \hat{X}_t^{\theta_0^V}, \text{ on } [x^*, \infty).$$

P_x-a.s. Since $\zeta(x)$ is non-decreasing in x and $\alpha'(x) = \sigma'(x) = 0$ for all $x \in E$, we have by Theorem 12 and the comparison theorem that on $(-\infty, x^*)$,

$$\begin{aligned}
v_{\kappa_0}(x) = V'(x) &= w_{\kappa_0}(x, \tau^*, \phi^*) \\
&= \mathsf{E}_x^{\theta_0^V}\left[\int_0^{\tau^*\wedge\phi^*} e^{-\rho t}\zeta(\hat{X}_t^{\theta_0^V})\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)e^{-\rho\tau^*}\ell\right]
\end{aligned}$$

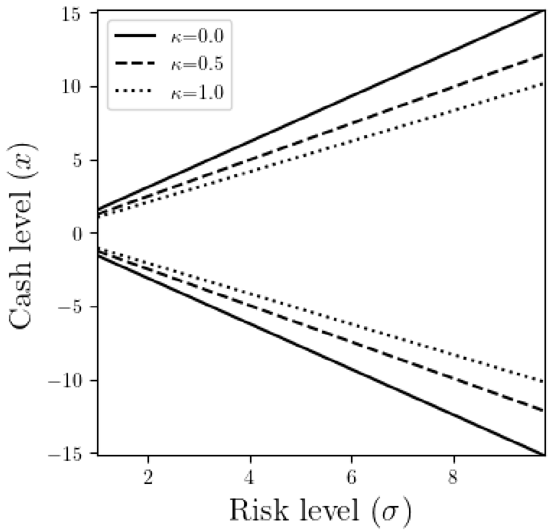

Figure 3.2: Control barriers as a function of σ in the symmetric case, using base-case parameters: $\alpha = 0$, $\rho = 0.1$, $\ell = u = 2$, and $\check{c} = \hat{c} = 1$. Each line stylesolid, dashed, and dottedrepresents control barriers for $\kappa \in \{0, 0.5, 1\}$, respectively. The upper line corresponds to the upper control barrier, $\overline{x}$, while the lower line represents the lower control barrier, $\underline{x}$. Given a fixed level of risk, when the cash level reaches $\overline{x}$ ($\underline{x}$), the DM must intervene to maintain the cash level below $\overline{x}$ (above $\underline{x}$), incurring a proportional cost u (ℓ).

$$
\begin{aligned}
&= \mathsf{E}_x\Big[\int_0^{\tau^*\wedge\phi^*} e^{-\rho t}\zeta(\hat{X}_t^0 - \kappa_0\sigma t)\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)e^{-\rho\tau^*}\ell\Big] \\
&\geq \mathsf{E}_x\Big[\int_0^{\tau^*\wedge\phi^*} e^{-\rho t}\zeta(\hat{X}_t^0 - \kappa_1\sigma t)\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)e^{-\rho\tau^*}\ell\Big] \\
&= \mathsf{E}_x^{\theta_1^V}\Big[\int_0^{\tau^*\wedge\phi^*} e^{-\rho t}\zeta(\hat{X}_t^{\theta_1^V})\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)e^{-\rho\tau^*}\ell\Big] \\
&= w_{\kappa_1}(x,\tau^*,\phi^*) = v_{\kappa_1}(x).
\end{aligned} \tag{3.5.6}
$$

Since this result holds for any $x < x^*$, it follows that $-\ell = v_{\kappa_0}(\underline{x}) \geq v_{\kappa_1}(\underline{x})$. Because $V_{\kappa_1}(x)$ is convex on E, we know that there exists $\delta > 0$ such that $v_{\kappa_1}(\underline{x} + \delta) = -\ell$. This means that $\underline{x}$ increases as κ increases. Since the parameters are symmetric, it follows from Lemma 1 that $v_{\kappa_1}(\underline{x} + \delta) = -v_{\kappa_1}(\overline{x} - \delta) = -u$, implying that $\overline{x}$ decreases as κ increases. Following a similar procedure, it is easy to show that the same result holds when $x > x^*$. In total, we conclude that $(\underline{x}, \overline{x})$ shrinks as the level of ambiguity rises. ■

Proposition 2 implies that a more ambiguous DM exerts control earlier and keeps a smaller cash inventory. This contrasts with the comparative statics of risk and is illustrated by Figure 3.2. This results reflects the pessimistic mindset of managers who lack full confidence in the drift of the inventory flow process. Mathematically speaking, κ-ignorance increases the drift when the value function is increasing and decreases it when the value function is decreasing. As a result, the manager faces higher expected holding costs near the control barriers. Thus, the optimal control policy implies that taking action earlier becomes less costly. However, this inevitably leads to a higher average frequency of control exertion, which, as previously discussed, contributes to an increase in the value function.

Next, we conduct a comparative static analysis for the remaining parameters under a fixed level of ambiguity κ. For simplicity, we start from the symmetric base case. We start by looking at changes in the discount rate.

Proposition 3 (Comparative statics of discounted rate)**.** *Suppose that the model parameters are symmetric. Then, the control barriers $(\underline{x}, \overline{x})$ expand as ρ increases. That is, $\rho \mapsto \underline{x}(\rho)$ is non-increasing and $\rho \mapsto \overline{x}(\rho)$ is non-decreasing.*

Proof of Proposition 3. Suppose that ρ_1 where $\rho_0 \triangleq \rho \leq \rho_1$ and denoted by $v_{\rho_i}(x)$ is the first derivative with respect to x of the value function $V(x; \rho_i)$, for each $\rho_i, i = 0, 1$. Applying the Itō's integration by parts to (3.4.6) yields

$$\begin{aligned}
v_{\rho_0}(x) \triangleq v(x) &= \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^* \wedge \phi^*} e^{-\rho_0 t}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)e^{-\rho_0\tau^*}\ell\Big] \\
&= \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^* \wedge \phi^*} (-\rho_0 v(\hat{X}_t^{\theta^V}) + \zeta(\hat{X}_t^{\theta^V}))\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)\ell\Big] \\
&\leq \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^* \wedge \phi^*} (-\rho_1 v(\hat{X}_t^{\theta^V}) + \zeta(\hat{X}_t^{\theta^V}))\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)\ell\Big] \\
&= \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^* \wedge \phi^*} e^{-\rho_1 t}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t - 1_{(0,\phi^*]}(\tau^*)e^{-\rho_1\tau^*}\ell\Big] \\
&= v_{\rho_1}(x)
\end{aligned}$$

for any $x \leq x^*$. The inequality holds because $v(x) \leq 0$ for all $x \leq x^*$. In turn, this implies that $-\ell = v_{\rho_0}(\underline{x}) \leq v_{\rho_1}(\underline{x})$. Since $v_{\rho_1}(x)$ is increasing in x, we can deduce that there exists $\delta > 0$ such that $v_{\rho_1}(\underline{x} - \delta) = -\ell$, meaning that $\underline{x}$ decreases as ρ increases. By the symmetry assumption, we deduce that $v_{\rho_1}(\underline{x} - \delta) = v_{\rho_1}(\overline{x} + \delta) = -u$, that is, $\overline{x}$ increases as ρ increases. This concludes that $(\underline{x}, \overline{x})$ expands as ρ increases. ■

This result implies that a more patient DM exerts less control. This happens because, from the DM's perspective, the worst-case prior is less bad, because expected future payments are discounted more.

The effect of a change in the reference drift is asymmetric: the inaction region shifts downwards.

Proposition 4 (Comparative statics of drift)**.** *Suppose that the model parameters are symmetric. Then the control barriers $(\underline{x}, \overline{x})$ shift downward as α increases. That is, $\alpha \mapsto \underline{x}(\rho)$ and $\alpha \mapsto \overline{x}(\rho)$ are non-increasing.*

Proof of Proposition 4. Suppose that α_1 where $\alpha_0 \triangleq \alpha \leq \alpha_1$ and denoted by $v_{\alpha_i}(x)$ is the first derivative with respect to x of the value function $V(x;\alpha_i)$, for each $\alpha_i, i = 0, 1$. Given that $\zeta(x)$ is non-decreasing in x and $\alpha'(x) = \sigma'(x) = 0$ for all $x \in E$, Theorem 12 implies that

$$\begin{aligned}
v_{\alpha_0}(x) &= \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^*\wedge\eta^*} e^{-\rho t}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t - 1_{(0,\eta^*)}(\tau^*)e^{-\rho\tau^*}\ell + 1_{(0,\tau^*)}(\eta^*)e^{-\rho\eta^*}u\Big] \\
&\leq \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^*\wedge\eta^*} e^{-\rho t}\zeta(\hat{X}_t^{\theta^V} + (\alpha_1 - \alpha_0)t)\mathrm{d}t - 1_{(0,\eta^*)}(\tau^*)e^{-\rho\tau^*}\ell + 1_{(0,\tau^*)}(\eta^*)e^{-\rho\eta^*}u\Big] \\
&= v_{\alpha_1}(x)
\end{aligned}$$

for any $x \in E$. In turn, this implies that $-\ell = v_{\alpha_0}(\underline{x}) \leq v_{\alpha_1}(\underline{x})$ and $u = v_{\alpha_0}(\overline{x}) \leq v_{\alpha_1}(\overline{x})$. Since $v_{\alpha_1}(x)$ is increasing in x, we can deduce that there exists $\delta, \delta' > 0$ such that $v_{\alpha_1}(\underline{x} - \delta) = -\ell$ and $v_{\alpha_1}(\overline{x} - \delta') = u$, meaning that $(\underline{x}, \overline{x})$ shift downward as α increases. ■

The intuition behind this result is as follows. An increasing reference drift makes it more likely, under the worst-case measure, that holding costs $\hat{c}$ need to be paid. That makes it more attractive to control the inventory on the upside, leading to a decrease in $\overline{x}$. Conversely, with higher α it is less likely that holding costs $\check{c}$ are incurred, which leads the DM to reduce the control barrier $\underline{x}$. The comparative statics for the holding costs follow a similar pattern.

Proposition 5 (Comparative statics of holding costs)**.** *Suppose that the model parameters are symmetric. Then the control barriers $(\underline{x}, \overline{x})$ shift upward as $\check{c}$ increases, but shift downward as $\hat{c}$ increases. That is, $\check{c} \mapsto \underline{x}(\check{c})$ and $\check{c} \mapsto \overline{x}(\check{c})$ are non-decreasing, while $\hat{c} \mapsto \underline{x}(\hat{c})$ and $\hat{c} \mapsto \overline{x}(\hat{c})$ are non-increasing.*

Proof of Proposition 5. Suppose that $\check{c}_1$ where $\check{c}_0 \triangleq \check{c} \leq \check{c}_1$ and denoted by $v_{\check{c}_i}(x)$ is the first derivative with respect to x of the value function $V(x;\check{c}_i)$, for each $\check{c}_i, i = 0, 1$. Then it follows from Theorem 12 that

$$\begin{aligned}
v_{\check{c}_0}(x) &= \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^*\wedge\eta^*} e^{-\rho t}(\hat{c}\cdot 1_{(0,\infty)}(\hat{X}_t^{\theta^V}) - \check{c}_0\cdot 1_{(-\infty,0)}(\hat{X}_t^{\theta^V}))\mathrm{d}t \\
&\qquad\qquad - 1_{(0,\eta^*]}(\tau^*)e^{-\rho\tau^*}\ell + 1_{(0,\tau^*)}(\eta^*)e^{-\rho\eta^*}u\Big] \\
&\geq \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^*\wedge\eta^*} e^{-\rho t}(\hat{c}\cdot 1_{(0,\infty)}(\hat{X}_t^{\theta^V}) - \check{c}_1\cdot 1_{(-\infty,0)}(\hat{X}_t^{\theta^V}))\mathrm{d}t \\
&\qquad\qquad - 1_{(0,\eta^*]}(\tau^*)e^{-\rho\tau^*}\ell + 1_{(0,\tau^*)}(\eta^*)e^{-\rho\eta^*}u\Big] \\
&= v_{\check{c}_1}(x),
\end{aligned}$$

for all $x \in E$, which implies that $-\ell = v_{\check{c}_0}(\underline{x}) \geq v_{\check{c}_1}(\underline{x})$ and $u = v_{\check{c}_0}(\overline{x}) \geq v_{\check{c}_1}(\overline{x})$. As a result, there is $\delta, \delta' > 0$ such that $v_{\check{c}_1}(\underline{x} + \delta) = -\ell$ and $v_{\check{c}_1}(\overline{x} + \delta') = u$, since $v_{\check{c}_1}(x)$ does not decrease in x. Therefore, $(\underline{x}, \overline{x})$ shift upward as $\check{c}$ increases. One can use the same argument to show that $(\underline{x}, \overline{x})$ shift downward as $\hat{c}$ increases. ■

Proposition 6 (Comparative statics of control costs)**.** *Suppose that the model parameters are symmetric. Then the control barriers $(\underline{x}, \overline{x})$ expand as ℓ or u increase. That is, $(\ell, u) \mapsto \underline{x}(\ell, u)$ is non-increasing, while $(\ell, u) \mapsto \overline{x}(\ell, u)$ is non-decreasing.*

Proof of Proposition 6. Suppose that ℓ_1 where $\ell_0 \triangleq \ell \leq \ell_1$ and denoted by $v_{\ell_i}(x)$ is the first derivative with respect to x of the value function $V(x; \ell_i)$, for each $\ell_i, i = 0, 1$. If $x \leq x^*$, then it follows from Corollary 1 that

$$\begin{aligned}
v_{\ell_0}(x) + \ell_0 &= \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^* \wedge \phi^*} e^{-\rho t}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t + 1_{(0,\phi^*]}(\tau^*)(1 - e^{-\rho\tau^*})\ell_0 + 1_{[\phi^*,\infty)}(\tau^*)\ell_0\Big] \\
&\leq \mathsf{E}_x^{\theta^V}\Big[\int_0^{\tau^* \wedge \phi^*} e^{-\rho t}\zeta(\hat{X}_t^{\theta^V})\mathrm{d}t + 1_{(0,\phi^*]}(\tau^*)(1 - e^{-\rho\tau^*})\ell_1 + 1_{[\phi^*,\infty)}(\tau^*)\ell_1\Big] \\
&= v_{\ell_1}(x) + \ell_1.
\end{aligned}$$

Then we have $x = \underline{x}$ that $0 = v_{\ell_0}(\underline{x}) + \ell_0 \leq v_{\ell_1}(\underline{x}) + \ell_1$. Because $v_{\ell_1}(x)$ increases in x, there exists $\delta > 0$ such that $v_{\ell_1}(\underline{x} - \delta) + \ell_1 = 0$. This means that $\ell \mapsto \underline{x}(\ell)$ is decreasing. Moreover, we have the symmetry assumption from Lemma 1 that $0 = v_{\ell_1}(\underline{x} - \delta) + \ell_1 = -v_{\ell_1}(-\underline{x} + \delta) - \ell_1 \leq -v_{\ell_1}(-\underline{x} + \delta) - u$. Therefore, One can infer the existence of $\delta' > 0$ such that $v_{\ell_1}(-\underline{x} + \delta + \delta') + u = 0$, implying that $\ell \mapsto \overline{x}(\ell)$ is increasing. This concludes that $(\underline{x}, \overline{x})$ expand as ℓ increases. A similar result holds in the case of increasing u, using an analogous argument. It is also clear that the proposition holds when $x > x^*$, after repeating the same procedure. Hence, the proof is complete. ■

When the control cost on one side is higher, the DM responses by delaying the exertion of control on that side, reducing the frequency of action to minimize the running cost. Furthermore, the DM also delays control on the opposite side, in expectation, to prevent the cash level from reaching a threshold where activating the more expensive control becomes necessary. This behavior leads to an expansion of the control barriers, as demonstrated in Proposition 6.

We next explore the effects arising from asymmetries, such as: non-zero drift, unequal holding costs for positive and negative cash balances, and unequal control costs, for different levels of ambiguity.

For simplicity, we suppose from now on that $\kappa_1 \geq \kappa_0$, where $\kappa_0 \triangleq \kappa$, and that denoted by $v_{\kappa_i}(x)$ the marginal value function (3.4.6) for each $\kappa_i, i = 0, 1$.

Proposition 7 (Comparative statics of ambiguity with non-zero drift)**.** *Suppose that α is non-zero while the other parameters remain symmetric. Then, the continuation region $(\underline{x}, \overline{x})$ shrinks as κ increases. That is, $\kappa \mapsto \underline{x}(\kappa)$ is non-decreasing and $\kappa \mapsto \overline{x}(\kappa)$ is non-increasing.*

Proof of Proposition 7. It is easy to see from Proposition 2 that (3.5.6) holds, even if α is non-zero. Therefore, it holds that $v_{\kappa_0}(x) \geq v_{\kappa_1}(x)$ for any $x < x^*$. Similarly, it also follows that $v_{\kappa_0}(x) \leq v_{\kappa_1}(x)$ on which $x > x^*$. This means that there are $\delta, \delta' > 0$ such that $v_{\kappa_1}(\underline{x} + \delta) = -\ell$

and $v_{\kappa_1}(\overline{x} - \delta') = u$. Hence, we conclude that $(\underline{x}, \overline{x})$ shrink as κ increases for any non-zero drift. ■

With a similar reasoning to the proof of Proposition 7, we obtain the following proposition.

Proposition 8 (Comparative statics of ambiguity with unequal holding costs)**.** *Suppose that $\check{c} \neq \hat{c}$, while the other parameters remain symmetric. Then, the control barriers $(\underline{x}, \overline{x})$ shrink as κ increases. That is, $\kappa \mapsto \underline{x}(\kappa)$ is non-decreasing and $\kappa \mapsto \overline{x}(\kappa)$ is non-increasing.*

Propositions 7 and 8 imply that the "shrink" effect caused by increasing ambiguity persists even in the presence of non-zero drift or unequal holding costs for positive and negative cash balances. However, this effect does not hold when unequal control costs are considered, resulting in a non-monotonicity established in the following proposition.

Proposition 9 (Comparative statics of ambiguity with unequal control costs)**.** *Suppose that $\ell \neq u$, while the other parameters remain symmetric. Then, the control barriers $(\underline{x}, \overline{x})$ shrink as κ increases if $|\ell - u|$ small enough. Otherwise, there exist $\epsilon > 0$ such that $(\underline{x}, \overline{x})$ shift downward if $\ell - u < \epsilon$, or shift upward if $\ell - u > \epsilon$, as κ increases.*

Proof of Proposition 9. Suppose that $\ell = u + \epsilon$ for some $\epsilon > 0$. By Lemma 1, we have $\ell = -v_{\kappa_0}(\underline{x}) = -v_{\kappa_1}(\underline{x} + \delta) = -v_{\kappa_1}(-\underline{x} - \delta) > \ell - \epsilon = u$ for some $\delta > 0$. Proposition 6 then implies that there exists $\delta' > 0$ such that $-v_{\kappa_1}(-\underline{x} - \delta + \delta') = u$. It is straightforward to observe that the mapping $\epsilon \mapsto \delta'(\epsilon)$ is increasing with $\delta'(0) = 0$. Consequently, there exist $\underline{\epsilon}, \overline{\epsilon} > 0$ such that $\overline{\epsilon} > \underline{\epsilon}$ and $\delta'(\underline{\epsilon}) < \delta < \delta'(\overline{\epsilon})$. This implies that if $u - \ell$ is sufficiently small, i.e., $\epsilon = \underline{\epsilon}$, then $(\underline{x}, \overline{x})$ shrink. Conversely, if $u - \ell$ is sufficiently large, i.e., $\epsilon = \overline{\epsilon}$, the control barriers lead to an upward shift. Similarly, it can be shown that $(\underline{x}, \overline{x})$ shift downward when $\ell - u$ becomes sufficiently large. ■

This result indicates that when the difference between upper and lower control costs becomes too large, the "shrink effect" from ambiguity no longer holds. As shown in Figures 3.2, 3.3a and 3.3b, when the cost of controlling the lower inventory becomes significantly higher than that of the upper, the upper barriers shift upward, contrary to the conventional expectation that more ambiguous DM acts earlier. This occurs due to the mitigation effect of singular control. In the symmetric case, the DM exerts control earlier at either the upper or lower barrier. However, when the cost of controlling one barrier is significantly lower than the other, the optimal policy suggests delaying the cheaper action to reduce the likelihood of triggering the more costly control. In other words, there exists a scenario where the optimal singular control policy offsets the impact of ambiguity, illustrating the duality between singular control and ambiguity.

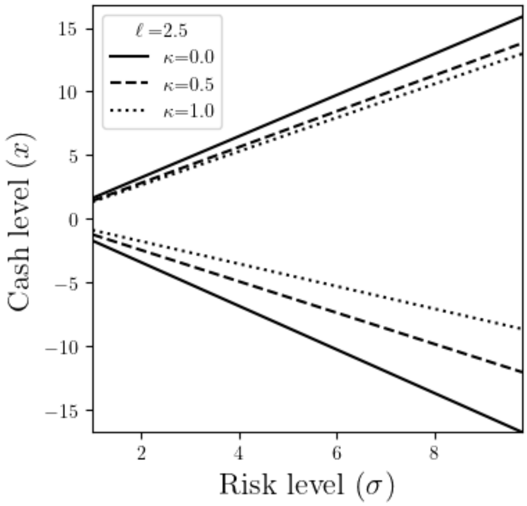

(a) Control barriers given $\ell = 2.5$ for $\kappa \in \{0, 0.5, 1\}$.

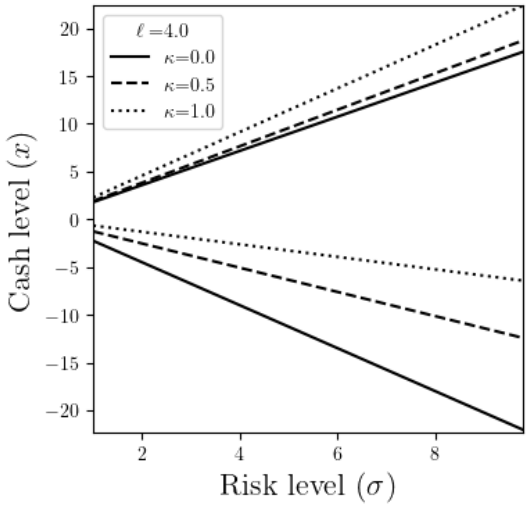

(b) Control barriers given $\ell = 4$ for $\kappa \in \{0, 0.5, 1\}$.

Figure 3.3: Control barriers as a function of σ for $\kappa \in \{0, 0.5, 1\}$, with fixed parameters $\alpha = 0$, $\rho = 0.1$, $u = 2$, $\check{c} = \hat{c} = 1$. Panels (3.3a) and (3.3b) displays the control barriers for $\ell = 2.5$ and $\ell = 4$, respectively.

3.6 MANAGERIAL IMPLICATION

According to the earlier argument, firms operating under the model with a presumably known prior tend to incur lower average costs for holding cash. This naturally raises a question: why should ambiguity be taken into account if doing so results in higher operational costs? From the standpoint of revenue optimization, relying on the classical model may initially seem more attractive, as it appears to minimize control and holding expenses.

However, this conclusion only holds if the assumed probability distribution accurately reflects the environment the firm faces. The key managerial insight emerges when a firm disregards ambiguity but then encounters unfavorable outcomes. In such situations, the cost of ignoring uncertainty can become significantly higher than anticipated. For example, consider Figure 3.4a. The blue line represents the standard model, with dotted control barriers indicating the optimal strategy under perceived known risk. In contrast, the red line shows the cash trajectory under the worst-case scenario, guided by dashed control barriers that account for ambiguity. These paths reveal that, under the worst-case prior, the firm's cash position tends to shift more rapidly, rising when the marginal cost of holding cash is high ($X_t > x^*$) and falling when it is low ($X_t \leq x^*$).

If the firm continues to rely on the standard model while facing a worst-case scenario, it will respond suboptimally. This leads to more frequent interventions and higher holding costs, as illustrated in the top-left and bottom-right panels of Figure 3.4a. In other words, while the standard model may appear cost-effective ex ante, as indicated by the solid blue line of Figure 3.4b, it can

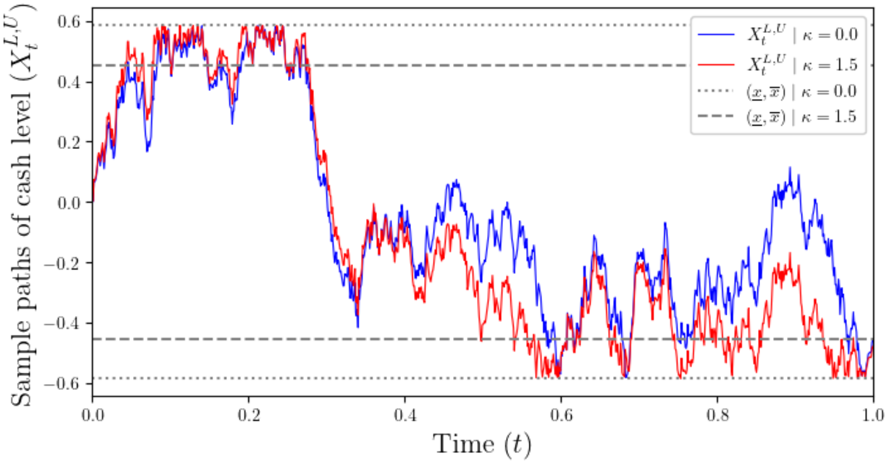

(a) Sample paths

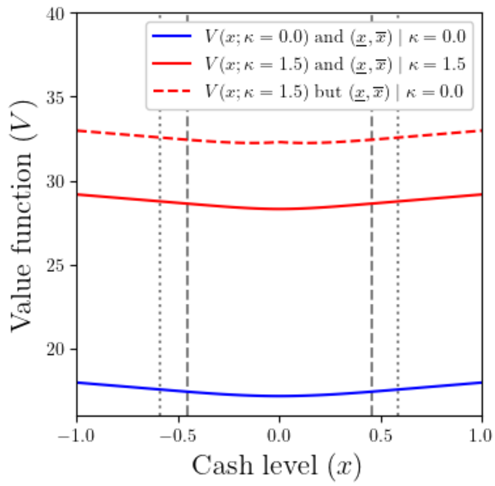

(b) Value functions

Figure 3.4: Panel (3.4a) shows sample paths of cash levels without ambiguity ($\kappa = 0$, blue line) and with ambiguity, ($\kappa = 1.5$, red line), using base-case parameters: $\alpha = 0$, $\rho = 0.1$, $\sigma = 0.5$, $\ell = u = 2$, and $\check{c} = \hat{c} = 1$. The value $X_0 = 0$ is fixed for both cases. The dotted and dashed gray lines correspond to the control barriers for $\kappa = 0$ and $\kappa = 1.5$, respectively. The associated value functions are displayed in Panel (3.4b), where the solid blue and red lines are for the cases $\kappa = 0$ and $\kappa = 1.5$, respectively. The dashed red line is the cost function when the worse-case happens ($\kappa = 1.5$), but the DM uses the control policy belonging to that of reference prior ($\kappa = 0$).

lead to greater losses when uncertainty proves more severe than expected, as depicted by the dashed red line in the same figure. By contrast, adopting an ambiguity-aware approach, such as the maxmin utility model, allows the firm to anticipate such a complication and respond more proactively. Although this requires earlier or more frequent interventions, as shown by the higher value function represented by the solid red line in Figure 3.4b, it acts as a form of precautionary adjustment. Importantly, the average running cost under this approach is lower than that of the standard model that ignores ambiguity. One can interpret this surplus as an *ambiguity premium*, a strategic cost paid upfront to mitigate future downside uncertainty.

This highlights a critical managerial implication: when ambiguity exerts a first-order effect on cash management, it is no longer a secondary consideration but a central determinant of optimal policy. Firms that incorporate ambiguity into their decision-making frameworks are better positioned to protect firm value in uncertain or poorly understood environments.

3.7 CONCLUSION

In this paper, we revisit the classical problem of two-sided singular control in the optimal cash reserve problem where the net cash position evolves according to an Itō diffusion. We extend this model to allow for managerial ambiguity within the multiple prior framework with κ-ignorance and maxmin expected utility. We establish a verification theorem for the minimal expected present value for holding and control costs under the worst-case prior, as well as an optimal cash holding policy. A Dynkin game approach is used to explore the effects of the level of ambiguity (and other parameters) on the optimal control of cash holdings.

From a managerial perspective, our most important observation is that ambiguity increases the frequency with which control is exerted. This is due to the fact that under the worst-case prior the manager expects higher holding costs, which makes exerting control relatively cheaper. This results, on average, in a smaller cash inventory, which is the opposite effect to an increase in risk. If risk, as measured by the variance of the net cash flow, increases, then the standard "option value of waiting" (cf., Dixit and Pindyck, 1994) increases which implies that, typically, control is exerted later. This results, in expectation, in a larger cash inventory.

Our work suggests several avenues for future research. First, we note that our findings are purely theoretical, and the numerical results presented serve only as illustrative examples. However, because our model is both analytically tractable and robust for any range of model parameters, comparable to the seminal work of Nishimura and Ozaki (2007), it can readily be used as a benchmark for empirical validation using real-world economic data. We believe that this would be a valuable avenue for future research.

Secondly, the assumption of proportional control costs, such as for equity issuance or dividend payments, may not fully reflect the complexity of real-world cost structures, which could involve both variable and fixed components. The existence of such fixed costs leads to a problem of *impulse control*. For an overview of the method, see, e.g., Harrison (2013, Chapter 7).

Thirdly, one of the assumptions of our model is that the manager does not learn about the set of priors. It is as if the manager is confronted with a new Ellsberg urn at every point in time. In some real-world situations it may be more realistic to assume that the manager is confronted with the *same* Ellsberg urn at every point in time. This then opens up the possibility of managerial learning about ambiguity.

Finally, our model of κ-ambiguity describes a fairly extreme version of cautious behavior. A more realistic version of the model would allow the manager to average over multiple priors. That would naturally lead to the smooth ambiguity model of Klibanoff et al. (2005).

3.8 APPENDIX

A PROOF OF PROPOSITION 11

First note that the constants A and B depend on $\underline{x}$, whereas the constants C and D depend on $\overline{x}$. In what follows we will make this dependence explicit by writing, say, the constant A, as a mapping $\underline{x} \mapsto A(\underline{x})$. In fact, for given $\underline{x}$ and $\overline{x}$, the systems of linear equations

$$\begin{bmatrix} \hat{V}'_{-\kappa}(\underline{x}+) & \check{V}'_{-\kappa}(\underline{x}+) \\ \hat{V}''_{-\kappa}(\underline{x}+) & \check{V}''_{-\kappa}(\underline{x}+) \end{bmatrix} \begin{bmatrix} A(\underline{x}) \\ B(\underline{x}) \end{bmatrix} = \begin{bmatrix} -\ell - R'_{-\kappa}(\underline{x}) \\ -R''_{-\kappa}(\underline{x}) \end{bmatrix}, \tag{A.1}$$

and

$$\begin{bmatrix} \hat{V}'_{+\kappa}(\overline{x}-) & \check{V}'_{+\kappa}(\overline{x}-) \\ \hat{V}''_{+\kappa}(\overline{x}-) & \check{V}''_{+\kappa}(\overline{x}-) \end{bmatrix} \begin{bmatrix} C(\overline{x}) \\ D(\overline{x}) \end{bmatrix} = \begin{bmatrix} u - R'_{+\kappa}(\overline{x}) \\ -R''_{+\kappa}(\overline{x}) \end{bmatrix}, \tag{A.2}$$

have unique solutions:

$$\begin{aligned}
A(\underline{x}) &= \frac{\check{V}''_{-\kappa}(\underline{x}+)(-\ell - R'_{-\kappa}(\underline{x})) + \check{V}'_{-\kappa}(\underline{x}+)R''_{-\kappa}(\underline{x})}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}, \\
B(\underline{x}) &= \frac{-\hat{V}''_{-\kappa}(\underline{x}+)(-\ell - R'_{-\kappa}(\underline{x})) - \hat{V}'_{-\kappa}(\underline{x}+)R''_{-\kappa}(\underline{x})}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}, \\
C(\overline{x}) &= \frac{\check{V}''_{+\kappa}(\overline{x}-)(u - R'_{+\kappa}(\overline{x})) + \check{V}'_{+\kappa}(\overline{x}-)R''_{+\kappa}(\overline{x})}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)}, \\
D(\overline{x}) &= \frac{-\hat{V}''_{+\kappa}(\overline{x}-)(u - R'_{+\kappa}(\overline{x})) - \hat{V}'_{+\kappa}(\overline{x}-)R''_{+\kappa}(\overline{x})}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)}.
\end{aligned}$$

Define the function φ on $(\underline{x}, \overline{x})$ as follows:

$$\varphi(x) = \begin{cases} R_{-\kappa}(x) + A(\underline{x})\hat{V}_{-\kappa}(x) + B(\underline{x})\check{V}_{-\kappa}(x) & \text{if } \underline{x} < x < x^* \\ R_{+\kappa}(x) + C(\overline{x})\hat{V}_{+\kappa}(x) + D(\overline{x})\check{V}_{+\kappa}(x) & \text{if } x^* \leq x < \overline{x}. \end{cases}$$

We establish that φ is convex on $(\underline{x}, \overline{x})$.

First, suppose that $x \in (\underline{x}, x^*)$ and that $x^* < 0$. It then holds that

$$\varphi''(x) = R''_{-\kappa}(x) + A(\underline{x})\hat{V}''_{-\kappa}(x) + B(\underline{x})\check{V}''_{-\kappa}(x)$$

$$
\begin{aligned}
&= R''_{-\kappa}(x) + \frac{\check{V}''_{-\kappa}(\underline{x}+)(-\ell - R'_{-\kappa}(\underline{x})) + \check{V}'_{-\kappa}(\underline{x}+)R''_{-\kappa}(\underline{x})}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\hat{V}''_{-\kappa}(x) \\
&\quad + \frac{-\hat{V}''_{-\kappa}(\underline{x}+)(-\ell - R'_{-\kappa}(\underline{x})) - \hat{V}'_{-\kappa}(\underline{x}+)R''_{-\kappa}(\underline{x})}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\check{V}''_{-\kappa}(x) \\
&= \check{c}f''_{-\kappa}(x) + \hat{E}_{-\kappa}\hat{V}''_{-\kappa}(x) \\
&\quad + \frac{\check{V}''_{-\kappa}(\underline{x}+)(-\ell - \check{c}f'_{-\kappa}(x)(\underline{x}) - \hat{E}_{-\kappa}\hat{V}'_{-\kappa}(\underline{x}+))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\hat{V}''_{-\kappa}(x) \\
&\quad + \frac{\check{V}'_{-\kappa}(\underline{x}+)(\check{c}f''_{-\kappa}(x)(\underline{x}) + \hat{E}_{-\kappa}\hat{V}''_{-\kappa}(\underline{x}+))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\hat{V}''_{-\kappa}(x) \\
&\quad - \frac{\hat{V}''_{-\kappa}(\underline{x}+)(-\ell - f'_{-\kappa}(\underline{x}) - \hat{E}_{-\kappa}\hat{V}'_{-\kappa}(\underline{x}+))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\check{V}''_{-\kappa}(x) \\
&\quad - \frac{\hat{V}'_{-\kappa}(\underline{x}+)(\check{c}f''_{-\kappa}(x)(\underline{x}) + \hat{E}_{-\kappa}\hat{V}''_{-\kappa}(\underline{x}+))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\check{V}''_{-\kappa}(x) \\
&= \left[\frac{\check{V}''_{-\kappa}(\underline{x}+)(-\ell - \check{c}f'_{-\kappa}(x)(\underline{x}))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\right]\hat{V}''_{-\kappa}(x) \\
&\quad - \left[\frac{\hat{V}''_{-\kappa}(\underline{x}+)(-\ell - \check{c}f'_{-\kappa}(x)(\underline{x}))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\right]\check{V}''_{-\kappa}(x) \\
&= \left[\frac{(-\ell - \check{c}f'_{-\kappa}(x)'(\underline{x}))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\right] \\
&\quad \times (\check{V}''_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(x) - \hat{V}''_{-\kappa}(\underline{x}+)\check{V}''_{-\kappa}(x)) \\
&\geq 0.
\end{aligned}
$$

Here, the last inequality holds due to (i) assumption (3.4.1), (ii)

$$\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)$$

is strictly positive and (iii) non-negativity of the term

$$\check{V}''_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(x) - \hat{V}''_{-\kappa}(\underline{x}+)\check{V}''_{-\kappa}(x)$$

on $(\underline{x}, x^*)$. This last part can be seen as follows; the term is zero at $\underline{x}+$, and is increasing in x on $(\underline{x}, x^*)$.

When $x^* \geq 0$, the same result holds on $(\underline{x}, 0)$, so in that case we only need to prove that φ is convex on $[0, x^*)$. Unlike the previous proof that the sign of $A(\underline{x})$ does not affect the convexity of φ, it does matter in this case. Since $B(\underline{x})$ is always non-positive, we thus separate the proof into two cases: $A(\underline{x}) \geq 0$ and $A(\underline{x}) < 0$. In what follows, we use the fact that on $[0, x^*)$ it holds that $R_{-\kappa}(x) = \hat{c}f_{-\kappa}(x) + \check{E}\check{V}_{-\kappa}(x)$.

Case 1: when $A(\underline{x}) \geq 0$, it holds that

$$\begin{aligned}
\varphi''(x) &= R''_{-\kappa}(x) + A(\underline{x})\hat{V}''_{-\kappa}(x) + B(\underline{x})\check{V}''_{-\kappa}(x) \\
&= \hat{c}f''_{-\kappa}(x) + \check{E}_{-\kappa}\check{V}''_{-\kappa}(x) + A(\underline{x})\hat{V}''_{-\kappa}(0)\frac{\hat{V}''_{-\kappa}(x)}{\hat{V}''_{-\kappa}(0)} + B(\underline{x})\check{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&\geq \check{E}_{-\kappa}\check{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} + A(\underline{x})\hat{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} + B(\underline{x})\check{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&= \hat{E}_{-\kappa}\hat{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&\quad + \frac{\check{V}''_{-\kappa}(\underline{x}+)(\ell - \check{c}f'_{-\kappa}(\underline{x}) - \hat{E}_{-\kappa}\hat{V}'_{-\kappa}(\underline{x}+)) + \check{V}'_{-\kappa}(\underline{x}+)\hat{E}_{-\kappa}\hat{V}''_{-\kappa}(\underline{x}+)}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)} \\
&\quad \times \hat{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&\quad - \frac{\hat{V}''_{-\kappa}(\underline{x}+)(\ell - \check{c}f'_{-\kappa}(\underline{x}) - \hat{E}_{-\kappa}\hat{V}'_{-\kappa}(\underline{x}+)) + \hat{V}'_{-\kappa}(\underline{x}+)\hat{E}_{-\kappa}\hat{V}''_{-\kappa}(\underline{x}+)}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)} \\
&\quad \times \check{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&= \hat{E}_{-\kappa}\hat{V}''_{-\kappa}(0)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&\quad + \left[\frac{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(0)(\ell - \check{c}f'_{-\kappa}(\underline{x}))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)} - \hat{E}_{-\kappa}\hat{V}''_{-\kappa}(0)\right]\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&\quad - \left[\frac{\hat{V}''_{-\kappa}(\underline{x}+)\check{V}''_{-\kappa}(0)(\ell - \check{c}f'_{-\kappa}(\underline{x}))}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}\right]\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&= \left[\frac{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(0) - \hat{V}''_{-\kappa}(\underline{x}+)\check{V}''_{-\kappa}(0)}{\check{V}''_{-\kappa}(\underline{x}+)\hat{V}'_{-\kappa}(\underline{x}+) - \check{V}'_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(\underline{x}+)}(\ell - \check{c}f'_{-\kappa}(\underline{x}))\right]\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)} \\
&\geq 0.
\end{aligned}$$

The first inequality holds since $\frac{\hat{V}''_{-\kappa}(x)}{\hat{V}''_{-\kappa}(0)} \geq \frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(0)}$ for all $x \geq 0$. The last equality obtains from the the fact that $f'_{-\kappa}$ is constant, together with (3.4.1), i.e. $\hat{E}_{-\kappa}\hat{V}''_{-\kappa}(0) = \check{E}_{-\kappa}\check{V}''_{-\kappa}(0)$. Condition (3.4.1) and the fact that $\check{V}''_{-\kappa}(\underline{x}+)\hat{V}''_{-\kappa}(0) - \hat{V}''_{-\kappa}(\underline{x}+)\check{V}''_{-\kappa}(0) \geq 0$ give the last inequality.

Case 2: when $A(\underline{x}) < 0$, it holds that

$$\begin{aligned}
\varphi''(x) &= R''_{-\kappa}(x) + A(\underline{x})\hat{V}''_{-\kappa}(x) + B(\underline{x})\check{V}''_{-\kappa}(x) \\
&= \hat{c}f''_{-\kappa}(x) + \check{E}_{-\kappa}\check{V}''_{-\kappa}(x) + A(\underline{x})\hat{V}''_{-\kappa}(x^*)\frac{\hat{V}''_{-\kappa}(x)}{\hat{V}''_{-\kappa}(x^*)} + B(\underline{x})\check{V}''_{-\kappa}(x^*)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)}
\end{aligned}$$

$$
\begin{aligned}
&\geq \check{E}_{-\kappa}\check{V}''_{-\kappa}(x^*)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} + A(\underline{x})\hat{V}''_{-\kappa}(x^*)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} + B(\underline{x})\check{V}''_{-\kappa}(x^*)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} \\
&= \left(\check{E}_{-\kappa}\check{V}''_{-\kappa}(x^*) + A(\underline{x})\hat{V}''_{-\kappa}(x^*) + B(\overline{x})\check{V}''_{-\kappa}(x^*)\right)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} \\
&= \left(\check{E}_{+\kappa}\check{V}''_{+\kappa}(x^*) + C(\overline{x})\hat{V}''_{+\kappa}(x^*) + D(\overline{x})\check{V}''_{+\kappa}(x^*)\right)\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} \\
&= \Bigg[\check{E}_{+\kappa}\check{V}''_{+\kappa}(x^*) \\
&\quad + \frac{\check{V}''_{+\kappa}(\overline{x}-)(u - \hat{c}f'_{+\kappa}(\overline{x}) - \check{E}_{+\kappa}\check{V}'_{+\kappa}(\overline{x})) + \check{V}'_{+\kappa}(\overline{x}-)\check{E}_{+\kappa}\check{V}''_{+\kappa}(\overline{x})}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)}\hat{V}''_{+\kappa}(x^*) \\
&\quad - \frac{\hat{V}''_{+\kappa}(\overline{x}-)(u - \hat{c}f'_{+\kappa}(\overline{x}) - \check{E}_{+\kappa}\check{V}'_{+\kappa}(\overline{x})) + \hat{V}'_{+\kappa}(\overline{x}-)\check{E}_{+\kappa}\check{V}''_{+\kappa}(\overline{x})}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)}\check{V}''_{+\kappa}(x^*)\Bigg] \\
&\quad \times \frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} \\
&= \Bigg[\check{E}_{+\kappa}\check{V}''_{+\kappa}(x^*) + \frac{\check{V}''_{+\kappa}(\overline{x}-)(u - \hat{c}f'_{+\kappa}(\overline{x}))}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)}\hat{V}''_{+\kappa}(x^*) \\
&\quad - \left(\frac{\hat{V}''_{+\kappa}(\overline{x}-)(u - \hat{c}f'_{+\kappa}(\overline{x}))}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)} + \check{E}_{+\kappa}\right)\check{V}''_{+\kappa}(x^*)\Bigg]\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} \\
&= \left[\frac{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(x^*) - \hat{V}''_{+\kappa}(\overline{x}-)\check{V}''_{+\kappa}(x^*)}{\check{V}''_{+\kappa}(\overline{x}-)\hat{V}'_{+\kappa}(\overline{x}-) - \check{V}'_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(\overline{x}-)}(u - \hat{c}f'_{+\kappa}(\overline{x}))\right]\frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)} \\
&\geq 0.
\end{aligned}
$$

Here, the first inequality follows from the fact that $\frac{\hat{V}''_{-\kappa}(x)}{\hat{V}''_{-\kappa}(x^*)} \leq \frac{\check{V}''_{-\kappa}(x)}{\check{V}''_{-\kappa}(x^*)}$ for all $x \geq 0$. The fourth equality follows from (3.4.4). From $\check{V}''_{+\kappa}(\overline{x}-)\hat{V}''_{+\kappa}(x^*) - \hat{V}''_{+\kappa}(\overline{x}-)\check{V}''_{+\kappa}(x^*) \geq 0$, together with (3.4.1), we obtain that φ'' is non-negative in this case too.

Hence, we conclude that φ is convex on $(\underline{x}, x^*)$. By a similar argument, it can be shown that φ is convex on $[x^*, \overline{x})$. This establishes convexity of φ on $(\underline{x}, \overline{x})$.

It, therefore, holds that φ is decreasing on $(\underline{x}, x^*)$ and increasing on $(x^*, \overline{x})$. Direct verification then shows that Condition 1 of Proposition 10 is satisfied. The second condition is satisfied by assumption. The third condition is satisfied by the choice of $A(\underline{x}), B(\underline{x}), C(\overline{x}), D(\overline{x})$, as given by (A.1) and (A.2), respectively. Conditions 4 and 5 are satisfied due to (3.4.1) and transversality (condition 6) is also trivially satisfied given requirement (3.2.2). Hence, all assumptions of Proposition 10 are satisfied and the conclusion follows.

— 4 —

Optimal Water Reservoir Management under Climate Uncertainty: A Singular Control Model with Smooth Ambiguity

Abstract

This paper studies optimal dam reservoir management under climate-driven ambiguity. We model inflows as a mean-reverting process with precipitation trends captured by a Bernoulli-distributed hidden variable, and ambiguity is incorporated through smooth ambiguity preferences. This framework unifies standard learning, smooth ambiguity, and maxmin utility within a single model, providing flexibility in representing decision makers attitudes toward uncertainty. We develop a numerical scheme based on Markov chain approximation and coordinate transformation to solve the associated Hamilton-Jacobi-Bellman equation. Comparative statics show that ambiguity aversion leads to earlier and more cautious interventions when information is scarce, reducing the risk of costly possible extreme drought or flood events. Unlike the maxmin framework, smooth ambiguity allows decision makers to revise and diminish aversion as more data are observed, maintaining a balance between robustness and learning. These findings highlight how climate-induced ambiguity reshapes reservoir operation policies and long-run outcomes. Our results suggest actionable guidelines for adaptive and resilient water management under deep climate uncertainty.
Keywords: Climate Change, Dam Reservoir Management, Ambiguity, Singular Stochastic Control, Filtering Theory.

4.1 INTRODUCTION

Climate change has profoundly reshaped the landscape of water resource management by altering weather patterns, intensifying droughts, and amplifying extreme rainfall. These shifts fundamentally undermine the long-standing *stationarity* assumption (Milly et al., 2008) under which decision makers (DMs) have traditionally relied on historical data and conventional models (e.g., the real options framework of Dixit and Pindyck, 1994) to guide operational policies. In essence, climate change introduces *ambiguity*:[1] a condition where DMs face incomplete knowledge of future climate scenarios and cannot assign precise probabilities to alternative outcomes. This type of uncertainty goes far beyond traditional notions of risk (volatility) and creates profound challenges for operational decision-making across disciplines ranging from the physical sciences to economics and sociology (Walker et al., 2013).

This challenge is acute in many areas, including dam reservoir management, which is directly exposed to climate variability and extremes. Ignoring climate-driven ambiguity therefore risks systematic under-preparation for extreme droughts or floods, potentially resulting in costly economic losses and severe ecological damage. Incorporating ambiguity into reservoir control models is, thus, not merely a theoretical refinement but a practical necessity: it provides managers with decision frameworks that remain robust under uncertain climate regimes, supporting policies that secure water supply, safeguard ecosystems, and strengthen long-term resilience.

Dam reservoirs are designed with an optimal storage capacity that sustains a long-run stream level, balancing social, economic, and ecological benefits (cf. Huang et al., 2019; Niu and Shah, 2021; Null et al., 2024). This trade-off leads us to assume that there is an optimal operating state that maximizes overall benefits. Any deviation from this state risks significant losses.

On the one hand, when water levels become excessive, several costs arise. Excess water reduces the reservoirs buffering capacity against seasonal or extreme precipitation, increasing the risk of flooding if not promptly addressed (Prakash et al., 2015). Persistently high levels also accelerate sedimentation, which not only diminishes storage capacity (Randle et al., 2021) but also traps nutrients vital to downstream ecosystems, contributing to biodiversity loss and ecological degradation (Schmutz and Moog, 2018). On the other hand, when water levels fall into shortage, the impacts on social welfare are immediate and severe: reduced access to drinking water, diminished irrigation supply, and decreased hydropower generation (World Commission on Dams, 2000).

These dual risks underscore the need for well-defined intervention rules: releasing water before levels exceed critical thresholds, or replenishing when levels approach scarcity. Yet, applying such rules is not without costs nor without risk. Releasing water, for instance, may entail an opportunity cost from forgoing future use, while delaying outflows can generate penalty costs through temporary shortages or reduced hydropower generation. Such additional cost structures fall precisely within the domain of inventory control models, which rigorously balance inflows, outflows, holding costs, and intervention costs under uncertainty. Viewing dam reservoir management through the lens of inventory control is therefore not only natural but also essential for designing policies that safeguard both economic efficiency and ecological sustainability.

[1]Some studies also refer to ambiguity as *Knightian uncertainty*, acknowledging Frank Knights original distinction between risk and uncertainty in Knight (1921).

4.1.1 RELATED LITERATURE

There is a rich literature on inventory control, which provides a natural foundation for modeling reservoir operations. The classical starting point is Arrow et al. (1951), who studied inventory as a countable sequence of random variables with fixed reflecting barriers, analogous to reservoir storage being constrained by minimum and maximum capacity levels. This setup is often referred to as an *obstacle problem*. Later, Eppen and Fama (1969) extended the framework to situations where reflecting barriers are not fixed but instead chosen by the decision maker to minimize total costs, including both holding costs (inventory or water storage) and adjustment costs (releases or replenishments). This approach is now known as *singular control*.

The continuous-time analog of the obstacle problem was developed by Bather (1966), Constantinides (1976), and Vial (1972), where inventory demand evolves as a Brownian-motion-driven diffusion. Building on this, Harrison (1978) provided the first formal treatment of singular control in a continuous-time inventory framework. This seminal work has since been extended to more general settings by, among others, Bar-Ilan and Sulem (1995), Bensoussan et al. (2005), Dai and Yao (2013a, 2013b), Harrison and Taksar (1983), and Karatzas (1983). For a comprehensive overview of this line of research, see Harrison (2013).

Empirical studies of reservoir inflow and outflow dynamics, most notably Lucia and Schwartz (2002), Figuerola-Ferretti et al. (2024), and references therein, consistently confirm that these dynamics evolve stochastically, with a pronounced tendency toward mean reversion. Importantly, both the mean level and the speed of mean reversion are typically non-stationary, shaped by seasonal variations and fluctuating electricity demand throughout the year. These features make it essential to incorporate a mean-reverting process when modeling reservoir levels.

The existing literature on singular control of mean-reverting processes, such as Cadenillas et al. (2010), Ferrari and Vargiolu (2020), and Jiang et al. (2022), offers compatible tools that can be adapted to our setting. However, these works abstract from ambiguity as an additional layer of uncertainty, assuming instead that the DM operates under a single, known probability measure. While this assumption preserves analytical tractability, it risks producing suboptimal control policies. Evidence for this comes from Brekke (2009), who show that climate-driven shifts in hydrological processes substantially diminish the effectiveness of traditional water management strategies. In other words, relying solely on past information is insufficient for predicting future water dynamics and designing robust policies.

One way to address this issue is to treat ambiguity as a hidden variable and incorporate a learning procedure. Fletcher et al. (2019), in their study of reservoir infrastructure investment in Mombasa, Kenya, show that adopting Bayesian learning to capture regional climate shifts enables planners to dynamically adjust policies, thereby enhancing both flexibility and cost-effectiveness. In other words, improving knowledge about climate dynamics (i.e., reducing ambiguity) allows for more accurate predictions of future uncertainty and hence better decision-making strategies.

A related singular control model that incorporates learning, which we can adapt to our setting, is developed by Federico et al. (2023). They employ sequential detection to update the unknown trend, interpreted as a proxy for ambiguity, in a two-sided singular control problem under arithmetic Brownian motion. While such approaches demonstrate the power of learning, they generally overlook the role of the DMs attitude toward ambiguity (e.g., aversion or preference).

This omission is critical. As argued by Brugnach et al. (2025), conventional approaches to uncertainty in water resource management, focused primarily on refining data and predictive accuracy, are no longer sufficient under climate change. Based on evidence from Mediterranean water basins in Spain, they emphasize that uncertainty is not purely scientific but also deeply social and relational, shaped by conflicting values, diverse lived experiences, and varying interpretations of past, present, and future conditions. In this sense, climate change not only alters the underlying dynamics to be learned but also reshapes how DMs perceive the reliability of the learning process itself. Hence, in addition to Bayesian learning, a new framework is needed, one that explicitly accounts for various ambiguity attitudes rather than assuming neutrality.

The behavioral study on attitude toward ambiguity was first highlighted by Ellsberg (1961) in the so-called Ellsberg urn experiment, which showed that people typically prefer bets with known probabilities over bets with unknown ones, clear evidence of ambiguity aversion. Years later, Gilboa and Schmeidler (1989) formalized this behavior into a theoretical framework within subjective expected utility, by which they introduce the maxmin utility model, where decision makers evaluate outcomes against the worst-case prior. While mathematically tractable and widely applied in economics, finance, and operations research (e.g., Archankul et al., 2025; Chen and Epstein, 2002; Cheng and Riedel, 2013; Hellmann and Thijssen, 2018; Nishimura and Ozaki, 2007; Thijssen, 2011), the maxmin approach is restrictive: it ignores the possibility that beliefs can be updated over time. For example, Archankul et al. (2025) show in a cash-reserve control problem that ambiguity aversion leads to earlier interventions, but their framework cannot capture how managers learn from new information.

The smooth ambiguity model of Klibanoff et al. (2005) was developed to overcome this limitation. It allows Bayesian learning while adjusting posterior beliefs according to a DMs' attitude toward ambiguity. Unlike maxmin preferences, which always lock onto the most pessimistic scenario, smooth ambiguity blends beliefs across different priors, creating a richer and more flexible representation. This feature is especially relevant for reservoir management, where decision makers must adapt policies as climate patterns unfold but still account for their own attitudes toward uncertainty.

Originally proposed in a static setting, smooth ambiguity was extended to dynamic problems by Klibanoff et al. (2009) using recursive utility of Epstein and Zin (1989), where beliefs are updated step by step through Bayesian learning. Challenges arise in continuous time, where ambiguity effects can fade out (Skiadas, 2013), but later work shows how they can be preserved. Gindrat and Lefoll (2011) introduced a time-dependent certainty-equivalence operator, while Hansen and

Sargent (2011) linked smooth ambiguity to a backward stochastic differential equation (BSDE) representation, assuming that the ambiguity resembles Gaussian-generated hidden variables which can be updated via the Kalman-Bucy filter (Liptser & Shiryaev, 2013b, Chapter 12). A logarithmic transformation (Fleming & Soner, 2006) then connects this framework back to recursive preferences, ensuring that ambiguity attitudes remain even in continuous-time decision-making.

4.1.2 CONTRIBUTION

This paper investigates the optimal management of a dam reservoir within a singular stochastic inventory control framework, extending the seminal work of Harrison and Taksar (1983) where the underlying dynamics follow an Ornstein-Uhlenbeck (OU) diffusion. Ambiguity arises because the DM cannot fully ascertain precipitation trends, shaped by climate variability, which we model through a Bernoulli-distributed hidden variable representing two possible regimes: drought or flood. To capture the DMs attitude toward such ambiguity, we adopt the smooth ambiguity preferences and develop the approach of Hansen and Sargent (2011) to allow recursive learning and updating of ambiguity attitude consistent with dynamic programming principles. The papers main objective is to examine how ambiguity attitudes interact with learning dynamics and the cost-benefit trade-offs inherent in reservoir management.

Our contributions to the literature are threefold. First, we develop a rigorous mathematical framework that incorporates finite-state, learnable ambiguity into singular stochastic control under smooth ambiguity preferences. Within this framework, we derive a forward-backward stochastic differential equation (FBSDE) and provide a verification theorem through Hamilton-Jacobi-Bellman (HJB) equations in the viscosity sense (Crandall & Lions, 1983), yielding an optimal control policy. To the best of our knowledge, this is the first attempt to embed this form of ambiguity into the stochastic control literature.

Second, we address a critical computational challenge. In our setting, the reservoir dynamics and the ambiguity process are driven by a perfectly correlated Brownian motion, which violates standard stability conditions (cf. Barles and Souganidis, 1991, Theorem 3.4) and makes conventional numerical approaches, such as Markov chain approximation (MCA), fail to converge. We resolve this by adapting the coordinate transformation technique of Johnson and Peskir (2017) (see also Basei et al., 2024; De Angelis, 2020; Federico et al., 2023), which reduces the diffusion dimension to a single stochastic driver while turning the ambiguity process into a bounded variation component. This transformation ensures convergence of the MCA scheme to the HJB solution and substantially improves computational efficiency. Crucially, while the representation of the dynamics changes, the managerial interpretation of reservoir control remains intact.

Third, and most importantly, we explicitly establish the link between three benchmark frameworks: (i) learning without ambiguity aversion (zero ambiguity attitude), (ii) smooth ambiguity (finite ambiguity attitude), and (iii) maxmin utility (infinite ambiguity attitude). Our analysis shows that decision makers with a finite degree of ambiguity aversion can revise and diminish their aversion as more information is revealed. Unlike the maxmin framework, where decision makers must remain fully ambiguity-averse at all times, incurring significantly higher operational cost (cf. Section 4.8), the smooth ambiguity approach allows learning to actively reduce ambiguity concerns. This insight provides a novel perspective on managing climate-related uncertainty: decision makers can remain cautious without being locked into extreme pessimism. We believe this contribution is particularly relevant for the management science and operations research communities, as it highlights how integrating learning with ambiguity preferences can yield more flexible, cost-effective, and realistic decision-support models for climate-sensitive resource management.

The remainder of the paper is organized as follows. Section 4.2 formulates the dam reservoir management problem as a singular control model with a learnable Bernoulli-generated prior. Section 4.3 introduces the optimization problem under smooth ambiguity preferences. Section 4.4 establishes the verification theorem for the value function. Section 4.6 presents the coordinate transformation technique, while Section 4.7 develops the Markov chain approximation scheme. This scheme is then applied in Section 4.8 to solve the HJB equation and conduct comparative statics. Section 4.9 concludes the paper.

4.2 MODEL FORMULATION

Let $(\Omega, \mathscr{F}, \mathbf{F} = (\mathscr{F}_t)_{t\geq 0}, \mathsf{P})$ be a filtered probability space, satisfying the usual conditions. For each $t \geq 0$, the associated conditional expectation operator under P is denoted by $\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t}$, where $\mathsf{E}^{\mathsf{P}} \triangleq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_0}$. Suppose that the dam reservoir level is given by the stochastic process $X \triangleq (X_t)_{t\geq 0}$, which is the unique strong solution of the Ornstein-Uhlenbeck (OU) process:

$$\mathrm{d}X_t = \beta(\tilde{x} - X_t)\mathrm{d}t + \sigma \mathrm{d}B_t, \quad X_0 = x > 0.$$

Here $\beta > 0, \tilde{x} > 0$ and $\sigma > 0$ are mean reversion speed, long term mean and volatility of X, respectively. The process $B = (B_t)_{t\geq 0}$ is an $\mathbf{F}$-adapted standard Brownian motion.

Suppose that $\underline{x}$ and $\overline{x}$ are the minimum and maximum level of the water level, where $0 \leq \underline{x} < \overline{x}$. The decision maker (DM) has two primary objectives. The first is to lower the water level before it reaches $\overline{x}$, in order, e.g., to increase flood buffering, especially in the event of major storms or excessive seasonal precipitation. This helps prevent damage to infrastructure, crops, homes, and livelihoods. The second objective is to maintain the water level above $\underline{x}$. There are several reasons for this. The first one is to prevent water scarcity, which can reduce agricultural yields, limit domestic and industrial use, and threaten ecosystem sustainability. A low reservoir level may impair electricity production, which poses a challenge for cities that rely heavily on renewable energy. This shortfall may force them to purchase electricity from the market or turn to more expensive and environmentally harmful fossil fuel sources.

This regulation is achieved through a so-called *control policy*, which consists of a pair of processes $A \triangleq (A^-, A^+)$ that are non-decreasing, non-negative, and $\mathbf{F}$-adapted. The reservoir process governed by A is denoted by $X^A \triangleq \left(X_t^A\right)_{t\geq 0}$, which satisfies the stochastic differential equation:

$$\mathrm{d}X_t^A = \beta(\tilde{x} - X_t^A)\mathrm{d}t + \sigma \mathrm{d}B_t + \mathrm{d}A_t^{\pm}, \quad \mathsf{P}\text{-a.s.}$$

where $A^{\pm} \triangleq A^+ - A^-$. In this framework, the process A^+ increases and $-A^-$ decreases whenever control is applied. We call X^A the *controlled reservoir process*. The control policy A is said to be *feasible* if

$$\sup_{t\geq 0} |X_t^A| \in (\underline{x}, \overline{x}), \quad \mathsf{P}\text{-a.s.}, \tag{4.2.1}$$

In other words, the feasible policy A ensures that X_t^A lies within $(\underline{x}, \overline{x})$ for all $t \geq 0$. The set of feasible control policy is denoted by $\mathscr{D}$.

In our model, ambiguity is captured by the fact that the DM cannot identify a unique probability measure to describe the fluctuations of the reservoir level. Instead, they consider a set of priors, denoted by $\mathscr{P}$, which consists of probability measures equivalent to the reference measure P. Each $\mathsf{Q} \in \mathscr{P}$ is generated by a set of density generator (model) Λ. That is, each $\lambda \in \Lambda$ is such that the process

$$\mathrm{d}\xi_t = -\lambda_t \xi_t \mathrm{d}B_t, \ \xi_0 = 1,$$

is a P-martingale. For any $t > 0$, Q is generated through the Radon-Nikodym derivative $\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}}|\mathscr{F}_t = \xi_t$, and $\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}}|\mathscr{F}_\infty = \xi_\infty$, where $\mathscr{F}_\infty = \sigma(\cup_{t=0}^{\infty}\mathscr{F}_t)$, the smallest σ-algebra that contains $\mathscr{F}_t$ for all $t \geq 0$. For notional convenience, we denote

$$\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} |_{\mathscr{F}_t}^{(t,s)} \triangleq \frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \mid \mathscr{F}_s \Big/ \frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \mid \mathscr{F}_t, \quad \text{for any } s \in [t, T].$$

Accordingly, by Girsanovs theorem, the process $B_t^\lambda \triangleq B_t + \int_0^t \lambda_s \mathrm{d}s$ is a Q-Brownian motion. Under Q, the controlled reservoir process is given by

$$\begin{aligned}\mathrm{d}X_t^{A,\lambda} &= \beta(\tilde{x} - X^{A,\lambda})\mathrm{d}t + \sigma \mathrm{d}B_t^\lambda + \mathrm{d}A_t^{\pm} \\ &= \beta\Big(\tilde{x} + \frac{\sigma}{\beta}\lambda_t - X_t^{A,\lambda}\Big)\mathrm{d}t + \sigma \mathrm{d}B_t + \mathrm{d}A_t^{\pm}.\end{aligned}$$

In other words, the DM becomes uncertain about the long term mean level of the reservoir in which the variation is determined by the term $\tilde{x} + \frac{\sigma}{\beta}\lambda_t$.

For simplicity in the comparative statics analysis, which will be presented in Section 4.8, we redefine the precipitation trend by setting[2] $\theta \triangleq -\lambda$, and correspondingly reparametrize the set of density generators by Θ.

[2]As shown later, the comparative statics rely on a coordinate transformation introduced to ensure numerical stability. This transformation reverses the sign of the precipitation trend, which complicates direct economic interpretation. To avoid this issue, we redefine the sign of the trend parameter here.

When considering the nature of ambiguity, it is reasonable to assume that, due to limited information or the potential unreliability of available data, the model $\theta \in \Theta$ is unobservable. However, it can still be inferred from a filtration $\mathbf{F}^X$, which is generated by what the DM can observe, i.e., the controlled reservoir process $X^{A,\theta}$. That is, we take $\mathbf{F}^X$ to be the filtration generated by $X^{A,\theta}$. Note that $\mathscr{F}_t^X \subseteq \mathscr{F}_t$ for all $t \geq 0$ and that, while θ_t is $\mathscr{F}_t$-measurable, it is not $\mathscr{F}_t^X$-measurable. By general filtering theory (see Øksendal, 2010, Chapter 6), the best estimation of θ_t is found by computing its conditional expectation with respect to the filtration $\mathbf{F}^X$, i.e., $\widehat{\theta_t^{\mathsf{P}}} \triangleq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}[\theta_t]$.

In this paper, we assume that the DM perceives ambiguity as uncertainty regarding the effects of climate change on annual precipitation. As a result, the average long-term mean level of the reservoir cannot be predetermined with certainty. Specifically, the DM considers two potential scenarios, i.e., $\theta \in \Theta \triangleq \{\overline{\theta}, \underline{\theta}\}$. Here, $\overline{\theta} < 0$ determines the long-term mean associated with excessive precipitation (flood). If this is the case, then the long-run mean level is $\tilde{x} + \frac{\sigma}{\beta}|\overline{\theta}|$. Meanwhile, the insufficient precipitation (drought) is represented by $\underline{\theta} > 0$, where if this scenario is revealed, long-run mean level is $\tilde{x} - \frac{\sigma}{\beta}|\underline{\theta}|$.

To ensure that the long term mean of $X^{A,\theta}$ lies within $(\underline{x}, \overline{x})$, we make the following assumption.

Assumption 2. The precipitation ambiguity levels $\overline{\theta}$ and $\underline{\theta}$ satisfy

$$\overline{\theta}_{\min} \triangleq \frac{\beta}{\sigma}(\tilde{x} - \underline{x}) < \overline{\theta} < 0 \quad \text{and} \quad 0 < \underline{\theta} < \frac{\beta}{\sigma}(\tilde{x} - \overline{x}) \triangleq \underline{\theta}_{\max}.$$

For simplicity, we let $\overline{\theta}$ and $\underline{\theta}$ be fractions of $\overline{\theta}_{\min}$ and $\underline{\theta}_{\max}$, respectively. That is, $\underline{\theta}(\underline{\rho}) = \underline{\rho}\underline{\theta}_{\max}$ and $\overline{\theta}(\overline{\rho}) = \overline{\rho}\overline{\theta}_{\min}$ for $\underline{\rho}, \overline{\rho} \in [0, 1)$. In other words, the values $\underline{\rho} \times 100\%$ and $\overline{\rho} \times 100\%$ determine the percentage at which the ambiguous long term mean reservoir level could deviate from the reference $\tilde{x}$ before reaching $\underline{x}$ and $\overline{x}$, respectively.

Furthermore, we assume that the DM has a prior belief $\mathsf{P}\theta_0 = \underline{\theta}] \triangleq \pi \in (0, 1)$. That is, θ has a *Bernoulli distribution*[3] with $\mathsf{P}[\theta_0 = \overline{\theta}] = 1 - \pi$. The posterior updating of this belief is a special case of the *Wonham filter* (cf. Liptser and Shiryaev, 2013a, Theorem 9.6). Specifically, the estimate of θ at time $t > 0$ is given by $\widehat{\theta_t^{\mathsf{P}}} = \underline{\theta}\Pi_t^{\mathsf{P}} + \overline{\theta}(1 - \Pi_t^{\mathsf{P}})$, where $\Pi^{\mathsf{P}} = \left(\Pi_t^{\mathsf{P}}\right)_{t \geq 0}$ is an Itō process solving

$$\mathrm{d}\Pi_t^{\mathsf{P}} = \eta\Pi_t^{\mathsf{P}}(1 - \Pi_t^{\mathsf{P}})\mathrm{d}\overline{B}_t, \quad \Pi_0^{\mathsf{P}} = \pi, \tag{4.2.2}$$

[3]This type of belief updating is widely used in decision-making under uncertainty because it offers a clear managerial interpretation (e.g., distinguishing between "good" and "bad" regimes), making results easier to communicate to policymakers. For instance, it can help justify whether a project is "profitable" or "unprofitable," whether a new technology is "effective" or "ineffective," or, in our case, whether climate change leads to "excessive" or "scarce" rainfall. Applications of such sequential detection methods to adoption and abandonment decisions in investment include Jensen (1982), Kwon and Lippman (2011), Ryan and Lippman (2003), and Thijssen et al. (2006). In healthcare technology assessment, see Bregantini et al. (2023) and Thijssen and Bregantini (2017), and for related work in environmental economics, see Dalby et al. (2018) and Kelly and Kolstad (1999).

which is the updated belief of θ_t being $\overline{\theta}$ at time $t > 0$, based on the observation of X_t^A. Here, $\eta = \frac{\underline{\theta}-\overline{\theta}}{\sigma}$ and $\overline{B} \triangleq \left(\overline{B}_t\right)_{t\geq 0}$ is the associated innovation process, which is $\mathscr{F}_t^X$–measurable, where the controlled reservoir process solves

$$\mathrm{d}X_t^{A,\Pi^{\mathsf{P}}} = \beta(\tilde{x} - \frac{\sigma}{\beta}(\Pi_t^{\mathsf{P}}\underline{\theta} + (1-\Pi_t^{\mathsf{P}})\overline{\theta}) - X_t^{A,\Pi^{\mathsf{P}}})\mathrm{d}t + \sigma\mathrm{d}\overline{B}_t + \mathrm{d}A_t^{\pm}. \tag{4.2.3}$$

Hence, the uncertainty is now characterized by the pair $(X^{A,\Pi^{\mathsf{P}}}, \Pi^{\mathsf{P}})$, which constitutes partially observed, time-homogeneous Markov processes defined on the filtered probability space $(\Omega, \mathscr{F}^X, \mathbf{F}^X, \mathsf{P})$. Moreover, we denote $\mathcal{E} \triangleq (\underline{x}, \overline{x}) \times (0,1)$

Note that the variance of Π^{P} diminishes as it approaches the boundary points $0, 1$. In other words, 0 and 1 are absorbing states for Π^{P} in the limit as $t \to \infty$. This convergence indicates that full information is gradually revealed as $X^{A,\Pi^{\mathsf{P}}}$ is observed over time.

4.3 OPTIMIZATION PROBLEM

We assume that the reservoir is deliberately designed to yield a maximum running benefit (or profit) of $b > 0$ when the water level reaches the long-term mean level $\tilde{x}$. This benefit represents an estimate of community welfare, incorporating both the utility of water usage and the profit generated from hydropower. However, this benefit decreases as the water level deviates from $\tilde{x}$. Specifically, when the water level exceeds $\tilde{x}$, a proportional cost of $\hat{c} > 0$ is incurred. This cost can be interpreted as, or attributed to, the estimated losses from reduced flood buffer capacity and increased maintenance demands on dam infrastructure. Conversely, when the water level falls below $\tilde{x}$, a proportional cost of $\check{c} > 0$ arises. This can reflect losses in productivity due to water shortages, as well as reduced profits from insufficient hydropower generation. Therefore, the instantaneous revenue from managing the reservoir is defined by the function

$$f(x) \triangleq b - \check{c}(\tilde{x} - x)^+ - \hat{c}(x - \tilde{x})^+, \quad x \in (\underline{x}, \overline{x})$$

where $x^+ \triangleq \max\{x, 0\}$.

When the upper control is applied, achieved by releasing water from the reservoir at a maximum effort to keep the water level below a certain threshold, a proportional cost of $u > 0$ is incurred. This cost can be interpreted as an opportunity loss, representing the forfeited potential for future electricity generation or agricultural yield from utilizing the retained water. Conversely, when the DM applies the lower control, by reducing the rate of water release or terminating water discharge temporarily, a proportional cost of $\ell > 0$ is incurred. This cost may reflect, for instance, the expense of purchasing electricity from the energy market to support the community, or the estimated immediate losses resulting from water shortages.

In summary, within the standard expected utility framework, the DM aims to maximize the (perpetual) discounted net profit generated from managing the reservoir level, as defined by the following value function.

$$V(x,\pi) = \sup_{A\in\mathscr{D}} \mathsf{E}^{\mathsf{P}}\left[\int_0^\infty e^{-rt}\Big(f\big(X_t^{A,\Pi^{\mathsf{P}}}\big)\mathrm{d}t - u\mathrm{d}A_t^- - \ell\mathrm{d}A_t^+\Big)\right], \tag{4.3.1}$$

where $r > 0$ is the DM's discounted factor. To ensure that the value function exists (at least in a viscosity sense), we make the following assumptions.

Assumption 3. $\ell \leq \frac{\check{c}}{r+\beta}$ and $u \leq \frac{\hat{c}}{r+\beta}$.

This assumption ensures that it is never optimal to do nothing, since at some point it is better to exert controls to prevent further losses. We demonstrate in the following section that this assumption is indeed necessary.

To explore the DM's attitude to ambiguity, we assume that the DM adopts the *smooth ambiguity* preference framework *à la* Klibanoff et al. (2005). In other words, the DM is concerned that the (reference) probability measure used to estimate the model may be misspecified. This concern arises either from the questionable reliability of historical information or from the lack of sufficient data to establish a prior distribution. As a result, the DM employs a certainty equivalence criterion to robustly and dynamically update the posterior distribution based on the current state of the world.

Mathematically, instead of relying on the standard expectation described in (4.3.1), the DM considers a certainty equivalence induced by a *convex function* ψ_h, which depends on a time interval $h > 0$. That is, ψ_h ensures that the recursive discounted profit flow functional $\mathbb{F}_h$ corresponding to each feasible policy $A \in \mathscr{D}$ is determined by

$$\mathbb{F}_h(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}}) \triangleq \psi_h^{-1}\Big(\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\Big[\psi_h\Big(F_{T,h}(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}})\Big)\Big]\Big) + \mathbb{F}_h(X_{t+h}^{A,\Pi^{\mathsf{P}}},\Pi_{t+h}^{\mathsf{P}}), \tag{4.3.2}$$

where $F_{T,h}$ is the running profit functional on interval $(t,\, t+h)$ satisfying

$$F_{T,h}(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}}) \triangleq \int_t^{t+h} e^{-r(s-t)}\Big(f\big(X_s^{A,\Pi^{\mathsf{P}}}\big)\mathrm{d}s - u\mathrm{d}A_s^- - \ell\mathrm{d}A_s^+\Big).$$

for $0 \leq t < t+h \leq T < \infty$.

It is evident that the convexity of ψ_h biases the standard expectation toward lower values when φ_h is *decreasing*, and toward higher values when it is *increasing*. To further refine the analysis, we assume that ψ follows a time-increment-dependent exponential "dis-utility" function, as proposed by Hansen and Sargent (2011), i.e.,

$$\psi_h(x) \triangleq \exp\left(\frac{\gamma x}{he^{-rh}}\right), \quad \text{for } x \in \mathbb{R}, h > 0.$$

where γ determines the attitude toward ambiguity. Given $h > 0$, if $\gamma < 0$, then φ_h is a decreasing function, indicating that the DM is *ambiguity averse*. Conversely, if $\gamma > 0$, the DM is *ambiguity seeking*. When $\gamma \to 0$, the recursive functional (4.3.2) converges to the standard model (4.3.1). In other word, the DM is said to be *ambiguity neutral*. The presence of h is essential for preserving the effect of ambiguity attitude in decision making over the interval $[t, t+h]$. Without it, ambiguity vanishes over time, causing (4.3.2) to a.s. converge to a linear conditional expectation for any feasible control policy, as shown by Skiadas (2013). Consequently, under smooth ambiguity preferences, the DM's objective is to identify a feasible control policy that maximizes the perpetual discounted running revenue over each (infinitesimal) decision-making interval. That is, the value function under smooth ambiguity is

$$V(x,\pi) = \lim_{\substack{h\downarrow 0 \\ T\to\infty}} \sup_{A\in\mathscr{D}} \mathbb{F}_h(x,\pi), \quad (X_0^{A,\Pi^{\mathsf{P}}}, \Pi_0^{\mathsf{P}}) = (x,\pi). \tag{4.3.3}$$

In the following sections, we solves (4.3.2) by re-expressing it into the form of robust control problem through logarithmic transformation (cf. Fleming and Soner, 2006), before making a connection with a forward-backward stochastic differential equation (FBSDEs). We then provide a verification theorem to find the value function (4.3.3) and the optimal control policy.

4.4 VERIFICATION THEOREM

In this section, we establish the robust control formulation of (4.3.2).

Lemma 2. *Let $t \in (0,T]$, and $\mathbb{F}_h(X_t^{A,\Pi^{\mathsf{P}}}, \Pi_t^{\mathsf{P}})$ be the recursive functional (4.3.2), where $h > 0$ is chosen small enough that $t+h < t+2h < \ldots \leq t+nh = T$, for some $n \in \mathbb{N}$. If we take $h \downarrow 0$, then*

$$J_t^A \triangleq \lim_{h\downarrow 0} \mathbb{F}_h(X_t^{A,\Pi^{\mathsf{P}}}, \Pi_t^{\mathsf{P}}) \leq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Bigg[\Bigg(\int_t^T e^{-r(s-t)}\Bigg(f\big(X_s^{A,\Pi^{\mathsf{P}}}\big)\mathrm{d}s - \frac{1}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,s)}_{\mathscr{F}_t^X}\right)\mathrm{d}s - u\mathrm{d}A_s^- - \ell\mathrm{d}A_s^+\Bigg) + e^{-(T-t)}J_T^A\Bigg)\Bigg], \tag{4.4.1}$$

for any $\mathsf{Q} \in \mathscr{P}$.

Proof of Lemma 2. Let $\mathsf{Q} \in \mathscr{P}$. Then,

$$\begin{aligned}
&\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\exp\left(-\frac{\gamma}{he^{-rh}}F_{T,h}(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}})\right)\right] \\
&= \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\left[\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,t+h)}_{\mathscr{F}_t^X}\exp\left(-\frac{\gamma}{he^{-rh}}F_{T,h}(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}})\right)\right] \\
&= \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\left[\exp\left(-\frac{\gamma}{he^{-rh}}\left(F_{T,h}(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}}) - \frac{h}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,t+h)}_{\mathscr{F}_t^X}\right)\right)\right)\right] \\
&\geq \exp\left(-\mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\left[\frac{\gamma}{he^{-rh}}\left(F_{T,h}(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}}) - \frac{h}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,t+h)}_{\mathscr{F}_t^X}\right)\right)\right]\right).
\end{aligned}$$

The inequality holds by the Jensen's inequality. Observe that the above inequality applies recursively on interval $(t+kh,\, t+(k+1)h)$ for $k=0,\dots,n-1$. Therefore, it follows that

$$\begin{aligned}\mathbb{F}_h(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}}) \le \inf_{\mathsf{Q}\in\mathscr{P}} \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Bigg[\sum_{k=0}^{n-1}\Bigg\{ & F_{T,h}(X_{t+kh}^{A,\Pi^{\mathsf{P}}},\Pi_{t+kh}^{\mathsf{P}}) \\ & -\frac{h}{\gamma}e^{-rh}\underbrace{\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t+kh,t+(k+1)h)}_{\mathscr{F}^X_{t+kh}}\right)}_{(*)}\Bigg\}\Bigg], \end{aligned} \tag{4.4.2}$$

The infimum holds since Q is arbitrary for any k. By the mean value theorem (MVT), for each k there exists $\delta_k\in(0,h]$ such that

$$e^{-r\delta_k}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t+kh,t+kh+\delta_k)}_{\mathscr{F}^X_{t+kh}}\right) = \frac{1}{h}\int_{t+kh}^{t+(k+1)h} e^{-r(s-(t+kh))}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t+kh,s)}_{\mathscr{F}^X_{t+kh}}\right)\mathrm{d}s. \tag{4.4.3}$$

Observe that the term $(*)$ converges to the LHS of (4.4.3) as one takes $h\to 0$ for any k. Hence, in limit as $h\to 0$, the RHS of (4.4.2) converges to that of (4.4.1), which completes the proof. ■

Remark 15. *Observe from Lemma 2 that if there exists $\mathsf{Q}\in\mathscr{P}$ for which equality holds in (4.4.1), then this measure Q attains the infimum, exhibiting the robust control formulation and solving (4.3.2). In what follows, we show that such measure Q can be explicitly obtained through solving an FBSDE.*

We now establish a unique representation of the cost functional (4.3.2) in terms of a FBSDE, and the associated multiple prior as $h\downarrow 0$ for any admissible control policy. Before proceeding, we introduce the associated solution spaces. Note that $|\cdot|$ is the Euclidean norm, i.e., $|x|=\sqrt{x\cdot x}$, $x=(x_1,\dots,x_n)\in\mathbb{R}^n$. Then, for $t\in[0,T]$, we define:

- $\mathbb{L}^2(\mathscr{F}^X,\mathsf{P})$: the set of $\mathscr{F}^X$-measurable random variable ξ such that $\mathsf{E}^{\mathsf{P}}\left[|\xi|^2\right]<\infty$.
- $\mathbb{L}^2_T(\mathbf{F}^X,\mathsf{P})$: the set of $\mathbf{F}^X$-adapted process Y such that $\mathsf{E}^{\mathsf{P}}\left[\sup_{0\le s\le T}|Y_s|^2\right]<\infty$.
- $\mathbb{H}^2_T(\mathbf{F}^X,\mathsf{P})$: the set of $\mathbf{F}^X$-adapted process Y such that $\mathsf{E}^{\mathsf{P}}\left[\int_0^T|Y_s|^2\mathrm{d}s\right]<\infty$.

Lemma 3. *Suppose that there exists a triple $(Y^A,Z,A)\in\mathbb{L}^2_T(\mathbf{F}^X,\mathsf{P})\times\mathbb{H}^2_T(\mathbf{F}^X,\mathsf{P})\times\mathbb{L}^2(\mathscr{F}^X_T,\mathsf{P})$ that solves the BSDE,*

$$\begin{aligned}\mathrm{d}Y_t^A &= -F(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}},Y_t^A,Z_t)\mathrm{d}t+\ell\mathrm{d}A_t^+ + u\mathrm{d}A_t^- \\ &\quad +Z_t\mathrm{d}\overline{B}_t,\quad Y_T^A=\xi(X_T^{A,\Pi^{\mathsf{P}}})\in\mathbb{L}^2(\mathscr{F}_T^X,\mathsf{P})\end{aligned} \tag{4.4.4}$$

where $X^{A,\Pi^{\mathsf{P}}}$ and Π^{P} solve (4.2.3) and (4.2.2), respectively, and

$$F(x,\pi,y,z)\triangleq -ry+f(x)+h(\pi,z),\ \text{and} \tag{4.4.5}$$

$$h(\pi,z)\triangleq +(\pi\underline{\theta}+(1-\pi)\overline{\theta})z-\frac{1}{\gamma}\log\left(\pi\exp\left(-\gamma\underline{\theta}z\right)+(1-\pi)\exp\left(-\gamma\overline{\theta}z\right)\right).$$

If there exists $\bar{z} < \infty$ such that $|Z_t| \le \bar{z}$, P-a.s., then the solution (Y^A, Z, A) is unique in $\mathbb{L}^2_T(\mathbf{F}^X, \mathsf{P}) \times \mathbb{H}^2_T(\mathbf{F}^X, \mathsf{P}) \times \mathbb{L}^2(\mathscr{F}^X_T, \mathsf{P})$.

Proof of Lemma 3. Adapting from El Karoui, Kapoudjian, et al. (1997, Theorem 5.2) and the standing assumption (4.2.1), it is sufficient to show that

A1) (X, Π^{P}) admits a unique strong solution in $\mathbb{H}^2_T(\mathbf{F}^X, \mathsf{P})$,

A2) F is uniformly continuous in (t, x, π, y) and there is a constant L such that for any $t \in [0, T],\ x_1, x_2 \in (\underline{x}, \overline{x}),\ \pi_1, \pi_2 \in (0, 1),\ y_1, y_2, z_1, z_2 \in \mathbb{R}$,

$$\begin{aligned}&|F(x_1, \pi_1, y_1, z_1) - F(x_2, \pi_2, y_2, z_2)| \\ &\qquad\qquad \le L(|x_1 - x_2| + |\pi_1 - \pi_2| + |y_1 - y_2| + |z_2 - z_1|)\end{aligned}$$

A3) ℓ and u are deterministic real values.

Condition A3) is immediate. Condition A1) holds since the drift and diffusion coefficients of (X, Π^{P}) satisfy global Lipschitz continuity and linear growth conditions (see, for instance, Karatzas and Shreve, 1991, Theorem 2.9). To verify Condition A2), it suffices to show that $h(\pi, z)$ has bounded first-order partial derivatives with respect to π and z for any $\gamma > 0$. A straightforward calculation gives

$$\begin{aligned}\left|\frac{\partial h}{\partial \pi}(\pi, z)\right| &= \left|\underline{\theta} - \overline{\theta} + \frac{1}{\gamma} \cdot \frac{\exp(-\gamma\overline{\theta}z) - \exp(-\gamma\underline{\theta}z)}{\pi\exp(-\gamma\overline{\theta}z) + (1-\pi)\exp(-\gamma\underline{\theta}z)}\right| \\ &\le \underline{\theta} - \overline{\theta} + \left|\frac{\exp(-\gamma\overline{\theta}z) - \exp(-\gamma\underline{\theta}z)}{\gamma\min\{\exp(-\gamma\overline{\theta}z), \exp(-\gamma\underline{\theta}z)\}}\right| \\ &= \underline{\theta} - \overline{\theta} + \begin{cases}\left|\frac{1}{\gamma}(\exp(-\gamma(\underline{\theta} - \overline{\theta})z) - 1)\right| & \text{if } \ z \le 0 \\ \left|\frac{1}{\gamma}(1 - \exp(-\gamma(\underline{\theta} - \overline{\theta})z))\right| & \text{if } \ z > 0\end{cases} \\ &\le \underline{\theta} - \overline{\theta} + \frac{C_z}{\gamma} \quad \text{for some } C_z < \infty, \text{ induced by } \bar{z},\end{aligned} \tag{4.4.6}$$

$$\left|\frac{\partial h}{\partial z}(\pi, z)\right| = \left|\frac{\underline{\theta}\pi\exp(-\gamma\underline{\theta}z) + \overline{\theta}(1-\pi)\exp(-\gamma\overline{\theta}z)}{\pi\exp(-\gamma\overline{\theta}z) + (1-\pi)\exp(-\gamma\underline{\theta}z)}\right| \le \underline{\theta} - \overline{\theta} \tag{4.4.7}$$

As a result, (4.4.5), (4.4.6), and (4.4.7) induce the Lipschitz constant $L = r + \check{c} + \hat{c} + \frac{C_z}{\gamma} + \underline{\theta} - \overline{\theta} + L_\xi$, where L_ξ is the Lipschitz constant of ξ. This concludes the proof. ∎

Since the function g is time-homogeneous, one can adapt the argument from Peng and Shi (2000, Theorem 6) to extend the result from Lemma 3 to the case of an infinite time horizon. The result is presented by the following lemma.

Lemma 4. *Suppose that (Y^A, Z, A) solves the BSDE (4.4.4). Then (Y^A, Z, A) has a unique solution in $\mathbb{L}^2_\infty(\mathbf{F}^X, \mathsf{P}) \times \mathbb{H}^2_\infty(\mathbf{F}^X, \mathsf{P}) \times \mathbb{L}^2(\mathscr{F}^X_\infty, \mathsf{P})$.*

In what follows, we show that the BSDE (4.4.4) solves the recursive functional inequality (4.4.1), where the quality holds for some $\mathsf{Q} \in \mathscr{P}$.

Theorem 13. *Suppose that the triple* (Y^A, Z, A) *solves* (4.4.4)*, then it satisfies* (4.3.2) *as* $h \downarrow 0$.

Proof of Theorem 13. From (4.4.4), one can infer that, for some $h > 0$ and $t + h < T$,

$$Y_t^A = \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\Big[\int_t^{t+h}\big(-rY_s^A + f(X_s^{A,\Pi^{\mathsf{P}}}) + h(\Pi_s^{\mathsf{P}}, Z_s)\big)\mathrm{d}s - \int_t^{t+h}\big(\ell\mathrm{d}A_s^+ - u\mathrm{d}A_s^-\big) + Y_{t+h}^A\Big]. \tag{4.4.8}$$

For any $\mathsf{Q} \in \mathscr{P}$, if we change the measure from P to Q, the standard Girsanov theorem gives that

$$\overline{B}_t = \widehat{B}_t - \int_0^t \left(\widehat{\theta_s^{\mathsf{Q}}} - \widehat{\theta_s^{\mathsf{P}}}\right)\mathrm{d}s \tag{4.4.9}$$

where $\widehat{\theta^{\mathsf{Q}}}_s \triangleq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_s^X}[\theta_s]$ and $\widehat{B} = \left(\widehat{B}_t\right)_{t\geq 0}$ is the innovation process under Q. By plugging (4.4.9) into (4.4.4), it implies that the inequality (4.4.8) yields

$$Y_t^A = \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Big[\int_t^{t+h}\Big(\big(-rY_s^A + f(X_s^{A,\Pi^{\mathsf{P}}}) - \left(\widehat{\theta_s^{\mathsf{Q}}} - \widehat{\theta_s^{\mathsf{P}}}\right)Z_s + h(\Pi_s^{\mathsf{P}}, Z_s)\big)\mathrm{d}s\Big) - \int_t^{t+h}\big(\ell\mathrm{d}A_s^+ + u\mathrm{d}A_s^-\big) + Y_{t+h}^A\Big].$$

Notice that $h(\Pi_s^{\mathsf{P}}, Z_s)$ is the log-function of the moment generating function of the process $-\gamma\theta_s Z_s$. Thus, for any $s \in (t, t+h]$, we have that

$$\begin{aligned} h(\Pi_s^{\mathsf{P}}, Z_s) &= -\widehat{\theta^{\mathsf{P}}} Z_s - \frac{1}{\gamma}\log\left(\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_s^X}[\exp(-\gamma\theta_s Z_s)]\right) \\ &\leq \left(\widehat{\theta_s^{\mathsf{Q}}} - \widehat{\theta_s^{\mathsf{P}}}\right)Z_s - \frac{1}{\gamma}\mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_s^X}\left[\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(s,t+h)}_{\mathscr{F}_s^X}\right)\right], \end{aligned} \tag{4.4.10}$$

The inequality holds by the logarithmic transformation (cf. Fleming and Soner, 2006, Chapter 6), and the fact that Z_s is $\mathscr{F}_s^X$-measurable. Consequently, it follows that

$$Y_t^A \leq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Big[\int_t^{t+h}\Big(\big(-rY_s^A + f(X_s^{A,\Pi^{\mathsf{P}}}) - \frac{1}{\gamma}\mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_s^X}\Big[\log\Big(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(s,t+h)}_{\mathscr{F}_s^X}\Big)\Big]\Big)\mathrm{d}s - \int_t^{t+h}\big(\ell\mathrm{d}A_s^+ + u\mathrm{d}A_s^-\big) + Y_{t+h}^A\Big]. \tag{4.4.11}$$

Since $\widehat{\theta^{\mathsf{Q}}}$ is $\mathbf{F}^X$-adapted, it is thus right continuous and has a left limit. Therefore, taking $h \downarrow 0$ implies $\mathscr{F}_s^X \downarrow \mathscr{F}_t^X$. As a results, we obtain

$$Y_t^A \leq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Big[\int_t^{T}\Big(\big(-rY_s^A + f(X_s^{A,\Pi^{\mathsf{P}}}) - \frac{1}{\gamma}\log\Big(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,s)}_{\mathscr{F}_t^X}\Big)\Big)\mathrm{d}s - \int_t^{T}\big(\ell\mathrm{d}A_s^+ + u\mathrm{d}A_s^-\big) + Y_T^A\Big]. \tag{4.4.12}$$

Note that if we choose, dynamically, $\mathsf{Q} = \mathsf{Q}^*$ where $\mathsf{Q}^* \in \mathscr{P}$ is such that

$$\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}^*} \Big|_{\mathscr{F}_t^X}^{(s,t+h)} = \frac{\exp(-\gamma\theta_s Z_s)}{\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_s^X}[\exp(-\gamma\theta_s Z_s)]}$$

for any $s \in (t, t+h]$, where $t \in [0, T]$, $h > 0$ and $t + h \leq T$, then we have equality in (4.4.10), and, thus, (4.4.11) and (4.4.12) (in the limit $h \downarrow 0$). According to Lemma 2, we can conclude that the BSDE (Y^A, Z) solves (4.3.2) as $h \downarrow 0$. ■

According to Theorem 13, the value function is given by

$$V(t, x, \pi) \triangleq \sup_{A \in \mathscr{D}} Y_t^{A;x,\pi}, \tag{4.4.13}$$

where $Y_t^{A;x,\pi} \triangleq Y_t^A \mid_{X_t^{A,\Pi^{\mathsf{P}}} = x, \Pi_t^{\mathsf{P}} = \pi, T \to \infty}$. We denote $V(x, \pi) \triangleq V(0, x, \pi)$. Now, we show that (4.4.13) is a concave function and locally Lipchitz continuous.

Proposition 10. *$V(t, x, \pi)$ is concave in $x \in (\underline{x}, \overline{x})$.*

Proof of Proposition 10. Suppose, by contradiction, that for any $\pi \in (0, 1)$, there exist $x_0, x_1 \in (\underline{x}, \overline{x})$ and $\alpha \in (0, 1)$ such that

$$\alpha V(t, x_0, \pi) + (1 - \alpha) V(t, x_1, \pi) \geq V(t, \overbrace{\alpha x_0 + (1 - \alpha) x_1}^{\triangleq x_\alpha}, \pi), \quad t \geq 0. \tag{4.4.14}$$

For simplicity, we denote by $X_s^{A,\Pi^{\mathsf{P}};x}$, $t \leq s$, the controlled process given by $X_t^{A,\Pi^{\mathsf{P}}} = x$. Then, from (4.4.4), we have

$$\begin{aligned} \alpha Y_t^{A;x_0,\pi} + (1-\alpha) Y_t^{A;x_1,\pi} &= \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^\infty \alpha F(X_s^{A,\Pi^{\mathsf{P}};x_0}, \Pi_s^{\mathsf{P}}, Y_s^{A;x_0,\pi}, Z_s)\mathrm{d}s\right] \\ &+ \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^\infty (1-\alpha) F(X_s^{A,\Pi^{\mathsf{P}};x_1}, \Pi_s^{\mathsf{P}}, Y_s^{A;x_1,\pi}, Z_s)\mathrm{d}s - \int_t^\infty (\ell \mathrm{d}A_s^+ + u\mathrm{d}A_s^-)\right] \end{aligned}$$

A straightforward calculation gives $\frac{\partial^2 h}{\partial \pi^2}(x, \pi) \leq 0$ and $\frac{\partial^2 h}{\partial z^2}(\pi, z) \leq 0$ in (π, z). Then it follows that $F(x, \pi, y, z)$ is a concave function in (x, π, y, z). Together with the linearity of $X_t^{A,\Pi^{\mathsf{P}};x}$ in x, it holds that

$$\begin{aligned} &\alpha Y_t^{A;x_0,\pi} + (1-\alpha) Y_t^{A;x_1,\pi} \\ &\quad \leq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^\infty F(X_s^{A,\Pi^{\mathsf{P}};x_\alpha}, \Pi_s^{\mathsf{P}}, \alpha Y_s^{A;x_0,\pi} + (1-\alpha) Y_s^{A;x_1,\pi}, Z_s)\mathrm{d}s - \int_t^\infty (\ell \mathrm{d}A_s^+ + u\mathrm{d}A_s^-)\right] \end{aligned}$$

Since this holds for any feasible control policy, it holds that for any $\varepsilon > 0$ there exists $\overline{A} \in \mathscr{D}$ such that

$$\alpha V(t, x_0, \pi) + (1 - \alpha) V(t, x_1, \pi)$$

$$\le \sup_{A\in\mathscr{D}} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^\infty F(X_s^{A,\Pi^{\mathsf{P}};x_\alpha},\Pi^{\mathsf{P}}_s,\alpha Y_s^{A;x_0,\pi}+(1-\alpha)Y_s^{A;x_1,\pi},Z_s)\mathrm{d}s - \int_t^\infty(\ell\mathrm{d}A_s^+ + u\mathrm{d}A_s^-)\right]$$

$$\le \mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^\infty F(X_s^{A,\Pi^{\mathsf{P}};x_\alpha},\Pi^{\mathsf{P}}_s,\alpha Y_s^{\overline{A};x_0,\pi}+(1-\alpha)Y_s^{\overline{A};x_1,\pi},Z_s)\mathrm{d}s - \int_t^\infty(\ell\mathrm{d}\overline{A}_s^+ + u\mathrm{d}\overline{A}_s^-)\right]+\varepsilon.$$

Since the second inequality holds for arbitrary $\varepsilon > 0$, we have

$$\begin{aligned}\alpha V(t,x_0,\pi)&+(1-\alpha)V(t,x_1,\pi)\\ &\le \sup_{A\in\mathscr{D}} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^\infty F(X_s^{A,\Pi^{\mathsf{P}};x_\alpha},\Pi^{\mathsf{P}}_s,\alpha V(s,x_0,\pi)+(1-\alpha)V(s,x_1,\pi),Z_s)\mathrm{d}s - \int_t^\infty(\ell\mathrm{d}A_s^+ + u\mathrm{d}A_s^-)\right].\end{aligned} \tag{4.4.15}$$

Recalling (4.4.14) and that $F(\cdot,\cdot,y,\cdot)$ is strictly decreasing in y, we conclude from (4.4.15) that $\alpha V(t,x_0,\pi)+(1-\alpha)V(t,x_1,\pi) < V(t,x_\alpha,\pi)$, which is a contradiction. Hence, $V(x,\cdot)$ is concave in x. ■

Proposition 11. *Let $\bar{c} = \max\{\check{c},\hat{c}\}$. The value function $V(\cdot,x,\pi)$ is Lipschitz continuous in x and π with modulus $\frac{\bar{c}}{r+\beta}$ and $\frac{\bar{c}\sigma(\underline{\theta}-\overline{\theta})}{r(r+\beta)}$, respectively. Moreover, if Assumption 3 holds, then there exists $(a,b)\subseteq(\underline{x},\overline{x})$ such that $-u \le \frac{V(t,x_0,\pi)-V(t,x_1,\pi)}{x_0-x_1} \le \ell$ for any $x_0,x_1,\in(a,b)$ that $x_0 > x_1$.*

Proof of Proposition 11. Take $x_0,x_1\in(\underline{x},\overline{x})$, such that $x_0 > x_1$. Note that

$$-\hat{c} \le \frac{f(x_0)-f(x_1)}{x_0-x_1} \le \check{c}$$

and

$$\begin{aligned}X_s^{A,\Pi^{\mathsf{P}};x_0} - X_s^{A,\Pi^{\mathsf{P}};x_1} &= x_0 - x_1 - \beta\int_t^s \left(X_u^{A,\Pi^{\mathsf{P}};x_0} - X_u^{A,\Pi^{\mathsf{P}};x_1}\right)\mathrm{d}u\\ &= (x_0-x_1)\exp(-\beta(s-t)).\end{aligned}$$

Then for any $A\in\mathscr{D}$, it follows from (4.4.4) that,

$$\begin{aligned}Y_t^{A;x_0,\pi} - Y_t^{A;x_1,\pi} &= \mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^\infty e^{-r(s-t)}(f(X_s^{A,\Pi^{\mathsf{P}};x_0}) - f(X_s^{A,\Pi^{\mathsf{P}};x_1}))\mathrm{d}s\right]\\ &\le \check{c}\,\mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^\infty e^{-r(s-t)}\left(X_s^{A,\Pi^{\mathsf{P}};x_0} - X_s^{A,\Pi^{\mathsf{P}};x_1}\right)\mathrm{d}s\right]\\ &\le \frac{\check{c}}{r+\beta}(x_0-x_1),\end{aligned} \tag{4.4.16}$$

Similarly, it holds that

$$Y_t^{A;x_0,\pi} - Y_t^{A;x_1,\pi} \geq -\frac{\hat{c}}{r+\beta}(x_0 - x_1), \tag{4.4.17}$$

Since (4.4.16) and (4.4.17) hold for any $A \in \mathscr{D}$, we have $-\frac{\hat{c}}{r+\beta} \leq \frac{V(t,x_0,\pi)-V(t,x_1,\pi)}{x_0-x_1} \leq \frac{\check{c}}{r+\beta}$. In other words, $V(\cdot,x,\cdot)$ is Lipschitz continuous with modulus $\frac{\bar{c}}{r+\beta}$. A similar result is obtained for $V(\cdot,\cdot,\pi)$ using an analogous technique.

Now, suppose that Assumption 3 holds. Because $V(\cdot,x,\cdot)$ is strictly concave, thanks to Proposition 10, it follows that $\frac{V(t,x_0,\pi)-V(t,x_1,\pi)}{x_0-x_1}$ is decreasing in x_0. Together with Assumption 3, this implies that there exists $(a,b) \subseteq (\underline{x},\overline{x})$ such that $-u \leq \frac{V(t,x_0,\pi)-V(t,x_1,\pi)}{x_0-x_1} \leq \ell$. ■

Under certain assumptions, as demonstrated in the following verification theorem, we can establish the existence of a "smooth" solution for the value function together with the conditions under which a unique optimal control policy is obtained. The proof is standard, (see, for example, Harrison and Taksar, 1983, Fleming and Soner, 2006, Chapter 8, or Archankul et al., 2025, Theorem 1), and is, thus, omitted for brevity.

Theorem 14 (Verification theorem)**.** *Suppose that $V \in C^2(\mathcal{E})$ such that*

$$V(x,\pi) = \sup_{A\in\mathscr{D}} \mathsf{E}^{\mathsf{P}}\left[\int_0^\infty e^{-rt}\left(\left(f(X_t^{A,\Pi^{\mathsf{P}}}) + h\left(\Pi_t^{\mathsf{P}}, \overline{\nabla}V(X_t^{A,\Pi^{\mathsf{P}}},\Pi_t^{\mathsf{P}})\right)\right)\mathrm{d}t - \ell\mathrm{d}A_t^+ - u\mathrm{d}A_t^-\right)\right].$$

where

$$\overline{\nabla}V(x,\pi) \triangleq \sigma\frac{\partial V}{\partial x}(x,\pi) + \eta\pi(1-\pi)\frac{\partial V}{\partial \pi}(x,\pi),$$

If for any $\phi \in C^2(\mathcal{E})$ satisfying

$$\mathscr{L}\phi(x,\pi) + F(x,\pi,\phi(x,\pi),\overline{\nabla}\phi(x,\pi)) \leq 0 \ \text{ and}$$
$$-u \leq \frac{\partial \phi}{\partial x}(x,\pi) \leq \ell,$$

where

$$\mathscr{L}\phi(x,\pi) \triangleq \beta\left(\tilde{x} - \frac{\sigma}{\beta}\left(\underline{\theta}\pi + (1-\pi)\overline{\theta}\right) - x\right)\frac{\partial\phi}{\partial x}(x,\pi) + \sigma\frac{\partial^2\phi}{\partial x^2}(x,\pi)$$
$$+ \eta^2\pi^2(1-\pi)^2\frac{\partial^2\phi}{\partial\pi^2}(x,\pi) + \sigma\eta\pi(1-\pi)\frac{\partial^2\phi}{\partial x\partial\pi}(x,\pi),$$

then $V(x,\pi) \geq \phi(x,\pi)$ for any $(x,\pi) \in \mathcal{E}$ and $A \in \mathscr{D}$.

Moreover, for each $\pi \in (0,1)$, if there exists a pair of free-boundary interval $\mathcal{X}(\pi) \subseteq (\underline{x},\overline{x})$ such that

$$\mathscr{L}\phi(x,\pi) + F(x,\pi,\phi(x,\pi),\overline{\nabla}\phi(x,\pi)) = 0 \qquad \text{on } \mathcal{X}(\pi),$$

$$\frac{\partial \phi}{\partial x}(a(\pi), \pi) = \ell, \tag{4.4.18}$$

$$\frac{\partial \phi}{\partial x}(b(\pi), \pi) = -u, \tag{4.4.19}$$

$$\frac{\partial^2 \phi}{\partial x^2}(a(\pi), \pi) = \frac{\partial^2 \phi}{\partial x^2}(b(\pi), \pi) = 0, \tag{4.4.20}$$

where

$$a(\pi) \triangleq \inf\{x : x \in \mathcal{X}(\pi)\},$$
$$b(\pi) \triangleq \sup\{x : x \in \mathcal{X}(\pi)\},$$

and $A^* \in \mathscr{D}$ *such that* A^{*-} *and* A^{*-} *only increase at* $a(x_2)+$ *and* $b(x_2)-$*, respectively,* P*-a.s., then* $V(x, \pi) = \phi(x, \pi)$*.*

Remark 16. *The set* $\mathcal{X}(\pi)$ *is commonly known as the* continuation region*, while* (4.4.18) *and* (4.4.19) *represent the* value-matching conditions*, and* (4.4.20) *imposes the* smooth-pasting condition.

Theorem 14 implies that V is a classical solution to the HJB equation

$$0 = \max\left\{\mathscr{L}V(x, \pi) + F(x, \pi, V(x, \pi), \overline{\nabla}V(x, \pi)), -\ell + \frac{\partial V}{\partial x}(x, \pi), -u - \frac{\partial V}{\partial x}(x, \pi)\right\}. \tag{4.4.21}$$

Moreover, the control barriers $\mathcal{X}(\pi)$ are unique for any $\pi \in (0, 1)$, since V is concave in x (cf. Proposition 10). That is,

$$a(\pi) \triangleq \sup\left\{x \in (\underline{x}, \overline{x}) : \frac{\partial V}{\partial x}(x, \pi) \geq \ell\right\}, \quad \text{and}$$
$$b(\pi) \triangleq \inf\left\{x \in (\underline{x}, \overline{x}) : \frac{\partial V}{\partial x}(x, \pi) \leq -u\right\}.$$

Moreover, Proposition 11 ensures that the first derivatives have at most polynomial growth in (x, π), i.e.,

$$\left|\frac{\partial V}{\partial x}(x, \pi)\right| + \left|\frac{\partial V}{\partial \pi}(x, \pi)\right| \leq C(1 + |x|^p + |\pi|^q), \text{ for some } p, q > 1. \tag{4.4.22}$$

Then the BSDE (4.4.4) with $Z = \overline{\nabla}V$ admits a unique solution, since $|\overline{\nabla}V| \leq \sigma \max\{\ell, u\}$ and thus $\overline{\nabla}V \in \mathbb{H}^2_\infty(\mathbf{F}^X, \mathsf{P})$.

4.5 VISCOSITY SOLUTION

Given the nonlinearity of the function $h(\cdot, \cdot)$, to the best of our knowledge, there does not exist a convex function in the class C^2 that solves the HJB equation (4.4.21) and satisfies the growth condition (4.4.22). Furthermore, the smoothness assumption may be overly restrictive This is because, first of all, V is only known to be continuous but not necessarily smooth. Secondly, $h(\cdot,\cdot)$ may fail to be differentiable, especially in the case when $\gamma \to \infty$. In fact,

$$\lim_{\gamma\to\infty} h(\pi,z) = -(\pi\underline{\theta} + (1-\pi)\overline{\theta})z + \min\{\underline{\theta}z, \overline{\theta}z\}. \tag{4.5.1}$$

In other words, in our smooth ambiguity setting, the extremely ambiguity-averse DM acts as if they consider the maxmin utility *à la* Gilboa and Schmeidler (1989). This prompts us to consider a weaker class of solutions, the so-called *viscosity solution* (cf. Crandall and Lions, 1983), where only the locally boundedness of the solution to the HJB equation (4.4.21) is assumed. The definition of viscosity solution in our setting is of the following type.

Definition 23 (Viscosity Solution)**.** Suppose that $V \in C(\mathcal{E})$, and let $\varphi \in C^2(\mathcal{E})$. Then

1) V is a viscosity *subsolution* of (4.4.21) if for all local maximum points $(x,\pi) \in \mathcal{E}$ of $V-\varphi$, i.e., $(V-\varphi)(x,\pi) = 0$ and $(V-\varphi)(x,\pi) \geq (V-\varphi)(x_1', x_2')$, $(x_1', x_2') \neq (x,\pi)$, it holds that

$$\max\left\{\mathscr{L}\varphi(x,\pi) + F(x,\pi,\varphi(x,\pi),\overline{\nabla}\varphi(x,\pi)), -\ell + \frac{\partial\varphi}{\partial x}(x,\pi), -u - \frac{\partial\varphi}{\partial x}(x,\pi)\right\} \geq 0,$$

2) V is a viscosity *supersolution* of (4.4.21) if for all local minimum points $(x,\pi) \in \mathcal{E}$ of $v-\varphi$, i.e., $(V-\varphi)(x,\pi) = 0$ and $(V-\varphi)(x,\pi) \leq (V-\varphi)(x_1', x_2')$, $(x_1', x_2') \neq (x,\pi)$, it holds that

$$\max\left\{\mathscr{L}\varphi(x,\pi) + F(x,\pi,\varphi(x,\pi),\overline{\nabla}\varphi(x,\pi)), -\ell + \frac{\partial\varphi}{\partial x}(x,\pi), -u - \frac{\partial\varphi}{\partial x}(x,\pi)\right\} \leq 0,$$

3) V is a *viscosity solution* of φ if it is both viscosity subsolution and supersolution.

Clearly, $Y_0^{A;x,\pi}$ is continuous in $\mathcal{E}$. Now, we show that $Y_0^{A;x,\pi}$ is a viscosity solution to the HJB equation (4.4.21).

Proposition 12. *Suppose that Assumption 3 holds. Then $V(x,\pi)$ is a viscosity solution to the HJB equation* (4.4.21)*.*

Proof of Proposition 12. Given $\varphi \in C^2(\mathcal{E})$. We first show that φ is a viscosity supersolution to (4.4.21). The idea of proof is adapted from that of the nonlinear Feynman-Kac formula given by Zhang (2017, Chapter 5).

Supersolution. Suppose that (x,π) is a local minimum point of $V-\varphi$. Suppose, in contradiction, that

$$\mathscr{L}\varphi(x,\pi) + F(x,\pi,\varphi(x,\pi),\overline{\nabla}\varphi(x,\pi)) > 0, \tag{4.5.2}$$

$$-\ell + \frac{\partial\varphi}{\partial x}(x,\pi) > 0, \quad \text{or} \tag{4.5.3}$$

$$-u - \frac{\partial\varphi}{\partial x}(x,\pi) > 0 \tag{4.5.4}$$

According to the proof of Proposition 11, together with Assumption 3, we obtain $-uh < Y_0^{A;x+h,\pi} - Y_0^{A;x,\pi} < \ell h$, for any $h > 0$ such that $\underline{x} \le x < x+h \le \overline{x}$. Since (x,π) is a local minimum point of $Y_0^A - \varphi$, it follows that $\varphi(x+h,\pi) - \varphi(x,\pi) < \ell h$. The arbitrariness of $h > 0$ and $A \in \mathscr{D}$ then contradicts (4.5.3). We can also conclude that (4.5.4) leads to a similar contradiction. Thus, we must have $-u \le \frac{\partial\varphi}{\partial x}(x,\pi) \le \ell$.

Now, it is left to show that (4.5.2) leads to a contradiction. Let us consider

$$\check{Y}_t^A = \varphi(X_t^{A,\Pi^{\mathsf{P}}}, \Pi_t^{\mathsf{P}}) \text{ and } \check{Z}_t = \overline{\nabla}\varphi(X_t^{A,\Pi^{\mathsf{P}}}, \Pi_t^{\mathsf{P}})$$

for some $\varphi \in C^2(\mathcal{E})$. Since $F(\cdot,\cdot,\cdot,\cdot)$ is continuous, the following stopping times are finite P-a.s.,

$$\overline{\tau}_t \triangleq \inf\{s > t : \mathscr{L}\varphi(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) + F(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, \check{Y}_s^A, \check{Z}_s)) = 0\}$$

Denote by $\tau_{t,h} \triangleq \overline{\tau}_t \wedge (t+h)$ for some $h > 0$, and by $\tau_h \triangleq \tau_{0,h}$, where $a \wedge b = \min\{a,b\}$. Then it follows from Itō's lemma that

$$\check{Y}_{\tau_h}^A = \check{Y}_0^A + \int_0^{\tau_h} \mathscr{L}\varphi(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}})\mathrm{d}s + \int_0^{\tau_h} \frac{\partial\varphi}{\partial x}(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}})\mathrm{d}A_s + \int_0^{\tau_h} \overline{\nabla}\varphi(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}})\mathrm{d}\overline{B}_s.$$

Now, recall from (4.4.4) that and suppose that $A^* = \arg\sup_{A\in\mathscr{D}} Y^A$. Then, we have $V(x,\pi) = Y_0^{A^*}$ and

$$Y_{\tau_h}^{A^*} = Y_0^{A^*} - \int_0^{\tau_h} F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, Y_s^{A^*}, Z_s)\mathrm{d}s + \int_0^{\tau_h} (\ell \mathrm{d}A_s^{*+} + u\mathrm{d}A_s^{*-}) + \int_0^{\tau_h} Z_s \mathrm{d}\overline{B}_s.$$

Let $\Delta Y_t \triangleq \check{Y}_t^{A^*} - Y_t^{A^*}$ and $\Delta Z_t \triangleq \check{Z}_t - Z_t$. Since $(X_0^{A^*,\Pi^{\mathsf{P}}}, \Pi_0^{\mathsf{P}}) = (x,\pi)$ is the local minimum point, we have $\Delta Y_0 = 0$. Therefore, the definition of τ_h ensures that

$$\begin{aligned}
\Delta Y_{\tau_h} = &\int_0^{\tau_h} \Big(\mathscr{L}\varphi(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) + F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, Y_s^{A^*}, Z_s)\Big)\,\mathrm{d}s + \int_0^{\tau_h} \Delta Z_s \mathrm{d}\overline{B}_s \\
&+ \int_0^{\tau_h} \left(\frac{\partial\varphi}{\partial x}(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) - \ell\right)\mathrm{d}A_s^{*+} + \int_0^{\tau_h} \left(-\frac{\partial\varphi}{\partial x}(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) - u\right)\mathrm{d}A_s^{*+} \\
> &\int_0^{\tau_h} \Big(F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, Y_s^{A^*}, Z_s) - F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, \check{Y}_s^{A^*}, \check{Z}_s))\Big)\,\mathrm{d}s + \int_0^{\tau_h} \Delta Z_s \mathrm{d}\overline{B}_s \\
&+ \int_0^{\tau_h} \left(\frac{\partial\varphi}{\partial x}(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) - \ell\right)\mathrm{d}A_s^{*+} + \int_0^{\tau_h} \left(-\frac{\partial\varphi}{\partial x}(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) - u\right)\mathrm{d}A_s^{*+} \\
= &\int_0^{\tau_h} (\delta_s^Y \Delta Y_s + \delta_s^Z \Delta Z_s)\mathrm{d}s + \int_0^{\tau_h} \Delta Z_s \mathrm{d}\overline{B}_s + \int_0^{\tau_h} \delta_s^+ \Delta Y_s \mathrm{d}A_s^{*+} + \int_0^{\tau_h} \delta_s^- \Delta Y_s \mathrm{d}A_s^{*-},
\end{aligned}$$

where

$$\begin{aligned}
\delta_s^Y &= \frac{F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, \check{Y}_s^{A^*}, Z_s) - F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, Y_s^{A^*}, Z_s))}{\Delta Y_s} 1_{\mathbb{R}\setminus\{0\}}(\Delta Y_s) \\
\delta_s^Z &= \frac{F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, Y_s^{A^*}, \check{Z}_s) - F(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, Y_s^{A^*}, Z_s))}{\Delta Z_s} 1_{\mathbb{R}\setminus\{0\}}(\Delta Z_s) \\
\delta_s^+ &= \frac{\frac{\partial \varphi}{\partial x}(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) - \ell}{\Delta Y_s} 1_{\mathbb{R}\setminus\{0\}}(\Delta Y_s) \\
\delta_s^- &= \frac{-\frac{\partial \varphi}{\partial x}(X_s^{A^*,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) - u}{\Delta Y_s} 1_{\mathbb{R}\setminus\{0\}}(\Delta Y_s)
\end{aligned}$$

Now, consider the processes $(\Gamma^{Y^{A^*}}, \Gamma^Z) = (\Gamma_t^{Y^{A^*}}, \Gamma_t^Z)_{t \geq 0}$ such that

$$\begin{aligned}
\Gamma_{\tau_h}^{Y^{A^*}} &= 1 - \int_0^{\tau_h} \Gamma_s^{Y^{A^*}} (\delta_s^Y \mathrm{d}s + \delta^+ \mathrm{d}A_s^{*+} + \delta^- \mathrm{d}A_s^{*-}) \\
\Gamma_{\tau_h}^Z &= 1 - \int_0^{\tau_h} \Gamma_s^Z \delta_s^Y \mathrm{d}\overline{B}_s.
\end{aligned}$$

Clearly, $\Gamma^{Y^{A^*}} \geq 0$ and $\Gamma^Z \geq 0$. Then, we obtain by Itō's integration by parts that

$$\begin{aligned}
\Gamma_{\tau_h}^{Y^{A^*}} \Gamma_{\tau_h}^Z \Delta Y_{\tau_h} &= \int_0^{\tau_h} \Gamma_s^Z \Delta Y_s \mathrm{d}\Gamma_s^{Y^{A^*}} + \int_0^{\tau_h} \Gamma_s^{Y^{A^*}} \Gamma_s^Z \mathrm{d}\Delta Y_s \\
&\quad - \int_0^{\tau_h} \delta_s^Z \Gamma_s^{Y^{A^*}} \Gamma_s^Z \Delta Y_s \mathrm{d}\overline{B}_s - \int_0^{\tau_h} \delta_s^Z \Gamma_s^{Y^{A^*}} \Gamma_s^Z \Delta Z_s \mathrm{d}s \\
&> - \int_0^{\tau_h} \Gamma_s^Z \Delta Y_s \Gamma_s^{Y^{A^*}} (\delta_s^Y \mathrm{d}s + \delta^+ \mathrm{d}A_s^{*+} + \delta^- \mathrm{d}A_s^{*-}) \\
&\quad + \int_0^{\tau_h} \Gamma_s^{Y^{A^*}} \Gamma_s^Z ((\delta_s^Y \Delta Y_s + \delta_s^Z \Delta Z_s) \mathrm{d}s + \Delta Z_s \mathrm{d}\overline{B}_s \\
&\qquad + \delta_s^+ \Delta Y_s \mathrm{d}A_s^{*+} + \delta_s^- \Delta Y_s \mathrm{d}A_s^{*-}) \\
&\quad - \int_0^{\tau_h} \delta_s^Z \Gamma_s^{Y^{A^*}} \Gamma_s^Z \Delta Y_s \mathrm{d}\overline{B}_s - \int_0^{\tau_h} \delta_s^Z \Gamma_s^{Y^{A^*}} \Gamma_s^Z \Delta Z_s \mathrm{d}s \\
&= \int_0^{\tau_h} \Gamma_s^{Y^{A^*}} \Gamma_s^Z (\Delta Z_s - \delta_s^Z \Delta Y_s) \mathrm{d}\overline{B}_s
\end{aligned} \tag{4.5.5}$$

Since $F(\cdot, \cdot, y, z)$ is Lipschitz continuous in (y, z), there exists a positive constant C' such that $|\delta_s^Y| + |\delta_s^Z| \leq C'$. Moreover, $\delta^{\pm} \in \mathbb{H}_{\tau_{t,h}}^2(\mathbf{F}^X, \mathsf{P})$, since φ is a C^2-function and $X^{A^*,\Pi^{\mathsf{P}}}, \Pi^{\mathsf{P}} \in \mathbb{H}_{\tau_{t,h}}^2(\mathbf{F}^X, \mathsf{P})$. These imply that $\Gamma^{Y^{A^*}}, \Gamma^Z \in \mathbb{H}_{\tau_h}^2(\mathbf{F}^X, \mathsf{P})$ (see, e.g., Jeanblanc et al., 2009, Exercise 9.4.3.5). Therefore, the integral (4.5.5) is a martingale (cf. Øksendal, 2010, Theorem 3.2.1), and it follows that

$$\mathsf{E}_{\mathscr{F}_t^X}^{\mathsf{P}}[\Gamma_{\tau_h}^{Y^{A^*}} \Gamma_{\tau_h}^Z \Delta Y_{\tau_h}] > \mathsf{E}_{\mathscr{F}_t^X}^{\mathsf{P}}\left[\int_0^{\tau_h} \Gamma_s^{Y^{A^*}} \Gamma_s^Z (\Delta Z_s - \delta_s^Z \Delta Y_s) \mathrm{d}\overline{B}_s\right] = 0,$$

which is a contradiction, since $(X_0^{A^*,\Pi^{\mathsf{P}}}, \Pi_0^{\mathsf{P}}) = (x, \pi)$ is the local minimum point, i.e., $\Delta Y_t \leq \Delta Y_0 = 0$ for any $t > 0$. Since this holds for any feasible policy, we conclude that $\varphi(x, \pi)$ is the viscosity supersolution to the HJB equation (4.4.21).

Subsolution. The proof of subsolution can be obtained similarly, using a contradiction argument. That is, if we assume that (x, π) is a local maximum point of $V - \varphi$ and suppose, in contrary, that

$$\begin{aligned}\mathscr{L}\varphi(x,\pi) + F(x,\pi,\varphi(x,\pi),\overline{\nabla}\varphi(x,\pi)) &< 0,\\ -\ell + \frac{\partial\varphi}{\partial x}(x,\pi) &< 0, \quad \text{and}\\ -u - \frac{\partial\varphi}{\partial x}(x,\pi) &< 0,\end{aligned}$$

and define the stopping time $\overline{\varrho}_t$, $\check{\varrho}_t$, $\hat{\varrho}_t$ by

$$\begin{aligned}\overline{\varrho}_t &\triangleq \inf\{s > t : \mathscr{L}\varphi(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) + F(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}, \check{Y}_s^A, \check{Z}_s)) = 0\}\\ \check{\varrho}_t &\triangleq \inf\{s > t : -\ell + \frac{\partial\varphi}{\partial x}(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) = 0\} \quad \text{and}\\ \hat{\varrho}_t &\triangleq \inf\{s > t : -u - \frac{\partial\varphi}{\partial x}(X_s^{A,\Pi^{\mathsf{P}}}, \Pi_s^{\mathsf{P}}) = 0\},\end{aligned}$$

then one can adapt the argument used in the second part of the supersolution proof, by replacing $\tau_{t,h}$ with $\varrho_{t,h} \triangleq \overline{\varrho}_t \wedge \check{\varrho}_t \wedge \hat{\varrho}_t \wedge (t+h)$, to derive a contradiction here. Hence, $\varphi(x, \pi)$ is the viscosity subsolution to the HJB equation (4.4.21), which completes the proof. ■

4.6 COORDINATE TRANSFORMATION

As mentioned previously, there is no known analytical solution to the HJB equation (4.4.21). Therefore, we must resort to numerical methods. In this paper, we develop a numerical scheme based on the *Markov Chain Approximation* (MCA) method (cf. Kushner and Martins, 1991) to address the singular stochastic control problem under smooth ambiguity. We adopt the MCA approach because it is among the most efficient methods for solving free boundary problems, such as our HJB equation, on a uniform, rectangular grid, which aligns well with the structure of our problem.[4] Our primary objective is to ensure that the proposed MCA scheme not only maintains numerical consistency with the HJB equation, but also has computational efficiency and yields implications that are intuitive from a managerial perspective.

However, such a numerical scheme cannot be obtained in a straightforward manner, as the HJB equation (4.4.21) involves a non-constant covariance structure, and the processes $X^{A,\Pi^{\mathsf{P}}}$ and Π^{P} are perfectly correlated (i.e., exhibit 100% dependence). These features introduce several challenges for numerical implementation.

First, in the MCA framework, multi-dimensional diffusion processes require solving large systems of (nonlinear) equations. Reducing the dimensionality of the problem can significantly improve computational efficiency. Second, the presence of a non-constant covariance matrix and

[4] See Kumar and Muthuraman (2004) for a method tailored to irregular grid domains in free boundary problems.

the perfect correlation between $X^{A,\Pi^{\mathsf{P}}}$ and Π^{P} complicates the construction of a numerical grid that guarantees convergence of the MCA to the solution of the HJB equation. See, Kushner and Dupuis (2001, p. 110) for an example of such complications in the presence of correlated diffusions.

To address these challenges, we adopt the coordinate transformation technique introduced by Johnson and Peskir (2017).[5] This transformation reduces the two-dimensional stochastic system $(X^{A,\Pi^{\mathsf{P}}}, \Pi^{\mathsf{P}})$ to a system involving a single diffusion process and two processes of bounded variations. We show that this approach not only overcomes the aforementioned difficulties but also retains strong managerial interpretability.

To proceed, we define the transformation $\mathcal{T} : \mathcal{E} \to \mathbb{R}$ by $\mathcal{T}(x,\pi) = x - \frac{\sigma}{\eta}\log(\frac{\pi}{1-\pi})$. Then, we denote $X^{(1)} \triangleq X^{A,\Pi^{\mathsf{P}}}$ and $X^{(2)} \triangleq \mathcal{T}(X^{(1)}, \Pi^{\mathsf{P}})$. By Itō's lemma, we have

$$\mathrm{d}X_t^{(1)} = \alpha^{(1)}\left(X_t^{(1)}, X_t^{(2)}\right)\mathrm{d}t + \mathrm{d}A_t + \sigma \mathrm{d}\overline{B}_t \tag{4.6.1}$$

$$\mathrm{d}X_t^{(2)} = \alpha^{(2)}\left(X_t^{(1)}, X_t^{(2)}\right)\mathrm{d}t + \mathrm{d}A_t \tag{4.6.2}$$

where $(X_0^{(1)}, X_0^{(2)}) = (x_1, x_2) = (x, x - \frac{\sigma}{\eta}\log(\frac{\pi}{1-\pi}))$,

$$\alpha^{(1)}(x_1,x_2) \triangleq \beta\left(\tilde{x} - x_1\right) + \frac{\sigma}{2}(\overline{\theta} + \underline{\theta}) - \frac{1}{2}\eta\sigma^2 \cdot \tanh\left(\frac{\eta}{2\sigma}\left(x_1 - x_2\right)\right),$$
$$\alpha^{(2)}(x_1,x_2) \triangleq \beta\left(\tilde{x} - x_1\right) + \frac{\sigma}{2}(\overline{\theta} + \underline{\theta}) - \frac{1}{2}\eta\sigma(\sigma+1) \cdot \tanh\left(\frac{\eta}{2\sigma}\left(x_1 - x_2\right)\right),$$

and $\tanh$ is the hyperbolic tangent function, i.e., $\tanh(x) = \frac{\exp(2x)-1}{\exp(2x)+1}$. The invertibility of $\mathcal{T}$ implies

$$\Pi_t^{\mathsf{P}} \triangleq \mathcal{T}^{-1}(X_t^{(1)}, X_t^{(2)}) = \frac{1}{2} + \frac{1}{2}\tanh\left(\frac{\eta}{2\sigma}\left(X_t^{(1)} - X_t^{(2)}\right)\right). \tag{4.6.3}$$

Observe that the transformed process (4.6.2) shares a similar structure with the controlled process (4.6.1), but without a diffusion term. Moreover, its evolution, when coupled with the dynamics of the controlled process, facilitates the identification of the belief process (4.6.3) over time. For this reason, we refer to $X^{(2)}$ as the *auxiliary controlled process*. Furthermore, the value function can be expressed as $V(x,\pi) = V(x_1, \mathcal{T}^{-1}(x_1,x_2))$. Hence, solving the transformed problem yields the same solution as that of the original value function.

Now notice that (4.6.1) and (4.6.2) imply that

$$\mathrm{d}\left(X_t^{(1)} - X_t^{(2)}\right) = \frac{1}{2}\eta\sigma \cdot \tanh\left(\frac{\eta}{2\sigma}\left(X_t^{(1)} - X_t^{(2)}\right)\right)\mathrm{d}t + \sigma\mathrm{d}\overline{B}_t. \tag{4.6.4}$$

In fact, $\tanh$ is an increasing function, with $\tanh(0) = 0$ and $\tanh(x) \in (-1,1)$ for any $x \in \mathbb{R}$, which ensures that Π^{P} defined in (4.6.3) lies in $(0,1)$. This also implies that the process (4.6.4)

[5] See also Basei et al. (2024), De Angelis (2020), and Federico et al. (2023) for further applications of this technique.

has a *mean-repulsion* from 0. That is, the process $X^{(1)} - X^{(2)}$ is unstable around 0 and once it moves away from 0, the drift reinforces the movement, contrary to the familiar effect of mean reversion. Since $\tanh$ is an increasing function bounded between $(-1, 1)$, a highly positive value of $X^{(1)} - X^{(2)}$ results in a drift of at most $\frac{1}{2}\eta\sigma$, while a highly negative value yields a drift as low as $-\frac{1}{2}\eta\sigma$. Consequently, the process (4.6.4) is *transient*, with $\pm \lim_{T\to\infty} \frac{1}{2}\eta\sigma T = \pm\infty$ representing the corresponding absorbing states. This implies, by (4.6.3), that $\Pi^{\mathsf{P}}_{\infty} \in \{0, 1\}$, consistent with its fundamental characteristic.

In practical terms, the previous argument implies that when the reservoir-controlled process is substantially larger than its auxiliary process, the belief process tends to increase, indicating a downward (drought) trend. Conversely, when the reservoir-controlled process is significantly lower than its auxiliary process, the belief process tends to increase, indicating an upward (flood) trend.

Under the coordinate transformation, the newly defined BSDE is

$$\mathrm{d}\widehat{Y}^A_t = -\widehat{F}(X^{(1)}_t, X^{(2)}_t, \widehat{Y}^A_t, \widehat{Z}_t)\mathrm{d}t + \ell \mathrm{d}A^+_t + u\mathrm{d}A^-_t + \widehat{Z}_t \mathrm{d}\overline{B}_t,$$

where $\widehat{F}(x_1, x_2, y, z) \triangleq F(x_1, \mathcal{T}^{-1}(x_1, x_2), y, z)$. Thus, $\widehat{Y}^A = Y^A$, by the comparison theorem (cf. Zhang (2017, Theorem 6.2.5)). Therefore, the denoted by $\widehat{V}(x_1, x_2) \triangleq \sup_{A\in\mathscr{D}} \widehat{Y}^{A;x_1,x_2}_0$ is the optimal singular control of the transformed BSDE, then it also holds that $\widehat{V}(x_1, x_2) = V(x_1, \mathcal{T}^{-1}(x_1, x_2))$. Now, we provide a verification theorem for the transformed value function $\widehat{V}$.

Theorem 15 (Verification theorem)**.** *Suppose that* $\widehat{V} \in C^2((\underline{x}, \overline{x}) \times \mathbb{R})$ *is such that*

$$\begin{aligned}\widehat{V}(x_1, x_2) \triangleq \sup_{A\in\mathscr{D}} \mathsf{E}^{\mathsf{P}}\Bigg[\int_0^\infty e^{-rt} f(X^{(1)}_t)\mathrm{d}t - \int_0^\infty (\ell \mathrm{d}A^+_t + u\mathrm{d}A^-_t) \\ + \int_0^\infty h\left(\mathcal{T}^{-1}(X^{(1)}_t, X^{(2)}_t), \sigma\frac{\partial \widehat{V}}{\partial x_1}(X^{(1)}_t, X^{(2)}_t)\right)\mathrm{d}t\Bigg].\end{aligned}$$

If $\phi \in C^2((\underline{x}, \overline{x}) \times \mathbb{R})$ *is such that*

$$\begin{aligned}\widehat{\mathscr{L}}\phi(x_1, x_2) + \widehat{F}(x_1, x_2, \phi(x, x_2), \sigma\frac{\partial\phi}{\partial x_1}(x_1, x_2)) \leq 0 \ \ \textit{and} \\ -u \leq \frac{\partial\phi}{\partial x_1}(x_1, x_2) + \frac{\partial\phi}{\partial x_2}(x_1, x_2) \leq \ell,\end{aligned}$$

where

$$\widehat{\mathscr{L}}\phi(x_1, x_2) \triangleq \alpha^{(1)}(x_1, x_2)\frac{\partial\phi}{\partial x_1}(x_1, x_2) + \alpha^{(2)}(x_1, x_2)\frac{\partial\phi}{\partial x_2}(x_1, x_2) + \sigma\frac{\partial^2\phi}{\partial x^2}(x_1, x_2),$$

then $\widehat{V}(x_1, x_2) \geq \phi(x_1, x_2)$ *for any* $(x_1, x_2) \in (\underline{x}, \overline{x}) \times \mathbb{R}$ *and* $A \in \mathscr{D}$.

Moreover, for each $x_2 \in \mathbb{R}$ if there exists a pair of free-boundary interval $\widehat{\mathcal{X}}(x_2) \subseteq (\underline{x}, \overline{x})$ such that

$$
\begin{aligned}
\widehat{\mathscr{L}}\phi(x_1,x_2) + \widehat{F}(x_1,x_2,\phi(x_1,x_2),\sigma\frac{\partial\phi}{\partial x_1}(x_1,x_2)) &= 0 \qquad \text{on } \widehat{\mathcal{X}}(x_2),\\
\frac{\partial\phi}{\partial x_1}(\hat{a}(x_2),x_2) + \frac{\partial\phi}{\partial x_2}(\hat{a}(x_2),x_2) &= \ell,\\
\frac{\partial\phi}{\partial x_1}(\hat{b}(x_2),x_2) + \frac{\partial\phi}{\partial x_2}(\hat{b}(x_2),x_2) &= -u,\\
\frac{\partial^2\phi}{\partial x_1^2}(\hat{a}(x_2),x_2) - \frac{\partial^2\phi}{\partial x_2^2}(\hat{a}(x_2),x_2) + 2\frac{\partial^2\phi}{\partial x_1\partial x_2}(\hat{a}(x_2),x_2) &= 0\\
= \frac{\partial^2\phi}{\partial x_1^2}(\hat{b}(x_2),x_2) - \frac{\partial^2\phi}{\partial x_2^2}(\hat{b}(x_2),x_2) + 2\frac{\partial^2\phi}{\partial x_1\partial x_2}(\hat{b}(x_2),x_2),&
\end{aligned}
$$

where

$$
\begin{aligned}
\hat{a}(x_2) &\triangleq \inf\{x_1 : x_1 \in \widehat{\mathcal{X}}(x_2)\}\\
\hat{b}(x_2) &\triangleq \sup\{x_1 : x_1 \in \widehat{\mathcal{X}}(x_2)\}
\end{aligned}
$$

and $\widehat{A}^* \in \mathscr{D}$ such that $\widehat{A}^{*-}$ and $\widehat{A}^{*+}$ only increase at $\hat{a}(x_2)+$ and $\hat{b}(x_2)-$, respectively, P-a.s., then $\widehat{V}(x_1,x_2) = \phi(x_1,x_2)$.

Similar to that of Theorem 14, we can infer of Theorem 15 that $\widehat{V}$ is a classical solution of the HJB equation:

$$
\begin{aligned}
0 = \max\Bigg\{&\widehat{\mathscr{L}}\widehat{V}(x_1,x_2) + \widehat{g}\left(x_1,x_2,\widehat{V}(x_1,x_2),\sigma\frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2)\right),\\
&-\ell + \frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2) + \frac{\partial\widehat{V}}{\partial x_2}(x_1,x_2), -u - \frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2) - \frac{\partial\widehat{V}}{\partial x_2}(x_1,x_2)\Bigg\}. \qquad (4.6.5)
\end{aligned}
$$

Note that the total derivatives of $\widehat{V}$ with respect to x_1 and x_2 are

$$
\begin{aligned}
\frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2) &= \frac{\partial V}{\partial x_1}(x_1,\pi) + \frac{\eta}{4\sigma}\left(1+\tanh^2\left(\frac{\eta}{2\sigma}(x_1-x_2)\right)\right)\frac{\partial V}{\partial \pi}(x_1,\pi), \qquad (4.6.6)\\
\frac{\partial\widehat{V}}{\partial x_2}(x_1,x_2) &= -\frac{\eta}{4\sigma}\left(1+\tanh^2\left(\frac{\eta}{2\sigma}(x_1-x_2)\right)\right)\frac{\partial V}{\partial \pi}(x_1,\pi)
\end{aligned}
$$

Therefore, the concavity of V in x_1 (cf. Proposition 10) implies that $\frac{\partial V}{\partial x_1} = \frac{\partial \widehat{V}}{\partial x_1} + \frac{\partial \widehat{V}}{\partial x_2}$ is strictly decreasing in x_1. This means that $(\hat{a},\hat{b})(x_2)$ is unique for any $x_2 \in \mathbb{R}$, where

$$
\hat{a}(x_2) \triangleq \sup\left\{x_1 \in (\underline{x},\overline{x}) : \frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2) + \frac{\partial\widehat{V}}{\partial x_2}(x_1,x_2) \geq \ell\right\}, \quad \text{and} \qquad (4.6.7)
$$

$$
\hat{b}(x_2) \triangleq \inf\left\{x_1 \in (\underline{x},\overline{x}) : \frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2) + \frac{\partial\widehat{V}}{\partial x_2}(x_1,x_2) \leq -u\right\}. \qquad (4.6.8)
$$

Moreover, we have from (4.6.6) that $\sigma\frac{\partial\widehat{V}}{\partial x_1} = \overline{\nabla}V(x_1, \mathcal{T}^{-1}(x_1,x_2))$. Hence, we conclude that BSDE $(\widehat{V}, \sigma\frac{\partial\widehat{V}}{\partial x_1}, \widehat{A}^*)$ has a unique solution, and identical to $(V, \overline{\nabla}V, A^*)$.

Since the transformation $\mathcal{T}$ is smooth and injective, we can straightforwardly adapt the proof from Proposition 12 to guarantee the viscosity property of $\widehat{V}(x_1,x_2)$.

Proposition 13. *Suppose that Assumption 3 holds. Then $\widehat{V}(x_1,x_2)$ is a viscosity solution to the transformed HJB equation (4.6.5).*

Remark 17. *Observe that if $\widehat{V}$ is a classical solution to (4.6.5), then the optimal controlled dynamics are*

$$\mathrm{d}X_t^{(1)} = \alpha_g^{(1)}\Big(X_t^{(1)}, X_t^{(2)}, \frac{\partial\widehat{V}}{\partial x_1}(X_t^{(1)}, X_t^{(2)})\Big)\mathrm{d}t + \sigma\mathrm{d}\overline{B}_t + \mathrm{d}\widehat{A}_t^*, \tag{4.6.9}$$

where $X^{(2)}$ solves (4.6.2), and

$$\alpha_g^{(1)}(x_1,x_2,y) \triangleq \begin{cases} \alpha^{(1)}(x_1,x_2) - \sigma G_\gamma(\pi,y) & \text{if } \gamma\in(0,\infty) \\ \alpha^{(1)}(x_1,x_2) - \sigma G_\infty(\pi,y) & \text{if } \gamma\to\infty \end{cases} \tag{4.6.10}$$

$$G_\gamma(\pi,y) \triangleq \pi(1-\pi)\Bigg(\frac{\gamma(\underline{\theta}-\overline{\theta})^2}{2!}y - \frac{\gamma^2(\underline{\theta}-\overline{\theta})^3}{3!}(1-2\pi)y^2 + \frac{\gamma^3(\underline{\theta}-\overline{\theta})^4}{4!}(1-6\pi(1-\pi))y^3 + \ldots + \mathcal{O}(y^N)\Bigg)$$

$$G_\infty(\pi,y) \triangleq \pi\underline{\theta} + (1-\pi)\overline{\theta} - \underline{\theta}\cdot 1_{(-\infty,0)}(y) - \overline{\theta}\cdot 1_{[0,\infty)}(y)$$

$$\pi \triangleq \mathcal{T}^{-1}(x_1,x_2),$$

$$y \triangleq \frac{\partial V}{\partial x_1}(x_1,x_2).$$

The term $G_\gamma(\pi,y)$ factored with y follows from the N^{th}-order Taylor expansion of the function $h(\cdot,z)$, while $G_\infty(\pi,y)$ is obtained from the limiting case (4.5.1). Here N determines how accurate the approximation of the drift (4.6.10) is when $\frac{\partial\widehat{V}}{\partial x_1}$ deviates from zero. Since $\frac{\partial\widehat{V}}{\partial x_1} \in \Big(-u - \frac{\bar{c}\sigma(\underline{\theta}-\overline{\theta})}{r(r+\beta)}, \ell + \frac{\bar{c}\sigma(\underline{\theta}-\overline{\theta})}{r(r+\beta)}\Big)$, it follows that for a given error tolerance, say $\varepsilon>0$, there exists at least $N_\varepsilon\in\mathbb{N}$ terms in the Taylor expansion such that $\mathcal{O}\big((\max\{u,\ell\}+\frac{\bar{c}\sigma(\underline{\theta}-\overline{\theta})}{r(r+\beta)})^{N_\varepsilon}\big)=\varepsilon$, providing a sufficiently accurate truncation of the drift (4.6.10), when γ is finite. Importantly, regardless of what γ or N_ε are, the process in (4.6.9) always admits a unique strong solution, owing to the Lipschitz continuity of $h(\cdot,z)$ (cf. (4.4.7) and Karatzas and Shreve, 1991, Theorem 2.5).

To illustrate how the ambiguity-adjusted drift (4.6.10) evolves, let us consider Figure 4.1, which plots the graphs of the functions $G_\infty(\pi,y)$ and $G_\gamma(\pi,y)$ for different levels of ambiguity aversion γ and marginal value functions $y=\frac{\partial\widehat{V}}{\partial x_1}$.

First, observe that for any finite γ, the function $G_\gamma(\pi,y)$ equals zero at the absorbing states $\pi\in\{0,1\}$. This indicates that the DM eventually becomes neutral to ambiguity once sufficient information is revealed.

Away from the absorbing states, however, ambiguity aversion alters the drift. As γ increases, $G_\gamma(\pi, y)$ tilts upward when $\pi \to 1$ (i.e., as $\theta \to \underline{\theta}$, the drought regime) provided that the marginal value $\frac{\partial \widehat{V}}{\partial x_1}$ is positive. Conversely, when $\frac{\partial \widehat{V}}{\partial x_1}$ is negative, the function tilts downward as $\pi \to 0$ (i.e., as $\theta \to \overline{\theta}$, the flood regime). As shown by the last panel of Figure 4.1, in the case where the drift follows a maxmin utility *specification ($\gamma \to \infty$), it always tilts toward the worst-case scenario at every states of ambiguity learning.*

Overall, this demonstrates that (4.6.10) *embodies* drift ambiguity[6]*: the reservoir process is distorted toward the worst-case drift whenever the marginal value function is increasing or decreasing. In our setting, this drift ambiguity, after some rearrangement, takes the form of a mean-reversion ambiguity: the reservoir tends to revert either toward a more severe flood regime (water level even higher than expected) or toward a more severe drought regime (water level even lower than expected).*

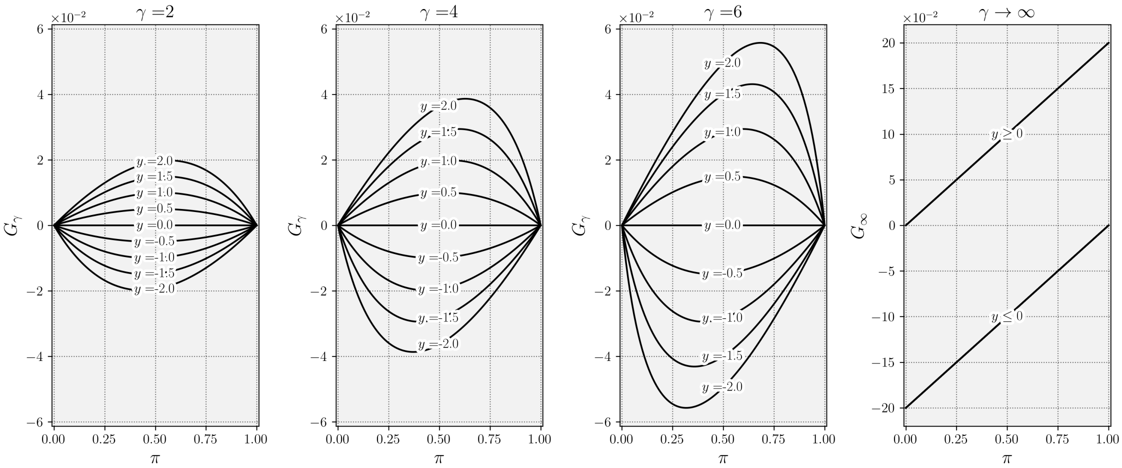

Figure 4.1: Simulation of the term G_γ and G_∞ for $y \triangleq \frac{\partial V}{\partial x_1}(x_1, x_2) \in [-2, 2]$ and $\gamma \in \{2, 4, 6\}$.

4.7 MARKOV CHAIN APPROXIMATION

The main idea of the MCA is to construct a suitable finite difference scheme on a discretized state space that (weakly) solves the transformed HJB equation (4.6.5). This scheme also ensures that the approximate sample paths of $(X^{(1)}, X^{(2)})$, generated via Markov transition processes on the same grid, have locally consistent transition probabilities.

We will later show that such a finite difference scheme exists in our setting, along with the corresponding locally consistent transition probabilities. To establish this, we must first define a

[6] This feature is common in models incorporating ambiguity aversion (or seeking) into continuous-time decision making under uncertainty. See, for example, Archankul et al. (2025), Cheng and Riedel (2013), Hellmann and Thijssen (2018), Nishimura and Ozaki (2007), and Thijssen (2011).

suitable discretized state space for the controlled processes $(X^{(1)}, X^{(2)})$. This step is non-trivial, as $X^{(2)}$ is unbounded and subject to Dirichlet-type far-field boundary conditions. In the following subsection, we introduce an assumption that allows us to address this issue in a tractable manner.

4.7.1 FAR-FIELD BOUNDARY CONDITIONS

Recall that $X^{(2)}$ is a transient process with absorbing states at $\pm\infty$. This implies that $\Pi^{\mathsf{P}}_\infty = 1$ or 0 whenever $\lim_{t\to\infty} X^{(2)}_t = -\infty$ or $+\infty$, respectively. To make numerical implementation feasible, we must impose a truncation: for some small $\varepsilon > 0$, the DM is assumed to believe that $\Pi^{\mathsf{P}}_t \in (\varepsilon, 1-\varepsilon)$ for all $t \geq 0$. Once Π^{P}_t hits either boundary ε or $1-\varepsilon$ at some time $t \geq 0$, the DM concludes that $\Pi^{\mathsf{P}}_s = 0$ or 1, respectively, for all $s > t$. In fact, this assumption is not far-fetched. According to the dynamics of Π^{P} in (4.2.2), when Π^{P} is close to 0 or 1, it is more likely to be absorbed at those values than to revert away. In other words, this assumption allows the DM to conclude the estimate value of their belief early once the states ε or $1-\varepsilon$ is hit. Under this approximation, the transformed process $X^{(2)}$ is assumed to be bounded within the interval:

$$(\underline{x}^\varepsilon_2, \overline{x}^\varepsilon_2) \triangleq \left(\tilde{x} - \frac{\sigma}{\eta}\log\left(\frac{1-\varepsilon}{\varepsilon}\right),\ \tilde{x} - \frac{\sigma}{\eta}\log\left(\frac{\varepsilon}{1-\varepsilon}\right)\right).$$

Therefore, the transformed value function at the boundaries $\underline{x}^\varepsilon_2$ and $\overline{x}^\varepsilon_2$ is

$$\widehat{V}(x_1, x_2) = \begin{cases} \phi_0(x_1) & \text{if } x_2 = \underline{x}^\varepsilon_2 \\ \phi_1(x_1) & \text{if } x_2 = \overline{x}^\varepsilon_2 \end{cases}$$

where $\phi_p(x_1),\ p = 0, 1$ is a classical solution to

$$0 = \max\left\{\mathscr{L}_p\phi_p(x_1) - r\phi_p(x_1) + f(x_1), -\ell + \frac{\partial\phi_p}{\partial x_1}(x_1), -u - \frac{\partial\phi_p}{\partial x}(x_1)\right\} \quad (4.7.1)$$

$$\mathscr{L}_p\phi_p(x_1) \triangleq \beta\left(\tilde{x} - \frac{\sigma}{\beta}\left(p\underline{\theta} + (1-p)\overline{\theta}\right) - x_1\right)\frac{\partial\phi_p}{\partial x_1}(x_1) + \sigma\frac{\partial^2\phi_p}{\partial x_1^2}(x_1).$$

Employing the argument from Archankul et al. (2025, Section 4), we obtain that there exists a concave C^2-function solving (4.7.1). In fact, for each $p = 0, 1$, there are unique inaction region $(a, b) \subseteq (\underline{x}, \overline{x})$ such that

$$\phi_p(x_1) = \begin{cases} -\ell(a - x_1) + F_p(a) + A\hat{\phi}_p(a) + B\check{\phi}_p(a) & \text{if } x_1 \in (\underline{x}, a] \\ F_p(x_1) + A\hat{\phi}_p(x_1) + B\check{\phi}_p(x_1) & \text{if } x_1 \in (a, b) \\ -u(x_1 - b) + F_p(b) + A\hat{\phi}_p(b) + B\check{\phi}_p(b) & \text{if } x_1 \in [b, \overline{x}) \end{cases}$$

where

$$\hat{\phi}_p(x_1) \triangleq \exp\left(\frac{\beta}{2\sigma^2}\alpha_p^2(x_1)\right)\mathcal{D}_{-\frac{r}{\beta}}\left(-\frac{\sqrt{2\beta}}{\sigma}\alpha_p(x_1)\right),$$

$$
\begin{aligned}
\check{\phi}_p(x_1) &\triangleq \exp\left(\frac{\beta}{2\sigma^2}\alpha_p^2(x_1)\right)\mathcal{D}_{-\frac{r}{\beta}}\left(\frac{\sqrt{2\beta}}{\sigma}\alpha_p(x_1)\right),\\
F_p(x_1) &\triangleq \begin{cases} \frac{b}{r} - \frac{\check{c}}{r+\beta}\alpha_p(x_1) + C\hat{\phi}_p(x_1) & \text{if } x_1 \in (\underline{x}, \tilde{x}) \\ \frac{b}{r} + \frac{\hat{c}}{r+\beta}\alpha_p(x_1) + D\check{\phi}_p(x_1) & \text{if } x_1 \in [\tilde{x}, \overline{x}), \text{ and} \end{cases}\\
\alpha_p(x_1) &\triangleq \tilde{x} - \frac{\sigma}{\beta}\left(p\underline{\theta} + (1-p)\overline{\theta}\right) - x_1.
\end{aligned}
$$

Here, $\mathcal{D}_z$ denotes the parabolic cylinder function of index z (see Jeanblanc et al., 2009, Chapter 5). The constants C and D are determined by applying the value matching and smooth pasting conditions on F_p at $\tilde{x}$, that is, by solving $F_p(\tilde{x}-) = F_p(\tilde{x}+)$ and $F_p'(\tilde{x}-) = F_p'(\tilde{x}+)$. Similarly, the constants A, B, a, and b are obtained by imposing value matching and smooth pasting conditions (specific to singular control) on ϕ_p at a and b, i.e., by solving $\phi_p'(a-) = \phi_p'(a+)$, $\phi_p'(b-) = \phi_p'(b+)$, $\phi_p''(a-) = \phi_p''(a+)$, and $\phi_p''(b-) = \phi_p''(b+)$.

Now, we are ready to apply the MCA to solve the transformed value function over the interval $(\underline{x}_2^\varepsilon, \overline{x}_2^\varepsilon)$.

4.7.2 Finite difference approximation

To construct the discretized grid, we partition the domains of $X^{(1)}$ and $X^{(2)}$ into $n \in \mathbb{N}$ equally spaced intervals. The corresponding grid nodes are given by:

$$
\begin{aligned}
\Omega_{x_1}^n &\triangleq \{\underline{x}, \underline{x}+h_1, \underline{x}+2h_1, \ldots, \underline{x}+(n-1)h_1, \overline{x}\}, \quad h_1 = \frac{\overline{x}-\underline{x}}{n}\\
\Omega_{x_2}^{n,\varepsilon} &\triangleq \{\underline{x}_2^\varepsilon, \underline{x}_2^\varepsilon+h_2, \underline{x}_2^\varepsilon+2h_2, \ldots, \underline{x}_2^\varepsilon+(n-1)h_2, \overline{x}_2^\varepsilon\}, \quad h_2 = \frac{\overline{x}_2^\varepsilon - \underline{x}_2^\varepsilon}{n},
\end{aligned}
$$

respectively, and denote $\mathcal{E}^{n,\varepsilon} \triangleq \Omega_{x_1}^n \times \Omega_{x_2}^{n,\varepsilon} \setminus \{\underline{x}_2^\varepsilon, \overline{x}_2^\varepsilon\}$. Then for each $(x_1, x_2) \in \mathcal{E}^{n,\varepsilon}$, the approximate transformed value function is obtained by the recursive mapping $\mathbb{T}^n : \mathbb{R} \to \mathbb{R}$:

$$
\mathbb{T}^n\widehat{V}(x_1,x_2) \triangleq \begin{cases} \max\{\widehat{V}_c(x_1,x_2), \widehat{V}_\ell(x_1,x_2), \widehat{V}_u(x_1,x_2)\} & \text{if } (x_1,x_2) \in \mathcal{E}^{n,\varepsilon} \\ \phi_0(x_1) & \text{if } x_1 \in \Omega_{x_1}^n \text{ and } x_2 = \underline{x}_2^\varepsilon \\ \phi_1(x_1) & \text{if } x_1 \in \Omega_{x_1}^n \text{ and } x_2 = \overline{x}_2^\varepsilon \end{cases} \tag{4.7.2}
$$

where

$$
\begin{aligned}
\widehat{V}_c(x_1,x_2) &\triangleq (1 - r\Delta t_{x_1,x_2}^{h_1,h_2})\left(\sum_{x_1',x_2' \in \mathcal{E}^{n,\varepsilon}} \bar{p}_{x_1,x_2}^{h_1,h_2}(x_1',x_2')\widehat{V}(x_1',x_2') + f(x_1)\Delta t_{x_1,x_2}^{h_1,h_2}\right)\\
\widehat{V}_\ell(x_1,x_2) &\triangleq \sum_{x_1',x_2' \in \mathcal{E}^{n,\varepsilon}} \check{p}_{x_1,x_2}^{h_1,h_2}(x_1',x_2')\widehat{V}(x_1',x_2') - \frac{\ell h_1 h_2}{h_1+h_2}\\
\widehat{V}_u(x_1,x_2) &\triangleq \sum_{x_1',x_2' \in \mathcal{E}^{n,\varepsilon}} \hat{p}_{x_1,x_2}^{h_1,h_2}(x_1',x_2')\widehat{V}(x_1',x_2') - \frac{u h_1 h_2}{h_1+h_2}.
\end{aligned}
$$

The terms $p^{h_1,h_2}_{x_1,x_2}, \check{p}^{h_1,h_2}_{x_1,x_2}$ and $\hat{p}^{h_1,h_2}_{x_1,x_2}$ are obtained by applying the upwind finite-difference scheme to the transform HJB equation (4.6.5). That is,

$$
\begin{aligned}
\bar{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') &\triangleq \frac{\tilde{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')}{\tilde{q}^{h_1,h_2}_{x_1,x_2}} \\
\tilde{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \\
&\triangleq \begin{cases}
\frac{1}{h_1}\Big(\alpha^{(1)}(x_1,x_2)^+ + \sigma\mathcal{T}^{-1}(x_1,x_2)(1-\mathcal{T}^{-1}(x_1,x_2)) \\
\qquad \times H(x_1,x_2,[\widehat{V}(x_1)-\widehat{V}(x_1')]^+/h_1)\Big) + \frac{\sigma^2}{2h_1^2} & \text{if } (x_1',x_2') = (x+h_1,x_2) \\
\frac{1}{h_1}\Big(\alpha^{(1)}(x_1,x_2)^- + \sigma\mathcal{T}^{-1}(x_1,x_2)(1-\mathcal{T}^{-1}(x_1,x_2)) \\
\qquad \times H(x_1,x_2,[\widehat{V}(x_1')-\widehat{V}(x_1)]^-/h_1)\Big) + \frac{\sigma^2}{2h_1^2} & \text{if } (x_1',x_2') = (x-h_1,x_2) \\
\frac{\alpha^{(2)}(x_1,x_2)^\pm}{h_2} & \text{if } (x_1',x_2') = (x_1,x_2\pm h_2) \\
0 & \text{otherwise,}
\end{cases} \\
\tilde{q}^{h_1,h_2}_{x_1,x_2} &\triangleq \frac{1}{h_1}\Big(\big|\alpha^{(1)}(x_1,x_2)\big| + \sigma\mathcal{T}^{-1}(x_1,x_2)(1-\mathcal{T}^{-1}(x_1,x_2)) \\
&\qquad \times H(x_1,x_2,\overline{\nabla}_{h_1}\widehat{V}(x_1,x_2))\Big) + \frac{\big|\alpha^{(2)}(x_1,x_2)\big|}{h_2} + \frac{\sigma^2}{h_1^2}, \\
\overline{\nabla}_{h_1}\widehat{V}(x_1,x_2) &\triangleq \frac{1}{h_1}([\widehat{V}(x_1,x_2)-\widehat{V}(x_1+h_1,x_2)]^+ + [\widehat{V}(x_1-h_1,x_2)-\widehat{V}(x_1,x_2)]^-) \\
\Delta t^{h_1,h_2}_{x_1,x_2} &\triangleq \frac{1}{\tilde{q}^{h_1,h_2}_{x_1,x_2}},
\end{aligned}
$$

$$
\check{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \triangleq \begin{cases}
\frac{h_2}{h_1+h_2} & \text{if } (x_1',x_2') = (x_1+h_1,x_2) \\
\frac{h_1}{h_1+h_2} & \text{if } (x_1',x_2') = (x_1,x_2+h_2) \\
0 & \text{otherwise,}
\end{cases} \tag{4.7.3}
$$

$$
\hat{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \triangleq \begin{cases}
\frac{h_2}{h_1+h_2} & \text{if } (x_1',x_2') = (x_1-h_1,x_2) \\
\frac{h_1}{h_1+h_2} & \text{if } (x_1',x_2') = (x_1,x_2-h_2) \\
0 & \text{otherwise.}
\end{cases} \tag{4.7.4}
$$

Now we suppose that the process $(X^{(1),h_1}, X^{(2),h_2}) = (X^{(1),h_1}_k, X^{(2),h_2}_k)_{k\in\mathbb{N}}$, is the discrete time Markov chain with transition probabilities:

$$
\begin{aligned}
&\mathsf{P}\left(\left(X^{(1),h_1}_{k+1}, X^{(2),h_2}_{k+1}\right) = (x_1',x_2') \mid \left(X^{(1),h_1}_{k}, X^{(2),h_2}_{k}\right) = (x_1,x_2) \in \mathcal{E}^{h_1,h_2}\right) \\
&\triangleq p^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \triangleq \begin{cases}
\bar{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_c(x_1,x_2) \\
\check{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_\ell(x_1,x_2) \\
\hat{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_u(x_1,x_2)
\end{cases}
\end{aligned}
$$

where $\Delta t^{h_1,h_2}_{x_1,x_2}$ is the time interval of the Markov chains. The associated conditional expectation operator is denoted by $\mathsf{E}^{h_1,h_2}_{x_1,x_2}$. Observe that

$$p^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \in (0,1), \quad \forall (x_1',x_2') \in \mathcal{E}^{h_1,h_2} \tag{4.7.5}$$
$$\sum_{(x_1',x_2')\in\mathcal{E}^{h_1,h_2}} p^{h_1,h_2}_{x_1,x_2}(x_1',x_2') = 1,$$

so that the transition probabilities are well-defined.

Note that when $\mathbb{T}\widehat{V}(x_1,x_2) = \widehat{V}_c(x_1,x_2)$, the state (x_1,x_2) lies within the inaction region, and it is optimal to continue with transition probabilities $\bar{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')$, for all $(x_1',x_2') \in \mathcal{E}^{h_1,h_2,\varepsilon}$. On the other hand, if $\mathbb{T}\widehat{V}(x_1,x_2)$ equals either $\widehat{V}_\ell(x_1,x_2)$ or $\widehat{V}_u(x_1,x_2)$, this indicates that it is optimal to exert control from below with probability $\check{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')$ or above with probability $\hat{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')$, respectively. In other words, the state (x_1,x_2) is instantaneously shifted to

$$(x_1 + \Delta A^{+,h_1,h_2}_k, x_2 + \Delta A^{+,h_1,h_2}_k) \text{ or} \tag{4.7.6}$$
$$(x_1 - \Delta A^{-,h_1,h_2}_k, x_2 - \Delta A^{-,h_1,h_2}_k), \tag{4.7.7}$$

respectively, where $\Delta A^{\pm,h_1,h_2}_k \triangleq \frac{h_1 h_2}{h_1+h_2}$ for all $k \in \mathbb{N}$.

Remark 18. *Observe that $(\underline{x}^\varepsilon_2, \overline{x}^\varepsilon_2)$ is getting significantly larger than $(\underline{x}, \overline{x})$ when ε is brought closer and closer to zero. Consequently, increasing n to a sufficiently large number means $h_2 \gg h_1$. That is, in limit $\varepsilon \to 0$ and $n \to \infty$, the reflecting state transition probabilities (4.7.3) and (4.7.4) reduce to*

$$\check{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \approx \begin{cases} 1 & \text{if } (x_1',x_2') = (x_1+h_1,x_2) \\ 0 & \text{otherwise,} \end{cases} \quad \textit{and}$$
$$\hat{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2') \approx \begin{cases} 1 & \text{if } (x_1',x_2') = (x_1-h_1,x_2) \\ 0 & \text{otherwise,} \end{cases}$$

respectively, consistent with the property of the pre-transformed processes $(X^{A,\Pi^{\mathsf{P}}}, \Pi^{\mathsf{P}})$.

The following result affirms that the recursive mapping $\mathbb{T}^n$ gives a unique solution for a sufficiently large $n \in \mathbb{N}$.

Theorem 16. *Suppose that $n \in \mathbb{N}$ is chosen large enough that*

$$\max_{x_1,x_2\in\mathcal{E}^{n,\varepsilon}} r\Delta t^{h_1,h_2}_{x_1,x_2} < 1. \tag{4.7.8}$$

Then the recursive mapping $\mathbb{T}^n\widehat{V}$ has a unique fixed point, which converges to the viscosity solution of the transformed HJB equation (4.4.21) as $n \to \infty$. Moreover, for each $x_2 \in \Omega^{n,\varepsilon}_{x_2}$ the optimal control barriers are approximated by $(\hat{a}^n(x_2), \hat{b}^n(x_2))$:

$$\hat{a}^n(x_2) \triangleq \max\left\{ x_1 \in \Omega^n_{x_1} : \widehat{V}(x_1,x_2) \le \frac{h_2}{h_1+h_2}\widehat{V}(x_1+h_1,x_2)\right.$$

$$
+ \frac{h_2}{h_1+h_2}\widehat{V}(x_1, x_2+h_2) - \frac{\ell h_1 h_2}{h_1+h_2}\Bigg\}, \quad \text{and}
$$

$$
\hat{b}^n(x_2) \triangleq \min\Bigg\{x_1 \in \Omega^n_{x_1} : \widehat{V}(x_1,x_2) \le \frac{h_2}{h_1+h_2}\widehat{V}(x_1-h_1, x_2) + \frac{h_2}{h_1+h_2}\widehat{V}(x_1, x_2-h_2) - \frac{u h_1 h_2}{h_1+h_2}\Bigg\}.
$$

Proof of Theorem 16. To establish that $\mathbb{T}^n$ admits a unique fixed point, it suffices to show that it is a contraction mapping. This is immediate when $x_2 \in \underline{x}_2^\epsilon, \overline{x}_2^\epsilon$. Therefore, it remains to verify the contraction property for $x_2 \in (\underline{x}_2^\epsilon, \overline{x}_2^\epsilon)$.

Suppose that $\widehat{V}_i$, for $i = 1, 2$, are solutions to (4.7.2), and define $\Delta\widehat{V} \triangleq \widehat{V}_1 - \widehat{V}_2$. It is straightforward to verify that the right-hand side of (4.7.8) is positive and strictly decreasing in n, since $\widehat{V}_i$ is concave and Lipschitz continuous. Hence, there exists $n^* \in \mathbb{N}$ such that (4.7.8) holds for all $n \ge n^*$. Then, for any $(x_1, x_2) \in \mathcal{E}^{n,\epsilon},\ n \ge n^*$, we have

$$
\begin{aligned}
&|\mathbb{T}^n\widehat{V}_1(x_1,x_2) - \mathbb{T}^n\widehat{V}_2(x_1,x_2)| \\
&\qquad \le \max\Bigg\{\Bigg|(1 - r\Delta t^{h_1,h_2}_{x_1,x_2}) \sum_{x_1',x_2' \in \mathcal{E}^{n,\varepsilon}} \bar{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')\Delta\widehat{V}(x_1',x_2')\Bigg|, \\
&\qquad \Bigg|\sum_{x_1',x_2' \in \mathcal{E}^{n,\varepsilon}} \check{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')\Delta\widehat{V}(x_1',x_2')\Bigg|, \Bigg|\sum_{x_1',x_2' \in \mathcal{E}^{n,\varepsilon}} \hat{p}^{h_1,h_2}_{x_1,x_2}(x_1',x_2')\Delta\widehat{V}(x_1',x_2')\Bigg|\Bigg\} \\
&\qquad \le \widehat{C}\Delta\widehat{V}(x_1',x_2'), \text{ for some } 0 < \widehat{C} < 1, \text{ for any } (x_1',x_2') \in \mathcal{E}^{n,\varepsilon}.
\end{aligned}
$$

The second inequality follows from assumption (4.7.8) and the transition probabilities (4.7.5). This implies that $\mathbb{T}^n\widehat{V}$ is a contraction mapping, which, therefore, has a unique fixed point (cf. Aliprantis and Border, 2006, Theorem 3.48). This means that the recursive mapping (4.7.2) weakly solves the transformed HJB equation (4.6.5), and converges in the sense of viscosity solution (cf. Definition 23) when $n \to \infty$. Moreover, the pair $(\hat{a}^n(x_2), \hat{b}^n(x_2))$ is obtained by discretizing (4.6.7) and (4.6.8), which exists uniquely thanks to Proposition 11. This completes the proof. ∎

Now, let $(\Delta X_k^{(1),h_1}, \Delta X_k^{(2),h_2}) \triangleq (X_{k+1}^{(1),h_1} - X_k^{(2),h_2}, X_{k+1}^{(2),h_1} - X_k^{(2),h_2})$, then for each $(x_1, x_2) \in \mathcal{E}^{n,\varepsilon}$,

$$
\mathsf{E}^{h_1,h_2}_{x_1,x_2}\left[\Delta X_k^{(1),h_1}\right]
$$

$$= \begin{cases} \alpha_g^{(1)}\Big(X_k^{(1),h_1}, X_k^{(2),h_2}, \\ \overline{\nabla}_{h_1}\widehat{V}(X_k^{(1),h_1}, X_k^{(2),h_2})\Big)\Delta t_{x_1,x_2}^{h_1,h_2} + \mathcal{O}(h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_c(x_1,x_2) \\ \Delta A_k^{+,h_1,h_2} + \mathcal{O}(h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_\ell(x_1,x_2) \\ -\Delta A_k^{-,h_1,h_2} + \mathcal{O}(h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_u(x_1,x_2) \end{cases} \tag{4.7.9}$$

$$\mathsf{E}_{x_1,x_2}^{h_1,h_2}\left[\Delta X_k^{(2),h_2}\right]$$

$$= \begin{cases} \alpha^{(2)}(X_k^{(1),h_1}, X_k^{(2),h_2})\Delta t_{x_1,x_2}^{h_1,h_2} + \mathcal{O}(h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_c(x_1,x_2) \\ \Delta A_k^{+,h_1,h_2} + \mathcal{O}(h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_\ell(x_1,x_2) \\ -\Delta A_k^{-,h_1,h_2} + \mathcal{O}(h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(x_1,x_2) = \widehat{V}_u(x_1,x_2) \end{cases}$$

$$\mathsf{E}_{x_1,x_2}^{h_1,h_2}\left[\left(\Delta X_k^{(1),h_1} - \mathsf{E}_{x_1,x_2}^{h_1,h_2}\left[\Delta X_k^{(1),h_1}\right]\right)^2\right] = \sigma^2\Delta t_{x_1,x_2}^{h_1,h_2} + \mathcal{O}(h_1,h_2)$$

$$\mathsf{E}_{x_1,x_2}^{h_1,h_2}\left[\left(\Delta X_k^{(2),h_2} - \mathsf{E}_{x_1,x_2}^{h_1,h_2}\left[\Delta X_k^{(2),h_2}\right]\right)^2\right] = \mathcal{O}(h_1,h_2) \tag{4.7.10}$$

where $\mathcal{O}(h_1,h_2)$ is the approximate error proportion to the size of h_1, h_2. According to Kushner and Dupuis (2001, Section 8.3), equations (4.7.9) through (4.7.10) establish the local consistency of the discretized processes $(X^{(1),h_1}, X^{(2),h_2})$ with the original processes $(X^{(1)}, X^{(2)})$. In particular, the discrete transition probabilities are constructed to match the drift and diffusion of the continuous model in the limit as $\Delta t_{x_1,x_2}^{h_1,h_2} \to 0$. As a result, our numerical scheme based on the MCA is consistent with the original stochastic singular control problem (with smooth ambiguity) in the sense that the value functions and optimal policies derived from the discretized model converge to those of the continuous problem as $h_1, h_2 \to 0$.

In the following section, we employ our numerical scheme to conduct a comparative statics analysis.

4.8 NUMERICAL COMPARATIVE STATICS

To illustrate the numerical examples, we assume $(\underline{x}, \overline{x}) = (0, 10)$ and fix the following parameters as a based case.

$$r = 0.2,\ \beta = 0.15,\ \tilde{x} = 5,\ \sigma = 0.5,\ \underline{\rho} = \overline{\rho} = 0.25,\ \ell = u = 2, \check{c} = \hat{c} = 2, b = 5.$$

The marginal error of the absorbing states is assumed to be $\varepsilon = 10^{-14}$. Therefore, in the based case, the corresponding boundary of the process $X^{(2)}$ is $(\underline{x}_2^\varepsilon, \overline{x}_2^\varepsilon) \approx (-6.51, 16.51)$, which can be changed according to the different model parameters. The refinement of the grid is determined by the number n that obeys Assumption (4.7.8). The general characteristics of the optimal control policy based on these parameters are demonstrated in Figure 4.2 with a detailed description below. To keep the argument concise, from now on, we denote $\widehat{\alpha}_g^{(1)}(x_1,x_2) \triangleq \alpha_g^{(1)}(x_1, x_2, \frac{\partial\widehat{V}}{\partial x_1}(x_1,x_2))$ for any $(x_1,x_2) \in (\underline{x}, \overline{x}) \times (\underline{x}_2^\varepsilon, \overline{x}_2^\varepsilon)$

4.8.1 GENERAL CHARACTERISTICS OF THE OPTIMAL CONTROL POLICY

The controlled process achieves its maximum value when it remains near the long-term mean level $\tilde{x}$, which, by construction, corresponds to the highest expected profit. This occurs around the point where $\widehat{\alpha}_g^{(1)}(X^{(1)}, X^{(2)}) = 0$, i.e., when $X^{(1)} = X^{(2)} = \tilde{x} + \frac{\sigma}{2\beta}(\overline{\theta} + \underline{\theta})$ and $\frac{\partial \widehat{V}}{\partial x_1}(X^{(1)}, X^{(2)}) = 0$.

Intuitively, when the reservoir level consistently trends toward either $\overline{\theta}$ or $\underline{\theta}$, the likelihood of incurring losses increases due to water shortage or excess, respectively. Therefore, maintaining the reservoir level near $\tilde{x}$, when the expected inflow trend is $\hat{\theta} = (\overline{\theta} + \underline{\theta})/2$, yields the highest long-run profit. Moreover, this maximum profit is attained most stably when the reservoir has the least tendency to deviate from these optimal states, i.e., when $\frac{\partial \widehat{V}}{\partial x_1}(X^{(1)}, X^{(2)}) = 0$. For this reason, we refer to the states satisfying $\widehat{\alpha}_g^{(1)}(X^{(1)}, X^{(2)}) = 0$ as the *maximum benefit states*.

Consequently, the optimal control policy is to exert maximal intervention to keep the controlled processes near these maximum benefit states. Combined with the control directions defined in (4.7.6) and (4.7.7), this leads to the narrowest inaction region, as illustrated in Figure 4.2. We refer to this region as the *target region*: the set of states the DM aims to maintain in order to achieve optimal profit.

However, when the process $X^{(2)}$ begins to drift to the left of $X^{(1)}$, i.e., when $X^{(1)} > X^{(2)}$, it indicates that the inflow trend is shifting away from the average level $\hat{\theta}$ toward t he lower bound $\overline{\theta}$, which is an absorbing state. This suggests, with high probability, that the reservoir is increasingly likely to trend toward a drought regime. Consequently, it becomes progressively more costly, on average, to maintain a low reservoir level. As a result, it is optimal to exert lower control early (i.e., release less to no water) to mitigate potential future losses, while delaying action on the upper side in expectation of reducing the likelihood of reaching the lower, more costly state.

A symmetric argument applies when the process drifts in the opposite direction toward the upper absorbing state $\underline{\theta}$ (flood), as reflected in the numerical results shown in Figure 4.2. These observations remain robust across varying degrees of ambiguity attitude and are especially pronounced near the average belief state $\hat{\theta}$. We explore this scenario in more detail in the discussion that follows.

4.8.2 COMPARATIVE STATICS OF AMBIGUITY ATTITUDE

According to Remark 17, when ambiguity aversion is finite and non-zero, the reservoir process is subject to drift ambiguity of the (temporary) mean-reverting level. In other words, when the belief process begins to deviate from the maximum benefit state $\tilde{x}$, indicating a drought (resp. flood) regime, i.e., when $X^{(1)} > X^{(2)}$ (resp. $X^{(1)} < X^{(2)}$), the DM expects the reservoir to hit the worst-case state sooner. This is because, by (4.6.10), the drift distortion is directly proportional to the product of ambiguity aversion and the marginal value function.

Thus, when the reservoir level is decreasing (resp. increasing), it is expected to fall (resp. rise)

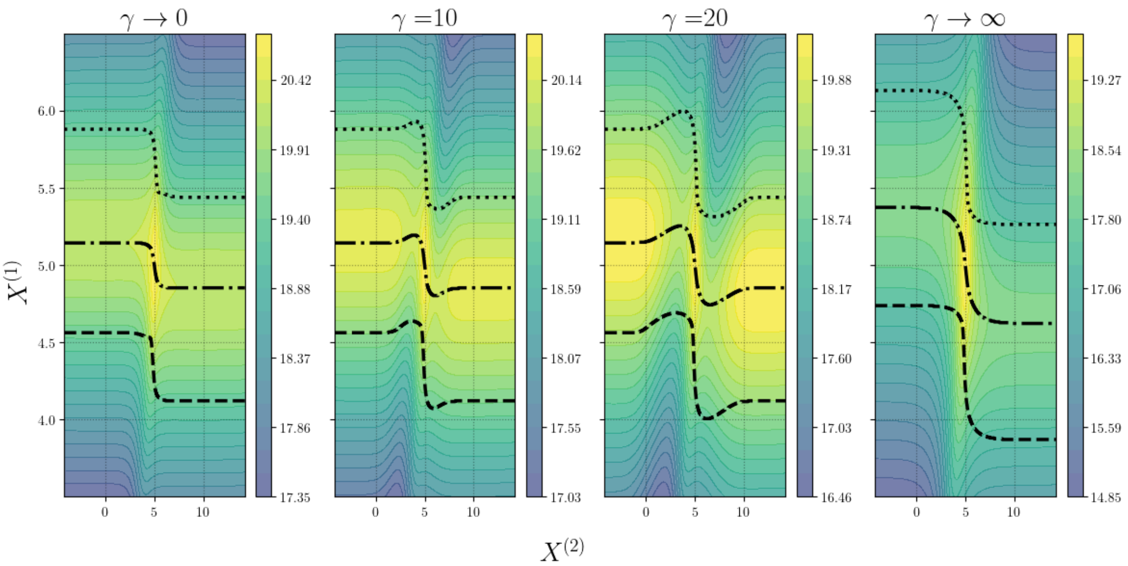

Figure 4.2: The contour plots display the value functions and the corresponding optimal control barriers under the base-case parameter set, for varying levels of ambiguity attitude $\gamma \in \{0+, 10, 20, \infty-\}$. The y-axis determines the reservoir levels, while the x-axis represents the *auxiliary* reservoir levels. The dashed and dotted lines represent the lower and upper control barriers, respectively, while the dash-dotted line indicates the state of maximum benefit for the controlled reservoir level $X^{(1)}$ at each $X^{(2)} \in (\underline{x}_2^\varepsilon, \overline{x}_2^\varepsilon)$. When the state $(X^{(1)}, X^{(2)})$ reaches either the lower or upper barrier, the control processes A^- or A^+ instantaneously reflect the state to $(X^{(1)} - \Delta A^{-,h_1,h_2}, X^{(2)} - \Delta A^{-,h_1,h_2})$ or $(X^{(1)} + \Delta A^{+,h_1,h_2}, X^{(2)} + \Delta A^{+,h_1,h_2})$, respectively, where $\Delta A^{\pm,h_1,h_2} = \frac{h_1 h_2}{h_1 + h_2}$.

faster toward the drought (resp. flood) regime. The optimal response is to exert the lower (resp. upper) control action promptly to prevent such extreme scenarios. Conversely, if the reservoir moves against the implied trend, e.g., $X^{(1)} > X^{(2)}$ but the level is rising, then it is optimal to delay intervention at the corresponding barrier. This is because, once the state enters a neighborhood of such a barrier, the likelihood of reverting to the worst-case outcome decreases, thereby mitigating potential losses.

Taken together, this generates a sudden upward (resp. downward) shift in the control boundary around the maximum benefit states, with the magnitude of the shift increasing proportionally to the degree of ambiguity aversion. Consequently, this results in a narrower target region as γ increases, consistent with the numerical results in the second and third panels of Figure 4.2. Moreover, according to the scales of the value functions in the same figure, higher ambiguity aversion leads to a lower value function. This is because, in expectation, interventions are exercised more frequently in a smaller inaction region, thereby incurring higher accumulated control costs.

However, as the belief process becomes more informative (through sufficient information recovery), the absorbing states become more conclusive, and the influence of ambiguity aversion diminishes. Ultimately, the reservoir process and the associated optimal policy converge to the standard model without ambiguity aversion.

Importantly, when ambiguity aversion tends to an extreme level ($\gamma \to \infty$), which is the limit case of the above phenomenon, the DM adopts maxmin preferences and thus remains perpetually uncertain about the true probability measure, regardless of how much information is revealed. In this case, the DM believes the drift always takes its worst possible form. That is, once the reservoir departs from the maximum benefit states, the perceived long-run mean immediately shifts to the worst regime under maxmin utility: the long-run reservoir level may turn out to be even lower or higher than previously estimated.

In summary, if the DM is ambiguity averse to climate patterns driving reservoir inflows, the optimal strategy differs markedly from both the ambiguity-free and max-min utility benchmarks. When the DM has not yet learned about the long-run trend, ambiguity induces earlier and more decisive interventions, the DM acts as soon as the water level begins to deviate from the reference maximum benefit state. This highlights the interplay between ambiguity and learning: when future conditions are most uncertain, the DM behaves conservatively to guard against unprecedented events. As more information is revealed and the long-run trend becomes clearer, the influence of ambiguity aversion diminishes, and control policies gradually converge to those of the standard model without ambiguity.

In what follows, we study comparative statics analysis of the remaining parameters under the base case model, starting from volatility of the reservoir fluctuation.

4.8.3 COMPARATIVE STATICS OF VOLATILITY

From Figure 4.3, one observes that the control barriers widen as volatility σ increases. This occurs because higher volatility raises the likelihood of extreme fluctuations in reservoir levels, thereby increasing the expected frequency of interventions. To mitigate these costs, it becomes optimal to *wait* by delaying interventions: under high volatility, the reservoir level is also more likely to revert naturally to the maximum benefit states after deviating, which increases the running benefit, on average.

4.8.4 COMPARATIVE STATICS OF DISCOUNT RATE

Figure 4.7 illustrates the comparative statics with respect to increasing discount rates. As shown, the control barriers widen as the discount rate rises, mirroring the effect of volatility. The intuition is that a higher discount rate lowers the present value of future reservoir running costs, which makes delaying intervention relatively more attractive. Consequently, the DM has stronger incentives to *wait*, since postponing control becomes more profitable than acting immediately when the discount rate is low.

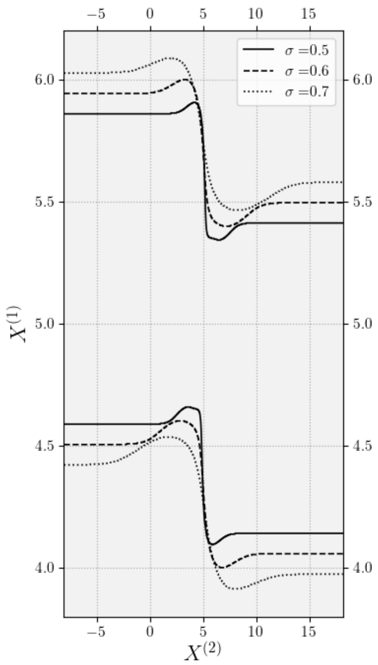

Figure 4.3: Optimal control policies of volatility $\sigma \in \{0.5, 0.6, 0.7\}$.

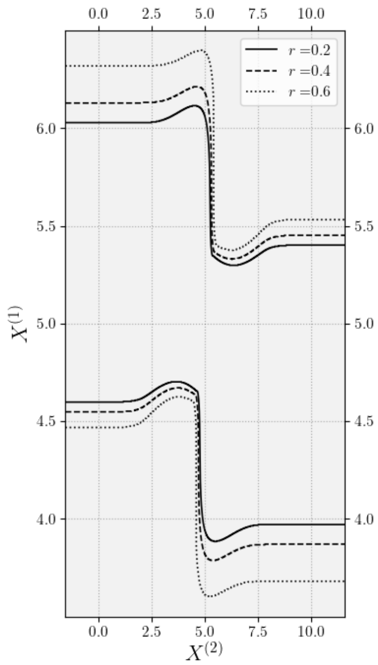

Figure 4.4: Optimal control policies of discounted rates $r \in \{0.2, 0.4, 0.6\}$.

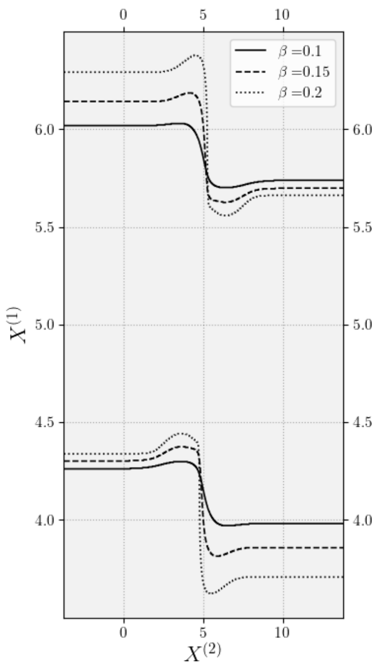

Figure 4.5: Optimal control policies of mean reversion speeds $\beta \in \{0.1, 0.15, 0.2\}$.

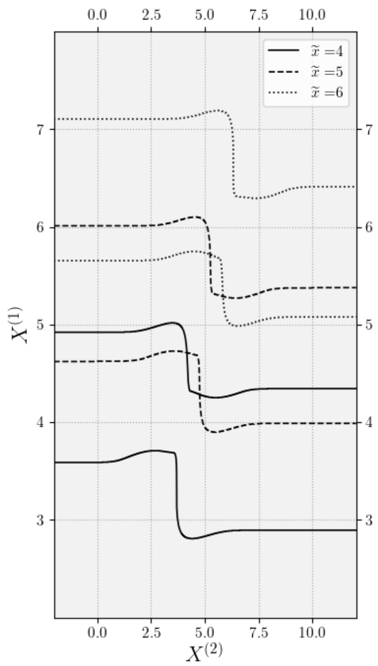

Figure 4.6: Optimal control policies of long run means $\tilde{x} \in \{4, 5, 6\}$.

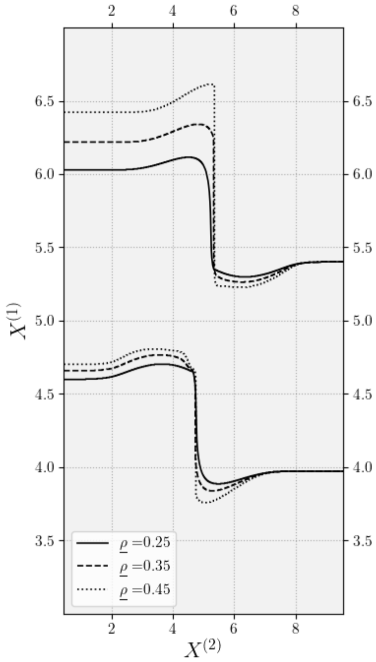

Figure 4.7: Optimal control policies of drought precipitation trends $\underline{\rho} \in \{0.25, 0.35, 0.45\}$.

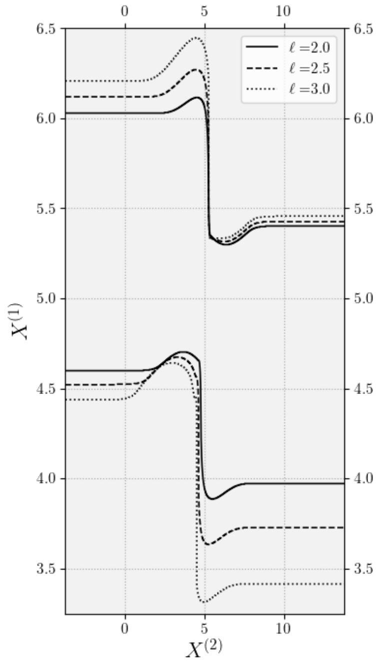

Figure 4.8: Optimal control policies of lower control costs $\ell \in \{2, 2.5, 3\}$.

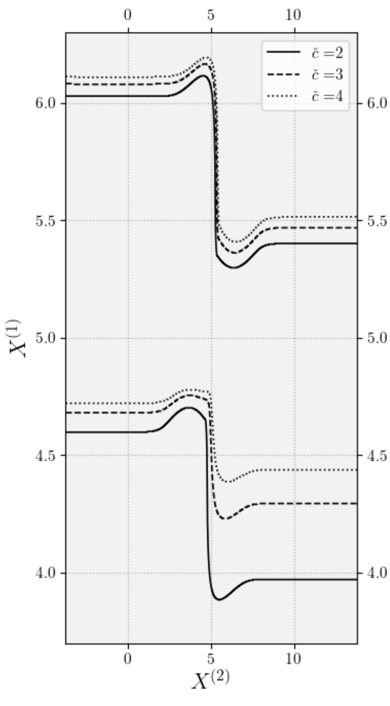

Figure 4.9: Optimal control policies of lower holding costs $\check{c} \in \{0.25, 0.35, 0.45\}$.

4.8.5 COMPARATIVE STATICS OF MEAN REVERSION SPEED

According to (4.6.9), when the reservoir and learning dynamics deviate from the maximum benefit state toward either extreme (a drought or flood long-run trend), a higher mean-reversion speed accelerates the movement toward these worst-case scenarios. Consequently, with faster mean reversion, the DM is expected to incur greater cumulative losses, either because the reservoir level falls more quickly under a perceived drought trend, or rises excessively under a perceived flood trend.

The optimal control policy is therefore to intervene earlier: applying lower control sooner in the drought case, and upper control sooner in the flood case. Conversely, as in other instances of optimal singular control, it remains advantageous to delay the opposite interventions in order to reduce running losses.

Overall, a higher mean-reversion speed shifts the control barriers upward when the learning dynamics suggest a drought regime, and downward when they indicate a flood regime. This feature is illustrated in Figure 4.5.

4.8.6 COMPARATIVE STATICS OF LONG-RUN MEAN

Recall that the maximum benefit states are characterized by $\alpha^{(1)}(X^{(1)}, X^{(2)}) \approx 0$. Thus, if $\tilde{x}$ increases, the definition of $\alpha^{(1)}$ requires that the maximum benefit states adjust so that the condition $\alpha^{(1)}(X^{(1)}, X^{(2)}) \approx 0$ continues to hold. To achieve this, two changes occur: first, the reservoir level $X^{(1)}$ must also increase, ideally remaining close to $\tilde{x}$; second, the state is most stable, yielding the longest expected maximum benefit, when the reservoir trend is perceived to lean slightly toward excess inflows.

This adjustment results in an upward shift of the maximum benefit states, particularly in the north-eastern direction, and consequently, the entire set of optimal control barriers shifts upward in the same direction, as demonstrated by the numerical results in Figure 4.6. Conversely, when $\tilde{x}$ decreases, a symmetric downward shift arises, also illustrated in the figure.

In managerial terms, when the long-run mean reservoir level $\tilde{x}$ increases, the DM should raise their intervention thresholds accordingly. In practice, this means tolerating higher reservoir levels before releasing water, since what previously counted as “excess” now represents the new normal. Conversely, when $\tilde{x}$ falls, thresholds should be lowered to avoid costly shortages.

4.8.7 COMPARATIVE STATICS OF PRECIPITATION TREND

Recall from Assumption 2 that $\underline{\rho}, \overline{\rho} \in [0, 1)$ determine the percentages by which the ambiguous long-term mean reservoir level may deviate from the reference $\tilde{x}$ before reaching the minimum $\underline{x}$ or maximum $\overline{x}$ capacity, respectively. For instance, if the DM *believes* that drought risk has a greater impact on reservoir management, then $\underline{\rho} > \overline{\rho}$. In this case, the numerical results shown in Figure 4.7, indicate that the control barriers shift upward when the drought trend is revealed.

This result is intuitive. A larger $\underline{\rho}$ increases the likelihood that the reservoir will drift toward the drought regime, which in turn raises the expected cost of managing persistent water shortages. The optimal control response is therefore to intervene earlier, for example, by postponing water releases, to preserve resources and prevent further losses. At the same time, delaying opposite actions (such as releasing additional water when the level is above $\tilde{x}$) becomes optimal, since this provides a safeguard against future drought scenarios. From a managerial perspective, this reflects a conservative yet profit-maximizing strategy: in the presence of drought risk, water managers should preserve capacity and defer outflows whenever possible, even at the expense of tolerating temporarily higher levels.

By contrast, when the flood trend is revealed under the same condition ($\underline{\rho} > \overline{\rho}$), the optimal control policy remains essentially unchanged. This is unsurprising, since ambiguity in our framework is represented by a Bernoulli-distributed random variable. Once a particular trend is revealed, the DM can commit to it, as the probability of the alternative trend materializing becomes negligible.

Symmetrically, when flood risk is perceived as more severe ($\overline{\rho} > \underline{\rho}$), the control barriers shift downward: the DM intervenes earlier by releasing water and delays opposite actions to hedge against flooding. As with drought, once the true trend is revealed, ambiguity vanishes and the policy converges to the standard regime-specific strategy.

4.8.8 COMPARATIVE STATICS OF CONTROL COSTS

We now examine how the control barriers adjust when lower control costs increase. Intuitively, higher lower-side costs raise the expected running cost, making it optimal to postpone intervention at the lower barrier. Furthermore, because exercising control on the lower side has become relatively more expensive than on the upper side, the DM also has an incentive to delay action at the upper barrier, thereby reducing the likelihood that the reservoir drifts into the costly lower region. This joint effect results in an outward expansion of the control barriers, as illustrated in Figure 4.8. The case of higher upper control costs is symmetric and can be interpreted analogously.

In other words, when the cost of replenishing or withholding water rises on one side of the system, the DM optimally tolerates wider fluctuations before intervening. In practice, this means reserving costly actions for more extreme states, and adjusting both upper and lower thresholds to reflect the asymmetry in intervention costs.

4.8.9 COMPARATIVE STATICS OF HOLDING COSTS

Figure 4.9 shows the control barrier configurations under increasing costs of holding a deficit water level. As this cost rises, maintaining a low reservoir becomes more expensive. The DM should therefore intervene earlier at the lower barrier to avoid large shortage losses. At the same time, it is optimal to delay action at the upper barrier, since holding excess water is relatively cheaper and also reduces the probability of drifting into the costly deficit region. The net effect is an upward shift of the control barriers. Symmetrically, when the cost of holding excessive water increases, the barriers shift downward.

For a key managerial takeaway, when the losses from maintaining insufficient or excessive water outweigh the other, the DM optimally shifts intervention toward the less costly side. For example, if extreme rainfall is anticipated, managers should lower the reservoir below its normal operating level to create buffer capacity, thereby preventing overflow and avoiding potentially catastrophic losses.

4.9 CONCLUSION

This paper studies optimal dam reservoir management as a singular stochastic control problem, where inflows follow an OU diffusion and ambiguity arises from climate uncertainty. Ambiguity is represented by a Bernoulli-distributed hidden variable distinguishing drought and flood regimes, while decision makers preferences are modeled using smooth ambiguity. In continuous time, this yields a recursive utility linked to a Wonham filtering problem, with the value function shown to satisfy a Hamilton-Jacobi-Bellman variational inequality in the viscosity sense. As closed-form solutions are unavailable, we propose an efficient numerical scheme based on Markov chain approximation and coordinate transformation.

A key contribution of our work is to unify standard learning, smooth ambiguity, and maxmin utility within a single framework. Our numerical results reveal that, unlike maxmin utility, where DMs remain permanently adhered in extreme ambiguity aversion, incurring higher operational costs, smooth ambiguity enables ambiguity attitudes to diminish as information accumulates. This flexibility allows managers to stay cautious when uncertainty is high, but adaptively relax interventions as learning progresses, producing more cost-effective and realistic water management strategies. These findings offer a novel perspective for environmental economics and operations research, showing how climate-induced ambiguity reshapes reservoir policies and long-run sustainability outcomes.

Several extensions are left for future research. The OU assumption, while analytically tractable, omits seasonal variability, electricity demand, evaporation (Figuerola-Ferretti et al., 2024; Jiang et al., 2022; Lucia and Schwartz, 2002), and engineering-specific operational mechanisms (e.g., turbine scheduling or spillway management; see Yoshioka and Yoshioka, 2019). Our framework could be extended to incorporate these features. Similarly, although the Bernoulli specification provides a transparent "drought vs. flood" interpretation, richer settings could allow for a neutral state (no ambiguity) or multiple rainfall regimes with possibility of transitions between them, leading to generalized Wonham filtering problems (cf. Liptser and Shiryaev, 2013a, Chapter 9).

— 5 —

Optimal Inventory Management under Net DemandSupply Uncertainty: A Singular Control Model with Smooth Ambiguity

Abstract

We study singular stochastic control in inventory management under ambiguity, where the decision-maker (DM) exhibits smooth ambiguity preferences over Gaussian-generated priors. We show that continuous-time smooth ambiguity arises as the infinitesimal limit of KalmanBucy filtering with recursive robust utility. The cost function is formulated through forwardbackward stochastic differential equations with quadratic growth and, under a sufficient condition, connected to a Hamilton-Jacobi-Bellman (HJB) variational inequality to determine the value function and optimal policy. Using a coordinate transformation, we reduce the diffusion dimension, reformulating the problem as a two-dimensional singular control, which enables a stable and efficient Markov chain approximation scheme for solving the HJB equation. Numerical experiments with an arithmetic Brownian motion inventory process reveal that when ambiguity is low, the optimal policy aligns with the standard model: higher risk delays action, while greater aversion to ambiguity triggers earlier intervention. In notable contrast, when ambiguity and ambiguity aversion are both high, increasing risk reverses this pattern, prompting earlier rather than later action. This reversal highlights a key managerial implication: in environments marked by extreme volatility and deep uncertainty, an ambiguity-averse DM benefits from acting sooner and adopting a more cautious approach to inventory management.
Keywords: Inventory Model, Singular Control, Smooth Ambiguity, Knightian Uncertainty, Kalman-Bucy Filtering

5.1 INTRODUCTION

Volatile supply chains, unpredictable demand patterns, and structural market shifts can drive inventory systems far from their normal operating conditions. In such environments, the decision-maker

(DM) must control inventory levels without knowing the underlying probability distribution of net demand or supply, an uncertainty that goes beyond stochastic variability and which is commonly referred to as *ambiguity*. We study a singular stochastic inventory control problem in which this ambiguity is modeled as an unobservable process, gradually inferred from observed inventory dynamics. To capture the DM's attitude toward such uncertainty, we employ the smooth ambiguity framework *à la* Klibanoff et al. (2005), embedding ambiguity aversion into a certainty-equivalent transformation of expected utility. Our work extends the seminal singular control formulation of Karatzas (1983) for Brownian storage systems in finite time horizons to accommodate ambiguity, enabling us to examine how the interplay between risk (volatility) and ambiguity attitude shapes optimal replenishment and depletion policies under related inventory cost structures.

Managing inventory becomes a concern when a DM needs to (i) offload excess inventory, which incurs high holding costs, or (ii) restock inventory when levels are too low, potentially leading to penalties for delayed shipments or credit losses. The former action can be executed, e.g., by offering the excess inventory at a promotional price, donating it, or eliminating it, each of which incurs a (proportional) cost. On the other hand, when the inventory is depleted fast due to high demand, the DM might need to, e.g., issue partial refunds to customers, or request orders restock from another factory, both of which incur costs relative to their volume. These in fact create a trade-off which prompts the DM to seek the optimal policy, taking form of the upper and lower triggers (also known as reflecting barriers) of continual inventory demand, to minimize the overall operational costs.

There is a rich literature on inventory control. The first paper that addresses the problem is Arrow et al. (1951), where inventory is modeled as a countable sequence of random variables with known reflecting barriers. This is also referred to in the literature as an *obstacle problem*. Later, Eppen and Fama (1969) developed the model to handle unknown reflecting barriers, allowing the DM to choose policies that minimize a cost criterion that includes holding costs as well as adjustment costs at the reflecting barriers. This is known as *singular control*. The continuous-time analogue of the obstacle problem was subsequently developed by Bather (1966), Constantinides (1976), and Vial (1972), where inventory demand follows a Brownian motion-dirven diffusion. Harrison (1978) was the first to address singular control of inventory in a continuous-time framework, which has since been extended to more generalized approaches by authors such as Bar-Ilan and Sulem (1995), Bensoussan et al. (2005), Dai and Yao (2013a, 2013b), Harrison and Taksar (1983), and Karatzas (1983), among others. For a detailed overview of related research, see Harrison (2013).

While prior literature has successfully applied inventory models in operations research, it is usually assumed that the DM operates under a known probability measure to guide inventory flow decisions. In practice, however, this assumption is frequently not satisfied, particularly when the demand distribution is unknown. This is evident in the newsvendor problem, a classic inventory model, where Ma and Aloysius (2022) show that ambiguity significantly increases decision error, as reflected by higher mean absolute percentage error from the normative benchmark (based on prior information) and lower expected profits. Ambiguity also leads to systematic under-ordering, especially for high-margin products. Crucially, the study reveals an interaction between risk and ambiguity: ambiguity has a stronger negative impact when stochastic variability (i.e., risk) is low, while high risk appears to dampen ambiguity's effect, highlighting the interplay between these two forms of uncertainty.

Further empirical support comes from Kocabıyıkoğlu et al. (2024), who examine how ambiguous demand and learning from observed demand affect inventory decisions. Their study shows that managers often default to normative models, resulting in suboptimal outcomes, capturing at most 84% of potential profit in high-margin settings, and only 51% (or losses) in low-margin ones. They also find that order deviations grow under ambiguity, increasing backorder risks and penalty costs. Consistent with Ma and Aloysius (2022), they observe that while ambiguity generally worsens performance, "demand chasing," updating decisions based on observed demand, can improve accuracy and profitability. Together, these findings highlight the critical impact of demand ambiguity and learning, emphasizing the need for rigorous decision making models in inventory management under ambiguity.

In this paper, we develop a framework for singular control of inventory under learnable multiple priors, where ambiguity attitudes are captured through a second-order expected utility represented by a certainty equivalence operator, that is, a subjective weighted average over a set of priors. This approach follows the smooth ambiguity model Klibanoff et al. (2005, 2009).

The concept of ambiguity was introduced by Knight (1921) as a crucial aspect of uncertainty measurement. Knight differentiates between uncertainties with known probability distributions, which he terms as *risk*, and those with unknown distributions, referred to as ambiguity[1]. The behavior of DMs under ambiguity was first studied by Ellsberg (1961) in what is now known as the Ellsberg urn experiment. This experiment involves making bets on two urns, each containing 100 red or blue balls. In the first urn, the proportion of red to blue balls is known, whereas in the second urn, this proportion is unknown. The results indicated that most people prefer to bet on the first urn, demonstrating an aversion to ambiguity. In other words, people favor a risky bet over an ambiguous one.

Several attempts have been made to incorporate ambiguity into decision support theory. Notably, Gilboa and Schmeidler (1989) introduced the concept of maxmin utility and Klibanoff et al. (2005) developed the smooth ambiguity model. Both models capture ambiguity preferences within the framework of subjective expected utility, albeit from different perspectives. The maxmin utility model shows that a rational DM can display ambiguity aversion by considering the worst-case prior out of a set of priors. Meanwhile, smooth ambiguity allows posterior learning under certainty equivalence to display ambiguity aversion. In terms of Ellsberg's urn experiment, a DM

[1]In some literature, ambiguity is also referred to as *Knightian uncertainty* in recognition of Knight's contributions.

using maxmin utility assumes the urn's composition changes with each draw. Conversely, a DM with smooth ambiguity preference believes that the urn's composition remains fixed and can be learned through a sequence of draws.

In the maxmin approach, ambiguity preferences are modeled through the "size" of the set of priors, offering significant mathematical tractability. This feature makes the framework particularly attractive for a wide range of applications in finance, economics, and operations research, especially in continuous-time settings, as developed by Chen and Epstein (2002). See, for instance, Cheng and Riedel (2013), Hellmann and Thijssen (2018), Nishimura and Ozaki (2007), Thijssen (2011), and Trojanowska and Kort (2010) for examples of such applications.

The first attempt to study (cash) inventory management under ambiguity is made by Archankul et al. (2025), who apply the maxmin utility framework to a singular control problem of firm cash reserves. They show analytically that a more ambiguity-averse DM, through a larger set of priors, tends to exert control earlier, leading to shorter cash hoarding durations. However, the maxmin approach cannot capture managerial learning. A related contribution by Federico et al. (2023) employs sequential detection to update the unknown trend, interpreted as a proxy for ambiguity, in a two-sided singular control problem under arithmetic Brownian motions. However, their model does not take into account the DM's attitude toward ambiguity, indicating the need for a framework that incorporates subjective utility alongside learning.

These limitations motivate the adoption of an alternative framework that accommodates learning under ambiguity. The smooth ambiguity model serves as a natural and flexible solution, as it allows for belief updating while accounting for ambiguity aversion. Unlike the maxmin approach, where the worst-case prior is always selected from the extreme boundary of the set, known as the *upper-rim generator* (cf. Chen and Epstein, 2002), the smooth ambiguity model enables a weighted average across priors, offering a richer and more flexible representation of subjective beliefs.

The concept of smooth ambiguity was initially introduced by Klibanoff et al. (2005) within a static framework. To adapt it for a dynamic setting, the same authors, in Klibanoff et al. (2009), employ a dynamic programming principle, by means of a recursive utility appearing in Epstein and Zin (1989), to evaluate the cost function in each discrete time period, with the likelihood updated one-step-ahead through the Bayesian framework. However, complications arise when considering smooth ambiguity in continuous time, as highlighted by Skiadas (2013). In continuous time, the role of smooth ambiguity diminishes over time. In other words, in the original model proposed by Klibanoff et al. (2009), short-term decision-making neglects ambiguity aversion.

Later, Gindrat and Lefoll (2011) demonstrated that it is possible to retain smooth ambiguity in continuous time by introducing a time increment dependence into a certainty equivalence operator between each one-step-ahead decision. In the same year, Hansen and Sargent (2011) established a connection between the model proposed by Gindrat and Lefoll and proposed a backward stochastic differential equation (BSDE) representation. They accomplished this by assuming that the probability density generator is an unobservable *Gaussian-generated* process and employed the Kalman-Bucy filter (cf. Liptser and Shiryaev, 2013b) to give the best estimate under the realized information. They then applied a logarithmic transformation (cf. Fleming and Soner, 2006) to convert Klibanoff et al.'s utility function into the BSDE.

In our contribution, we build on the idea of Hansen and Sargent (2011) and provide a rigorous foundation by establishing the existence and uniqueness of the associated BSDE. This BSDE has a distinctive form, with a driver exhibiting quadratic growth in the diffusion term. Drawing on the theory for such BSDEs (Kobylański, 2000; Zhang, 2017), we show that, under suitable regularity conditions, it can be integrated into a singular control framework. We then derive a verification theorem in form of a Hamilton-Jacobi-Bellman (HJB) equation, to characterize the optimal policy in the smooth ambiguity setting.

To the best of our knowledge, no analytical solution exists for the corresponding HJB equation. We, therefore, turn to numerical methods for sensitivity analysis. Specifically, we adapt the Markov chain approximation (MCA) approach for singular stochastic control (Kushner & Martins, 1991) to the smooth ambiguity framework. This adaptation is non-trivial: in our setting, the KalmanBucy filter produces two perfectly correlated diffusions, one for the controlled inventory process and one for the posterior estimate, causing instability in the MCA recursion and preventing convergence (Kushner & Dupuis, 2001, see Page 110).

To overcome this, we employ the coordinate transformation technique of Johnson and Peskir (2017) (see also Basei et al., 2024; De Angelis, 2020; Federico et al., 2023), which reduces the diffusion dimension to a single term associated with the controlled inventory process, while transforming the other into a bounded variation process. Under mild assumptions, this yields an MCA scheme that converges to the HJB solution and improves computational efficiency. Although the transformation alters the representation of the inventory dynamics, we demonstrate that the managerial interpretation remains consistent with the pre-transformation model.

The structure of our paper is as follows. In Section 5.2, we formulate the inventory model, transitioning from singular control to the KalmanBucy filter. Section 5.3 introduces an optimization problem incorporated with smooth ambiguity. Section 5.5 is dedicated to the viscosity solution proof of the value function. Section 5.6 discusses the coordinate transformation technique. In Section 5.7, we develop Markov chain approximation scheme before utilizing it to solve the HJB equation and perform comparative statics in Section 5.8. The conclusion of the paper is given by Section 5.9.

5.2 MODEL FORMULATION

Fix $T < \infty$. Let $(\Omega, \mathscr{F}, \mathbf{F} = (\mathscr{F}_t)_{t\in[0,T]}, \mathsf{P})$ be a filtered probability space, satisfying the usual conditions with $\mathscr{F} \supseteq \mathscr{F}_T$. For $t \in [0, T]$, the associated conditional expectation operator under P

is denoted by $\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t}$, where $\mathsf{E}^{\mathsf{P}} \triangleq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_0}$. We assume that the *inventory process* follows an arithmetic Brownian motion, $X \triangleq (X_t)_{t\geq 0}$, i.e., X is the unique strong solution to the stochastic differential equation (SDE),

$$\mathrm{d}X_t = \alpha \mathrm{d}t + \sigma \mathrm{d}B_t, \quad X_0 = x, \quad \mathsf{P}\text{-a.s.}, \tag{5.2.1}$$

for some $\alpha \in \mathbb{R}$ and $\sigma > 0$, where $B = (B_t)_{t\geq 0}$ is a standard P-Brownian motion, adapted to $\mathbf{F}$.

A *control policy* is a pair of bounded variation processes $A \triangleq (A^-, A^+)$, where A^- and A^+ are $\mathbf{F}-$adapted, bounded, non-decreasing, and non-negative. These processes are associated with increases and decreases, respectively, of X at times when control is exerted. We assume that the inventory can hold the minimum backlog $\underline{x} \in (-\infty, 0)$ and the maximum excess $\overline{x} \in (0, \infty)$. Then, a control policy A is said to be *feasible* if for all $x \in \mathbb{R}$,

1) there exists a unique X^A that strongly solves the SDE

$$\mathrm{d}X^A_t = \alpha \mathrm{d}t + \sigma \mathrm{d}B_t + \mathrm{d}A_t \quad X_0 = x \in [\underline{x}, \overline{x}], \quad \mathsf{P}\text{-a.s.}, \tag{5.2.2}$$

where $\mathrm{d}A_t = \mathrm{d}A^-_t - \mathrm{d}A^+_t$, and

2) $\mathsf{P}\left(\sup_{t\in[0,T]} X^A_t < \overline{x}, \inf_{t\in[0,T]} X^A_t > \underline{x}\right) = 1.$

The set of feasible control policies is denoted by $\mathscr{D}$. We call X^A a *controlled inventory process*. Notice that $X^0 = X$ for the uncontrolled process.

Ambiguity arises when a decision maker (DM) is unable to establish a single probability model. This inability might stem from a lack of information about the historical inventory flow evolution, or even if such information exists, the DM cannot agree on a single probability measure. Therefore, the DM contemplates a *set of priors* denoted by $\mathscr{P}$, which is a set of probability measures that are all equivalent to the *reference prior* P. The baseline prior represents the DM's best subjective estimate of the model.

In this paper, we assume that $\mathscr{P}$ is generated by a set of density generators denoted by Λ. That is to say, each $\lambda = (\lambda_t)_{t\geq 0} \in \Lambda$ is assumed to be an $\mathbf{F}$-adapted process such that the process

$$\mathrm{d}\xi_t = -\lambda_t \xi_t \mathrm{d}B_t,\ \xi_0 = 1,$$

is a P-martingale. We refer to λ as the *model*. Then, each $\mathsf{Q} \in \mathscr{P}$ is generated through the Radon-Nikodym derivative $\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \mid \mathscr{F}_T = \xi_T$. In addition, we denote for notional convenience by

$$\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \Big|^{(t,s)}_{\mathscr{F}_t} \triangleq \frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \mid \mathscr{F}_s \Big/ \frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \mid \mathscr{F}_t, \quad \text{for any } s \in [t, T].$$

By the Girsanov theorem (cf. Karatzas and Shreve, 1991, Theorem 5.1 in Chapter 3.5), the process $B^\lambda = \left(B^\lambda_t\right)_{t\geq 0}$ defined by

$$B^\lambda_t \triangleq B_t + \int_0^t \lambda_s \mathrm{d}s$$

is a Q-Brownian motion. Thus, under the measure Q, the controlled inventory process $X^{\lambda,A}$ follows the SDE

$$\mathrm{d}X_t^{\lambda,A} = \alpha \mathrm{d}t + \sigma \mathrm{d}B_t^{\lambda} + \mathrm{d}A_t = (\alpha + \sigma\lambda_t)\,\mathrm{d}t + \sigma \mathrm{d}B_t + \mathrm{d}A_t,\ X_0^{\lambda,A} = x \in [\underline{x}, \overline{x}],\ \mathsf{Q}\text{-a.s.}$$

and admits a unique strong solution. Moreover, the fact that Q and P share the same null set ensures that $A \in \mathscr{D}$ under Q. In other words, A is also adapted under ambiguity. Observe that $X^{0,A} = X^A$ if there is no ambiguity.

For convenience in the comparative statics analysis presented in Section 5.8, from now on, we reparametrize the inventory trend by defining[2] $\theta \triangleq -\lambda$, and accordingly redefine the associated set of density generators as Θ.

In this paper, we assumes that the DM treats the model θ as an unobservable process. Nevertheless, the model can be inferred from a stream of information $\mathbf{F}^X$, generated by an observable process, which, in this context, is the controlled inventory process $X^{\theta,A}$. Here, $\mathbf{F}^X = \left(\mathscr{F}_t^X\right)_{t\geq 0}$ is such that $\mathscr{F}_t^X \subseteq \mathscr{F}_t$ for all $t \in [0,T]$. In other words, θ_t is $\mathscr{F}_t$-measurable but not $\mathscr{F}_t^X$-measurable. According to filtering theory (see Liptser and Shiryaev, 2013b, Chapter 12), θ_t can be estimated by its conditional expectation with respect to $\mathbf{F}^X$. Under the reference measure P, we represent this estimator as a $\mathscr{F}_t^X$-measurable process

$$M_t^{\mathsf{P}} \triangleq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}[\theta_t]_{t\in[0,T]},$$

where the DM evaluates the noise using the mean square error, given by

$$S_t \triangleq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}[(\theta_t - M_t^{\mathsf{P}})^2]_{t\in[0,T]}.$$

As a result, $\mathsf{Q} \in \mathscr{P}$ under $\mathbf{F}^X$ is now generated through the Radon-Nikodym derivative $\frac{\mathrm{d}\mathsf{Q}}{\mathrm{d}\mathsf{P}} \mid \mathscr{F}_T^X = \overline{\xi}_T$, where $\overline{\xi}$ solves

$$\mathrm{d}\overline{\xi}_t = -M_t^{\mathsf{P}}\overline{\xi}_t \mathrm{d}B_t,\ \overline{\xi}_0 = 1,$$

which is a P-martingale.

Throughout this paper we consider the special case where θ is a *constant*, i.e., $\mathrm{d}\theta_t = 0$,, P-a.s., and has predetermined (believed) to be Gaussian distributed with mean $M_0^{\mathsf{P}} = m \in \mathbb{R}$ and variance $S_0 = s > 0$, i.e. $\theta_0 \sim \mathsf{N}(M_0^{\mathsf{P}}, S_0)$. By these assumptions, it is thus implied by Kalman-Bucy filtering (cf. Liptser and Shiryaev, 2013b, Theorem 12.1) that there exists a strong unique solution to the process $(X^{M^{\mathsf{P}},A}, M^{\mathsf{P}}, S) = (X_t^{M^{\mathsf{P}},A}, M_t^{\mathsf{P}}, S_t)_{t\in[0,T]}$ where

$$\begin{aligned} \mathrm{d}X_t^{M^{\mathsf{P}},A} &= (\alpha - \sigma M_t^{\mathsf{P}})\mathrm{d}t + \sigma \mathrm{d}\overline{B}_t + \mathrm{d}A_t, & X_0^{M^{\mathsf{P}},A} &= x, \\ \mathrm{d}M_t^{\mathsf{P}} &= S_t \mathrm{d}\overline{B}_t, & M_0^{\mathsf{P}} &= m, \end{aligned} \tag{5.2.3}$$

[2] As discussed later, the comparative statics are conducted under a coordinate transformation introduced to ensure numerical stability. This transformation reverses the sign of the inventory trend, which obscures direct economic interpretation. The present redefinition avoids this complication.

$$\mathrm{d}S_t = -S_t^2 \mathrm{d}t, \qquad\qquad S_0 = s,$$

P-a.s., and $\overline{B} \triangleq \left(\overline{B}_t\right)_{t\geq 0}$ is the so-called *innovation process*, which solves

$$\overline{B}_t = \frac{1}{\sigma}\int_0^t \left(X_r^{\theta,A} - (\alpha - \sigma M_r^{\mathsf{P}})\right)\mathrm{d}r$$

where $\overline{B}$ is a P–Brownian motion. We call $X^{M^{\mathsf{P}},A}$, M^{P} and S the *filtered controlled inventory process, belief process* and *belief variance*, respectively.

Note that the variance of the belief process is given by $S_t = \frac{s}{1+st}$, $t \in [0,T]$, i.e., $S \in (0,s]$. This expression reveals that the variance of the belief increment decreases monotonically over time and tends to zero as $t \to \infty$. In other words, the longer the learning process continues, the more precise the belief becomes, implying that the hidden information is eventually revealed. Consequently, the terminal belief M_∞^{P} becomes an *absorbing* state of the belief process.

Furthermore, the tail probability of the belief process can be derived explicitly, as the following lemma shows.

Lemma 5. *Let M be the solution of (5.2.3) given $M_0 = m$. Then, for any $t > 0$,*

$$\mathsf{P}\left(|M_t - m| > h\right) = 2\left(1 - \Phi\left(h\sqrt{\frac{1+st}{s^2 t}}\right)\right)$$

where Φ is the standard normal cumulative distribution function.

Proof of Lemma 5. Since the process $M_t - m$ is symmetric around m, we have

$$\mathsf{P}\left(|M_t - m| > h\right) = \mathsf{P}(M_t - m > h) + \mathsf{P}(M_t - m < -h) = 2\mathsf{P}(M_t - m > h).$$

Now, let σ_{M_t} denote the variance of M_t. Using the Itō's isometry gives $\sigma_{M_t} = \int_0^t S_t^2 \mathrm{d}t = \frac{s^2 t}{1+st}$. Consequently, we have $\frac{M_t - m}{\sqrt{\sigma_{M_t}}} \sim \mathsf{N}(0,1)$. Therefore, $\mathsf{P}\left(\frac{M_t-m}{\sqrt{\sigma_{M_t}}} > \frac{h}{\sqrt{\sigma_{M_t}}}\right) = 1 - \Phi\left(h\sqrt{\frac{1+st}{s^2t}}\right)$, completing the proof. ■

Lemma 5 indicates that the likelihood of the belief deviating significantly from its reference point $M_0^{\mathsf{P}} = m$ decreases as h increases. Thus, it becomes increasingly unlikely for such tail events to be the absorbing states of the belief process.

Based on this observation, we take, as a primitive of the model, the DM's *confidence interval*, denoted by $[\underline{m}, \overline{m}] = [m-h, m+h] \subset \mathbb{R}$, which we assume contains the absorbing state with high probability. Accordingly, when the belief process reaches the boundaries $\underline{m}$ or $\overline{m}$, the DM endogenously intervenes to confine the belief within the interval $[\underline{m}, \overline{m}]$.

Formally, this intervention is modeled via a bounded variation process $D \triangleq D^- + D^+$, where D^+ and D^- are $\mathbf{F}^X$-adapted, bounded, non-decreasing, non-negative. They increase only when the belief process M_t^{P} hits the lower boundary $\underline{m}$ or upper boundary $\overline{m}$, respectively, for any

$t \in [0, T]$. In other words, D ensures that

$$\mathsf{P}\left(\sup_{t\in[0,T]} M_t^{\mathsf{P}} \le \overline{m},\ \inf_{t\in[0,T]} M_t^{\mathsf{P}} \ge \underline{m}\right) = 1.$$

As a result, the inventory flow satisfies

$$\begin{aligned} \mathrm{d}X_t^{M^{\mathsf{P}},A} &= (\alpha - \sigma M_t^{\mathsf{P}})\mathrm{d}t + \sigma \mathrm{d}\overline{B}_t + \mathrm{d}A_t, & X_0^{M^{\mathsf{P}},A} &= x \\ \mathrm{d}M_t^{\mathsf{P}} &= S_t \mathrm{d}\overline{B}_t + \mathrm{d}D_t, & M_0^{\mathsf{P}} &= m \\ \mathrm{d}S_t &= -S_t^2 \mathrm{d}t, & S_0 &= s. \end{aligned} \tag{5.2.4}$$

P-a.s., where $\mathrm{d}D_t = \mathrm{d}D_t^- - \mathrm{d}D_t^+$. For notional convenience, we denote $\Omega_K^T \triangleq [\underline{x}, \overline{x}] \times [\underline{m}, \overline{m}]$, the domain of the processes $K_t^{A,\mathsf{P}} \triangleq (X_t^{M^{\mathsf{P}},A}, M_t^{\mathsf{P}})_{t\in[0,T]}$.

5.3 OPTIMIZATION PROBLEM

We now introduce the discounted cumulative cost of inventory holding from time t to T:

$$\begin{aligned} H(t, X_t^{M^{\mathsf{P}},A}) \triangleq \int_t^T e^{-\rho(r-t)} \left(f(X_r^{M^{\mathsf{P}},A})\mathrm{d}r + \ell \mathrm{d}A_r^+ + u \mathrm{d}A_r^- \right) \\ + e^{-\rho(T-t)} \overline{H}(X_T^{M^{\mathsf{P}},A}), \end{aligned} \tag{5.3.1}$$

where $f : \mathbb{R} \to \mathbb{R}^+$ and $\overline{H} : \mathbb{R} \to \mathbb{R}^+$ are Borel-measurable functions representing the running holding cost and the terminal cost, respectively. The constants $\ell > 0$ and $u > 0$ denote the proportional costs associated with positive and negative control actions, and $\rho > 0$ is the DM's discount rate.

To facilitate the analysis, we impose the following assumptions[3]:

1) *Piecewise Linear Holding Cost.* The running cost is piecewise linear in inventory levels:

$$f(x) \triangleq \check{c} \cdot x^- + \hat{c} \cdot x^+, \quad \text{where } x^- \triangleq -\min\{x, 0\},\ x^+ \triangleq \max\{x, 0\}. \tag{5.3.2}$$

Here, $\check{c}, \hat{c} > 0$ denote the instantaneous holding costs for negative and positive inventory levels, respectively.

2) *Intervention Cost Bound.* The intervention costs satisfy $\ell \le \check{c}/\rho$ and $u \le \hat{c}/\rho$. This ensures that a strategy of never intervening is suboptimal, since eventual intervention becomes less costly in the long run.

3) *No Terminal Cost.* The terminal cost is zero: $\overline{H} \equiv 0$. This reflects that T only represents the DM's subjective inventory observation horizon (i.e., longer horizons allow more precise trend estimation in (5.2.4)).

[3]These assumptions can be relaxed to more general settings. See, for example, Ferrari and Vargiolu (2020, Assumption 2.7).

We assume that the DM displays attitudes toward ambiguity through the smooth ambiguity framework introduced by Klibanoff et al. (2005, 2009). As a result, the DM evaluates the cumulative cost criterion using a certainty equivalence operator, which distorts the standard expected value of (5.3.1) based on the information available at each point in time. This is captured by a convex function $\psi_h : \mathbb{R} \to \mathbb{R}$, through which the DM recursively evaluates the cost associated with any admissible policy $A \in \mathscr{D}$ via the recursive functional,

$$H^h(t, X_t^{M^{\mathsf{P}},A}) \triangleq \psi_h^{-1} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X} \left[\psi_h \left(\Pi_{t+h} H(t, X_t^{M^{\mathsf{P}},A}) \right) \right] + H^h(t+h, X_{t+h}^{M^{\mathsf{P}},A}) \quad (5.3.3)$$

where $\Pi_{t+h}H$ is the running profit functional on interval $(t,\, t+h)$ satisfying

$$\Pi_{t+h} H(t, X_t^{M^{\mathsf{P}},A}) \triangleq \int_t^{t+h} e^{-\rho(r-t)} \left(f(X_r^{M^{\mathsf{P}},A}) \mathrm{d}r + \ell \mathrm{d}A_r^+ + u \mathrm{d}A_r^- \right) \quad (5.3.4)$$

for $0 \leq t < t+h \leq T < \infty$. The convexity of the function ψ_h distorts the standard expectation: when ψ_h is decreasing, it biases the evaluation toward lower values; when increasing, it biases toward higher values. To facilitate a more tractable analysis, we adopt a time-increment-dependent exponential disutility specification for φ_h, as proposed by Hansen and Sargent (2011) (see also, Hansen and Miao, 2018, 2022 for a similar treatment). Specifically, we take

$$\psi_h(x) \triangleq \exp\left(\frac{\gamma x}{h e^{-\rho h}} \right), \quad \text{for } x \in \mathbb{R}, h > 0,$$

where the parameter γ governs the DM's attitude toward ambiguity. For a fixed $h > 0$, the function φ_h is decreasing when $\gamma < 0$, reflecting *ambiguity aversion*; increasing when $\gamma > 0$, reflecting *ambiguity seeking*; and approaches linearity as $\gamma \to 0$, corresponding to *ambiguity neutrality*. The presence of h is crucial to maintain the impact of ambiguity in a decision making within an interval $[t, t+h]$. Without it, ambiguity diminishes with time, which means that (5.3.3) almost surely becomes a linear conditional expectation of (5.3.1), as demonstrated by Skiadas (2013). In the remainder of the paper we focus on the practically most relevant case, i.e., where the DM is *ambiguity averse*.

We show in the following section that in the limit $h \downarrow 0$ one can express the cost function (5.3.3) by the robust control problem3:

$$\sup_{\mathsf{Q} \in \mathscr{P}} \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X} \left[H(t, X_t^{M^{\mathsf{P}},A}) + \Psi_{t,T}(\mathsf{P} \parallel \mathsf{Q}) \right], \quad (5.3.5)$$

where $\Psi_{t,T}(\mathsf{P} \parallel \mathsf{Q})$ represents the distance between the probability measures P and Q at time T, given the information known at time t, known in the literature as *relative entropy* or *Kullback-Leibler divergence*.

Hence, within the framework of a singular control with smooth ambiguity, the DM's objective is to determine an optimal policy A that minimizes the cost function (5.3.5). Therefore, at time t given T, the *value function*, the process is defined as:

$$V(t,x,m) \triangleq \inf_{A\in\mathscr{D}} \sup_{\mathsf{Q}\in\mathscr{P}} \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}^X_t}\left[H(t,X_t^{M_t^{\mathsf{P}},A}) + \Psi_{t,T}(\mathsf{P}\parallel\mathsf{Q})\right], \tag{5.3.6}$$

where at $t\in[0,T]$, $(X_t^{M_t^{\mathsf{P}},A}, M_t^{\mathsf{P}}) = (x,m)$. Thus, in short, the value function reads $V(t,k)$ where $k\triangleq(x,m)$.

We provide a verification theorem of the value function in the following section.

5.4 VERIFICATION THEOREM

To analyze this value function (5.3.6), we take the following steps:

1) deriving an explicit form of the relative entropy;
2) demonstrate that solving the problem is equivalent to solving a system of forward-backward stochastic differential equations (FBSDEs);
3) establish sufficient conditions for the optimal conditions of the singular control;
4) establish a verification theorem to provide a unique viscosity solution of a Hamilton-Jacobi-Bellman (HJB) equation which gives a deterministic representation of the system of FBSDEs.

The first step is established in the following lemma.

Lemma 6 (Logarithmic Transformation)**.** *For all $\mathsf{Q}\in\mathscr{P}$ The conditional expectation (5.3.3) satisfies*

$$H^h(t,X_t^{M^{\mathsf{P}},A}) \geq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}^X_t}\left[\Pi_{t+h}H(t,X_t^{M^{\mathsf{P}},A}) + \frac{1}{\gamma}e^{-\rho h}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,t+h)}_{\mathscr{F}^X_t}\right)\right],$$

where $h>0$ is chosen to be small enough that $t+h<t+2h<\ldots\leq t+nh=T$, for some $n\in\mathbb{N}$. Moreover, in the limit $h\to 0$,

$$\begin{aligned}
Y_t^A &\triangleq \lim_{h\downarrow 0} H^h(t,X_t^{M^{\mathsf{P}},A}) \\
&\geq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}^X_t}\left[H(t,X_t^{M^{\mathsf{P}},A}) + \int_t^T e^{-\rho(r-t)}\frac{1}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,r)}_{\mathscr{F}^X_t}\right)\mathrm{d}r\right].
\end{aligned} \tag{5.4.1}$$

Proof of Lemma 6. The lemma is analogous to Lemma 2, so the proof is omitted here. ■

Let $|\cdot|$ be a Euclidean norm, i.e., $|x|=\sqrt{x\cdot x},\ x=(x_1,\ldots,x_n)\in\mathbb{R}^n$. Now, we introduce function spaces that existence and uniqueness of solutions to the FBSDEs can be established. To that end, we let $t\in[0,T]$ and denote by:

- $\mathbb{L}^\infty(\mathscr{F}^X,\mathsf{P})$ the set of the almost surely bounded $\mathscr{F}^X$-measurable random variable ξ, i.e., $|\xi|<\infty$, $\mathsf{P}-$a.s.

- $\mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P})$ the set of the almost surely bounded $\mathbf{F}^X$-adapted process Y, i.e., $|Y_t| < \infty$, P-a.s., for any $0 \leq t \leq T$,
- $\mathbb{H}_T^2(\mathbf{F}^X, \mathsf{P})$ the set of $\mathbf{F}^X$-adapted process Y such that $\mathsf{E}^{\mathsf{P}}\left[\int_0^T |Y_s|^2 \mathrm{d}s\right] < \infty$.

Definition 24. Given that $t \in [0, T]$, $r = (x, \dots, x_d) \in \mathbb{R}^d$ and $(y_i, z_i) \in \mathbb{R}^2$ for $i = 1, 2$. The random variable ξ and $\mathbb{R}$-valued mapping $(t, r, y_i, z_i) \mapsto F(t, r, y_i, z_i)$ are said to be *proper* if

1) $\xi \in \mathbb{L}^\infty(\mathscr{F}^X, \mathsf{P})$.
2) $F(t, r, \cdot, \cdot)$ is a continuous function.
3) There exists $C_1 < \infty$ such that

$$
\begin{aligned}
|F(t, k, y_i, z_i)| &\leq C_1(1 + |y_i| + |z_i|^2) \\
|F(t, k, y_1, z_1) - F(t, k, y_2, z_2)| &\leq C_1\Big(|y_1 - y_2| \\
&\qquad + (1 + |y_1| + |y_2| + |z_1| + |z_2|)|z_1 - z_2|\Big).
\end{aligned}
$$

We call a random variable ξ and a mapping F, *terminal value* and *driver*, respectively.

The existence and uniqueness result for FBSDEs satisfying Assumption 24 are given by the following comparison theorem. For further detail, we refer to Kobylański (2000), Kobylański et al. (2002) or Zhang (2017).

Lemma 7 (Comparison Theorem)**.** *Suppose that the process*

$$
\left(X^{(i)}, A^{(i)}, Y^{(i)}, Z^{(i)}\right) \in \left(\mathbb{H}_T^2(\mathbf{F}^X, \mathsf{P})\right)^d \times (\mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P}))^{d+1} \times \mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P}) \times \left(H_T^2(\mathbf{F}^X, \mathsf{P})\right)^{d+1},
$$

is a solution to

$$
-\mathrm{d}Y_t^{(i)} = F_i\left(t, X_t^{(i)}, Y_t^{(i)}, Z_t^{(i)}\right)\mathrm{d}t + \mathrm{d}A_t^{(i)} - Z_t^{(i)}.\mathrm{d}W_t, \quad Y_T^{(i)} = \xi^{(i)}
$$

where W is $(d+1)$-Brownian motions under P, $X^{(i)}$ solves (5.2.1), $A^{(i)} \in \mathscr{D}^{d+1}$, $\xi^{(i)}$ and F_i are proper, for $i = 1, 2$. Suppose, furthermore, that:

1) $A^{(1)} - A^{(2)} \geq \mathbf{0}$ *a.s.,*
2) $\xi^{(1)} - \xi^{(2)} \geq 0$ *a.s. and*
3) $F_1 - F_2 \geq 0$ *a.e.,*

where denoted by $\mathbf{0}$ is the $(d+1)$-dimensional null-vector. Then $Y_t^{(1)} - Y_t^{(2)} \geq 0$ P-a.s. for $t \in [0, T]$.

Now we are ready to establish the FBSDE representation for the value function (5.3.6).

Theorem 17. *There exists a unique solution to*

$$
(K^{A,\mathsf{P}}, A, Y^A, Z) \in \left(\mathbb{H}_T^2(\mathbf{F}^X, \mathsf{P})\right)^2 \times \mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P}) \times \mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P}) \times \mathbb{H}_T^2(\mathbf{F}^X, \mathsf{P}),
$$

where $K^{A,\mathsf{P}} \triangleq (X^{M^{\mathsf{P}},A}, M^{\mathsf{P}})$ *solve* (5.2.4) *through* (5.2.3) *and such that*

$$\mathrm{d}Y_t^A = -F(t, K_t^{A,\mathsf{P}}, Y_t^A, Z_t)\mathrm{d}t - \ell \mathrm{d}A_t^+ - u\mathrm{d}A_t^- + Z_t \mathrm{d}\overline{B}_t, \quad Y_T^A = 0 \tag{5.4.2}$$

where

$$F(t,k,y,z) \triangleq -\rho y + f(x) + \frac{\gamma S_t}{2} z^2, \quad k = (x,m).$$

Moreover, Y^A *satisfies the inequality* (5.4.1)*, where equality holds for which*

$$\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}^*} \,|_{\mathscr{F}_t^X}^{(r,t+h)} = \frac{\exp\left(\gamma(\theta_r - M_r^{\mathsf{P}})Z_r\right)}{\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_r^X}\left[\exp\left(\gamma(\theta_r - M_r^{\mathsf{P}})Z_r\right)\right]},$$

for any $r \in [t, t+h)$.

Proof of Theorem 17. Observe that

$$\begin{aligned}
|F(t,k,y,z)| &\le C_f \left(|y| + \frac{\gamma|s|}{2} z^2\right) \\
|F(t,k,y_1,z_2) - F(t,k,y_2,z_2)| &\le C_f \left(|y_1 - y_2| + \frac{\gamma|s|}{2}(z_1^2 - z_2^2)\right) \\
&= C_f \left(|y_1 - y_2| + \frac{\gamma|s|}{2}(z_1 + z_2)(z_1 - z_2)\right)
\end{aligned} \tag{5.4.3}$$

where C_f is the Lipschitz constant of f. Therefore, if we choose $C_1 = C_f + \frac{\gamma|s|}{2}$, then F is proper. Together with the boundedness of Y_T^A, we conclude by Lemma 7 that there is a unique solution to $(X^{M^{\mathsf{P}},A}, M^{\mathsf{P}}, A, Y^A, Z)$ in $\left(\mathbb{H}_T^2(\mathbf{F}^X, \mathsf{P})\right)^2 \times \mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P}) \times \mathbb{L}_T^\infty(\mathbf{F}^X, \mathsf{P}) \times \mathbb{H}_T^2(\mathbf{F}^X, \mathsf{P})$

Next, we show that Y_t^A satisfies the inequality (5.4.1). Note that (5.4.2) implies

$$Y_t^A = \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^{t+h} \left[\left(-\rho J_r^A + f(X^{M^{\mathsf{P}},A}) + \frac{\gamma S_r}{2} Z_r^2\right)\mathrm{d}r + \ell \mathrm{d}A_r^+ + u\mathrm{d}A_r^-\right] + Y_{t+h}^A\right],$$

for some $h > 0$ such that $t + h \le T$. Now we fix $\mathsf{Q} \in \mathscr{P}$. Consequently, for all $r \in [t, t+h)$, one can infer that

$$\begin{aligned}
\frac{\gamma S_r}{2} Z_r^2 &= \frac{1}{\gamma} \log\left(\exp\left(\gamma \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_r^X}[\theta_r - M_r^{\mathsf{P}}]Z_r + \frac{\gamma^2}{2}\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_r^X}[(\theta_r - M_r^{\mathsf{P}})^2]Z_r^2\right)\right) \\
&= \frac{1}{\gamma} \log \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_r^X}\left[\exp\left(\gamma(\theta_r - M_r^{\mathsf{P}})Z_r\right)\right] \\
&\ge \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_r^X}\left[(\theta_r - M_r^{\mathsf{P}})Z_r + \frac{1}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}} \,|_{\mathscr{F}_r^X}^{(r,t+h)}\right)\right] \\
&= -(M_r^{\mathsf{P}} - M_r^{\mathsf{Q}})Z_r + \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_r^X}\left[\frac{1}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}} \,|_{\mathscr{F}_r^X}^{(r,t+h)}\right)\right].
\end{aligned} \tag{5.4.4}$$

The second equality follows from the moment generating function of $\theta_s - M_s^{\mathsf{P}}$, while the definition of logarithmic transformation (similar to Lemma 6) gives rise to inequality in the third line. Thus, we have

$$Y_t^A \geq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\Bigg[\int_t^{t+h}\Bigg[\Bigg(-\rho J_r^A + f(X^{M^{\mathsf{P}},A}) - (M_r^{\mathsf{P}} - M_r^{\mathsf{Q}})Z_r \\ + \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_r^X}\left[\frac{1}{\gamma}\left(\log\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(r,t+h)}_{\mathscr{F}_r^X}\right)\right]\Bigg)\mathrm{d}r + \ell\mathrm{d}A_r^+ + u\mathrm{d}A_r^-\Bigg] + Y_{t+h}^A\Bigg].$$

The change of measure from P to Q gives

$$\overline{B}_t = \widehat{B}_t - \int_0^t (M_r^{\mathsf{Q}} - M_r^{\mathsf{P}})\mathrm{d}r, \quad t \in [0, T],$$

where $\widehat{B}$ is a Q–innovation process. As a result, we obtain that

$$Y_t^A \geq \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Bigg[\int_t^{t+h}\Bigg[\left(-\rho J_r^A + f(X^{M^{\mathsf{P}},A}) + \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_r^X}\left[\frac{1}{\gamma}\left(\log\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(r,t+h)}_{\mathscr{F}_r^X}\right)\right]\right)\mathrm{d}r \\ + \ell\mathrm{d}A_r^+ + u\mathrm{d}A_r^-\Bigg] + Y_{t+h}^A\Bigg].$$

Since M^{P} is càdlàg, taking $h \downarrow 0$ gives that $\mathscr{F}_r^X \downarrow \mathscr{F}_t^X$. Combining with the arbitrariness of $\mathsf{Q} \in \mathscr{P}$ and Lemma 6, we have

$$Y_t^A \geq \sup_{\mathsf{Q}\in\mathscr{P}} \mathsf{E}^{\mathsf{Q}}_{\mathscr{F}_t^X}\Bigg[\int_t^{T}\Bigg[\left(-\rho J_r^A + f(X^{M^{\mathsf{P}},A}) + \frac{1}{\gamma}\log\left(\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}}\,|^{(t,r)}_{\mathscr{F}_t^X}\right)\right)\mathrm{d}r \\ + \ell\mathrm{d}A_r^+ + u\mathrm{d}A_r^-\Bigg] + \xi(X_T^{M^{\mathsf{P}},A})\Bigg], \tag{5.4.5}$$

By Itō's lemma, we obtain (5.4.1).

Finally, we choose $\mathsf{Q}^* \sim \mathsf{P}$ such that the corresponding Radon-Nikodym derivative takes the form

$$\frac{\mathrm{d}\mathsf{P}}{\mathrm{d}\mathsf{Q}^*}\,|^{(r,t+h)}_{\mathscr{F}_r^X} = \frac{\exp\left(\gamma(\theta_r - M_r^{\mathsf{P}})Z_r\right)}{\mathsf{E}^{\mathsf{P}}_{\mathscr{F}_r^X}\left[\exp\left(\gamma(\theta_r - M_r^{\mathsf{P}})Z_r\right)\right]}. \tag{5.4.6}$$

By plugging (5.4.6) into (5.4.4), it is easy to see that equation (5.4.4) achieves the equality. This infers that the equality holds for (5.4.5) and Q^* gives rise to the supremum of the robust control problem (in limit), which is the completion of Theorem 17. ■

Following Theorem 17, the value function is

$$V(t,k) \triangleq \inf_{A\in\mathscr{D}} Y_t^{A;k} \tag{5.4.7}$$

where $Y_t^{A;k} = Y_t^A\,|_{K_t^{A,\mathsf{P}}=k}$. The following results show the convexity and the locally Lipschitz continuity of the value function in x.

Proposition 14. *$V(t, x, \cdot, \cdot)$ is convex in $x \in [\underline{x}, \overline{x}]$.*

Proof of Proposition 14. We suppose by contradiction that for any $t, m \in [0, T] \times [\underline{m}, \overline{m}]$, there exist $x_0, x_1 \in [\underline{x}, \overline{x}]$ and $\alpha \in (0, 1)$ such that

$$\alpha V(t, x_0, m) + (1-\alpha)V(t, x_1, m) \leq V(t, \overbrace{\alpha x_0 + (1-\alpha)x_1}^{\triangleq x_\alpha}, m), \quad t \in [0, T]. \tag{5.4.8}$$

For simplicity, we denote by $X_r^{A;x,m}$, $t \leq r$, the controlled process given by $X_t^{A;m} = x$; and $K_r^{A,\mathsf{P};x} \triangleq (r, X_r^{A;x,m}, M_r^{\mathsf{P}})$. Then, from (5.4.2), we have

$$\begin{aligned}\alpha Y_t^{A;x_0,m} + (1-\alpha)Y_t^{A;x_1,m} &= \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^T \alpha F(r, K_r^{A,\mathsf{P};x_0}, Y_r^{A;x_0,m}, Z_r)\mathrm{d}r\right] \\ &+ \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^T (1-\alpha)F(r, K_r^{A,\mathsf{P};x_1}, Y_r^{A;x_1,m}, Z_r)\mathrm{d}r + \int_t^T (\ell \mathrm{d}A_r^+ + u\mathrm{d}A_r^-)\right].\end{aligned}$$

Since F is strictly convex in x, the linearity of $X_t^{A;x,t,m}$ in x gives

$$\begin{aligned}&\alpha Y_t^{A;x_0,m} + (1-\alpha)Y_t^{A;x_1,m} \\ &\quad > \mathsf{E}^{\mathsf{P}}\left[\int_t^T F(r, K_r^{A,\mathsf{P};x_\alpha}, \alpha Y_r^{A;x_0,m} + (1-\alpha)Y_r^{A;x_1,m}, Z_r)\mathrm{d}r + \int_t^T (\ell \mathrm{d}A_r^+ + u\mathrm{d}A_r^-)\right]\end{aligned}$$

Since this holds for any feasible control policy, it holds that for any $\varepsilon > 0$ there exists $\overline{A} \in \mathscr{D}$ such that

$$\begin{aligned}&\alpha V(t, x_0, m) + (1-\alpha)V(t, x_1, m) \\ &> \inf_{A \in \mathscr{D}} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^T F(r, K_r^{A,\mathsf{P};x_\alpha}, \alpha Y_r^{A;x_0,m} + (1-\alpha)Y_r^{A;x_1,m}, Z_r)\mathrm{d}r + \int_t^T (\ell \mathrm{d}A_r^+ + u\mathrm{d}A_r^-)\right] \\ &\geq \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^T F(r, K_r^{A,\mathsf{P};x_\alpha}, \alpha Y_r^{\overline{A};x_0,m} + (1-\alpha)Y_r^{\overline{A};x_1,m}, Z_r)\mathrm{d}r + \int_t^T (\ell \mathrm{d}\overline{A}_r^+ + u\mathrm{d}\overline{A}_r^-)\right] - \varepsilon.\end{aligned}$$

Since the second inequality holds for arbitrary $\varepsilon > 0$, we have

$$\begin{aligned}&\alpha V(t, x_0, m) + (1-\alpha)V(t, x_1, m) \\ &> \inf_{A \in \mathscr{D}} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X}\left[\int_t^T F(r, K_r^{A,\mathsf{P};x_\alpha}, \alpha V(r, x_0, m) + (1-\alpha)V(r, x_1, m), Z_r)\mathrm{d}r\right. \\ &\qquad\left. + \int_t^T (\ell \mathrm{d}A_r^+ + u\mathrm{d}A_r^-)\right].\end{aligned} \tag{5.4.9}$$

Recalling (5.4.8) and that F is strictly decreasing in y, we conclude from (5.4.9) that $\alpha V(t, x_0, m) + (1-\alpha)V(t, x_1, m) > V(t, x_\alpha, m)$, which is a contradiction. Hence, $V(\cdot, x, \cdot, \cdot)$ is convex in x. This completes the proof. ■

Proposition 15. *For any $t \in [0,T]$, $V(t,x,m)$ is Lipschitz continuous in x and m with Lipschitz constant $\max\{\check{c},\hat{c}\}/\rho$ and $\sigma\max\{\check{c},\hat{c}\}/\rho$, respectively. Moreover, if Assumption 2) holds, then there exists $(a,b) \subseteq [\underline{x},\overline{x}]$ such that*

$$-\ell \leq \frac{V(t,x_0,m) - V(t,x_1,m)}{x_0 - x_1} \leq u$$

for any $x_0, x_1, \in (a,b)$ such that $x_0 > x_1$.

Proof of Proposition 15. Let $x_0, x_1 \in \mathbb{R}$ with $x_0 > x_1$. Observe that

$$-\check{c} \leq \frac{f(x_0) - f(x_1)}{x_0 - x_1} \leq \hat{c},$$

which follows directly from the piecewise linear structure of the holding cost function f defined in (5.3.2).

Then, for any admissible control $A \in \mathscr{D}$, it follows from (5.4.2) that

$$\begin{aligned} Y_t^{A;x_0,m} - Y_t^{A;x_1,m} &= \mathsf{E}^{\mathsf{P}}\left[\int_t^T e^{-\rho(r-t)}\left(f(X_r^{A;x_0,m}) - f(X_r^{A;x_1,m})\right)\mathrm{d}r\right] \\ &\leq \hat{c}\,\mathsf{E}^{\mathsf{P}}\left[\int_t^T e^{-\rho(r-t)}\left(X_r^{A;x_0,m} - X_r^{A;x_1,m}\right)\mathrm{d}r\right] \leq \frac{\hat{c}}{\rho}(x_0 - x_1). \end{aligned} \quad (5.4.10)$$

Similarly, one obtains the lower bound:

$$Y_t^{A;x_0,m} - Y_t^{A;x_1,m} \geq -\frac{\check{c}}{\rho}(x_0 - x_1). \quad (5.4.11)$$

Since inequalities (5.4.10) and (5.4.11) hold uniformly for all $A \in \mathscr{D}$, we conclude that

$$-\frac{\check{c}}{\rho} \leq \frac{V(t,x_0,m) - V(t,x_1,m)}{x_0 - x_1} \leq \frac{\hat{c}}{\rho}.$$

Hence, the value function $V(t,x,m)$ is Lipschitz continuous in x with Lipschitz constant at least $\max\{\check{c},\hat{c}\}/\rho$. A similar argument establishes Lipschitz continuity in the belief mean m.

Furthermore, since $V(\cdot,x,\cdot)$ is strictly convex in x (by Proposition 14), the finite difference quotient $(V(t,x_0,m) - V(t,x_1,m))/(x_0 - x_1)$ is increasing in x_0. Combined with Assumption 2), this implies that there exists an interval $(a,b) \subseteq [\underline{x},\overline{x}]$ such that

$$-\ell \leq \frac{V(t,x_0,m) - V(t,x_1,m)}{x_0 - x_1} \leq u.$$

This completes the proof. ■

Under certain assumptions, as demonstrated in the following verification theorem, we can establish the existence of a "smooth" solution for the value function together with the conditions under which a unique optimal control policy is obtained. The proof is standard, as can be seen in Theorem 18 or in seminal works of Harrison and Taksar (1983) or Fleming and Soner (2006, chapter 8), so omitted for brevity.

Theorem 18 (Verification theorem)**.** *Suppose that* $V \in [0,T] \times C^{1,2}([0,T] \times \Omega_K^T)$ *such that*

$$V(t,k) = \sup_{A \in \mathscr{D}} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}_t^X} \left[\int_t^T \left(F(r, K_r^{A,\mathsf{P}}, V(r, K_r^{A,\mathsf{P}}), \overline{\nabla} V(r, K_r^{A,\mathsf{P}}))\mathrm{d}r + \ell \mathrm{d}A_t^+ + u \mathrm{d}A_t^- \right) \right].$$

where

$$\overline{\nabla} V(t,k) \triangleq \sigma \frac{\partial V}{\partial x}(t,k) + S_t \frac{\partial V}{\partial m}(t,k),$$

If for any $\phi \in C^{1,2}([0,T] \times \Omega_K^T)$ *satisfying*

$$\mathscr{L}\phi(t,k) + F(k, \phi(t,k), \overline{\nabla}\phi(t,k)) \geq 0 \ \text{ and} - \ell \leq \frac{\partial \phi}{\partial x}(t,k) \leq u,$$

where

$$\mathscr{L}\phi(t,k) \triangleq \frac{\partial V}{\partial t}(t,k) + \alpha \frac{\partial \phi}{\partial x}(t,k) + \frac{\sigma^2}{2}\frac{\partial^2 \phi}{\partial x^2}(t,k) + \frac{S_t^2}{2}\frac{\partial^2 \phi}{\partial m^2}(t,k) + \sigma S_t \frac{\partial^2 \phi}{\partial x \partial m}(t,k).$$

Then $V(t,k) \geq \phi(t,k)$ *for any* $(t,k) \in [0,T] \times \Omega_K^T$ *and* $A \in \mathscr{D}$*.*

Moreover, for each $(t,m) \in [0,T] \times [\underline{m}, \overline{m}]$*, if there exists an open free-boundary interval* $\mathcal{X}(t,m) \subseteq [\underline{x}, \overline{x}]$ *such that*

$$\mathscr{L}\phi(t,k) + F(k, \phi(t,k), \overline{\nabla}\phi(t,k)) = 0 \qquad \text{on } \mathcal{X}(t,m),$$

$$\frac{\partial \phi}{\partial x}(t, a(t,m), m) = -\ell, \tag{5.4.12}$$

$$\frac{\partial \phi}{\partial x}(t, b(t,m), m) = u, \tag{5.4.13}$$

$$\frac{\partial^2 \phi}{\partial x^2}(t, a(t,m), m) = \frac{\partial^2 \phi}{\partial x^2}(t, b(t,m), m) = 0, \tag{5.4.14}$$

where

$$a(t,m) = \inf\{x : x \in \mathcal{X}(t,m)\},$$
$$b(t,m) = \sup\{x : x \in \mathcal{X}(t,m)\},$$

and $A^* \in \mathscr{D}$ *such that* A^{*-} *and* A^{*-} *only increase at* $a(t,m)+$ *and* $b(t,m)-$*, respectively,* P*-a.s., then* $V(k) = \phi(t,k)$*.*

Remark 19. *The set* $\mathcal{X}(t,m)$ *is usually called the* continuation region*, while* (5.4.12) *and* (5.4.13) *represent the* value-matching conditions*, and* (5.4.14) *imposes the* smooth-pasting condition.

Theorem 18 implies that if V is a classical solution to the HJB equation

$$0 = \min\left\{ \mathscr{L}V(t,k) + F(t,k,V(t,k), \overline{\nabla}V(t,k)),\ \ell + \frac{\partial V}{\partial x}(t,k),\ u - \frac{\partial V}{\partial x}(t,k) \right\} \tag{5.4.15}$$

then, since V is convex in x (cf. Proposition 14), there is a uniquely determined continuation region $\mathcal{X}(t,m)$ for any $(t,m) \in [0,T] \times [\underline{m}, \overline{m}]$. That is,

$$a(t,m) \triangleq \sup\left\{x \in [\underline{x}, \overline{x}] : \frac{\partial V}{\partial x}(t,k) \leq -\ell\right\}, \quad \text{and}$$

$$b(t,m) \triangleq \inf\left\{x \in [\underline{x}, \overline{x}] : \frac{\partial V}{\partial x}(t,k) \geq u\right\}.$$

Moreover, the first derivatives have at most polynomial growth in (t,k), thanks to Proposition 15. Therefore, $\overline{\nabla} V \in \mathbb{H}^2_T(\mathbf{F}^X, \mathsf{P})$ and so the BSDE (5.4.2) with $Z = \overline{\nabla} V$ admits a unique solution.

Remark 20. *If a classical solution to* (5.4.15) *exists, then under our smooth ambiguity framework the (optimal) controlled inventory process satisfies*

$$\mathrm{d}X_t^{M^{\mathsf{P}},A^*} = \left(\alpha - \sigma\left(M_t^{\mathsf{P}} - \frac{\gamma}{2} S_t \overline{\nabla} V\left(t, X_t^{M^{\mathsf{P}},A^*}, M_t^{\mathsf{P}}\right)\right)\right)\mathrm{d}t + \sigma \mathrm{d}\overline{B}_t + \mathrm{d}A_t^*, \quad (5.4.16)$$

$$\mathrm{d}M_t^{\mathsf{P}} = \frac{\gamma}{2} S_t^2 \overline{\nabla} V\left(t, X_t^{M^{\mathsf{P}},A^*}, M_t^{\mathsf{P}}\right)\mathrm{d}t + S_t \mathrm{d}\overline{B}_t + \mathrm{d}D_t. \quad (5.4.17)$$

By Propositions 14 and 15, the value function V is locally Lipschitz continuous in (x,m) for all $t \in [0,T]$. Consequently, the SDE system (5.4.16)(5.4.17) *admits a unique strong solution (cf. Karatzas and Shreve, 1991, Theorem 2.5).*

This formulation highlights the presence of drift ambiguity, *a common feature in models where ambiguity aversion (or seeking) is incorporated into decision-making under uncertainty driven by continuous-time stochastic processes.*[4]

5.5 VISCOSITY SOLUTION

Given the previous argument, we cannot immediately say that the FBSDE (5.4.7) is the solution the fully nonlinear PDE (5.4.15) (at least in a classical sense) since Y^A is only known to be continuous, but not necessary smooth, opposing to the assumption stated in Theorem 18. To deal with is issue, we employ the notion of viscosity solution introduced by Crandall and Lions (1983), in which they generalize the solution of fully nonlinear PDEs to equip in a weaker sense, where the only locally boundedness is assumed. The development of a viscosity solution to the second order PDEs thanks to Crandall et al. (1987), and our following argument will rely on this paper. In other words, we show that the FBSDE (5.4.7) is a viscosity solution to (5.4.15). Before doing so, let us first state the definition of viscosity solution.

Definition 25 (Viscosity Solution)**.** For all $v \in C([0,T] \times \Omega_K^T)$ and $\varphi \in C^{1,2}([0,T] \times \Omega_K^T)$, define the operator $\mathbb{L}$, for any $(t,k) \in [0,T] \times \Omega_K^T$, by

$$\mathbb{L}\varphi(t,k) \triangleq \mathscr{L}\varphi(t,k) + F\left(t, k, \varphi(t,k), \sigma \frac{\partial \varphi(t,k)}{\partial x} + S_t \frac{\partial \varphi(t,k)}{\partial m}\right).$$

[4] See, for example, Archankul et al. (2025), Cheng and Riedel (2013), Hellmann and Thijssen (2018), Nishimura and Ozaki (2007), and Thijssen (2011).

1) The function v is a *viscosity subsolution* of (5.4.15) if for all local maximum points $(t,k) \in [0,T]\times\Omega_K^T$ of $v-\varphi$, i.e., $(v-\varphi)(t,k)=0$ and $(v-\varphi)(t,k) \geq (v-\varphi)(t',k'),$ for any $(t',k') \neq (t,k)$, it holds that

$$\min\left\{\mathbb{L}\varphi(t',k'),\ \ \ell + \frac{\partial\varphi(t',k')}{\partial x},\ \ u - \frac{\partial\varphi(t',k')}{\partial x}\right\} \geq 0, \qquad \varphi(T,k')=0.$$

2) The function v is a *viscosity supersolution* of (5.4.15) if for all local minimum points $(t,k) \in [0,T]\times\Omega_K^T$ of $v-\varphi$, i.e., $(v-\varphi)(t,k)=0$ and $(v-\varphi)(t,k) \leq (v-\varphi)(t',k'),$ for any $(t',k') \neq (t,k)$, it holds that

$$\min\left\{\mathbb{L}\varphi(t',k'),\ \ \ell + \frac{\partial\varphi(t',k')}{\partial x},\ \ u - \frac{\partial\varphi(t',k')}{\partial x}\right\} \leq 0, \qquad \varphi(T,k')=0,$$

3) $v(t,k) \in C([0,T]\times\Omega_K^T)$ is a *viscosity solution* of (5.4.15) if it both viscosity subsolution and supersolution.

Now, we show that $V(t,k)$ is a viscosity solution of (5.4.15).

Proposition 16. *Suppose that Assumption 2) holds. Then $V(t,k)$ is a viscosity solution to the HJB equation* (5.4.15).

Proof of Proposition 16. Given $\varphi \in C^{1,2}([0,T]\times\Omega_K^T)$. We first show that φ is a viscosity supersolution to (5.4.15). The idea of proof is adapted from that of the nonlinear Feynman-Kac formula given by Zhang (2017, Chapter 5).
Subsolution. Let (t,k) be a local maximum point of $V-\varphi$. Suppose, in contrary, that

$$\mathscr{L}\varphi(t,k) + F(t,k,\varphi(t,k),\overline{\nabla}\varphi(t,k)) < 0, \tag{5.5.1}$$

$$\ell + \frac{\partial\varphi}{\partial x}(t,k) < 0, \quad \text{or} \tag{5.5.2}$$

$$u - \frac{\partial\varphi}{\partial x}(t,k) < 0 \tag{5.5.3}$$

According to the proof of Proposition 15, together with Assumption 2), we obtain $-\ell h \leq V(t,x+h,m) - V(t,x,m) \leq uh$, for any $h>0$ such that $\underline{x} \leq x < x+h \leq \overline{x}$. Since (t,k) is a local maximum point of $V-\varphi$, it follows that $-\ell h \leq \varphi(t,x+h,m)-\varphi(t,k)$. The arbitrariness of $h>0$ then contradicts (5.5.2). We can also conclude that (5.5.3) leads to a similar contradiction. Thus, we must have $-\ell \leq \frac{\partial\varphi}{\partial x}(t,k) \leq u$.

Now, it is left to show that (5.5.1) leads to a contradiction. Let us consider

$$\check{Y}_t^A = \varphi(t, K_t^{A,\mathsf{P}}) \text{ and } \check{Z}_t = \overline{\nabla}\varphi(t, K_t^{A,\mathsf{P}})$$

for some $\varphi \in C^{1,2}([0,T]\times\Omega_K^T)$. Since F is continuous, the following stopping times are finite P-a.s.,

$$\overline{\tau}_t \triangleq \inf\{r > t : \mathscr{L}\varphi(r, K_r^{A,\mathsf{P}}) + F(r, K_r^{A,\mathsf{P}}, \check{Y}_r^A, \check{Z}_r)) = 0\}$$

Denote by $\tau_{t,h} \triangleq \overline{\tau}_t \wedge (t+h)$ for some $h>0$, and by $\tau_{t,h} \triangleq \tau_{0,h}$, where $a\wedge b = \min\{a,b\}$. Then it follows from Itō's lemma that

$$\check{Y}_{\tau_{t,h}}^A = \check{Y}_t^A + \int_t^{\tau_{t,h}} \mathscr{L}\varphi(r, K_r^{A,\mathsf{P}})\mathrm{d}s + \int_t^{\tau_{t,h}} \frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}})\mathrm{d}A_s + \int_t^{\tau_{t,h}} \overline{\nabla}\varphi(r, K_r^{A,\mathsf{P}})\mathrm{d}\overline{B}_r.$$

Recall from (5.4.2) that

$$Y_{\tau_{t,h}}^A = Y_t^A - \int_t^{\tau_{t,h}} F(r, K_r^{A,\mathsf{P}}, Y_r^A, Z_r)\mathrm{d}r - \int_t^{\tau_{t,h}} (\ell \mathrm{d}A_r^+ + u\mathrm{d}A_r^-) + \int_t^{\tau_{t,h}} Z_r \mathrm{d}\overline{B}_r.$$

Let $\Delta Y_t^A \triangleq \check{Y}_t^A - Y_t^A$ and $\Delta Z_t \triangleq \check{Z}_t - Z_t$. Then, the definition of $\tau_{t,h}$ ensures that

$$\begin{aligned}
\Delta Y_{\tau_{t,h}}^A - \Delta Y_t^A &= \int_t^{\tau_{t,h}} \Big(\mathscr{L}\varphi(r, K_r^{A,\mathsf{P}}) + F(r, K_r^{A,\mathsf{P}}, Y_r^A), Z_r)\Big)\,\mathrm{d}s + \int_t^{\tau_{t,h}} \Delta Z_r \mathrm{d}\overline{B}_r \\
&\quad + \int_t^{\tau_{t,h}} \Big(\frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) + \ell\Big)\,\mathrm{d}A_r^+ + \int_t^{\tau_{t,h}} \Big(-\frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) + u\Big)\,\mathrm{d}A_r^+ \\
&< \int_t^{\tau_{t,h}} \Big(F(r, K_r^{A,\mathsf{P}}, Y_r^A), Z_r) - F(r, K_r^{A,\mathsf{P}}, \check{Y}_r^A, \check{Z}_r))\Big)\,\mathrm{d}r + \int_t^{\tau_{t,h}} \Delta Z_r \mathrm{d}\overline{B}_r \\
&\quad + \int_t^{\tau_{t,h}} \Big(\frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) + \ell\Big)\,\mathrm{d}A_r^+ + \int_t^{\tau_{t,h}} \Big(-\frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) + u\Big)\,\mathrm{d}A_r^+ \\
&= \int_t^{\tau_{t,h}} (\delta_r^Y \Delta Y_r^A + \delta_r^Z \Delta Z_r)\mathrm{d}r + \int_t^{\tau_{t,h}} \Delta Z_r \mathrm{d}\overline{B}_r + \int_t^{\tau_{t,h}} \delta_r^+ \Delta Y_r^A \mathrm{d}A_r^+ + \int_t^{\tau_{t,h}} \delta_r^- \Delta Y_r^A \mathrm{d}A_r^-,
\end{aligned} \tag{5.5.4}$$

where

$$\begin{aligned}
\delta_r^Y &= \frac{F(r, K_r^{A,\mathsf{P}}, \check{Y}_r^A, Z_r) - F(r, K_r^{A,\mathsf{P}}, Y_r^A, Z_r))}{\Delta Y_r^A} 1_{\mathbb{R}\setminus\{0\}}(\Delta Y_r^A) \\
\delta_r^Z &= \frac{F(r, K_r^{A,\mathsf{P}}, Y_r^A, \check{Z}_r) - F(r, K_r^{A,\mathsf{P}}, Y_r^A, Z_r))}{\Delta Z_r} 1_{\mathbb{R}\setminus\{0\}}(\Delta Z_r) \\
\delta_r^+ &= \frac{\frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) + \ell}{\Delta Y_r^A} 1_{\mathbb{R}\setminus\{0\}}(\Delta Y_r^A) \\
\delta_r^- &= \frac{-\frac{\partial\varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) + u}{\Delta Y_r^A} 1_{\mathbb{R}\setminus\{0\}}(\Delta Y_r^A)
\end{aligned}$$

Now, consider the processes $(\Gamma^Y, \Gamma^Z) = (\Gamma_t^Y, \Gamma_t^Z)_{t\in[0,T]}$ such that

$$\begin{aligned}
\Gamma_{\tau_{t,h}}^Y &= 1 - \int_t^{\tau_{t,h}} \Gamma_r^Y (\delta_r^Y \mathrm{d}r + \delta^+ \mathrm{d}A_r^+ + \delta^- \mathrm{d}A_r^-) \\
\Gamma_{\tau_{t,h}}^Z &= 1 - \int_t^{\tau_{t,h}} \Gamma_r^Z \delta_r^Y \mathrm{d}\overline{B}_r.
\end{aligned}$$

Clearly, $\Gamma^Y \geq 0$ and $\Gamma^Z \geq 0$. Then, we obtain by Itō's integration by parts that

$$\begin{aligned}
\Gamma^Y_{\tau_{t,h}}\Gamma^Z_{\tau_{t,h}}(\Delta Y^A_{\tau_{t,h}} - \Delta Y^A_t) &= \int_t^{\tau_{t,h}} \Gamma^Z_r \Delta Y^A_r \mathrm{d}\Gamma^Y_r + \int_t^{\tau_{t,h}} \Gamma^Y_r \Gamma^Z_r \mathrm{d}\Delta Y^A_r \\
&\quad - \int_t^{\tau_{t,h}} \delta^Z_r \Gamma^Y_r \Gamma^Z_r \Delta Y^A_r \mathrm{d}\overline{B}_r - \int_t^{\tau_{t,h}} \delta^Z_r \Gamma^Y_r \Gamma^Z_r \Delta Z_r \mathrm{d}r \\
&\underset{(5.5.4)}{<} - \int_t^{\tau_{t,h}} \Gamma^Z_r \Delta Y^A_r \Gamma^Y_r (\delta^Y_r \mathrm{d}r + \delta^+ \mathrm{d}A^+_r + \delta^- \mathrm{d}A^-_r) \\
&\quad + \int_t^{\tau_{t,h}} \Gamma^Y_r \Gamma^Z_r ((\delta^Y_r \Delta Y^A_r + \delta^Z_r \Delta Z_r)\mathrm{d}r + \Delta Z_r \mathrm{d}\overline{B}_r + \delta^+_r \Delta Y^A_r \mathrm{d}A^+_r + \delta^-_r \Delta Y^A_r \mathrm{d}A^-_r) \\
&\quad - \int_t^{\tau_{t,h}} \delta^Z_r \Gamma^Y_r \Gamma^Z_r \Delta Y^A_r \mathrm{d}\overline{B}_r - \int_t^{\tau_{t,h}} \delta^Z_r \Gamma^Y_r \Gamma^Z_r \Delta Z_r \mathrm{d}r \\
&= \int_t^{\tau_{t,h}} \Gamma^Y_r \Gamma^Z_r (\Delta Z_r - \delta^Z_r \Delta Y^A_r)\mathrm{d}\overline{B}_r. \qquad (5.5.5)
\end{aligned}$$

From (5.4.3), we have that $|\delta^Y| < C_f$. Moreover, the assumption that $\varphi \in C^2$, together with that $K^{A,\mathsf{P}} \in \mathbb{H}^2_{\tau_{t,h}}(\mathbf{F}^X, \mathsf{P})$, ensure that both δ^+ and δ^- are in $\mathbb{H}^2_{\tau_{t,h}}(\mathbf{F}^X, \mathsf{P})$. Since $\Delta Z \in \mathbb{H}^2_T(\mathscr{F}^X, \mathsf{P})$, it follows from (5.4.3) that $\delta^Z \in \mathbb{H}^2_T(\mathscr{F}^X, \mathsf{P})$ as well. Taken together, Jeanblanc et al. (2009, Exercise 9.4.3.5) implies that $\Gamma^Y, \Gamma^Z \in \mathbb{H}^2_T(\mathscr{F}^X, \mathsf{P})$. As a result, the stochastic integral in (5.5.5) is a martingale (cf. Øksendal, 2010, Theorem 3.2.1). Therefore,

$$\begin{aligned}
&\mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\Gamma^Y_{\tau_{t,h}}\Gamma^Z_{\tau_{t,h}}(\Delta Y^A_{\tau_{t,h}} - \Delta Y^A_t)\right] \\
&< \mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^{\tau_{t,h}} \Gamma^Y_r \Gamma^Z_r \left(\Delta Z_r - \delta^Z_r \Delta Y^A_r\right) \mathrm{d}\overline{B}_r\right] = 0, \qquad (5.5.6)
\end{aligned}$$

Since $(t, K^{A,\mathsf{P}}t) = (t,k)$ is assumed to be a local maximum point and the control policy $A \in \mathscr{D}$ is arbitrary, it follows that $0 = V(t,k) - \varphi(t,k) = -\Delta Y^{A^*}_t \geq -\Delta Y^{A^*}_{\tau_{t,h}}$, where $A^* = \operatorname{arginf}_{A\in\mathscr{D}} V(t,\hat{k})$, which contradicts (5.5.6). Hence, we conclude that $\varphi(t,k)$ is a viscosity sub-solution of the HJB equation (5.4.15).

Supersolution. The proof of supersolution can be obtained similarly, using a contradiction argument. That is, if we assume that (t,k) is a local minimum point of $V - \varphi$ and suppose, in contrary, that

$$\begin{aligned}
\mathscr{L}\varphi(t,k) + F(t,k,\varphi(t,k), \overline{\nabla}\varphi(t,k)) &> 0, \\
\ell + \frac{\partial \varphi}{\partial x}(t,k) &> 0, \quad \text{and} \\
u - \frac{\partial \varphi}{\partial x}(t,k) &> 0,
\end{aligned}$$

and define the stopping time $\overline{\varrho}_t, \check{\varrho}_t, \hat{\varrho}_t$ by

$$\begin{aligned}
\overline{\varrho}_t &\triangleq \inf\{r > t : \mathscr{L}\varphi(r, K^{A,\mathsf{P}}_r) + F(r, K^{A,\mathsf{P}}_r, \varphi(t, K^{A,\mathsf{P}}_r), \overline{\nabla}\varphi(t, K^{A,\mathsf{P}}_r)) = 0\} \\
\check{\varrho}_t &\triangleq \inf\{r > t : \ell + \frac{\partial \varphi}{\partial x}(r, K^{A,\mathsf{P}}_r) = 0\} \quad \text{and}
\end{aligned}$$

$$\hat{\varrho}_t \triangleq \inf\{r > t : u - \frac{\partial \varphi}{\partial x}(r, K_r^{A,\mathsf{P}}) = 0\},$$

then one can adapt the argument used in the second part of the subsolution proof, by replacing $\tau_{t,h}$ with $\varrho_{t,h} \triangleq \overline{\varrho}_t \wedge \check{\varrho}_t \wedge \hat{\varrho}_t \wedge (t+h)$, to derive a contradiction here. Hence, $\varphi(t,k)$ is the viscosity supersolution to the HJB equation (5.4.15), which completes the proof. ■

5.6 COORDINATION TRANSFORM

As previously noted, there is no known closed-form solution to the HJB equation (5.4.15). Consequently, we turn to numerical methods. In this paper, we develop a scheme based on the *Markov Chain Approximation* (MCA) method (cf. Kushner and Martins, 1991) to solve the singular stochastic control problem in finite time horizon under Gaussian-generated smooth ambiguity.

The MCA framework is particularly well suited to our setting because it efficiently handles free boundary problems on uniform, rectangular grids, a structure that closely matches the characteristics of our HJB equation.[5] Our goal is to construct an MCA scheme that not only preserves numerical consistency with the HJB equation, but also offers computational tractability and delivers results with intuitive managerial interpretations.

However, implementing such a scheme is non-trivial. The HJB equation (5.4.15) features a non-constant covariance structure, and the processes $X^{A,M^{\mathsf{P}}}$ and M^{P} are perfectly correlated, posing significant challenges for numerical approximation.

First, within the MCA framework, solving high-dimensional diffusion problems typically involves large nonlinear systems. Reducing the dimensionality is therefore essential for achieving computational efficiency. Second, the combination of a non-constant covariance matrix and perfect correlation complicates the construction of a numerical grid that guarantees convergence of the approximation to the solution of the HJB equation. See Kushner and Dupuis (2001, Page 110) or Aliprantis and Border (2006) for discussion of such difficulties in correlated diffusion settings. Lastly, in our singular control setting, the Markov chain transition probabilities are binary, taking values only 0 or 1, which can, in turn, cause the numerical scheme to diverge.

To overcome these issues, we employ the coordinate transformation technique proposed by Johnson and Peskir (2017).[6] This method transforms the original two-dimensional stochastic system $(X^{A,M^{\mathsf{P}}}, M^{\mathsf{P}})$ into an equivalent formulation involving a single diffusion process and two bounded variation processes. We demonstrate that this transformation not only resolves the numerical challenges outlined above but also preserves the economic and managerial interpretability of the model.

[5] For approaches designed for irregular domains, see Kumar and Muthuraman (2004).
[6] See also Basei et al. (2024), De Angelis (2020), and Federico et al. (2023) for further developments.

To do this, we define a transformation $\mathscr{T} : [0,T] \times \Omega_K^T \to \mathbb{R}$ by

$$\mathscr{T}(t,k) \triangleq x - \frac{\sigma}{S_t} m \triangleq x_2, \ k = (x,m)$$

where the inverse is given by

$$\mathscr{T}^{-1}(t,\hat{k}) \triangleq \frac{S_t}{\sigma}(x_1 - x_2) \triangleq m, \ \ \hat{k} = (t, x_1, x_2).$$

Now, define $X^{(2)} \triangleq \left(\mathscr{T}(t, K_t^{A,\mathsf{P}})\right)_{t\in[0,T]}$, and let $X^{(1)} \triangleq X^{M^{\mathsf{P}},A}$. Then it follows from Itō's lemma that the new system of the forward SDEs $\widehat{K}^{A,\mathsf{P}} \triangleq (X^{(1)}, X^{(2)})$ solve

$$\begin{aligned}
\mathrm{d}X_t^{(1)} &= \alpha^{(1)}(t, X_t^{(1)}, X_t^{(2)})\mathrm{d}t + \mathrm{d}A_t + \sigma \mathrm{d}\overline{B}_t, && X_0^{(1)} = x, && (5.6.1)\\
\mathrm{d}X_t^{(2)} &= \alpha^{(2)}(t, X_t^{(1)}, X_t^{(2)})\mathrm{d}t + \mathrm{d}A_t - \frac{\sigma}{S_t}\mathrm{d}D_t, && X_0^{(2)} = x - \frac{\sigma}{S_0}m, && (5.6.2)\\
\mathrm{d}S_t &= -S_t^2 \mathrm{d}t, && S_0 = s.
\end{aligned}$$

where

$$\alpha^{(1)}(t,x_1,x_2) \triangleq \alpha + S_t(x_2 - x_1) \ \text{ and } \ \alpha^{(2)}(t,x_1,x_2) \triangleq \alpha + 2S_t(x_2 - x_1).$$

The *transformed* value function is defined by

$$\widehat{V}(t,\hat{k}) \triangleq V(t, x_1, \mathscr{T}^{-1}(t,\hat{k})).$$

By applying the coordinate transformation, the process D now increases only when $X_t^{(2)} = X_t^{(1)} - \frac{\sigma}{S_t}\underline{m}$ or $X_t^{(1)} - \frac{\sigma}{S_t}\overline{m}$. The widest admissible range of $X^{(2)}$ is therefore

$$[\underline{x}_2, \overline{x}_2] \triangleq \left[\overline{x} - \frac{\sigma}{S_T}\underline{m}, \ \underline{x} - \frac{\sigma}{S_T}\overline{m}\right]. \tag{5.6.3}$$

We assume there exist bounded variation processes $\overline{D}^-$ and $\overline{D}^+$, which increase when $X^{(2)}$ hits $\underline{x}_2$ or $\overline{x}_2$, respectively, where $\overline{D}^{\pm}$ are defined analogously to D. Using $\overline{D}^{\pm}$ in place of D is consistent with the confidence interval imposed on (5.2.3), since it produces the same or a more accurate approximation of the value function by imposing a larger confidence interval on the belief process (after the inverse transformation).

From (5.3.4), no cost is incurred by adjustments via D or $\overline{D}$. Therefore, solving (5.6.2) is equivalent to

$$\mathrm{d}X_t^{(2)} = \alpha^{(2)}(t, X_t^{(1)}, X_t^{(2)})\mathrm{d}t + \mathrm{d}A_t + \mathrm{d}\overline{D}_t, \tag{5.6.4}$$

which will be used in the remainder of our analysis. We refer to the domain of $\widehat{K}_t^{A,\mathsf{P}}, t \in [0,T]$ as $\Omega_{\widehat{K}}^T \triangleq [\underline{x}, \overline{x}] \times [\underline{x}_2, \overline{x}_2]$.

Remark 21. *Note that $X^{(2)}$ resembles $X^{(1)}$ but carries no risk exposure. Moreover, eliciting M^{P} from the environment can be achieved by observing $X^{(2)}$, which is inferred dynamically from $X^{(1)}$. For this reason, we refer to $X^{(2)}$ as an* auxiliary controlled inventory process.

Remark 22. *Combining* (5.6.1) *and* (5.6.4)*, we obtain*

$$\mathrm{d}(X_t^{(1)} - X_t^{(2)}) = S_t(X_t^{(1)} - X_t^{(2)})\mathrm{d}t + \sigma \mathrm{d}\overline{B}_t - \mathrm{d}\overline{D}_t, \tag{5.6.5}$$

which implies a mean-repulsion *tendency from the long-run mean of zero, with repulsion rate S_t. Typically, a mean-repulsion process diverges to $\pm\infty$ over time due to instability around its mean. However, this is not the case here because S_t is positive and monotonically decreasing over time. As a result, the increment of* (5.6.5) *eventually converges to a Gaussian process with a constant mean. Consequently, the process $\mathscr{T}^{-1}(t, \widehat{K}_t^{A,\mathsf{P}})$ becomes nearly deterministic as the noise diminishes (See, Figure 5.1), implying that the DM eventually recovers full information about the controlled inventory process $X^{(1)}$ and its drift through sufficiently long observation.*

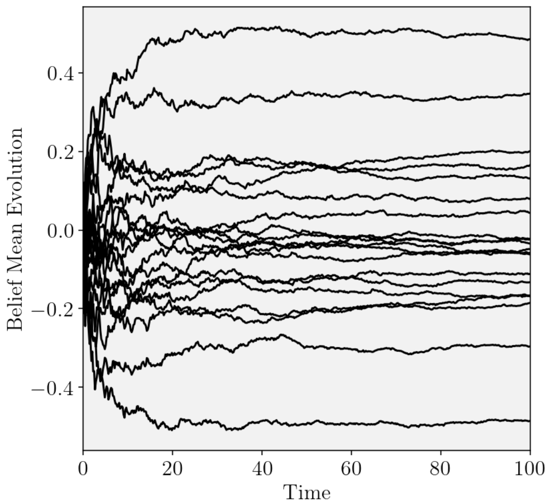

Figure 5.1: Sample paths of the belief mean process $M_t^{\mathsf{P}} \triangleq \mathscr{T}^{-1}(t, \widehat{K}_t^{A,\mathsf{P}})$.

Under the coordinate transformation, the corresponding BSDE becomes

$$\mathrm{d}\widehat{Y}_t^A = -\widehat{F}(t, \widehat{K}_t^{A,\mathsf{P}}, \widehat{Y}_t^A, \widehat{Z}_t)\mathrm{d}t - \ell \mathrm{d}A_t^+ - u\mathrm{d}A_t^- + \widehat{Z}_t \mathrm{d}\overline{B}_t,$$

where $\widehat{F}(t, \hat{k}, y, z) \triangleq F(t, x_1, \mathscr{T}^{-1}(t, \hat{k}), y, z)$, with $\hat{k} = (t, x_1, x_2)$. By the comparison theorem (cf. Lemma 7), we have $\widehat{Y}^A = Y^A$. Therefore, $\widehat{V}(t, \hat{k}) \triangleq \sup_{A \in \mathscr{D}} \widehat{Y}_t^{A;\hat{k}}$, $\widehat{K}_t^{A,\mathsf{P}} = \hat{k}$, is the optimal singular control value function for the transformed BSDE.

We now present a verification theorem for the transformed value function $\widehat{V}$.

Theorem 19 (Verification theorem)**.** *Suppose that $\widehat{V} \in C^{1,2}([0,T] \times \Omega^T_{\widehat{K}})$ such that*

$$\widehat{V}(t,\hat{k}) \triangleq \sup_{A \in \mathscr{D}} \mathsf{E}^{\mathsf{P}}_{\mathscr{F}^X_t}\left[\int_t^T \widehat{F}\left(r, \widehat{K}^{A,\mathsf{P}}_r, \widehat{V}(r, \widehat{K}^{A,\mathsf{P}}_r), \sigma \frac{\partial \widehat{V}}{\partial x_1}(t, \widehat{K}^{A,\mathsf{P}}_r)\right)\mathrm{d}t + \int_t^T (\ell \mathrm{d}A^+_t + u\mathrm{d}A^-_t)\right].$$

If $\phi \in C^{1,2}([0,T] \times \Omega^T_K)$ is such that

$$\widehat{\mathscr{L}}\phi(t,\hat{k}) + \widehat{F}(t,\hat{k},\phi(t,\hat{k}),\sigma\frac{\partial\phi}{\partial x_1}(t,\hat{k})) \le 0 \ \textit{ and} -\ell \le \frac{\partial\phi}{\partial x_1}(t,\hat{k}) + \frac{\partial\phi}{\partial x_2}(t,\hat{k}) \le u,$$

where

$$\widehat{\mathscr{L}}\phi(t,\hat{k}) \triangleq \alpha^{(1)}(t,\hat{k})\frac{\partial\phi}{\partial x_1}(t,\hat{k}) + \alpha^{(2)}(t,\hat{k})\frac{\partial\phi}{\partial x_2}(t,\hat{k}) + \sigma\frac{\partial^2\phi}{\partial x_1^2}(t,\hat{k})$$

Then $\widehat{V}(t,\hat{k}) \ge \phi(t,\hat{k})$ for any $(t,\hat{k}) \in [0,T] \times \Omega^T_{\widehat{K}}$ and $A \in \mathscr{D}$.

Moreover, for each $(t,x_2) \in [0,T] \times [\underline{x}_2, \overline{x}_2]$ if there exists an open free-boundary interval $\widehat{\mathcal{X}}(t,x_2) \subseteq [\underline{x},\overline{x}]$ such that

$$\begin{aligned}
\mathscr{L}\phi(t,\hat{k}) + \widehat{F}(t,\hat{k},\phi(t,\hat{k}),\sigma\frac{\partial\phi}{\partial x_1}(t,\hat{k})) &= 0 \qquad \text{on } \widehat{\mathcal{X}}(t,x_2),\\
\frac{\partial\phi}{\partial x_1}(t,\hat{a}(t,x_2),x_2) + \frac{\partial\phi}{\partial x_2}(t,\hat{a}(t,x_2),x_2) &= -\ell,\\
\frac{\partial\phi}{\partial x_1}(t,\hat{b}(t,x_2),x_2) + \frac{\partial\phi}{\partial x_2}(t,\hat{b}(t,x_2),x_2) &= u,\\
\frac{\partial^2\phi}{\partial x_1^2}(t,\hat{a}(t,x_2),x_2) - \frac{\partial^2\phi}{\partial x_2^2}(t,\hat{a}(t,x_2),x_2) + 2\frac{\partial^2\phi}{\partial x_1\partial x_2}(t,\hat{a}(t,x_2),x_2) &=\\
\frac{\partial^2\phi}{\partial x_1^2}(t,\hat{b}(t,x_2),x_2) - \frac{\partial^2\phi}{\partial x_2^2}(t,\hat{b}(t,x_2),x_2) + 2\frac{\partial^2\phi}{\partial x_1\partial x_2}(t,\hat{b}(t,x_2),x_2) &= 0,
\end{aligned}$$

where

$$\begin{aligned}
\hat{a}(t,x_2) &\triangleq \inf\{x_1 : x_1 \in \widehat{\mathcal{X}}(x_2)\},\\
\hat{b}(t,x_2) &\triangleq \sup\{x_1 : x_1 \in \widehat{\mathcal{X}}(x_2)\},
\end{aligned}$$

and $\widehat{A}^* \in \mathscr{D}$ such that $\widehat{A}^{*-}$ and $\widehat{A}^{*+}$ only increase at $\hat{a}(x_2)+$ and $\hat{b}(x_2)-$, respectively, P-a.s., then $\widehat{V}(t,\hat{k}) = \phi(t,\hat{k})$.

In the same spirit as Theorem 18, Theorem 19 shows that $\widehat{V}$ is a classical solution to the HJB equation

$$\begin{aligned}
0 = \min\Bigg\{&\widehat{\mathscr{L}}\widehat{V}(t,\hat{k}) + \widehat{F}\left(t,\hat{k},\widehat{V}(t,\hat{k}),\sigma\frac{\partial\widehat{V}}{\partial x_1}(t,\hat{k})\right),\\
&\ell + \frac{\partial\widehat{V}}{\partial x_1}(t,\hat{k}) + \frac{\partial\widehat{V}}{\partial x_2}(t,\hat{k}), u - \frac{\partial\widehat{V}}{\partial x_1}(t,\hat{k}) - \frac{\partial\widehat{V}}{\partial x_2}(t,\hat{k})\Bigg\}.
\end{aligned} \tag{5.6.6}$$

The total derivatives of $\widehat{V}$ with respect to x_1 and x_2 are

$$\begin{aligned}\frac{\partial \widehat{V}}{\partial x_1}(t,\hat{k}) &= \frac{\partial V}{\partial x}(t,k) + \frac{S_t}{\sigma}\frac{\partial V}{\partial m}(t,k), \text{ and} \\ \frac{\partial \widehat{V}}{\partial x_2}(t,\hat{k}) &= -\frac{S_t}{\sigma}\frac{\partial V}{\partial m}(t,k),\end{aligned} \tag{5.6.7}$$

respectively. From this, the convexity of V in x (cf. Proposition 14) implies that $\frac{\partial V}{\partial x} = \frac{\partial \widehat{V}}{\partial x_1} + \frac{\partial \widehat{V}}{\partial x_2}$ is strictly decreasing in x_1. Therefore, for each $t, x_2 \in [0,T] \times \Omega^T_{\widehat{K}}$, the pair $(\hat{a}, \hat{b})(t, x_2)$ is unique and is determined by

$$\hat{a}(t,x_2) \triangleq \sup\left\{x_1 \in [\underline{x}, \overline{x}] : \frac{\partial \widehat{V}}{\partial x_1}(t,\hat{k}) + \frac{\partial \widehat{V}}{\partial x_2}(t,\hat{k}) \le -\ell\right\}, \text{ and} \tag{5.6.8}$$

$$\hat{b}(t,x_2) \triangleq \inf\left\{x_1 \in [\underline{x}, \overline{x}] : \frac{\partial \widehat{V}}{\partial x_1}(t,\hat{k}) + \frac{\partial \widehat{V}}{\partial x_2}(t,\hat{k}) \ge u\right\}. \tag{5.6.9}$$

Furthermore, from (5.6.7) we have $\sigma\frac{\partial \widehat{V}}{\partial x_1} = \overline{\nabla} V(x_1, \mathscr{T}^{-1}(t,\hat{k}))$. This shows that the BSDE $(\widehat{V}, \sigma\frac{\partial \widehat{V}}{\partial x_1}, \widehat{A}^*)$ admits a unique solution, which coincides with $(V, \overline{\nabla} V, A^*)$.

Since the transformation $\mathscr{T}$ is smooth and injective, the argument in Proposition 16 can be directly applied to confirm that $\widehat{V}(t,\hat{k})$ is a viscosity solution to the transformed HJB equation.

Remark 23. *Analogous to Remark 20, if $\widehat{V}$ is a classical solution to (5.6.6), then the optimal controlled inventory dynamics are given by*

$$\mathrm{d}X^{(1)}_t = \alpha^{(1)}_g(t, X^{(1)}_t, X^{(2)}_t)\mathrm{d}t + \sigma \mathrm{d}\overline{B}_t + \mathrm{d}\widehat{A}^*_t, \tag{5.6.10}$$

where $X^{(2)}$ solves (5.6.4) and

$$\alpha^{(1)}_g(t,x_1,x_2) \triangleq \alpha^{(1)}(t,x_1,x_2) + \frac{\sigma^2\gamma S_t}{2}\frac{\partial \widehat{V}}{\partial x_1}(t,x_1,x_2).$$

For a similar reason as given in Remark 20, the process (5.6.10) has a strong unique solution.

Note that, under the coordinate transformation, ambiguity enters only through the drift of the controlled inventory process. Thus, the transformation not only reduces the dimensionality of the diffusion component, but also channels the ambiguity adjustment into a single coordinate, which subsequently, as shown in the next section, further improves computational efficiency.

Proposition 17. *Suppose that Assumption 2) holds. Then $\widehat{V}(t,\hat{k})$ is a viscosity solution to the transformed HJB equation (5.6.6).*

5.7 MARKOV CHAIN APPROXIMATION

To implement the MCA, we first partition the time-state domain $[0,T] \times \Omega_{\widehat{K}}^{T}$ into $n \in \mathbb{N}$ equal intervals. Namely,

$$\begin{aligned}
\Omega_T^n &\triangleq \{0, \delta, 2\delta, \ldots, (n-1)\delta, T\}, \quad \delta = \frac{T}{n} \\
\Omega_{X^{(1)}}^n &\triangleq \{\underline{x}, \underline{x}+h_1, \underline{x}+2h_1, \ldots, \underline{x}+(n-1)h_1, \overline{x}\}, \quad h_1 = \frac{\overline{x}-\underline{x}}{n} \\
\Omega_{X^{(2)}}^n &\triangleq \{\underline{x}_2, \underline{x}_2+h_2, \underline{x}_2+2h_2, \ldots, \underline{x}_2+(n-1)h_2, \overline{x}_2\}, \quad h_2 = \frac{\overline{x}_2-\underline{x}_2}{n}.
\end{aligned}$$

Then for each $(t, x_1, x_2) \in \Omega_T^n \times \Omega_{X^{(1)}}^n \times \Omega_{X^{(2)}}^n \triangleq \Omega_{\widehat{K}}^{n,T}$, the approximate transformed value function is obtained by the recursive mapping $\mathbb{T}^n : \Omega_{\widehat{K}}^{n,T} \to \mathbb{R}$:

$$\mathbb{T}^n \widehat{V}(t,x_1,x_2) \triangleq \begin{cases} \min\big\{\widehat{V}_c(t,x_1,x_2), \\ \quad \widehat{V}_\ell(t,x_1,x_2), \widehat{V}_u(t,x_1,x_2)\big\} & \text{if } (t,x_1) \in \Omega_T^n \times \Omega_{X^{(1)}}^n \text{ and } x_2 \notin \{\underline{x}_2, \overline{x}_2\} \\ \widehat{V}(t,x_1,x_2+h_2) & \text{if } (t,x_1) \in \Omega_T^n \times \Omega_{X^{(1)}}^n \text{ and } x_2 = \underline{x}_2 \\ \widehat{V}(t,x_1,x_2-h_2) & \text{if } (t,x_1) \in \Omega_T^n \times \Omega_{X^{(1)}}^n \text{ and } x_2 = \overline{x}_2, \end{cases} \tag{5.7.1}$$

where

$$\begin{aligned}
\widehat{V}_c(t,x_1,x_2) &\triangleq (1-\rho\Delta t_{t,x_1,x_2}^{\delta,h_1,h_2}) \Bigg(\sum_{(t',x_1',x_2') \in \Omega_{\widehat{K}}^{n,T}} \bar{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') \widehat{V}(t',x_1',x_2') \\
&\qquad + f(x_1)\Delta t_{t,x_1,x_2}^{\delta,h_1,h_2} \Bigg) \\
\widehat{V}_\ell(t,x_1,x_2) &\triangleq \sum_{(t',x_1',x_2') \in \Omega_{\widehat{K}}^{n,T}} \check{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') \widehat{V}(t',x_1',x_2') + \frac{\ell h_1 h_2}{h_1+h_2} \\
\widehat{V}_u(t,x_1,x_2) &\triangleq \sum_{(t',x_1',x_2') \in \Omega_{\widehat{K}}^{n,T}} \hat{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') \widehat{V}(t',x_1',x_2') + \frac{u h_1 h_2}{h_1+h_2}.
\end{aligned}$$

The terms $\bar{p}_{t,x_1,x_2}^{\delta,h_1,h_2}$, $\check{p}_{t,x_1,x_2}^{\delta,h_1,h_2}$ and $\hat{p}_{t,x_1,x_2}^{\delta,h_1,h_2}$ are obtained by applying the upwind implicit finite-difference scheme[7] to the transform HJB equation (5.6.6). That is,

$$\bar{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') \triangleq \frac{\widetilde{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2')}{\widetilde{q}_{x_1,x_2}^{h_1,h_2}}$$

[7]The reason we use the implicit scheme is the fact that the control barriers in our setting is time-invariant, so that we must treat time as another (virtual) state variable, namely ζ^δ, to have the forward equation (5.7.1) locally consistent with the corresponding HJB equation. See, Kushner and Dupuis (2001, Section 12.4), for the general implementation of this scheme.

$$\tilde{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2')$$
$$\triangleq \begin{cases} \frac{\sigma^2}{2h_1^2}+\frac{1}{h_1}\Big(\alpha^{(1)}(t,x_1,x_2)^+ \\ \quad +\frac{\sigma^2\gamma S_t}{2h_1}\left[\overline{V}(t,x_1',x_2)-\overline{V}(t,x_1,x_2)\right]^+\Big) & \text{if } (t',x_1',x_2')=(t,x_1+h_1,x_2) \\ \frac{\sigma^2}{2h_1^2}+\frac{1}{h_1}\Big(\alpha^{(1)}(t,x_1,x_2)^- \\ \quad +\frac{\sigma^2\gamma S_t}{2h_1}\left[\overline{V}(t,x_1,x_2)-\overline{V}(t,x_1',x_2)\right]^-\Big) & \text{if } (t',x_1',x_2')=(t,x_1-h_1,x_2) \\ \frac{\alpha^{(2)}(t,x_1,x_2)^{\pm}}{h_2} & \text{if } (t',x_1',x_2')=(t,x_1,x_2\pm h_2) \\ \frac{2}{\delta} & \text{if } (t',x_1',x_2')=(t+\delta,x_1,x_2) \\ 0 & \text{otherwise,} \end{cases} \tag{5.7.2}$$

$$\begin{aligned} \tilde{q}^{\delta,h_1,h_2}_{t,x_1,x_2} &\triangleq \frac{2}{\delta}+\frac{\sigma^2}{h_1^2}\bigg(1+\frac{\gamma S_t}{2}\Big(\left[\overline{V}(t,x_1+h_1,x_2)-\overline{V}(t,x_1,x_2)\right]^+ \\ &\quad +\left[\overline{V}(t,x_1,x_2)-\overline{V}(t,x_1-h_1,x_2)\right]^-\Big)\bigg) \\ &\quad +\frac{|\alpha^{(1)}(t,x_1,x_2)|}{h_1}+\frac{|\alpha^{(2)}(t,x_1,x_2)|}{h_2}, \\ \Delta t^{\delta,h_1,h_2}_{t,x_1,x_2} &\triangleq \frac{1}{\tilde{q}^{\delta,h_1,h_2}_{t,x_1,x_2}}, \end{aligned}$$

$$\check{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2') \triangleq \begin{cases} \frac{h_2}{h_1+h_2} & \text{if } (t',x_1',x_2')=(t,x_1+h_1,x_2) \\ \frac{h_1}{h_1+h_2} & \text{if } (t',x_1',x_2')=(t,x_1,x_2+h_2) \\ 0 & \text{otherwise,} \end{cases} \text{ and,} \tag{5.7.3}$$

$$\hat{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2') \triangleq \begin{cases} \frac{h_2}{h_1+h_2} & \text{if } (t',x_1',x_2')=(t,x_1-h_1,x_2) \\ \frac{h_1}{h_1+h_2} & \text{if } (t',x_1',x_2')=(t,x_1,x_2-h_2) \\ 0 & \text{otherwise.} \end{cases} \tag{5.7.4}$$

Now we suppose that the process $(\zeta^\delta, X^{(1),h_1}, X^{(2),h_2}) = (\zeta^\delta_k, X^{(1),h_1}_k, X^{(2),h_2}_k)_{0\le k\le n-1}$, are the discrete time Markov chains on $\Omega^{n,T}_{\widehat{K}}$ with transition probabilities:

$$\begin{aligned} &\mathsf{P}\left(\left(k+1, X^{(1),h_1}_{k+1}, X^{(2),h_2}_{k+1}\right)=(t',x_1',x_2') \mid \left(k, X^{(1),h_1}_k, X^{(2),h_2}_k\right)=(t,x_1,x_2)\in\Omega^{n,T}_{\widehat{K}}\right) \\ &\triangleq p^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2') \triangleq \begin{cases} \bar{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2') & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_c(t,x_1,x_2) \\ \check{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2') & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_\ell(t,x_1,x_2) \\ \hat{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2') & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_u(t,x_1,x_2) \end{cases} \end{aligned} \tag{5.7.5}$$

where $\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2}$ is the time interval of the Markov chains.

The associated conditional expectation operator is denoted by $\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}$. Observe that

$$p^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2')\in(0,1),\quad \forall (t',x_1',x_2')\in\Omega^{n,T}_{\widehat{K}}$$

$$\sum_{(t',x_1',x_2')\in\Omega_{\widehat{K}}^{n,T}} p_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') = 1.$$

Therefore, (5.7.5) is well-defined.

Importantly, when $\mathbb{T}\widehat{V}(t,x_1,x_2) = \widehat{V}_c(t,x_1,x_2)$, the state (x_1,x_2) lies within the inaction region, and it is optimal to continue with transition probabilities $\bar{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2')$, for all $(x_1',x_2') \in \Omega_{\widehat{K}}^{n,T}$. On the other hand, if $\mathbb{T}\widehat{V}(x_1,x_2)$ is equal to $\widehat{V}_\ell(x_1,x_2)$ or $\widehat{V}_u(x_1,x_2)$, this indicates that it is optimal to exert the lower control with probability $\check{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2')$ or the upper control with probability $\hat{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2')$, respectively. In other words, the state (x_1,x_2) is instantaneously shifted to

$$(x_1 + \Delta A_k^{+,h_1,h_2}, x_2 + \Delta A_k^{+,h_1,h_2}) \text{ or}$$
$$(x_1 - \Delta A_k^{-,h_1,h_2}, x_2 - \Delta A_k^{-,h_1,h_2}),$$

respectively, where $\Delta A_k^{\pm,h_1,h_2} \triangleq \frac{h_1 h_2}{h_1+h_2}$ for all $k \in \mathbb{N}$.

Remark 24. *From (5.6.3), one can see that the boundary of the process $X^{(2)}$ expands as the confidence interval widens, and becomes unbounded in the limit. In particular, when the confidence interval $[\underline{m},\overline{m}]$ is very large, the range $[\underline{x}_2,\overline{x}_2]$ can grow much larger than $[\underline{x}_1,\overline{x}_1]$. In such cases, we have $h_2 \gg h_1$.*

In the limiting case $\underline{m} \to -\infty$ and $\overline{m} \to \infty$, the reflecting state transition probabilities (5.7.3) and (5.7.4) reduce to

$$\check{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') \approx \begin{cases} 1 & \text{if } (t',x_1',x_2') = (t,x_1+h_1,x_2), \\ 0 & \text{otherwise,} \end{cases}$$
$$\hat{p}_{t,x_1,x_2}^{\delta,h_1,h_2}(t',x_1',x_2') \approx \begin{cases} 1 & \text{if } (t',x_1',x_2') = (t,x_1-h_1,x_2), \\ 0 & \text{otherwise,} \end{cases}$$

which is consistent with the behavior of the pre-transformed processes $(X^{A,M^{\mathsf{P}}}, M^{\mathsf{P}})$.

The following result affirms that the recursive mapping $\mathbb{T}^n$ gives a unique solution for a sufficiently large $n \in \mathbb{N}$.

Lemma 8. *Suppose that $a_i, b_i \in \mathbb{R}$ for $i = 1,2,\ldots,N$, for some $N \in \mathbb{N}$. Then the following inequality holds.*

$$\left| \min_{i=1,2,\ldots,N} a_i - \min_{i=1,2,\ldots,N} b_i \right| \leq \min_{i=1,2,\ldots,N} |a_i - b_i|. \tag{5.7.6}$$

Proof of Lemma 8. Suppose that that $i = 1,2$. Then, it holds that

$$|\min\{a_1,a_2\} - \min\{b_1,b_2\}|$$

$$= \begin{cases} |a_1 - b_1| & \text{if } \min\{a_1, a_2\} = a_1 \geq \min\{b_1, b_2\} = b_1 \\ |a_1 - b_2| \leq |a_2 - b_2| & \text{if } \min\{a_1, a_2\} = a_1 \geq \min\{b_1, b_2\} = b_2 \\ |a_2 - b_1| \leq |a_1 - b_1| & \text{if } \min\{a_1, a_2\} = a_2 \geq \min\{b_1, b_2\} = b_1 \\ |a_2 - b_2| & \text{if } \min\{a_1, a_2\} = a_2 \geq \min\{b_1, b_2\} = b_2. \end{cases}$$

In summary, we have $|\min\{a_1, a_2\} - \min\{b_1, b_2\}| \leq \min\{|a_1 - b_1|, |a_2 - b_2|\}$. Similarly, it follows that (5.7.6) holds for any $N \in \mathbb{N}$, therefore, completing the proof. ■

Theorem 20. *Suppose that there exists $n \in \mathbb{N}$, such that*

$$\max_{(t,x_1,x_2) \in \Omega^{m,T}_{\widehat{K}}} \rho \Delta t^{\delta,h_1,h_2}_{t,x_1,x_2} < 1, \textit{ for any } m \geq n. \tag{5.7.7}$$

Then the recursive mapping $\mathbb{T}^n \widehat{V}$ has a unique solution, which converges to the viscosity solution of the transformed HJB equation (5.4.15) as $n \to \infty$. Moreover, for each $(t, x_2) \in \Omega^n_T \times \Omega^{n,\varepsilon}_{X(2)}$ the optimal control barriers are approximated by $(\hat{a}^n(x_2), \hat{b}^n(x_2))$:

$$\hat{a}^n(x_2) \triangleq \max \Bigg\{ x_1 \in \Omega^n_{x_1} : \widehat{V}(t, x_1, x_2) \geq \frac{h_2}{h_1 + h_2} \widehat{V}(t, x_1 + h_1, x_2) + \frac{h_1}{h_1 + h_2} \widehat{V}(t, x_1, x_2 + h_2) + \frac{\ell h_1 h_2}{h_1 + h_2} \Bigg\}, \text{ and}$$

$$\hat{b}^n(x_2) \triangleq \min \Bigg\{ x_1 \in \Omega^n_{x_1} : \widehat{V}(t, x_1, x_2) \geq \frac{h_2}{h_1 + h_2} \widehat{V}(t, x_1 - h_1, x_2) + \frac{h_1}{h_1 + h_2} \widehat{V}(t, x_1, x_2 - h_2) + \frac{u h_1 h_2}{h_1 + h_2} \Bigg\}.$$

Proof of Theorem 20. To establish that $\mathbb{T}^n$ admits a unique fixed point, it suffices to show that it is a contraction mapping.

Suppose that $\widehat{V}_i$, for $i = 1, 2$, are solutions to (5.7.1), and define $\Delta \widehat{V} \triangleq \widehat{V}_1 - \widehat{V}_2$. Utilizing Proposition 15, it is straightforward to verify that the right-hand side of (5.7.7) is positive and strictly decreasing in n. Hence, there exists $n \in \mathbb{N}$ such that (5.7.7) holds for all $m \geq n$. Then, for any $(t, x_1, x_2) \in \Omega^{n,T}_{\widehat{K}}$, we have

$$\begin{aligned} &|\mathbb{T}^n \widehat{V}_1(t, x_1, x_2) - \mathbb{T}^n \widehat{V}_2(t, x_1, x_2)| \\ &\underset{\text{(Lemma 8)}}{\leq} \min \Bigg\{ \Bigg| (1 - r\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2}) \sum_{(t',x_1',x_2') \in \Omega^{n,T}_{\widehat{K}}} \bar{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t', x_1', x_2') \Delta \widehat{V}(t', x_1', x_2') \Bigg|, \\ &\qquad \Bigg| \sum_{(t',x_1',x_2') \in \Omega^{n,T}_{\widehat{K}}} \breve{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t', x_1', x_2') \Delta \widehat{V}(t', x_1', x_2') \Bigg|, \end{aligned}$$

$$\left|\sum_{(t',x_1',x_2')\in\Omega^{n,T}_{\widehat{K}}} \hat{p}^{\delta,h_1,h_2}_{t,x_1,x_2}(t',x_1',x_2')\Delta\widehat{V}(t',x_1',x_2')\right|\Bigg\}$$
$$\leq \widehat{C}\Delta\widehat{V}(t',x_1',x_2'),\ \text{for some } 0<\widehat{C}<1, \text{for any } (t',x_1',x_2')\in\Omega^{n,T}_{\widehat{K}}.$$

This implies that $\mathbb{T}^n\widehat{V}$ is a contraction mapping, which, therefore, has a unique fixed point (cf. Aliprantis and Border, 2006, Theorem 3.48). This means that the recursive mapping (5.7.1) weakly solves the transformed HJB equation (5.6.6), and converges in the sense of viscosity solution (cf. Definition 25) when $n\to\infty$. Moreover, the pair $(\hat{a}^n(t,x_2),\hat{b}^n(t,x_2))$ is obtained by discretizing (5.6.8) and (5.6.9), which exists uniquely thanks to Proposition 15. This completes the proof. ■

Now, let $(\Delta\zeta^\delta_k,\Delta X^{(1),h_1}_k,\Delta X^{(2),h_2}_k)\triangleq(\zeta^\delta_{k+1}-\zeta^\delta_k, X^{(1),h_1}_{k+1}-X^{(2),h_2}_k, X^{(2),h_1}_{k+1}-X^{(2),h_2}_k)$, then for each $(k,x_1,x_2)\in\Omega^{n,T}_{\widehat{K}}$,

$$\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\Delta\zeta^\delta_k\right]=2\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2}+\mathcal{O}(\delta,h_1,h_2) \tag{5.7.8}$$

$$\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\Delta X^{(1),h_1}_k\right]$$
$$=\begin{cases}\alpha^{(1)}(k,X^{(1),h_1}_k,X^{(2),h_2}_k)\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2} & \\ +\frac{\sigma^2\gamma S_k}{2}\overline{\nabla}_{h_1}\widehat{V}(k,X^{(1),h_1}_k,X^{(2),h_2}_k)\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2} & \\ +\mathcal{O}(\delta,h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_c(t,x_1,x_2)\\ \Delta A^{+,h_1,h_2}_k+\mathcal{O}(\delta,h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_\ell(t,x_1,x_2)\\ -\Delta A^{-,h_1,h_2}_k+\mathcal{O}(\delta,h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_u(t,x_1,x_2)\end{cases}$$

$$\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\Delta X^{(2),h_2}_k\right]$$
$$=\begin{cases}\alpha^{(2)}(X^{(1),h_1}_k,X^{(2),h_2}_k)\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2}+\mathcal{O}(\delta,h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_c(t,x_1,x_2)\\ \Delta A^{+,h_1,h_2}_k+\mathcal{O}(\delta,h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_\ell(t,x_1,x_2)\\ -\Delta A^{-,h_1,h_2}_k+\mathcal{O}(\delta,h_1,h_2) & \text{if } \mathbb{T}^n\widehat{V}(t,x_1,x_2)=\widehat{V}_u(t,x_1,x_2)\end{cases}$$

$$\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\left(\Delta\zeta^\delta_k-\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\Delta\zeta^\delta_k\right]\right)^2\right]=\mathcal{O}(\delta,h_1,h_2)$$
$$\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\left(\Delta X^{(1),h_1}_k-\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\Delta X^{(1),h_1}_k\right]\right)^2\right]=\sigma^2\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2}+\mathcal{O}(\delta,h_1,h_2)$$
$$\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\left(\Delta X^{(2),h_2}_k-\mathsf{E}^{\delta,h_1,h_2}_{t,x_1,x_2}\left[\Delta X^{(2),h_2}_k\right]\right)^2\right]=\mathcal{O}(\delta,h_1,h_2) \tag{5.7.9}$$

where $\overline{\nabla}_{h_1}\widehat{V}(k,x_1,x_2)\triangleq\left[\widehat{V}(k,x_1+h)-\widehat{V}(k,x_1)\right]^+-\left[\widehat{V}(k,x_1)-\widehat{V}(k,x_1-h_1)\right]^-$, and $\mathcal{O}(\delta,h_1,h_2)$ is the approximate error proportion to the size of δ,h_1,h_2. Note that the factor of 2 on the right-hand side of (5.7.8) arises from the fact that the stochastic system involves two time-dependent variables: S_t and t itself. According to Kushner and Dupuis (2001, Section 8.3), equations (5.7.8) through (5.7.9) establish the local consistency of the discretized processes $(\zeta^{\delta}, X^{(1),h_1}, X^{(2),h_2})$ with the original processes $(t, X^{(1)}, X^{(2)})$. Specifically, the discrete transition probabilities are constructed so that, in the limit $\Delta t^{\delta,h_1,h_2}_{t,x_1,x_2} \to 0$, they match the drift and diffusion characteristics of the continuous model under drift (smooth) ambiguity, as described in Remark 23.

Consequently, our MCA-based numerical scheme is consistent with the original stochastic singular control problem with smooth ambiguity, in the sense that the value functions and optimal policies obtained from the discretized model converge to those of the continuous problem as $\delta, h_1, h_2 \to 0$.

In the next section, we apply this numerical scheme to perform a comparative statics analysis.

5.8 NUMERICAL COMPARATIVE STATIC ANALYSIS

Now, we illustrate numerical results, obtained from the MCA to investigate the impact of ambiguity levels as well as other model parameters on the optimal control policy. Unless otherwise state, we use

$$\rho = 0.2,\ \alpha = 0,\ \sigma = 0.2, s = 0.2,\ \check{c} = \hat{c} = 2,\ \ell = u = 4,\ ,\tau = 0,\ T = 10,\ \text{and}\ \gamma = 5,$$

as baseline parameters. Suppose that the controlled inventory domain is $(\underline{x}, \overline{x}) = [-10, 10]$ and $(\underline{m}, \overline{m})$ represents the $100(1-d)\%$ confidence interval of M_T. That is,

$$(\underline{m}, \overline{m}) = \left(m - z_{\frac{d}{2}} \sqrt{\frac{s^2 T}{1 + sT}}, m + z_{\frac{d}{2}} \sqrt{\frac{s^2 T}{1 + sT}} \right),$$

where $z_{\frac{d}{2}}$ is the standard normal quantile function. For simplicity, we assume henceforth that the decision maker operates with a 99% (i.e., $z_{0.005} \approx 2.326$) over the zero mean reference prior, i.e., $m = 0$.

Therefore, under the baseline parameters, the boundary of the auxiliary process $X^{(2),n}$ is

$$(\underline{x}_2, \overline{x}_2) = \left(\overline{x} - \frac{\sigma}{S_T} \underline{m},\ \underline{x} - \frac{\sigma}{S_T} \overline{m} \right) \approx (-12.55, 12.55).$$

which can be subject to changes depending on the variation of parameters imposed in the following comparative static analysis.

5.8.1 GENERAL CHARACTERISTIC OF THE CONTROL POLICY

From Figure 5.2 and the dynamics (5.6.1), one observes that the inventory levels $X^{(1)}$ minimizing the value function for each auxiliary inventory level $X^{(2)}$ occur where $\alpha^{(1)}(t, X^{(1)}, X^{(2)}) = 0$, with the global minimum attained at $X^{(1)} = 0$. These are precisely the states where

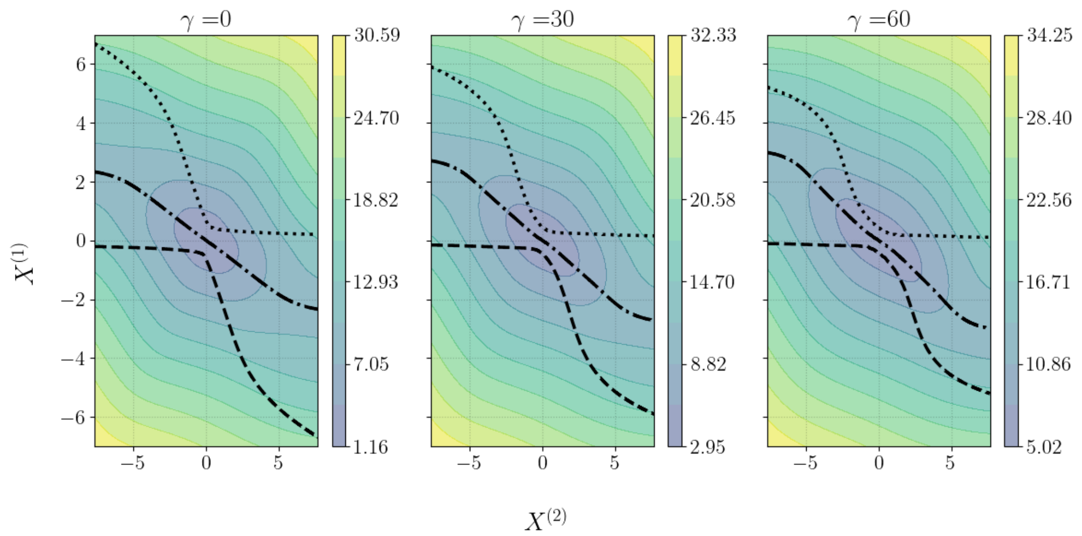

Figure 5.2: Contour plots of the value functions under the baseline parameters with ambiguity aversion $\gamma \in \{0, 30, 60\}$. The *dashed* and *dotted* lines determine lower and upper control barriers, respectively. A area between this two lines is the *inaction region*. The *dash-dotted* lines at the middle indicate of each figure the minimum inventory levels of the value function for each $X^{(2)}$.

the inventory exhibits no tendency to deviate from zero, and where the running cost is minimized. We call the set of $(X^{(1)}, X^{(2)})$ that satisfies this condition the *target lines*, which in particular, gives a minimum value function when the instantaneous holding cost satisfies $f(X^{(1)}) = 0$. Consequently, the optimal control policy is to intervene in the inventory dynamics with maximal effort to maintain the system near these states. Observe that the narrowest inaction region, as illustrated in Figure 5.2, is where $f(X^{(1)}) = 0$ and it is irrespective of the ambiguity attitude. We refer to this region as the *target region*, where the DM aims to keep the system in order to minimize inventory-related costs.

As $X^{(2)}$ deviates from the target region, the inaction region widens asymmetrically. From (5.6.1), fixing $X^{(2)}$ yields that $X^{(1)}$ evolves like an OrnsteinUhlenbeck (OU) process with (temporary) mean $\alpha + SX^{(2)}$ and mean reversion speed proportional to $X^{(2)}$. In our baseline model (i.e., $\alpha = 0$), negative values of $X^{(2)}$ imply that $X^{(1)}$ tends to revert to a negative inventory level. This leads to persistently higher lower holding costs relative to the case where $X^{(2)}$ is positive. Therefore, the optimal control policy is to promptly apply the lower control to avoid the downside risk, while simultaneously deferring the upper control. This delay ensures that if $X^{(1)}$ enters the upper inaction region, it lowers the likelihood of inventory shortfall and reduces the frequency of exerting the lower barrier.

Moreover, as $X^{(2)}$ becomes increasingly negative, the likelihood of maintaining a negative inventory, and thus incurring further losses increases. In response to this, it becomes optimal to further delay upper interventions to minimize such losses. This behavior results in an increasingly wider inaction region as $X^{(2)}$ becomes more negative for any $\gamma \in \{0, 30, 60\}$. A similar feature applies when in the opposite manner when $X^{(2)}$ becomes increasingly positive as shown by Figure 5.2.

Now we investigate the shape of these regions in the light of ambiguity.

5.8.2 COMPARATIVE STATICS OF AMBIGUITY ATTITUDE

As is evident from Figure 5.2, it shows that ambiguity aversion reduces the size of the inaction region. In other words, a more ambiguity-averse DM tends to intervene sooner, on average. Furthermore, the contour levels in Figure (5.2) indicate that increased ambiguity aversion leads to a higher value function. Taken together, this suggests that a more ambiguity-averse DM is expected to intervene more frequently, thereby incurring a higher average control cost.

From a technical perspective, the first and second lines of (5.7.2) imply that ambiguity amplifies perceived trends: the inventory level $X^{(1)}$ rises (or falls) more rapidly in response to upward (or downward) shifts, reflecting a worst-case posterior adjustment in estimating the model parameter θ under the reference measure.[8]

In what follows, we investigate the comparative statics of risk under different levels of ambiguity.

5.8.3 COMPARATIVE STATICS OF RISK

We begin by examining the optimal control policy under increasing levels of risk in the absence of ambiguity, as illustrated in Figure 5.3a. One observes that the control barriers expand as the level of risk increases. This is because higher risk amplifies the probability of extreme inventory fluctuations, thereby increasing the expected frequency of control interventions. As a result, it becomes more advantageous to delay these interventions. This observation is consistent with well-established results in the literature regarding the impact of risk on control barriers; see, for example, the seminal works of Archankul et al. (2025), Dixit and Pindyck (1994), Ferrari and Rodosthenous (2020), Ferrari and Vargiolu (2020), and Matomäki (2012).

However, complications arise when ambiguity aversion is present, particularly under a smooth ambiguity preference with a Gaussian-generated prior. On the one hand, when the confidence interval of $X^{(2)}$ is narrow (i.e., near the target region), the expansion effect of risk persists. On the other hand, this pattern may break down as the confidence interval widens, potentially giving rise to non-monotonic behavior.

As shown in Figure 5.3, for any level of risk, the control barriers shrink as ambiguity aversion γ increases, in line with the earlier comparative statics. For small values of γ, the expansion effect of risk remains visible, as in Figures 5.3a and 5.3b. However, when γ becomes sufficiently

[8]This trend is consistent with the findings of Archankul et al. (2025) in the case of maxmin utility, further validating the singular control decision-making under ambiguity.

large (e.g., $\gamma = 3$ or 5), the control barriers contract at an increasingly steep rate, proportional to the level of risk, especially when $X^{(2)}$ lies far from the target region. This contraction is strong enough to offset the expansion effect induced by risk. In other words, there exists a threshold at which ambiguity aversion begins to dominate the effect of risk on the control policy, and vice versa.

To understand the mechanism behind this unusual result, recall from Subsection 5.8.1 that when $X^{(1)} > X^{(2)}$ (resp. $X^{(1)} < X^{(2)}$), the inventory trend $\alpha_g^{(1)}(t, X^{(1)}, X^{(2)}) < 0$ (resp. > 0), and the control barriers widen to allow $X^{(1)}$ to explore further into the positive (resp. negative) region. This helps mitigate losses associated with holding a shortage (resp. surplus) of inventory.

In the presence of smooth ambiguity adjustment, we observe from the transition probability (5.7.2) that inventory trends are directly influenced not only by the ambiguity aversion parameter γ, but also by the level of risk σ and the absolute magnitude of the marginal value function. Since the marginal value function is strictly increasing, by Proposition 14, increasing (resp. decreasing) $X^{(1)}$ from the target line such that the inventory trend becomes positive (resp. negative) further amplifies this upward (resp. downward) trend. This amplification is aggravated at higher levels of σ. That is, greater diffusion increases the likelihood of $X^{(1)}$ deviating from its current state, which in turn raises the marginal value function and intensifies the inventory trend. As a result, when this scenario happens, the optimal control policy prompts the DM to act earlier in order to avoid potentially higher holding costs.

In summary, while the DM must account for trend uncertainty through learning, particularly given the potential for adverse inventory dynamics (i.e., excessively high or low trends), unlike the standard model, it is not always optimal to *wait* when the risk is perceptibly high. In some cases, delaying actions may themselves lead to even worse outcomes. In other words, under conditions of high volatility and limited information, it is optimal for an ambiguity-averse DM (within our model setting) to behave with extreme caution and pessimism.

In what follows, we study the comparative static of the remaining parameters under the baseline model, starting from the observation times.

5.8.4 COMPARATIVE STATICS OF OBSERVATION TIME

Figure 5.4 illustrates the shapes of the control barriers for observation times $\tau \triangleq T - t \in \{6, 8, 10\}$ under the baseline parameters. According to (5.6.5), the process $X^{(1)} - X^{(2)}$ exhibits a repelling tendency towards positive values when the drift is positive (i.e., when $X^{(1)} > X^{(2)}$). This repulsion intensifies when S_t is large, which corresponds to shorter observation periods (i.e., smaller τ).

In such cases, when $X^{(2)}$ is sufficiently negative, signalling an upward trend in (5.6.5), i.e., $S_t(X^{(1)} - X^{(2)}) > 0$, the process $X^{(1)} - X^{(2)}$ tends to accelerate further in the positive direction, pushing the inventory level $X^{(1)}$ upward. This leads to increased upper holding costs. As a result,

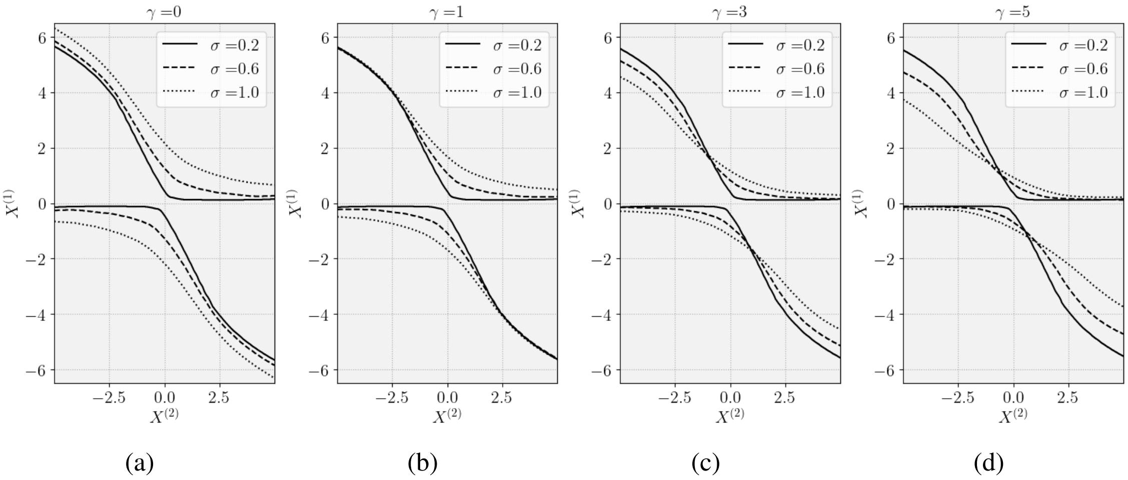

Figure 5.3: Panels 5.3a to 5.3d display the optimal control barriers for risk levels $\sigma \in \{0.2, 0.6, 1.0\}$ under the baseline parameter setting, with ambiguity aversion levels $\gamma = 0, 1, 3,$ and 5, respectively.

it becomes optimal to apply the upper control earlier. To mitigate the risk of $X^{(1)}$ reverting to costly positive levels, the optimal policy is to postpone the lower control once the process $X^{(1)}$ entering the negative region. Overall, this leads to an upward shift in the control barriers as shown on the left side of Figure 5.4. A similar reasoning applies when $X^{(1)} < X^{(2)}$, leading instead to a downward shift in the control barriers, as demonstrated in right side of the same figure.

Interpreted through the lens of the KalmanBucy filtering problem, shorter observation horizons (i.e., small τ) correspond to high levels of uncertainty (noise), which induce instability in inventory dynamics and increase the likelihood of extreme deviations. Consequently, it is optimal to apply upper or lower controls earlier when the inferred trend becomes excessively high or low, respectively, to avoid further divergence.

Conversely, when the observation period is sufficiently long, the noise is filtered out and the estimated trend becomes more reliable. In such a case, the DM faces lower uncertainty regarding future trajectories, and thus the inaction region broadens. This reflects a reduced need for precautionary intervention under more stable and predictable conditions.

5.8.5 COMPARATIVE STATICS OF BELIEF VARIANCE

The corresponding control barriers under the baseline parameters for belief variance $s \in 0.2, 0.25, 3$ are shown in Figure 5.5. When $s = 0$, the process M is constant, implying the absence of both ambiguity and learning. In this case, the control barrier corresponds to the classical singular control problem for the inventory process (5.2.2), with drift $\alpha - \sigma m$ and diffusion σ, where $m \in (\underline{m}, \overline{m})$. See, for instance, Archankul et al. (2025, Proposition 4) for the corresponding analytical comparative statics.

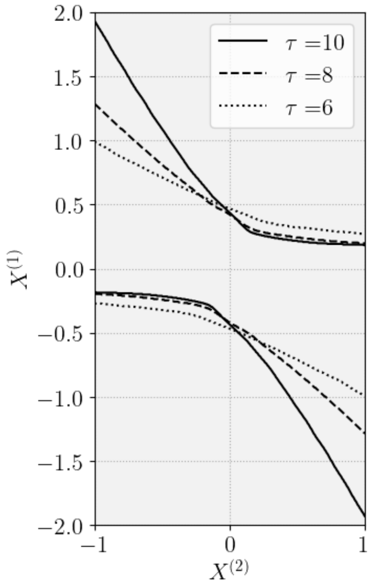

Figure 5.4: Optimal control policies of observation times $\tau \in \{6, 8, 10\}$.

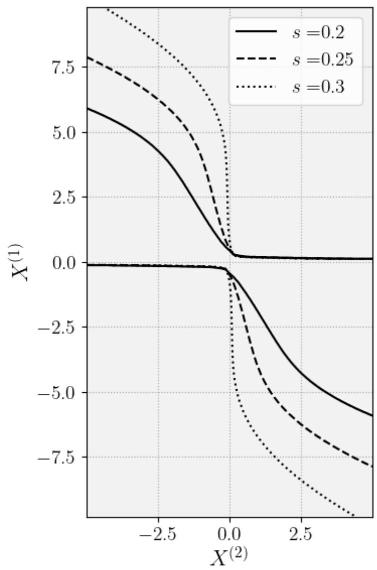

Figure 5.5: Optimal control policies of belief variance $s \in \{0.2, 0.25, 0.3\}$.

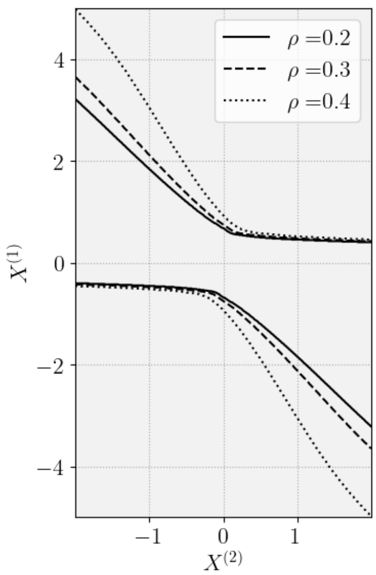

Figure 5.6: Optimal control policies of discounted rates $\rho \in \{0.2, 0.3, 0.4\}$.

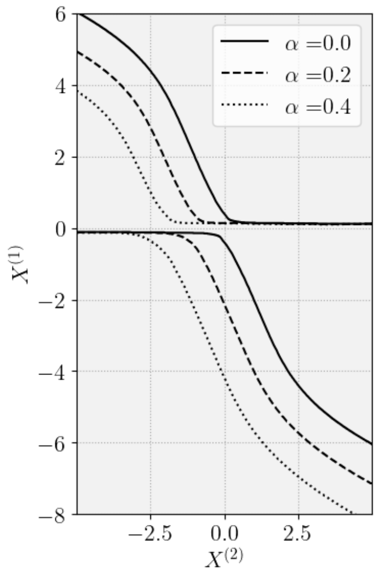

Figure 5.7: Optimal control policies of drifts $\alpha \in \{0, 0.2, 0.3\}$.

As the belief variance s increases, we observe that interventions at the upper (resp. lower) control barrier are postponed when $X^{(2)}$ induces a negative (resp. positive) trend. This behavior can be explained similarly to the mechanism discussed in Subsection 5.8.1. Specifically, since $s \mapsto S$ is strictly increasing, a higher s magnifies both the absolute (temporary) mean and the reversion

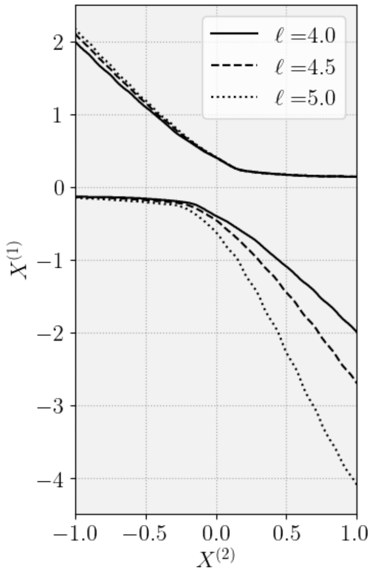

Figure 5.8: Optimal control policies of lower control costs $\ell \in \{4, 4.5, 5\}$.

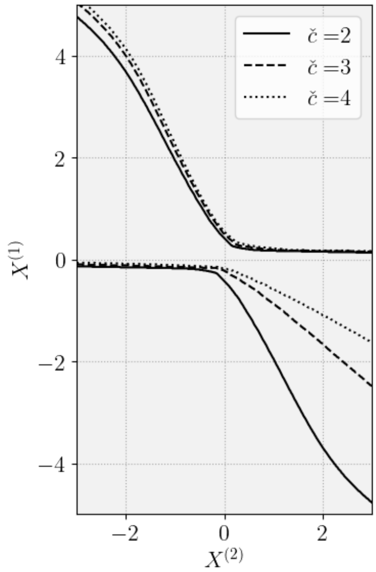

Figure 5.9: Optimal control policies of lower holding costs $\check{c} \in \{2, 3, 4\}$.

speed of the OU-interpreted process $X^{(1)}$. In other words, increasing the belief variance raises the likelihood of extreme inventory trends, either excessively high or excessively low. Consequently, the decision-maker has an incentive to delay applying the upper (resp. lower) control when the trend is strongly negative (resp. positive).

5.8.6 COMPARATIVE STATICS OF DISCOUNTED RATE

The comparative statics of the discount rate are illustrated in Figure 5.6, in line with the baseline parameters where the discount rate is varied by $\rho \in \{0.2, 0.3, 0.4\}$. As the figure shows, the control barriers expand as the discount rate increases. This occurs because a higher discount rate reduces the present value of future inventory holding costs, making it relatively less costly to delay intervention. As a result, the DM has a greater incentive to wait, knowing that postponing control is, in relative terms, more cost-effective than acting immediately under a lower discount rate.

5.8.7 COMPARATIVE STATICS OF DRIFT

We now investigate the optimal control policy under the baseline parameters with varying drift coefficients $\alpha \in \{0, 0.2, 0.3\}$. As shown in Figure 5.7, a higher drift increases the likelihood that the inventory process incurs the upper holding cost. To mitigate this, the DM should intervene earlier on the upper barrier, while delaying actions at the lower barrier. This allows the inventory more time to potentially return to the upper region, avoiding unnecessary intervention on the upper side.

Furthermore, an increase in α causes the target region to shift to the left side, as previously discussed in Section 5.8.1, since that is where $\alpha_g^{(1)}(t, X^{(1)}, X^{(2)}) = 0$. Together, these effects lead to a shift to the left of the entire control barrier structure, as illustrated in Figure 5.7. The case in which α decreases from 0 can be described similarly, resulting in a right shift of the control barriers.

5.8.8 COMPARATIVE STATICS OF CONTROL COST

We now analyze the shape of the control barriers under varying lower control costs, with $\ell \in \{2, 3, 4\}$ and under the baseline parameters. The case of higher upper control costs can be interpreted analogously. Intuitively, an increase in the lower control cost raises the expected running cost. As a result, it becomes optimal to delay intervention at the lower barrier. Moreover, since exercising control on the lower side is now relatively more costly than on the upper side, the DM should also defer action on the upper barrier to reduce the likelihood of the inventory drifting into the lower region. This interpretation is confirmed by the patterns observed in Figure 5.8.

5.8.9 COMPARATIVE STATICS OF HOLDING COST

Figure 5.9 displays the control barrier configurations for varying lower holding costs, with $\check{c} \in \{1, 2, 3\}$ under the baseline parameters. As the lower holding cost increases, it becomes more expensive to maintain a negative inventory level. Consequently, the DM should exert the lower control earlier to avoid incurring higher holding costs. Simultaneously, the DM should delay action on the upper control barrier, which is associated with a lower cost. Once the inventory reaches this upper region, not only is the holding cost relatively cheaper, but the probability of the inventory drifting into the costly negative region also decreases, thereby reducing the expected holding cost overall.

5.9 CONCLUSION

In this paper, we explore the application of singular control in inventory management under smooth ambiguity preference, where the model involves an unobservable parameter assumed to follow a Gaussian distribution. We establish a connection between the value function under smooth ambiguity and a forward-backward stochastic differential equation with quadratic growth. A verification theorem is provided in terms of the Hamilton-Jacobi-Bellman (HJB) equation to derive the optimal control policy. We then apply the Markov chains approximation based on the coordinate transform to provide an efficient numerical method that solve the HJB equation. Then we conduct comparative statics to analyze the effects of smooth ambiguity on the optimal control policy, assuming the inventory level follows an arithmetic Brownian motion.

Within this framework, one of our crucial findings is that when ambiguity is low, inferring that the observed inventory trend remains close to the reference trend, the optimal control policy aligns with the standard model: higher risk delays action (cf. Dixit and Pindyck, 1994), whereas greater ambiguity prompts earlier intervention (cf. Nishimura and Ozaki, 2007), regardless of the DM's ambiguity attitude. However, the opposite phenomenon holds when ambiguity is high, i.e., when the inventory trend deviates sharply from its initial estimate and the DM is strongly ambiguity averse. In such cases, increasing risk reverses the standard logic: instead of delaying action, it accelerates it. This inversion highlights a critical insight, when the environment is extremely volatile and the demand or supply trend is highly uncertain, an ambiguity-averse DM should act more promptly and conservatively, adopting a more pessimistic attitude in managing inventory.

It is important to highlight that while the treatment of smooth ambiguity in continuous time is based on the framework of singular control, our theory can be extended to a broader range of problems in stochastic control, provided the HJB equations are applicable. For instance, impulse control represents a direct extension of singular control, where actions cause not only continuous reflection but also jumps of a certain magnitude. This issue becomes particularly relevant when there are fixed and proportional costs associated with inventory intervention. Another area that could consider ambiguity, aside from barrier control, is, for example, the McDonald-Siegel-type model, which presents a stylized optimal stopping problem in investment timing decisions. There is extensive literature on these topics, and we refer to Harrison (2013) and Pham (2009) for a comprehensive overview.

— 6 —

Concluding Remarks and Future Research

This thesis presents a comprehensive study of singular stochastic control under ambiguity, extending the traditional risk-based framework of real options to contexts in which probabilities are uncertain. Across three chapters, the analysis traced a progression from restrictive to more flexible models of ambiguity, beginning with maxmin utility and advancing to smooth ambiguity with both finite- and continuous-state hidden variables. The unifying feature of all models is their focus on inventory-type problems: cash, water, and inventory management, where control is exercised through threshold-based interventions. These problems naturally lead to nonlinear expectations that cannot be addressed within classical, risk-based frameworks. To resolve them, the analysis relies on suitable FBSDEs and their associated HJB formulations.

The first chapter established a conservative benchmark by embedding maxmin preferences with κ-ignorance into a cash management model. The second chapter enriched the framework by introducing smooth ambiguity with finite-state beliefs in a reservoir management setting, and developed new numerical tools for computing viscosity solutions of singular-control HJB equations. The third chapter further broadened the scope by considering Gaussian hidden variables in inventory management, linking smooth ambiguity to quadratic FBSDEs and uncovering novel reversals in comparative statics. Taken together, these contributions extend both the theory of singular control and its methodological toolbox, while also generating insights into how ambiguity reshapes optimal policies in practice. In summary, the models developed here quantify the additional strategic costs required to manage uncertainty in environments where risk-based frameworks are insufficient.

Several avenues naturally emerge from this work. First, while the models presented are theoretical and rigorously analyzed, they provide a foundation for empirical validation. The cash management model could be calibrated against corporate finance data on cash holdings, the reservoir model against hydrological and climate series, and the inventory model against supply chain datasets. Such empirical studies would test the predictive power of ambiguity preferences and shed light on their relevance for managerial behavior.

Second, on the theoretical side, further extensions are possible. Beyond singular control, impulse control frameworks could be developed to capture fixed adjustment costs. More complex ambiguity structures, such as multi-state Wonham filters or Gaussian priors with regime-switching, would allow richer dynamics at the cost of computational complexity. Another promising direction is to investigate singular control with smooth ambiguity through the lens of Dynkin games. To the best of our knowledge, the complexity of smooth ambiguity currently prevents such a formulation. Overcoming this obstacle would enable a more rigorous analysis of optimal policies and broaden the range of parameters suitable for empirical validation.

Finally, this work points to policy and managerial implications. Ambiguity aversion consistently leads to earlier, more cautious interventions, but its quantitative impact depends on the specification of beliefs. Understanding how firms, regulators, or policymakers perceive and act upon ambiguity remains an open empirical question with significant economic importance.

In summary, the thesis provides a rigorous theoretical foundation for singular stochastic control under ambiguity, offers new analytical and numerical tools, and opens the door to both empirical validation and further theoretical generalizations.

References

Aliprantis, C. D., & Border, K. C. (2006, May). *Infinite dimensional analysis: A hitchhiker's guide* (3rd ed.). Springer.

Alvarez, L. (2003). On the properties of r-excessive mappings for a class of diffusions. *Annals of Applied Probability*, *13*, 1517–1533.

Anderson, E. W., Hansen, L. P., & Sargent, T. J. (2003). A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, *1*(1), 68–123.

Archankul, A., Ferrari, G., Hellmann, T., & Thijssen, J. J. (2025). Singular control in a cash management model with ambiguity. *European Journal of Operational Research*, *327*(2), 500–514.

Arrow, K. J., Harris, T., & Marschak, J. (1951). Optimal inventory policy. *Econometrica*, *19*(3), 250–272.

Asano, T., & Osaki, Y. (2021). Optimal investment under ambiguous technology shocks. *European Journal of Operational Research*, *293*(1), 304–311.

Balter, A. G., Mahayni, A., & Schweizer, Nikolaus. (2021). Time-consistency of optimal investment under smooth ambiguity. *European Journal of Operational Research*, *293*(2), 643–657.

Balter, A. G., & Pelsser, A. (2020). Pricing and hedging in incomplete markets with model uncertainty. *European Journal of Operational Research*, *282*(3), 911–925.

Bar-Ilan, A., & Sulem, A. (1995). Explicit solution of inventory problems with delivery lags. *Mathematics of Operations Research*, *20*(3), 709–720.

Barles, G., & Souganidis, P. E. (1991). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic analysis*, *4*(3), 271–283.

Basei, M., Ferrari, G., & Rodosthenous, N. (2024). Uncertainty over uncertainty in environmental policy adoption: Bayesian learning of unpredictable socioeconomic costs. *Journal of Economic Dynamics and Control*, *161*, 104841.

Bather, J. A. (1966). A continuous time inventory model. *Journal of Applied Probability*, *3*(2), 538–549.

Baumol, W. J. (1952). The transactions demand for cash: An inventory theoretic approach. *The Quarterly Journal of Economics*, *66*(4), 545–556.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society of London*, *53*, 370–418.

Bensoussan, A., Liu, R., & Sethi, S. P. (2005). Optimality of an (s,s) policy with compound poisson and diffusion demands: A quasi-variational inequalities approach. *SIAM journal on control and optimization*, *44*(5), 1650–1676.

Borgonovo, E., & Marinacci, M. (2015). Decision analysis under ambiguity. *European Journal of Operational Research*, *244*(3), 823–836.

Bregantini, D., Schmitt, L. H. M., & Thijssen, J. (2023). A bayesian change-point detection approach to the economic evaluation of risky projects : An application to health-care technology assessment. *Journal of the Royal Statistical Society, Series A: Statistic in Society*, 23.

Brekke, L. D. (2009). *Climate change and water resources management: A federal perspective*. Diane Publishing.

Breuer, W., Rieger, M. O., & Soypak, K. C. (2017). Corporate cash holdings and ambiguity aversion. *Review of Finance*, *21*(5), 1933–1974.

Brugnach, M., Ruiz, I., Zafra-Calvo, N., Vivas, L. D., & Sanz, M. J. (2025). Advancing decision-making amid uncertainty in water governance and management under climate change. *PLOS Climate*, *4*(5), e0000617.

Cadenillas, A., Lakner, P., & Pinedo, M. (2010). Optimal control of a mean-reverting inventory. *Operations Research*, *58*(6), 1697–1710.

Chakraborty, P., Cohen, A., & Young, V. R. (2023). Optimal dividends under model uncertainty. *SIAM Journal on Financial Mathematics*, *14*(2), 497–524.

Chen, Z., & Epstein, L. (2002). Ambiguity, risk, and asset returns in continuous time. *Econometrica*, *70*(4), 1403–1443.

Cheng, X., & Riedel, F. (2013). Optimal stopping under ambiguity in continuous time. *Mathematics and Financial Economics*, *7*(1), 29–68.

Constantinides, G. M. (1976). Stochastic cash management with fixed and proportional transaction costs. *Management Science*, *22*(12), 1320–1331.

Crandall, M. G., Ishii, H., & Lions, P.-L. (1987). Uniqueness of viscosity solutions of hamilton-jacobi equations revisited. *Journal of the Mathematical Society of Japan*, *39*(4), 581–596.

Crandall, M. G., & Lions, P.-L. (1983). Viscosity solutions of hamilton-jacobi equations. *Transactions of the American mathematical society*, *277*(1), 1–42.

Dai, J. G., & Yao, D. (2013a). Brownian inventory models with convex holding cost, part 1: Average-Optimal controls. *Stochastic Systems*, *3*(2), 442–499.

Dai, J. G., & Yao, D. (2013b). Brownian inventory models with convex holding cost, part 2: Discount-Optimal controls. *Stochastic Systems*, *3*(2), 500–573.

Dalby, P. A., Gillerhaugen, G. R., Hagspiel, V., Leth-Olsen, T., & Thijssen, J. J. (2018). Green investment under policy uncertainty and bayesian learning. *Energy*, *161*, 1262–1281.

De Angelis, T. (2020). Optimal dividends with partial information and stopping of a degenerate reflecting diffusion. *Finance and Stochastics*, *24*(1), 71–123.

Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton Univ. Press.

Driouchi, T., Trigeorgis, L., & So, R. H. Y. (2020). Individual antecedents of real options appraisal: The role of national culture and ambiguity. *European Journal of Operational Research*, *286*(3), 1018–1032.

Duffie, D., & Epstein, L. G. (1992). Stochastic differential utility. *Econometrica*, *60*(2), 353–394.

El Karoui, N., Kapoudjian, C., Pardoux, E., Peng, S., & Quenez, M. C. (1997). Reflected solutions of backward SDE's, and related obstacle problems for PDE's. *Annals of Probability*, *25*(2), 702–737.

El Karoui, N., Peng, S., & Quenez, M. (1997). Backward stochastic differential equations in finance. *Mathematical Finance*, *7*(1), 1–71.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, *75*, 643–669.

Eppen, G. D., & Fama, E. F. (1969). Cash balance and simple dynamic portfolio problems with proportional costs. *International Economic Review*, *10*(2), 119–133.

Epstein, L. G., & Ji, S. (2013). Ambiguous volatility and asset pricing in continuous time. *The Review of Financial Studies*, *26*(7), 1740–1786.

Epstein, L. G., & Wang, T. (1994). Intertemporal asset pricing under knightian uncertainty. *Econometrica*, *62*(2), 283–322.

Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, *57*(4), 937–969.

Federico, S., Ferrari, G., & Rodosthenous, N. (2023). Two-sided singular control of an inventory with unknown demand trend. *SIAM Journal on Control and Optimization*, *61*(5), 3076–3101.

Ferrari, G., Li, H., & Riedel, F. (2022). A Knightian irreversible investment problem. *Journal of Mathematical Analysis and Applications*, *507*(1), 125744.

Ferrari, G., & Rodosthenous, N. (2020). Optimal control of debt-to-gdp ratio in an n-state regime switching economy. *SIAM Journal on Control and Optimization*, *58*(2), 755–786.

Ferrari, G., & Vargiolu, T. (2020). On the singular control of exchange rates. *Annals of Operations Research*, *292*(2), 795–832.

Figuerola-Ferretti, I., Schwartz, E., & Segarra, I. (2024). Optimal operation of a hydropower plant in a stochastic environment. *Quantitative Finance*, *24*(5), 521–539.

Fleming, W. H., & Soner, H. M. (2006). *Controlled markov processes and viscosity solutions* (Vol. 25). Springer Science & Business Media.

Fletcher, S., Lickley, M., & Strzepek, K. (2019). Learning about climate change uncertainty enables flexible water infrastructure planning. *Nature communications*, *10*(1), 1782.

Fouque, J.-P., Pun, C. S., & Wong, H. Y. (2016). Portfolio optimization with ambiguous correlation and stochastic volatilities. *SIAM Journal on Control and Optimization*, *54*(5), 2309–2338.

Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, *18*(2), 141–153.

Gindrat, R., & Lefoll, J. (2011). Smooth ambiguity aversion and the continuous-time limit. *Available at SSRN 1690240*.

Goodell, J. W., Goyal, A., & Urquhart, A. (2021). Uncertainty of uncertainty and firm cash holdings. *Journal of Financial Stability*, *56*(100922), 100922.

Guo, X., & Tomecek, P. (2008). Connections between singular control and optimal switching. *SIAM Journal on Control and Optimization*, *47*(1), 421–443.

Hansen, L. P., & Sargent, T. J. (2010). Fragile beliefs and the price of uncertainty. *Quantitative Economics*, *1*(1), 129–162.

Hansen, L. P., & Miao, J. (2018). Aversion to ambiguity and model misspecification in dynamic stochastic environments. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(37), 9163–9168.

Hansen, L. P., & Miao, J. (2022). Asset pricing under smooth ambiguity in continuous time. *Economic Theory*, *74*(2), 335–371.

Hansen, L. P., & Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, *91*(2), 60–66.

Hansen, L. P., & Sargent, T. J. (2011). Robustness and ambiguity in continuous time. *Journal of Economic Theory*, *146*(3), 1195–1223.

Hansen, L. P., Sargent, T. J., Turmuhambetova, G., & Williams, N. (2006). Robust control and model misspecification. *Journal of Economic Theory*, *128*(1), 45–90.

Harrison, J. M. (1978). Optimal control of brownian storage system. *Stochastic Processes and Their Applications*, *6*, 179–194.

Harrison, J. M. (2013). *Brownian models of performance and control*. Cambridge University Press, Cambridge.

Harrison, J. M., & Taksar, M. I. (1983). Instantaneous control of brownian motion. *Mathematics of Operations research*, *8*(3), 439–453.

Hellmann, T., & Thijssen, J. J. J. (2018). Fear of the market or fear of the competitor? ambiguity in a real options game. *Operations Research*, *66*(6), 1744–1759.

Hu, Y., Imkeller, P., & Müller, M. (2005). Utility maximization in incomplete markets. *The Annals of Applied Probability*, *15*(3), 1691–1712.

Huang, L., Li, X., Fang, H., Yin, D., Si, Y., Wei, J., Liu, J., Hu, X., & Zhang, L. (2019). Balancing social, economic and ecological benefits of reservoir operation during the flood season: A case study of the three gorges project, china. *Journal of Hydrology*, *572*, 422–434.

Jeanblanc, M., Yor, M., & Chesney, M. (2009). *Mathematical methods for financial markets* (2009th ed.). Springer.

Jeanblanc-Picqué, M., & Shiryaev, A. N. (1995). Optimization of the flow of dividends. *Russian Mathematical Surveys*, *50*(2), 257–277.

Jensen, R. (1982). Adoption and diffusion of an innovation of uncertain profitability. *Journal of Economic Theory*, *27*(1), 182–193.

Jiang, H., Gibson, N. L., & Chen, Y. (2022). A stochastic model for the optimal allocation of hydropower flexibility in renewable energy markets. *Stochastic Models*, *38*(2), 288–307.

Jin, H., & Yu Zhou, X. (2015). Continuous-time portfolio selection under ambiguity. *Mathematical Control and Related Fields*, *5*(3), 475–488.

Johnson, P., & Peskir, G. (2017). Quickest detection problems for bessel processes.

Kalman, R. E., & Bucy, R. S. (1961). New results in linear filtering and prediction theory. *Journal of Basic Engineering*, *83*(1), 95–108.

Karatzas, I. (1983). A class of singular stochastic control problems. *Advances in Applied Probability*, *15*(2), 225–254.

Karatzas, I., & Shreve, S. (1991). *Brownian motion and stochastic calculus* (Vol. 113). Springer Science & Business Media.

Kelly, D. L., & Kolstad, C. D. (1999). Bayesian learning, growth, and pollution. *Journal of economic dynamics and control*, *23*(4), 491–518.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, *73*(6), 1849–1892.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2009). Recursive smooth ambiguity preferences. *Journal of Economic Theory*, *144*(3), 930–976.

Knight, F. H. (1921). *Risk, uncertainty and profit* (Vol. 31). Houghton Mifflin.

Kobylański, M., Lepeltier, J.-P., Quenez, M.-C., & Torres, S. (2002). Reflected bsde with superlinear quadratic coefficient. *Probability and Mathematical Statistics*, *22*.

Kobylański, M. (2000). Backward stochastic differential equations and partial differential equations with quadratic growth. *The Annals of Probability*, *28*(2), 558–602.

Kocabıyıkoğlu, A., Önkal, D., Göğüş, C. I., & Gönül, M. S. (2024). Newsvendor decisions under incomplete information: Behavioural experiments on information uncertainty. *IMA Journal of Management Mathematics*, *35*(3), 427–462.

Kumar, S., & Muthuraman, K. (2004). A Numerical Method for Solving Singular Stochastic Control Problems. *Operations research*, *52*(4), 563–582.

Kushner, H. J., & Martins, L. F. (1991). Numerical methods for stochastic singular control problems. *SIAM Journal on Control and Optimization*, *29*(6), 1443–1475.

Kushner, H., & Dupuis, P. (2001). *Numerical methods for stochastic control problems in continuous time*. Springer.

Kwon, H. D., & Lippman, S. A. (2011). Acquisition of project-specific assets with bayesian updating. *Operations research*, *59*(5), 1119–1130.

Laplace, P.-S. (1812). *Théorie analytique des probabilités*. Courcier.

Lin, Q., & Riedel, F. (2021). Optimal consumption and portfolio choice with ambiguous interest rates and volatility. *Economic Theory*, *71*(3), 1189–1202.

Liptser, R. S., & Shiryaev, A. N. (2013a). *Statistics of random processes I: General theory* (Vol. 5). Springer Science & Business Media.

Liptser, R. S., & Shiryaev, A. N. (2013b). *Statistics of random processes II: Applications* (Vol. 6). Springer Science & Business Media.

Lucia, J. J., & Schwartz, E. S. (2002). Electricity prices and power derivatives: Evidence from the nordic power exchange. *Review of derivatives research*, *5*(1), 5–50.

Luo, P., & Tian, Y. (2022). Investment, payout, and cash management under risk and ambiguity. *Journal of Banking & Finance*, *141*, 106551.

Ma, S., & Aloysius, J. (2022). Inventory control under different forms of uncertainty: Ambiguity and stochastic variability. *Decision Science.*

Maenhout, P. J. (2004). Robust portfolio rules and asset pricing. *Review of financial studies*, *17*(4), 951–983.

Mania, M., & Schweizer, M. (2005). Dynamic exponential utility indifference valuation. *The Annals of Applied Probability*, *15*(3), 2113–2143.

Matomäki, P. (2012). On solvability of a two-sided singular control problem. *Mathematical Methods of Operations Research*, *76*(3), 239–271.

Milly, P. C., Betancourt, J., Falkenmark, M., Hirsch, R. M., Kundzewicz, Z. W., Lettenmaier, D. P., & Stouffer, R. J. (2008). Stationarity is dead: Whither water management? *Science*, *319*(5863), 573–574.

Neamtiu, M., Shroff, N., White, H. D., & Williams, C. D. (2014). The impact of ambiguity on managerial investment and cash holdings. *Journal of Business Finance & Accounting*, *41*(7-8), 1071–1099.

Nishimura, K. G., & Ozaki, H. (2007). Irreversible investment and knightian uncertainty. *Journal of Economic Theory*, *136*(1), 668–694.

Niu, Y., & Shah, F. A. (2021). Economics of optimal reservoir capacity determination, sediment management, and dam decommissioning. *Water Resources Research*, *57*(7).

Null, S. E., Zeff, H., Mount, J., Gray, B., Sturrock, A. M., Sencan, G., Dybala, K., & Thompson, B. (2024). Storing and managing water for the environment is more efficient than mimicking natural flows. *Nature Communications*, *15*(1), 5462.

Øksendal, B. (2010). *Stochastic differential equations: An introduction with applications* (6th ed.). Springer Berlin, Heidelberg.

Øksendal, B., & Sulem, A. (2019). Stochastic differential games. In *Applied stochastic control of jump diffusions* (pp. 157–210). Springer International Publishing.

Pardoux, E., & Peng, S. (1990). Adapted solution of a backward stochastic differential equation. *Systems & Control Letters*, *14*(1), 55–61.

Pardoux, E., & Peng, S. (1992). Backward stochastic differential equations and quasilinear parabolic partial differential equations. *Stochastic Partial Differential Equations and Their Applications*, 200–217.

Peng, S., & Shi, Y. (2000). Infinite horizon forward–backward stochastic differential equations. *Stoch. Process. Their Appl.*, *85*(1), 75–92.

Pham, H. (2009). *Continuous-time stochastic control and optimization with financial applications* (Vol. 61). Springer Science & Business Media.

Prakash, O., Srinivasan, K., & Sudheer, K. (2015). Adaptive multi-objective simulation–optimization framework for dynamic flood control operation in a river–reservoir system. *Hydrology Research*, *46*(6), 893–911.

Protter, P. (2010). *Stochastic integration and differential equations* (2nd ed.). Springer.

Randle, T. J., Morris, G. L., Tullos, D. D., Weirich, F. H., Kondolf, G. M., Moriasi, D. N., Annandale, G. W., Fripp, J., Minear, J. T., & Wegner, D. L. (2021). Sustaining united states reservoir storage capacity: Need for a new paradigm.

Ryan, R., & Lippman, S. A. (2003). Optimal exit from a project with noisy returns. *Probability in the Engineering and Informational Sciences*, *17*(4), 435–458.

Schmutz, S., & Moog, O. (2018). Dams: Ecological impacts and management. In S. Schmutz & J. Sendzimir (Eds.), *Riverine ecosystem management: Science for governing towards a sustainable future* (pp. 111–127). Springer International Publishing.

Skiadas, C. (2013). Smooth ambiguity aversion toward small risks and Continuous-Time recursive utility. *Journal of Political Economy*, *121*(4), 775–792.

Suzuki, M. (2018). Continuous-time smooth ambiguity preferences. *Journal of Economic Dynamics and Control*, *90*, 30–44.

Taksar, M. I. (1985). Average optimal singular control and a related stopping problem. *Mathematics of Operations Research*, *10*(1), 63–81.

Tanaka, H. (1979). Stochastic differential equations with reflecting. *Stochastic Processes: Selected Papers of Hiroshi Tanaka*, *9*, 157.

Thijssen, J. J. J., & Bregantini, D. (2017). Costly sequential experimentation and project valuation with an application to health technology assessment. *Journal of Economic Dynamics and Control*.

Thijssen, J. J. J. (2011). Incomplete markets, ambiguity, and irreversible investment. *Journal of Economic Dynamics and Control*, *35*(6), 909–921.

Thijssen, J. J. J., Huisman, K. J., & Kort, P. M. (2006). The effects of information on strategic investment and welfare. *Economic Theory*, *28*(2), 399–424.

Tobin, J. (1956). The Interest-Elasticity of transactions demand for cash. *The Review of Economics and Statistics*, *38*(3), 241–247.

Trojanowska, M., & Kort, P. M. (2010). The worst case for real options. *Journal of Optimization Theory and Applications*, *146*(3), 709–734.

Vial, J.-P. (1972, January). *A continuous time model for the cash balance problem* (LIDAM Reprints CORE No. 127). Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Walker, W. E., Lempert, R. J., & Kwakkel, J. H. (2013). Deep uncertainty. In *Encyclopedia of operations research and management science* (pp. 395–402). Springer.

Wonham, W. M. (1965). Some applications of stochastic differential equations to optimal nonlinear filtering. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, *2*(3), 347–369.

World Commission on Dams. (2000). *Dams and development: A new framework for decision-making: The report of the world commission on dams*. Earthscan.

Yoshioka, H., & Yoshioka, Y. (2019). Modeling stochastic operation of reservoir under ambiguity with an emphasis on river management. *Optimal Control Applications and Methods*, *40*(4), 764–790.

Zhang, J. (2017). *Backward stochastic differential equations: From linear to fully nonlinear theory* (Vol. 86). Springer New York.

Zhao, Z., Liu, Y., & Pan, Q. (2023). Cash holdings, ambiguity aversion, and investment puzzles. *Economic Letter*, *229*(111192), 111192.