

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Methods for Optimal Stochastic Control and Optimal Stopping Problems Featuring Time-Inconsistency

Permalink

<https://escholarship.org/uc/item/56f5715d>

Author

Miller, Christopher Wells

Publication Date

2016

Peer reviewed|Thesis/dissertation

**Methods for Optimal Stochastic Control and Optimal Stopping Problems
Featuring Time-Inconsistency**

by

Christopher Wells Miller

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Applied Mathematics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Lawrence Craig Evans, Chair

Professor Fraydoun Rezakhanlou

Professor Claire Tomlin

Fall 2016

**Methods for Optimal Stochastic Control and Optimal Stopping Problems
Featuring Time-Inconsistency**

Copyright 2016
by
Christopher Wells Miller

Abstract

Methods for Optimal Stochastic Control and Optimal Stopping Problems Featuring
Time-Inconsistency

by

Christopher Wells Miller

Doctor of Philosophy in Applied Mathematics

University of California, Berkeley

Professor Lawrence Craig Evans, Chair

This thesis presents novel methods for computing optimal pre-commitment strategies in time-inconsistent optimal stochastic control and optimal stopping problems. We demonstrate how a time-inconsistent problem can often be re-written in terms of a sequential optimization problem involving the value function of a time-consistent optimal control problem in a higher-dimensional state-space. In particular, we obtain optimal pre-commitment strategies in a non-linear optimal stopping problem, in an optimal stochastic control problem involving conditional value-at-risk, and in an optimal stopping problem with a distribution constraint on the admissible stopping times. In each case, we relate the original problem to auxiliary time-consistent problems, the value functions of which may be characterized in terms of viscosity solutions of a Hamilton-Jacobi-Bellman equation.

Contents

Contents	i
List of Figures	ii
1 Introduction	1
1.1 Background	2
1.2 Optimal Control in Mathematical Finance	3
1.3 Overview of Time-Inconsistency	5
1.4 Outline of Results	6
2 A Time-Inconsistent Optimal Stopping Problem	8
2.1 Introduction	9
2.2 Equivalent Sequential Time-Consistent Problem	13
2.3 Properties of the Auxiliary Value Function	18
3 A Time-Inconsistent Optimal Stochastic Control Problem	29
3.1 Introduction	30
3.2 Equivalent Sequential Time-Consistent Problem	34
3.3 Properties of the Auxiliary Value Function	38
3.4 Application to Mean-CVaR Portfolio Optimization	47
4 Distribution-Constrained Optimal Stopping	53
4.1 Introduction	54
4.2 Main Results	56
4.3 Application to Superhedging with a Volatility Outlook	68
4.4 Proof of Lemmas 4.3–4.5	74
Bibliography	85

List of Figures

3.1	The efficient frontier in dynamic Mean-CVaR portfolio optimization.	49
3.2	The cumulative distribution function of an optimal dynamic strategy.	50
3.3	Sample path when following an optimal investment strategy (a).	50
3.4	Sample path when following an optimal investment strategy (b).	51
4.1	Comparison of model-free superhedging value to model-based prices (a).	71
4.2	Comparison of model-free superhedging value to model-based prices (b).	72
4.3	Probability density estimate of optimally-stopped terminal value.	73

Acknowledgments

This thesis is the culmination of many years of great mentorship and collaboration. I am forever indebted to my advisor, Craig Evans, whose encouragement and willingness to let me explore my various interests made my graduate school experience unforgettable.

I am fortunate to have had many informal mentors invest their time in me. I send my special thanks to Erhan Bayraktar, Peter Carr, Irene Gamba, John Griffin, Andrew Lesniewski, Mihai Sirbu, Peter Vinella, and Gordon Žitković. This thesis may not exist were it not for the many contributions by my wonderful co-author, Insoon Yang. Similarly, I would not be graduating were it not for the selfless efforts of Per-Olof Persson, Fraydoun Rezakhanlou, and Claire Tomlin serving on my qualifying exam and dissertation committees.

To my friends and family, I send my thanks for the constant encouragement. I particularly thank my parents for allowing me to explore my interest in mathematics from a young age. I cannot emphasize the credit that my fiancé, Megan, is due for supporting me over several years of laborious research.

It goes without saying that I am grateful to the National Science Foundation for their sponsorship through the Graduate Research Fellowship Program (Grant DGE 1106400). This research would not have been possible without this critical funding throughout my doctoral program.

Chapter 1

Introduction

1.1 Background

Stochastic analysis and partial differential equations (PDEs) are two broad sub-fields of mathematics whose interplay has proven fruitful in many financial and engineering applications. These tools provide a framework and calculus for modeling and understanding random processes and their related functions of interest with widespread application. In this thesis, we focus particularly on their use in dynamic optimization under uncertainty.

Optimal stochastic control deals with dynamic selection of inputs to a non-deterministic system with the goal of optimizing some pre-defined objective function. This is a natural extension of deterministic optimal control theory, but the introduction of uncertainty immediately opens countless applications in financial mathematics. An important sub-class of stochastic control is optimal stopping, where the user selects a time to perform a given action.

Historically, there have been two main approaches to solving optimal stochastic control problems – variational methods and Bellman’s dynamic programming principle [Bel52]. In a variational method, we generally obtain necessary conditions for an optimal control by focusing on small variations from the optimal point. In deterministic control, this leads to a system of ordinary differential equations via the celebrated Pontryagin maximum principle [PBG64]. The stochastic control analogue, which is often referred to as the stochastic maximum principle, leads to a system of forward-backward stochastic differential equations, which are often difficult to solve in practice.

While variational methods are simple to apply, their main drawbacks are the assumption that an optimal control exists and the difficulty in obtaining necessary conditions for global optimality. While these issues can be relaxed in problems with extra structure (e.g., convexity), many practical optimal control problems simply do not have an optimal control (instead, there exist maximizing sequences of controls).

In contrast, when applying Bellman’s dynamic programming principle (DPP) to an optimal control problem, we focus instead on the value function, which is defined as the supremum (or infimum) of the objective function over all admissible controls. This function is well-defined even without assuming existence of optimal controls. The DPP is then a technique which obtains a functional equation which encodes information about how to recursively compute the value function in terms of simpler sub-problems. In particular, this procedure assumes neither the existence of optimal controls nor substantial regularity of the value function.

The study of optimal control overlapped fruitfully with the theory of non-linear PDEs in the development of the notion of viscosity solutions by Crandall and Lions in the early 1980’s [CL83]. The theory of viscosity solutions provides a weak notion of solution to a PDE, which requires no assumed differentiability. More importantly, this theory provided a proof of uniqueness for many practical non-linear elliptic PDEs (see Ishii [Ish89], Jensen [Jen89], Crandall-Ishii [CI90], and Crandall-Ishii-Lions [CIL92]).

The remarkable connection between optimal stochastic control and PDEs is that the value function for a stochastic optimal control problem is, in fact, the unique viscosity solution of

an associated PDE, called the Hamilton-Jacobi-Bellman (HJB) equation (see Lions [Lio83a; Lio83b; Lio83c]). With this connection, it became possible to define the value function without any assumptions of existence of optimal controls, obtain a connection to PDEs via the DPP, obtain further regularity of the value function using PDE techniques, then, in many cases, use the further regularity to extract optimal controls for the original problem. The most powerful feature of this procedure is that it assumes essentially no special structure in the original optimal control problem (e.g., convexity).

1.2 Optimal Control in Mathematical Finance

The earliest applications of stochastic analysis to finance are generally attributed to the use of Brownian motion as a model of equity prices by Bachelier [Bac00]. It was not until later in the century that such models became mainstream, most notably through Merton's work on optimal portfolio selection [Mer69] and the acclaimed option pricing formulas of Black and Scholes [BS73].

In the following decades, applications of stochastic control and analysis in mathematical finance have becoming very wide-reaching. Examples of application areas include optimal portfolio selection (see Dumas-Luciano [DL91], Zhou-Li [ZL00], Li-Zhou-Lim [LZL02], Ou-Yang [OY03], Zhou-Yin [ZY03], Zhu [Zhu10], and Soner-Touzi [ST13]), option pricing under stochastic volatility models (see Heston [Hes93], Hagan-Kumar-Lesniewski-Woodward [HKLW02], and Fouque-Papanicolaou-Sircar-Sølna [FPSS11]), robust price bounds under model uncertainty (see Bonnans-Tan [BT13], Galichon-Henry-Labordère-Touzi [GHLT14], and Cox-Källblad [CK15]), models for price impact and illiquidity (see Çetin-Soner-Touzi [CST10], Gökay-Roch-Soner [GRS11], and Bayraktar-Ludkovski [BL14]), valuation of real options (see Trigeorgis [Tri95], Barrera-Esteve et al. [BEBD06], and Thompson-Davison-Rasmussen [TDR09]), as well as the dynamics and pricing of real-time electricity markets (see Humphrey-McClain [HM98], Rajagopal et al. [RBWV12], and Yang-Callaway-Tomlin [YCT14; YCT15]).

There have been two major trends in these applications of stochastic control in mathematical finance. First, there has been a transition from relying on linear dynamics and explicit formulas to allowing non-linear equations and the use of numerical methods. Second, we have seen a shift from simplistic models and objective functions (e.g. geometric Brownian motion with simple utility functions over terminal wealth) to more robust model specifications (e.g. stochastic volatility or super-replication problems) with various non-standard risk measures (e.g. value-at-risk, probability of ruin, et cetera).

As increasingly sophisticated applications of non-linear PDEs and stochastic control have permeated the mathematical finance literature, it is worth mentioning the parallel development of numerical schemes for approximating viscosity solutions. Because the value function of a stochastic control problem is not generally smooth, convergence results for traditional finite-difference schemes do not apply. Instead, it was demonstrated in the seminal paper of Barles and Souganidis [BS91] that so-called *monotone schemes* can be used to approximate

the viscosity solution of a PDE. This led the development of an array of numerical schemes for non-linear PDEs (e.g. Bonnans-Zidani [BZ03], Pooley-Forsyth-Vetzal [PFV03], Oberman [Obe08], Fahim-Touzi-Warin [FTW11], and Falcone-Ferretti [FF14]). Nevertheless, all currently-available general schemes suffer, in some form, from the so-called *curse of dimensionality*. Nevertheless, a promising literature in types of non-linear Monte Carlo methods has been making progress in relaxing these constraints (see McKean [McK75], Bouchard-Touzi [BT04], and Henry-Labordere et al. [HLOTTW16]).

In the spirit of the split between DPP and variational methods, many new models and applications in mathematical finance can be characterized by the use of Markov versus non-Markov dynamics. Early extensions of Black-Scholes dynamics included extra state variables for volatility (e.g., the Heston model, SABR, etcetera) and could still be investigated with PDE methods. However, as models became more complicated, with more state variables, the curse of dimensionality increasingly became an issue.

An alternative modeling framework which extends to non-Markov models is the use of backwards stochastic differential equations (BSDE), which naturally arise in variational approaches to stochastic control problems via the stochastic maximum principle (see Bismut [Bis73], Pardoux-Peng [PP90], and Peng [Pen90]). In principle, these techniques can model path-dependent payoffs and non-Markov dynamics without the introduction of extra state variables. Unfortunately, the most effective numerical solution of BSDE relies on the equivalence of certain BSDE to quasi-linear PDE, which reintroduces the curse of dimensionality (see Ma-Protter-Yong [MPY94]).

Recently, the BSDE modeling framework has been generalized to what are called second-order backwards stochastic differential equations (2BSDE) (see Cheridito-Soner-Touzi-Victoir [CSTV07] and Soner-Touzi-Zhang [STZ12]), which have been fruitfully applied to optimal stochastic control problems with uncertain volatility (see Nutz-Soner [NS12], Matoussi-Possamaï-Zhou [MPZ15], and Possamaï-Tan-Zhou [PTZ15]) and target constraints (see Soner-Touzi-Zhang [STZ+13] and Touzi [Tou10]), as well as to principal-agent problems (see Cvitanic-Possamaï-Touzi [CPT14; CPT15]). However, these are generally reduced to a corresponding fully non-linear PDE in some form for numerical solutions in practice.

Recently, there has been a push towards obtaining model-independent no-arbitrage price bounds for exotic derivatives, generally subject to prices of various vanilla derivatives which are assumed tradeable in the market. In an incomplete market with multiple risk-neutral measures, we generally define the super- (sub-)replication value of a derivative as the supremum (infimum) of the expected payoff over some collection of admissible pricing measures (see Avellaneda-Levy-Parás [ALP95] and Lyons [Lyo95]). In many circumstances, this is equivalent to the dual problem of the cheapest hedging portfolio which has a non-negative (non-positive) terminal payoff (see Possamaï et al. [PRT+13] and Acciaio et al. [ABPS13]).

Remarkably, there are many connections between super-replication value and optimal stochastic control problems (see Galichon-Henry-Labordere-Touzi [GHLT14] and Henry-Labordere et al. [HLOST+16]). For many derivatives with payoff depending on the realized-variance over some time period, the super-replication value can be related to a stopping time for a Brownian motion (see Bonnans-Tan [BT13] and Bayraktar-Miller [BM16]). There is

significant value and interest in computing model-free prices and hedging strategies for exotics, as the model-risk introduced by elaborate stochastic volatility models for pricing is in general large and often under-appreciated.

1.3 Overview of Time-Inconsistency

Time-inconsistent optimal stochastic control (and optimal stopping) problems are characterized by the failure of standard dynamic programming arguments to apply. Early study on these problems dates back to Strotz [Str55] and has developed into a sizeable literature within the economics community (see Tversky-Kahneman [TK85], Hoch-Loewenstein [HL91], Loewenstein-Prelec [LP92], Laibson [Lai97], O’Donoghue-Rabin [OR99], and Frederick et al. [FLO02]). From a mathematical perspective, we saw in the previous section that dynamic programming has been a powerful tool in solving many applied problems, so its failure and subsequent potential extension is certainly of interest.

Intuitively, time-inconsistency means that an optimal strategy today may not be optimal in the future. The most common way in which this is broken can be understood by interpreting the DPP as roughly saying “the optimal control does not depend upon the initial state.” There are three main ways in which time-inconsistency is often introduced in practice: the use of hyperbolic discounting in inter-temporal choice, an objective function featuring non-linear functions of an expected payoff, and direct dependence on the initial conditions (e.g. in endogenous habit formation models). For more detail on these three examples, we refer the reader to the excellent paper by Björk-Murgoci [BM14].

There are two common approaches to dealing with time-inconsistency in the literature. The first is known as solving for a “pre-commitment strategy,” and refers to solving the problem as stated at some initial time, assuming the optimizing agent has the ability to commit to a strategy he may later regret. The second is to reformulate the problem in game theoretic terms. Roughly, a dynamic optimization problem is viewed as a sequential game between your current self and your future self (who has potentially different preferences). While the latter has an extensive and interesting literature (see Ekeland-Lazrak [EL10], Hu-Jin-Jin [HJZ12], Yong [Yon12], and Björk-Murgoci [BM14] for example), in this thesis, we primarily take the former point of view as we examine optimization problems which may arise from specific applications which does not warrant an examination of the more behavioral aspects of the notion of solution.

From the perspective of applications, time-inconsistency often appears when solving optimal portfolio selection problems (see Zhou-Li [ZL00], Li-Zhou-Lim [LZL02], and Pedersen-Peskir [PP13]), with a classic example being dynamic mean-variance optimization. This is in direct contrast to early study of optimal portfolio selection via dynamic programming, which generally relied on exponential time-preference and specific objective functions based on terminal wealth. Once we move to objective functions based on realistic risk-measures, the problems quickly become time-inconsistent.

1.4 Outline of Results

The unifying theme of this thesis is that we can develop new methods for approaching time-inconsistent optimal stochastic control and optimal stopping problems by re-writing a time-inconsistent problem in terms of an iterated optimization problem involving the value functions of time-consistent problems. Rather than employ a game-theoretic interpretation of time-inconsistent problems, we aim to compute optimal pre-commitment strategies directly.

In Chapter 2 and Chapter 3, we consider the case of a time-inconsistent stochastic optimal stopping and optimal stochastic control problem, respectively. In each of these problems we are able to define an auxiliary value function corresponding to a time-consistent problem in a higher-dimensional state-space. We then demonstrate that we can obtain an optimal pre-commitment strategy by first solving an optimization problem over the starting value of the additional state, then computing an optimal strategy in the time-consistent problem. In Chapter 4, we consider a type of optimal stopping problem which features a constraint on the distribution of the stopping time. We demonstrate that this problem may be re-written as a sequence of iterated time-consistent optimal stochastic control problems.

In the following, we outline the specific results of each chapter.

Chapter 2 is based on Miller [Mil16], which presents a novel method for solving a class of time-inconsistent optimal stopping problems by reducing them to a family of standard stochastic optimal control problems.

In particular, we convert an optimal stopping problem with a non-linear function of the expected stopping time in the objective into optimization over an auxiliary value function for a standard stochastic control problem with an additional state variable. This approach differs from the previous literature which primarily employs Lagrange multiplier methods or relies on exact solutions. In contrast, we characterize the auxiliary value function as the unique viscosity solution of a non-linear elliptic PDE which satisfies certain growth constraints and investigate basic regularity properties. We demonstrate a connection between optimal stopping times for the original problem and optimal controls of the auxiliary control problem.

More broadly within the scope of this thesis, this chapter lays out a way of thinking about obtaining a pre-commitment solution of time-inconsistent problems by re-writing the problem as a sequence of optimization in extra state variables.

Chapter 3 is based on Miller-Yang [MY15]. In this chapter, we consider continuous-time stochastic optimal control problems featuring Conditional Value-at-Risk (CVaR) in the objective. Again, the major difficulty in these problems arises from time-inconsistency, which prevents us from directly using dynamic programming. To resolve this challenge, we convert to an equivalent bilevel optimization problem in which the inner optimization problem is standard stochastic control.

Furthermore, we provide conditions under which the outer objective function is convex and compute the outer objective's value via a Hamilton-Jacobi-Bellman equation. The key observation is that we can then solve the outer optimization problem via a gradient descent algorithm. The significance of this result is that we provide an efficient dynamic programming-based algorithm for optimal control of CVaR without lifting the state-space.

We refer the interested reader to Miller-Yang [MY15] for additional analysis of differentiability of the outer objective function and methods for computing its gradient.

Lastly, Chapter 4 is based Bayraktar-Miller [BM16], which considers optimal stopping featuring a novel distribution constraint. While there is no reason to expect this problem to be time-consistent, we convert it to an equivalent sequence of standard stochastic control problems.

In particular, we solve the problem of optimal stopping of a Brownian motion subject to the constraint that the stopping time's distribution is a given measure consisting of finitely-many atoms. We show that this problem can be converted to a finite sequence of state-constrained optimal control problems with additional states corresponding to the conditional probability of stopping at each possible terminal time. The proof of this correspondence relies on a new variation of the dynamic programming principle for state-constrained problems which avoids measurable selection. We emphasize that distribution constraints lead to novel and interesting mathematical problems on their own, but also demonstrate an application in mathematical finance to model-free superhedging with an outlook on volatility.

Chapter 2

A Time-Inconsistent Optimal Stopping Problem

2.1 Introduction

The following chapter is based upon the results in Miller [Mil16]. We present a simple example of a time-inconsistent optimal stopping problem which can be solved by converting to a maximization problem over the value function of a time-consistent optimal stochastic control problem in an additional state variable. More broadly within the scope of this thesis, we outline a way of thinking about obtaining a pre-commitment solution to a time-inconsistent problem by re-writing the problem as a sequence of optimization problems featuring additional state variables.

2.1.1 Overview of results

In this chapter, we consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which supports a standard Brownian motion W . We let $\mathbb{F} := \{\mathcal{F}_t\}_{t \geq 0}$ denote a filtration, which is assumed to be right-continuous and have all \mathbb{P} -negligible sets contained in \mathcal{F}_0 . We denote by \mathcal{T} the collection of all \mathbb{F} -stopping times such that $\mathbb{E}[\tau^2] < \infty$. We remind the reader that a random variable $\tau : \Omega \rightarrow \mathbb{R}^+$ is a \mathbb{F} -stopping time if $\{\tau \leq t\}$ is \mathcal{F}_t -measurable for all $t \geq 0$.

Let $f, g : \mathbb{R} \rightarrow \mathbb{R}$ be fixed continuous functions. In this chapter, we take f to be Lipschitz continuous, though this is extended to quadratic growth in Miller [Mil16]. For some fixed $x_0 \in \mathbb{R}$, we define the main problem considered in this chapter.

Definition 2.1. *The time-inconsistent optimal stopping problem is to compute*

$$p^* := \sup_{\tau \in \mathcal{T}} [\mathbb{E}[f(x_0 + W(\tau))] + g(\mathbb{E}[\tau])] \quad (2.1)$$

and to find a stopping time $\tau^* \in \mathcal{T}$ which attains the supremum.

At this point, it is unclear whether p^* is finite or τ^* exists without making additional assumptions on f and g . Necessary and sufficient conditions will be made clear throughout the chapter.

As usual, we define a corresponding value function.

Definition 2.2. *We define the value function for the time-inconsistent optimal stopping problem, $v : \mathbb{R} \rightarrow \mathbb{R}$, as*

$$v(x) := \sup_{\tau \in \mathcal{T}} [\mathbb{E}[f(x + W(\tau))] + g(\mathbb{E}[\tau])]$$

for each $x \in \mathbb{R}$.

Notice that $v(x_0) = p^*$.

Time-inconsistent optimization problems are characterized in general by the failure of standard dynamic programming arguments to apply to the value function v . In the class of problems considered in this chapter, the time-inconsistency stems from non-linearity of

g. The two common approaches to dealing with time-inconsistency in the literature are to reformulate the problem as a time-consistent problem, while possibly changing the value function, or to employ a “pre-commitment strategy” which may need to be recomputed for each new initial condition. The main result of this paper is to show that we can convert the computation of a pre-commitment strategy to a bi-level optimization problem whose lower-level consists of a time-consistent stochastic control problem.

In particular, let

$$\mathcal{A} := \left\{ \alpha : \Omega \times [0, \infty) \rightarrow \mathbb{R} \mid \alpha \text{ is progressively-measurable and } \mathbb{E} \int_0^\infty \alpha_t^2 dt < +\infty \right\}.$$

For each control $\alpha \in \mathcal{A}$, we consider the controlled stochastic differential equation

$$\begin{cases} dX_t^\alpha = dW_t \\ dY_t^\alpha = -dt + \alpha_t dW_t. \end{cases} \quad (2.2)$$

For any choice of $(x, y) \in \mathbb{R} \times [0, \infty)$ and $\alpha \in \mathcal{A}$, we write $\{(X_s^x, Y_s^{y,\alpha}) \mid s \geq 0\}$ to denote the solution of (2.2) with initial conditions $X_0^x = x$ and $Y_0^{y,\alpha} = y$. Furthermore, we write $\tau^{y,\alpha}$ to denote the stopping time

$$\tau^{y,\alpha} := \inf \{t \geq 0 \mid Y_t^{t,\alpha} = 0\}. \quad (2.3)$$

Next, we define the value function for an auxiliary control problem.

Definition 2.3. *Define the value function of an auxiliary stochastic control problem as*

$$w(x, y) := \sup_{\alpha \in \mathcal{A}} \mathbb{E} [f(X_{\tau^{y,\alpha}}^x)] \quad (2.4)$$

for every $(x, y) \in \mathbb{R} \times [0, \infty)$.

The main observation of this chapter is then that

$$p^* = \sup_{y \geq 0} [w(x_0, y) + g(y)]. \quad (2.5)$$

Furthermore, we will show that w can be characterized as the unique viscosity solution of the following Hamilton-Jacobi-Bellman (HJB) PDE:

$$\begin{cases} u_y - \sup_{a \in \mathbb{R}} \left[\frac{1}{2} u_{xx} + a u_{xy} + \frac{1}{2} a^2 u_{yy} \right] = 0 & \text{in } \mathbb{R} \times (0, \infty) \\ u = f & \text{on } \mathbb{R} \times \{y = 0\}. \end{cases} \quad (2.6)$$

Some amount of work will be put into making sense of (2.6) when the coefficient in front of the diffusion is unbounded¹. The upside to this approach is that one can then compute

¹In the original paper, we must additionally put particular asymptotic growth constraints on w to pin down uniqueness because we allow f to have super-linear growth. In this simplified presentation, the analysis is more straightforward.

w via standard methods, solve the optimization problem in (2.5), then obtain a solution to (2.1).

In the remainder of this chapter, we emphasize the following three-step strategy for dealing with time-inconsistency:

1. Condition on the time-inconsistent feature to obtain a constrained problem,
2. Embed the constrained problem in a time-consistent problem in a higher-dimensional state-space, and
3. Construct an optimal control of the original problem by starting at an optimal choice of the new state variables.

This procedure allows us to compute the value function of a pre-commitment strategy, and construct optimal stopping times under suitable regularity. While the first and last step are well-understood in the previous literature (See Pedersen-Peskir [PP16a; PP13]), the main contribution of this chapter is our second step, which is typically replaced by an application of Lagrange multipliers or an appeal to exact solutions.

2.1.2 Overview of previous literature

The particular type of time-inconsistency featured in this chapter is a non-linear function of the expected stopping time appearing in the objective function. Following the ideas developed in Pedersen-Peskir [PP16a], we condition on the expected value of the stopping time to obtain an expectation-constrained optimal stopping problem. In contrast with the previous literature, we embed the constrained optimal stopping problem into a time-consistent control problem with one extra state variable rather than employing the method of Lagrange multipliers. We characterize this auxiliary value function as the viscosity solution of a degenerate-elliptic Hamilton-Jacobi-Bellman (HJB) PDE subject to certain growth constraints.

The overarching idea of our approach is inspired by that reported in Pedersen-Peskir [PP16a]. In that paper, the authors solve a mean-variance stopping problem with a similar non-linearity by conditioning on the time-inconsistent feature and solving the resulting constrained stopping problem with free-boundary techniques and verification arguments along the lines of Peskir-Shiryaev [PS06]. In comparison, the reader may view the main contribution of this chapter as a novel solution of the expectation-constrained optimal stopping problem by embedding in a time-consistent stochastic control problem.

The investigation of constrained optimal stopping problems is not new, but most previous work focuses on Lagrange multiplier approaches to the constraint (See Kennedy [Ken82], López-San Miguel-Sanz [LSMS95], Horiguchi [Hor01], and Makasu [Mak09]). In contrast, we identify the expectation-constrained auxiliary value function as the unique viscosity solution of degenerate-elliptic Hamilton-Jacobi-Bellman equation, subject to certain growth constraints. The main advantage of our approach is that it depends neither on the specific

form of non-linearity in the problem, nor on the availability of analytic solutions. When analytic solutions are not available, a solution via the method of Lagrange multipliers generally relies on a numerical optimization of the Lagrange dual problem. While effective, this approach can be unstable in practice since we have no regularity estimates on the dual function apart from convexity, and computation of the sub-gradient is often subject to truncation error. In contrast, the non-linear elliptic PDE featured in this chapter has established numerical approximation schemes with guaranteed convergence and, more importantly, stability (See Oberman [Obe08]).

The ideas developed in this chapter can be extended to deal with other types of time-inconsistent features in optimal stopping and optimal stochastic control problems. We primarily emphasize the case of optimal stopping due to its relative technical and pedagogical simplicity. We briefly discuss extensions to diffusion processes and more general types of time-inconsistencies in the final section, albeit formally. The ideas are related to recent solutions of mean-variance portfolio optimization, optimal control under certain non-linear risk-measures, and distribution-constrained optimal stopping (See Pedersen-Peskir [PP13], Miller-Yang [MY15], and Bayraktar-Miller [BM16]).

Our approach is similar to the dynamic approach of Karnam-Ma-Zhang [KMZ16], wherein the authors introduce extra state variables to remove time-inconsistency introduced by a system of controlled backwards stochastic differential equations. Similar to this chapter, the authors convert a problem without an immediate dynamic nature to a dynamic problem with additional state variables. There appear to be additional analogies with the formal generalizations provided in Miller [Mil16]. In terms of a focus on time-inconsistent optimal stopping problems, this chapter is similar to Xu-Zhou [XZ13], where the authors consider a non-linear functional of a stopped process as the objective function. However, the method of solution differs entirely. Whereas these authors relate their problem to optimization over the distribution of the stopped process using Skorokhod embedding, we convert to a dynamic problem in a larger state-space. Since the original circulation of Miller [Mil16], the same elliptic PDE has been obtained independently by Ankirchner-Klein-Kruse [AKK15] in the direct analysis of a related expectation-constrained optimal stopping problem.

Some other notable works in the literature on time-inconsistent problems include Björk-Murgoci [BM14], Hu-Jin-Zhou [HJZ12], Yong [Yon12], and Ekeland-Lazrak [EL10]. Most of the previous literature focuses either on specific examples of time-inconsistency (often arising from non-exponential discounting or mean-variance optimization) or on notions of equilibrium strategies, which view time-inconsistent problems as a sequential game against one's future self. In an equilibrium strategy, the optimal equilibrium strategy can be characterized as the solution to an “extended HJB” system. In general, these systems feature multiple solutions and exhibit values strictly less than the value function of a pre-committed strategy. The price we must pay in our approach is that the entire value function must be recomputed if we change the initial conditions. This point is related to the notions of static and dynamic optimality which are explored in depth by Pedersen-Peskir [PP16a].

2.2 Equivalent Sequential Time-Consistent Problem

Our goal is to convert this time-inconsistent optimal stopping problem into a sequential optimization problem involving a time-consistent control problem. Our general approach to time-inconsistency is the following:

- Step 1: Condition on any time-inconsistent features in the problem to generate a family of optimal-stopping problems with constraints,
- Step 2: Enlarge the state-space to embed the constrained problems in a time-consistent problem, and
- Step 3: Construct an optimal stopping time for the time-inconsistent problem by picking an optimal value of the time-inconsistent feature and generating an optimal solution to the time-consistent problem starting from that choice.

In our particular problem, this will involve adding a new state variable to track the expectation of the optimal stopping time. As the system evolves, we expect this variable to be a super-martingale as the expected time until stopping drifts downward. However, it is possible to allow this expected stopping time to increase along certain paths so-long as it is compensated by a decrease along other paths.

2.2.1 Conditioning on time-inconsistent features

For any $y \in [0, \infty)$, consider the following subset of stopping times:

$$\mathcal{T}_y := \{\tau \in \mathcal{T} \mid \mathbb{E}[\tau] = y\} \subset \mathcal{T}.$$

Furthermore, consider the following family of expectation-constrained optimal stopping problems.

Definition 2.4. *Define the value function corresponding to an expectation-constrained optimal stopping problem as*

$$\tilde{w}(x, y) := \sup_{\tau \in \mathcal{T}_y} \mathbb{E}[f(x + W(\tau))] \tag{2.7}$$

for each $(x, y) \in \mathbb{R} \times [0, \infty)$.

We first claim that we can reformulate the time-inconsistent optimal stopping problem (2.1) as a sequential optimization problem involving these constrained optimal stopping problems.

Theorem 2.1 (Pedersen-Peskir [PP16a]). *For any $x \in \mathbb{R}$, we have*

$$v(x) = \sup_{y \geq 0} [\tilde{w}(x, y) + g(y)].$$

Proof. The key is to note that since all stopping times in \mathcal{T} are assumed to have finite variance, they also have finite expectation. Then

$$\mathcal{T} = \bigcup_{y \geq 0} \mathcal{T}_y.$$

Then it is simple to check that

$$\begin{aligned} v(x) &= \sup_{\tau \in \mathcal{T}} [\mathbb{E} [f(x + W(\tau))] + g(\mathbb{E} [\tau])] \\ &= \sup_{y \geq 0} \sup_{\tau \in \mathcal{T}_y} [\mathbb{E} [f(x + W(\tau))] + g(\mathbb{E} [\tau])] \\ &= \sup_{y \geq 0} \sup_{\tau \in \mathcal{T}_y} [\mathbb{E} [f(x + W(\tau))] + g(y)] \\ &= \sup_{y \geq 0} [\tilde{w}(x, y) + g(y)]. \end{aligned}$$

□

2.2.2 Equivalence with time-consistent optimal control

Next, we reformulate the constrained optimal stopping problem as a time-consistent stochastic control problem. The benefit of this will be that we can write down an HJB PDE associated with the time-consistent stochastic control problem. Let \mathcal{A} be the set of all real-valued, progressively-measurable, and square-integrable processes. We state a lemma which identifies $\tau \in \mathcal{T}_y$ with a control in \mathcal{A} .

Lemma 2.1. *Fix $y \in [0, \infty)$.*

1. *For any $\tau \in \mathcal{T}_y$, there exists $\alpha \in \mathcal{A}$ such that $\tau = \tau^{y, \alpha}$ almost-surely.*
2. *For any $\alpha \in \mathcal{A}$, we have $\tau^{y, \alpha} \in \mathcal{T}_y$.*

Proof. 1. Because $\tau \in \mathcal{T}_y$ is a square-integrable random variable with expectation y , there exists $\alpha \in \mathcal{A}$ such that

$$\tau = \mathbb{E} [\tau] + \int_0^\infty \alpha_s dW_s = y + \int_0^\infty \alpha_s dW_s,$$

almost-surely, by the Martingale Representation Theorem (See Section 3.4 in Karatzas-Shreve [KS91]). However, τ is \mathcal{F}_τ -measurable, so if we take conditional expectations and use the martingale property of the Ito integral, we see

$$\tau = y + \mathbb{E} \left[\int_0^\infty \alpha_s dW_s \mid \mathcal{F}_\tau \right] = y + \int_0^\tau \alpha_s dW_s, \tag{2.8}$$

almost-surely.

Recall the definition of $\tau^{y,\alpha}$ from (2.3). From (2.8), we immediately deduce $Y_\tau^{y,\alpha} = 0$ almost-surely. This implies that $\tau^{y,\alpha} \leq \tau$ almost-surely.

Next, we take expectations of (2.8) conditional upon $\mathcal{F}_{\tau^{y,\alpha}}$ and using the martingale property of the Ito integral and note that

$$\mathbb{E}[\tau \mid \mathcal{F}_{\tau^{y,\alpha}}] = y + \int_0^{\tau^{y,\alpha}} \alpha_s dW_s.$$

Subtracting $\tau^{y,\alpha}$ from both sides, we see

$$\mathbb{E}[\tau - \tau^{y,\alpha} \mid \mathcal{F}_{\tau^{y,\alpha}}] = Y_{\tau^{y,\alpha}}^{y,\alpha} = 0, \quad (2.9)$$

using the definition of $\tau^{y,\alpha}$.

Taking an unconditional equation of (2.9), we see $\mathbb{E}[\tau] = \mathbb{E}[\tau^{y,\alpha}]$. This, together with $\tau^{y,\alpha} \leq \tau$ almost-surely, implies that $\tau^{y,\alpha} = \tau$ almost-surely.

2. Fix $\alpha \in \mathcal{A}$ and recall the definition of $\tau^{y,\alpha}$ from (2.3). We first claim that $\tau^{y,\alpha} < +\infty$ almost-surely. To that end, we investigate the random variable $Y_t^{y,\alpha}$ for large t .

It is clear that $\mathbb{E}[Y_t^{y,\alpha}] = y - t$ and

$$\begin{aligned} \text{Var}[Y_t^{y,\alpha}] &= \mathbb{E}\left[(Y_t^{y,\alpha} - y + t)^2\right] \\ &= \mathbb{E}\left[\left(\int_0^t \alpha_s dW_s\right)^2\right] \\ &= \mathbb{E}\left[\int_0^t \alpha_s^2 ds\right] \end{aligned}$$

by Ito's Isometry (See Section 3.2 in Karatzas-Shreve [KS91]). But because $\alpha \in \mathcal{A}$, there exists $M > 0$ such that

$$\text{Var}[Y_t^{y,\alpha}] = \mathbb{E}\left[\int_0^t \alpha_s^2 ds\right] \leq \mathbb{E}\left[\int_0^\infty \alpha_s^2 ds\right] \leq M.$$

The goal is to bound the probability that $Y_t^{y,\alpha}$ is non-negative. We first compute

$$\begin{aligned} \mathbb{P}[Y_t^{y,\alpha} \geq 0] &\leq \mathbb{P}[|Y_t^{y,\alpha} - \mathbb{E}[Y_t^{y,\alpha}]| \geq y - t] \\ &\leq \mathbb{P}\left[|Y_t^{y,\alpha} - \mathbb{E}[Y_t^{y,\alpha}]| \geq (y - t)M^{-1/2}\sqrt{\text{Var}[Y_t^{y,\alpha}]}\right]. \end{aligned}$$

However, we can now apply Chebyshev's Inequality (See Section 1.6 in Durrett [Dur10]) to conclude

$$\mathbb{P}[Y_t^{y,\alpha} \geq 0] \leq M(t - y)^{-2}.$$

But taking $t \rightarrow \infty$ and noting that $Y_t^{y,\alpha} \leq 0$ implies $\tau^{y,\alpha} \leq t$, this inequality contradicts the claim that $\mathbb{P}[\tau^{y,\alpha} = +\infty] > 0$.

Then because $\tau^{y,\alpha} < +\infty$ almost-surely, we conclude by the definition of $\tau^{y,\alpha}$ that

$$y - \tau^{y,\alpha} + \int_0^{\tau^{y,\alpha}} \alpha_s dW_s = Y_{\tau^{y,\alpha}}^{y,\alpha} = 0, \quad (2.10)$$

almost-surely. Taking expectations on both sides of (2.10), we see $\mathbb{E}[\tau^{y,\alpha}] = y$. Similarly, by re-arranging and squaring both sides of (2.10) then taking expectations, we see

$$\mathbb{E}[(\tau^{y,\alpha})^2] = y^2 + \mathbb{E}\left[\int_0^{\tau^{y,\alpha}} \alpha_s^2 ds\right] \leq y^2 + \mathbb{E}\left[\int_0^\infty \alpha_s^2 ds\right] < +\infty.$$

Then $\tau^{y,\alpha} \in \mathcal{T}_y$. □

The key to the main result is to convert between stopping times in \mathcal{T}_y and controls in \mathcal{A} via Lemma 2.1 and instead view w as the value function of a stochastic optimal control problem.

Recall the value function for a stochastic control problem from (2.4).

Theorem 2.2. *We have the equivalence*

$$w(x, y) = \tilde{w}(x, y)$$

for any $(x, y) \in \mathbb{R} \times [0, \infty)$.

Proof. Fix any $(x, y) \in \mathbb{R} \times [0, \infty)$.

1. Let $\alpha \in \mathcal{A}$ be an arbitrary control. By Lemma 2.1, we know $\tau^{y,\alpha} \in \mathcal{T}_y$. Then

$$\begin{aligned} \mathbb{E}[f(X_{\tau^{y,\alpha}}^x)] &= \mathbb{E}[f(x + W_{\tau^{y,\alpha}})] \\ &\leq \tilde{w}(x, y). \end{aligned}$$

Because $\alpha \in \mathcal{A}$ was arbitrary, we conclude

$$w(x, y) \leq \tilde{w}(x, y).$$

2. Let $\tau \in \mathcal{T}_y$ be an arbitrary stopping time. By Lemma 2.1 there exists a control $\alpha \in \mathcal{A}$ such that $\tau^{y,\alpha} = \tau$ almost-surely. Then

$$\begin{aligned} \mathbb{E}[f(x + W_\tau)] &= \mathbb{E}[f(x + W_{\tau^{y,\alpha}})] \\ &= \mathbb{E}[f(X_{\tau^{y,\alpha}}^x)] \\ &\leq w(x, y). \end{aligned}$$

Because $\tau \in \mathcal{T}_y$ was arbitrary, we conclude

$$\tilde{w}(x, y) \leq w(x, y). \quad \square$$

In the remainder of the chapter, we will refer to the auxiliary value function w in terms of the stochastic control value function (2.4) or the expectation-constrained optimal stopping value function (2.7) as is convenient.

2.2.3 Construction of optimal stopping times

We have shown from Theorem 2.1 and Theorem 2.2 that we can recover p^* in the time-inconsistent optimal stopping problem (2.1) by maximizing over choice of $y \geq 0$ and control $\alpha \in \mathcal{A}$. However, it remains to be shown that we can construct a corresponding optimal stopping time for the time-inconsistent problem.

The next theorem relates nearly-optimal choices of $(y, \alpha) \in [0, \infty) \times \mathcal{A}$ with nearly-optimal choices of stopping times \mathcal{T} .

Theorem 2.3. *For any $\epsilon \geq 0$, let $y \geq 0$ satisfy*

$$v(x_0) \leq w(x_0, y) + \epsilon$$

and let $\alpha \in \mathcal{A}$ satisfy

$$w(x_0, y) \leq \mathbb{E}[f(X_{\tau^{y,\alpha}}^{x_0})] + \epsilon.$$

Then $\tau^{y,\alpha} \in \mathcal{T}_y \subset \mathcal{T}$ satisfies

$$p^* - 2\epsilon \leq \mathbb{E}[f(x_0 + W_{\tau^{y,\alpha}})] \leq p^*.$$

Proof. Recall from Definition 2.2 that $p^* = v(x_0)$. Then by combining this with the two assumed inequalities and the definition of $\{X_s^{x_0} \mid s \geq 0\}$, we have

$$\begin{aligned} p^* &= v(x_0) \\ &\leq w(x_0, y) + \epsilon \\ &\leq \mathbb{E}[f(X_{\tau^{y,\alpha}}^{x_0})] + 2\epsilon \\ &= \mathbb{E}[f(x_0 + W_{\tau^{y,\alpha}})] + 2\epsilon. \end{aligned}$$

Because $\tau^{y,\alpha} \in \mathcal{T}_y \subset \mathcal{T}$, we also have

$$\mathbb{E}[f(x_0 + W_{\tau^{y,\alpha}})] \leq p^*.$$

□

Then we can record the following corollary regarding obtaining an optimal control.

Corollary 2.1. *Let $y^* \geq 0$ satisfy*

$$w(x_0, y^*) = \max_{y \geq 0} w(x_0, y)$$

and let $\alpha^ \in \mathcal{A}$ satisfy*

$$\mathbb{E}[f(X_{\tau^{y^*,\alpha^*}}^{x_0})] = w(x_0, y^*).$$

Then $\tau^{y^,\alpha^*} \in \mathcal{T}_{y^*} \subset \mathcal{T}$ is an optimal stopping time for (2.1). That is,*

$$p^* = \mathbb{E}[f(x_0 + W_{\tau^{y^*,\alpha^*}})].$$

This is straightforward from Theorem 2.3 with $\epsilon = 0$.

2.3 Properties of the Auxiliary Value Function

In this section, we investigate various properties of the auxiliary value function w , which is defined in (2.4). We remind the reader that we can equivalently consider the definition given by (2.7) as is convenient because of the equivalence given in Theorem 2.2.

The main result of this section will be to characterize w as the unique uniformly continuous viscosity solution of (2.6) which satisfies certain growth conditions to be specified.

2.3.1 Analytical properties of the auxiliary value function

We begin by observing a trivial boundary condition for w .

Proposition 2.1. *For each $x \in \mathbb{R}$, we have $w(x, 0) = f(x)$.*

Proof. This follows immediately from the stochastic control interpretation of w in (2.4). \square

Next, we prove a more subtle continuity result. In particular, we show that for a fixed control $\alpha \in \mathcal{A}$, we have Hölder continuity of the expected payoff when varying the initial conditions $(x, y) \in \mathbb{R} \times [0, \infty)$.

Lemma 2.2. *For any $(x_1, y_1), (x_2, y_2) \in \mathbb{R} \times [0, \infty)$ and $\alpha \in \mathcal{A}$, we have*

$$|\mathbb{E}[f(X_{\tau^{y_1, \alpha}}^{x_1})] - \mathbb{E}[f(X_{\tau^{y_2, \alpha}}^{x_2})]| \leq L(|x_1 - x_2| + |y_1 - y_2|^{1/2}),$$

where $L > 0$ is the Lipschitz constant of f .

Proof. Without loss of generality, we can assume $y_1 \leq y_2$. It is then clear from (2.2) and (2.3) that $\tau^{y_1, \alpha} \leq \tau^{y_2, \alpha}$ almost-surely. Furthermore, by Lemma 2.1 we see that $\mathbb{E}[\tau^{y_1, \alpha}] = y_1$ and $\mathbb{E}[\tau^{y_2, \alpha}] = y_2$. We can then compute

$$|\mathbb{E}[f(X_{\tau^{y_1, \alpha}}^{x_1})] - \mathbb{E}[f(X_{\tau^{y_2, \alpha}}^{x_2})]| \leq L|x_1 - x_2| + L\mathbb{E}[|W_{\tau^{y_1, \alpha}} - W_{\tau^{y_2, \alpha}}|],$$

where $L > 0$ is the Lipschitz constant for f .

Using Jensen's inequality, we have

$$\mathbb{E}[|W_{\tau^{y_1, \alpha}} - W_{\tau^{y_2, \alpha}}|] = \mathbb{E}\left[\sqrt{(W_{\tau^{y_1, \alpha}} - W_{\tau^{y_2, \alpha}})^2}\right] \leq \sqrt{\mathbb{E}[(W_{\tau^{y_1, \alpha}} - W_{\tau^{y_2, \alpha}})^2]}.$$

Because $\tau^{y_1, \alpha} \leq \tau^{y_2, \alpha}$ almost-surely, we can use the Markov property of Brownian motion to see

$$\begin{aligned} \mathbb{E}[(W_{\tau^{y_1, \alpha}} - W_{\tau^{y_2, \alpha}})^2] &= \mathbb{E}[\mathbb{E}[(W_{\tau^{y_1, \alpha}} - W_{\tau^{y_2, \alpha}})^2 \mid \mathcal{F}_{\tau^{y_1, \alpha}}]] \\ &= \mathbb{E}[\mathbb{E}[\tau^{y_2, \alpha} - \tau^{y_1, \alpha} \mid \mathcal{F}_{\tau^{y_1, \alpha}}]] \\ &= \mathbb{E}[\tau^{y_2, \alpha}] - \mathbb{E}[\tau^{y_1, \alpha}] \\ &= y_2 - y_1. \end{aligned}$$

Putting these three inequalities together, the claimed result follows. \square

We next extend this result to Hölder continuity of the auxiliary value function.

Proposition 2.2. *For any $(x_1, y_1), (x_2, y_2) \in \mathbb{R} \times [0, \infty)$, we have*

$$|w(x_1, y_1) - w(x_2, y_2)| \leq L (|x_1 - x_2| + |y_1 - y_2|^{1/2}),$$

where $L > 0$ is the Lipschitz constant of f .

Proof. Let $\alpha \in \mathcal{A}$ be an arbitrary control. By Lemma 2.2, we have

$$\begin{aligned} \mathbb{E}[f(X_{\tau y_1, \alpha}^{x_1})] &\leq \mathbb{E}[f(X_{\tau y_2, \alpha}^{x_2})] + L (|x_1 - x_2| + |y_1 - y_2|^{1/2}) \\ &\leq w(x_2, y_2) + L (|x_1 - x_2| + |y_1 - y_2|^{1/2}), \end{aligned}$$

where $L > 0$ is the Lipschitz constant for f . Because $\alpha \in \mathcal{A}$ was arbitrary, we then conclude

$$w(x_1, y_2) \leq w(x_2, y_2) + L (|x_1 - x_2| + |y_1 - y_2|^{1/2}).$$

Reversing the roles of (x_1, y_1) and (x_2, y_2) , the stated result follows. \square

An immediate corollary of this result is that w is uniformly continuous and has linear asymptotic growth in x and sub-linear asymptotic growth in y .

Corollary 2.2. *The auxiliary value function w is uniformly continuous and there $C > 0$, which depends only on f , such that*

$$|w(x, y)| \leq C (1 + |x| + \sqrt{y})$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$.

Proof. The uniform continuity of w follows immediately from the Hölder continuity in Proposition 2.2. The claimed growth bound follows from Proposition 2.1 and Proposition 2.2, because

$$\begin{aligned} |w(x, y)| &\leq |w(0, 0)| + |w(x, y) - w(0, 0)| \\ &\leq |f(0)| + L (|x| + \sqrt{y}). \end{aligned}$$

Then the result holds with $C \equiv |f(0)| + L$. \square

Lastly, we state an important functional equality for the auxiliary value function – a Dynamic Programming Principle. We will later pass this equality to smooth test functions to show that w is a viscosity solution of (2.6).

Proposition 2.3 (Dynamic Programming Principle). *Fix $(x, y) \in \mathbb{R} \times [0, \infty)$ and let $\{\theta^\alpha\}$ be any collection of stopping times indexed by $\alpha \in \mathcal{A}$. Then we have*

$$w(x, y) = \sup_{\alpha \in \mathcal{A}} \mathbb{E}[w(X_{\theta^\alpha \wedge \tau y, \alpha}^x, Y_{\theta^\alpha \wedge \tau y, \alpha}^{y, \alpha})].$$

The upside to showing that w is equivalent to a stochastic control problem is that, while technical, the proof of this proposition is standard. In the following we will provide a formal sketch of the proof. For full details, see Fleming-Soner [FS06] or Touzi [Tou13].

Sketch of Proof. Fix a collection of stopping times $\{\theta^\alpha\}$ and define $\phi : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ by

$$\phi(x, y) := \sup_{\alpha \in \mathcal{A}} \mathbb{E} [w(X_{\theta^\alpha \wedge \tau^{y, \alpha}}^x, Y_{\theta^\alpha \wedge \tau^{y, \alpha}}^{y, \alpha})]$$

for each $(x, y) \in \mathbb{R} \times [0, \infty)$. We aim to show that $\phi = w$.

1. Fix $(x, y) \in \mathbb{R} \times [0, \infty)$ and let $\alpha \in \mathcal{A}$ be an arbitrary control. For notational convenience, in the remainder of this step we denote $X_{\theta^\alpha \wedge \tau^{y, \alpha}}^x$ by \tilde{X} , $Y_{\theta^\alpha \wedge \tau^{y, \alpha}}^{y, \alpha}$ by \tilde{Y} , and $s \mapsto \alpha(s + \tau^{y, \alpha})$ by $s \mapsto \tilde{\alpha}(s)$. By the Markov property of solutions of (2.2), we note that

$$X_{\tau^{\tilde{Y}, \tilde{\alpha}}}^{\tilde{X}} = X_{\tau^{y, \alpha}}^x,$$

almost-surely. By the tower property of conditional expectation, we have

$$\begin{aligned} \mathbb{E} [f(X_{\tau^{y, \alpha}}^x)] &= \mathbb{E} \left[\mathbb{E} \left[f \left(X_{\tau^{\tilde{Y}, \tilde{\alpha}}}^{\tilde{X}} \right) \mid \mathcal{F}_{\theta^\alpha \wedge \tau^{y, \alpha}} \right] \right] \\ &\leq \mathbb{E} \left[w \left(\tilde{X}, \tilde{Y} \right) \right] \\ &\leq \phi(x, y). \end{aligned}$$

But because $\alpha \in \mathcal{A}$ was arbitrary, we conclude

$$w(x, y) \leq \phi(x, y).$$

2. To prove the reverse inequality, fix $(x, y) \in \mathbb{R} \times [0, \infty)$ and let $\alpha \in \mathcal{A}$ be an arbitrary control. Again, for notational convenience, in the remainder of this step we denote $X_{\theta^\alpha \wedge \tau^{y, \alpha}}^x$ by \tilde{X} and $Y_{\theta^\alpha \wedge \tau^{y, \alpha}}^{y, \alpha}$ by \tilde{Y} . For any $\epsilon > 0$, let $\alpha_\epsilon \in \mathcal{A}$ be an ϵ -suboptimal control starting from (\tilde{X}, \tilde{Y}) . That is, it satisfies

$$w \left(\tilde{X}, \tilde{Y} \right) - \epsilon \leq \mathbb{E} \left[f \left(X_{\tau^{\tilde{Y}, \alpha_\epsilon}}^{\tilde{X}} \right) \mid \mathcal{F}_{\theta^\alpha \wedge \tau^{y, \alpha}} \right]. \quad (2.11)$$

But then if we define a new control $\bar{\alpha} \in \mathcal{A}$ as

$$\bar{\alpha}(s) := \begin{cases} \alpha(s) & 0 \leq s < \theta^\alpha \wedge \tau^{y, \alpha} \\ \alpha_\epsilon(s - \theta^\alpha \wedge \tau^{y, \alpha}) & \theta^\alpha \wedge \tau^{y, \alpha} \leq s, \end{cases}$$

then, again by the Markov property of solutions of (2.2), we deduce that

$$X_{\tau^{\tilde{Y}, \alpha_\epsilon}}^{\tilde{X}} = X_{\tau^{y, \bar{\alpha}}}^x,$$

almost-surely. Then taking unconditional expectations of (2.11), we see

$$\begin{aligned} \mathbb{E} [w (X_{\theta^\alpha \wedge \tau^{y, \alpha}}^x, Y_{\theta^\alpha \wedge \tau^{y, \alpha}}^{y, \alpha})] - \epsilon &= \mathbb{E} \left[w \left(\tilde{X}, \tilde{Y} \right) \right] \\ &\leq \mathbb{E} \left[f \left(X_{\tau^{\tilde{Y}, \alpha \epsilon}}^{\tilde{X}} \right) \right] \\ &= \mathbb{E} [f (X_{\tau^{y, \bar{\alpha}}}^x)] \\ &\leq w(x, y). \end{aligned}$$

Because $\alpha \in \mathcal{A}$ and $\epsilon > 0$ were both arbitrary, we conclude

$$\phi(x, y) \leq w(x, y).$$

□

Remark 2.1. While there are several abuses of notation in the sketch above, the main difficulty lies in the assumption that we can construct an ϵ -suboptimal control starting from any point in $\mathbb{R} \times [0, \infty)$. In general, it is not obvious that this can be done in a measurable way. A typical complete proof of this result generally uses an open covering of the state-space and the Lindelöf Covering Theorem to obtain a countable open covering by neighborhoods which each correspond to a single ϵ -suboptimal control. The fact that a single control can locally be ϵ -suboptimal requires the use a continuity result like Lemma 2.2. For more details on how this process proceeds, we refer the interested reader to Touzi [Tou13] or a similar proof in Chapter 4 of this thesis.

2.3.2 Viscosity solution characterization

We next claim that the auxiliary value function w is the unique uniformly continuous viscosity solution of (2.6) which satisfies the growth condition from Corollary 2.2. Because of the unbounded term in front of the diffusion in y , this is not a completely standard task. In the following we begin by recalling a working definition of viscosity solution for this equation.

Definition 2.5. Let $u : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ be a continuous function.

1. We say that u is a viscosity supersolution of (2.6) if

- a) $u(x, 0) \geq f(x)$ for all $x \in \mathbb{R}$, and
- b) For any $(x_0, y_0) \in \mathbb{R} \times (0, \infty)$, any smooth function $\phi : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ for which $(u - \phi)$ attains a local minimum at (x_0, y_0) , and for each $a \in \mathbb{R}$, we have

$$\phi_y(x_0, y_0) - \left(\frac{1}{2} \phi_{xx}(x_0, y_0) + a \phi_{xy}(x_0, y_0) + \frac{1}{2} a^2 \phi_{yy}(x_0, y_0) \right) \geq 0.$$

2. We say that u is a viscosity subsolution of (2.6) if

- a) $u(x, 0) \leq f(x)$ for all $x \in \mathbb{R}$, and
 b) For any $(x_0, y_0) \in \mathbb{R} \times (0, \infty)$, any smooth function $\phi : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ for which $(u - \phi)$ attains a local maximum at (x_0, y_0) , and any $\epsilon > 0$, there exists $a \in \mathbb{R}$ such that

$$\phi_y(x_0, y_0) - \left(\frac{1}{2} \phi_{xx}(x_0, y_0) + a \phi_{xy}(x_0, y_0) + \frac{1}{2} a^2 \phi_{yy}(x_0, y_0) \right) \leq \epsilon.$$

3. We say that u is a viscosity solution of (2.6) if it is both a viscosity supersolution and a viscosity subsolution.

Remark 2.2. The difficulty with (2.6) is that the Hamiltonian is not continuous with respect to the Hessian D^2u because of the unbounded supremum. Following Da Lio-Ley [DLL11], we have modified the definition of viscosity solution slightly to avoid technical difficulties in cases where

$$\sup_{a \in \mathbb{R}} \left[\frac{1}{2} \phi_{xx}(x_0, y_0) + a \phi_{xy}(x_0, y_0) + \frac{1}{2} a^2 \phi_{yy}(x_0, y_0) \right] = +\infty.$$

Functionally, however, essentially nothing in the theory has changed. We will still be able to prove a comparison theorem for uniformly continuous super- and subsolutions, subject to certain growth assumptions, which provides a uniqueness result.

Proposition 2.4. *The auxiliary value function w is a viscosity solution of (2.6).*

Again, the upside to showing that w is equivalent to a stochastic control problem is that, up to some technicalities, the proof of this proposition is standard. In the following we provide an essentially complete sketch of the proof to emphasize how the modified definition of viscosity solution is used. For full details, see Fleming-Soner [FS06] or Touzi [Tou13].

Sketch of Proof. In the following we provide a sketch of the proof. By Proposition 2.1, we have $w(x, 0) = f(x)$ for all $x \in \mathbb{R}$, so we only need to consider the viscosity solution properties at interior points of $\mathbb{R} \times (0, \infty)$. The idea is to pass the Dynamic Programming Equality from Proposition 2.3 to smooth test functions via a careful choice of controls and stopping times.

1. Fix $(x_0, y_0) \in \mathbb{R} \times (0, \infty)$ and let $\phi : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ be a smooth function for which $(w - \phi)$ has a local minimum at (x_0, y_0) . Without loss of generality, we can assume $w(x_0, y_0) = \phi(x_0, y_0)$. Fix $\delta > 0$ small enough that

$$(w - \phi)(x, y) \geq (w - \phi)(x_0, y_0) = 0$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$ such that $|x_0 - x| + |y_0 - y| \leq \delta$.

For each $a \in \mathbb{R}$, associate it with a square-integrable control defined by

$$\alpha(t) := a \exp(-t)$$

for all $t \geq 0$. Then for small $h > 0$, define a stopping time

$$\theta_h := h \wedge \inf \{t \geq 0 \mid (X_t^{x_0}, Y_t^{y_0, \alpha}) \notin [x_0 - \delta, x_0 + \delta] \times [y_0 - \delta, y_0 + \delta]\}.$$

Then by the Dynamic Programming Principle (Proposition 2.3), we have

$$w(x_0, y_0) \geq \mathbb{E} \left[w \left(X_{\theta_h \wedge \tau^{y_0, \alpha}}^{x_0}, Y_{\theta_h \wedge \tau^{y_0, \alpha}}^{y_0, \alpha} \right) \right].$$

Then using the previous two inequalities, and applying Itô's change of variable formula (See Section 3.3 in Karatzas-Shreve [KS91]) to the smooth function ϕ , we compute

$$\begin{aligned} \phi(x_0, y_0) &= w(x_0, y_0) \\ &\geq \mathbb{E} \left[w \left(X_{\theta_h \wedge \tau^{y_0, \alpha}}^{x_0}, Y_{\theta_h \wedge \tau^{y_0, \alpha}}^{y_0, \alpha} \right) \right] \\ &\geq \mathbb{E} \left[\phi \left(X_{\theta_h \wedge \tau^{y_0, \alpha}}^{x_0}, Y_{\theta_h \wedge \tau^{y_0, \alpha}}^{y_0, \alpha} \right) \right] \\ &= \phi(x_0, y_0) + \mathbb{E} \left[\int_0^{\theta_h \wedge \tau^{y_0, \alpha}} (\mathcal{L}^\alpha \phi) (s, X_s^{x_0}, Y_s^{y_0, \alpha}) ds \right], \end{aligned}$$

where

$$\mathcal{L}^\alpha \phi(u, x, y) := \frac{1}{2} \phi_{xx}(x, y) + a e^{-s} \phi_{xy}(x, y) + \frac{1}{2} a^2 e^{-2s} \phi_{yy}(x, y) - \phi_y(x, y).$$

Re-arranging the inequality above and examining convergence as $h \rightarrow 0^+$, we see

$$\begin{aligned} 0 &\geq \liminf_{h \rightarrow 0^+} \mathbb{E} \left[\frac{1}{h} \int_0^{\theta_h \wedge \tau^{y_0, \alpha}} (\mathcal{L}^\alpha \phi) (s, X_s^{x_0}, Y_s^{y_0, \alpha}) ds \right] \\ &\geq \mathbb{E} \left[\liminf_{h \rightarrow 0^+} \frac{1}{h} \int_0^{\theta_h \wedge \tau^{y_0, \alpha}} (\mathcal{L}^\alpha \phi) (s, X_s^{x_0}, Y_s^{y_0, \alpha}) ds \right] \\ &= (\mathcal{L}^\alpha \phi) (0, x_0, y_0) \\ &= \frac{1}{2} \phi_{xx}(x_0, y_0) + a \phi_{xy}(x_0, y_0) + \frac{1}{2} a^2 \phi_{yy}(x_0, y_0) - \phi_y(x_0, y_0). \end{aligned}$$

But then we conclude that w is a viscosity supersolution of (2.6).

2. Fix $(x_0, y_0) \in \mathbb{R} \times (0, \infty)$ and let $\phi : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ be a smooth function for which $(w - \phi)$ has a local maximum at (x_0, y_0) . Without loss of generality, we can assume $w(x_0, y_0) = \phi(x_0, y_0)$. Fix $\delta > 0$ small enough that

$$(w - \phi)(x, y) \leq (w - \phi)(x_0, y_0) = 0$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$ such that $|x_0 - x| + |y_0 - y| \leq \delta$.

For each $\alpha \in \mathcal{A}$ and each $h > 0$, define a stopping time

$$\theta_h^\alpha := h \wedge \inf \{t \geq 0 \mid (X_t^{x_0}, Y_t^{y_0, \alpha}) \notin [x_0 - \delta, x_0 + \delta] \times [y_0 - \delta, y_0 + \delta]\}.$$

Then by the Dynamic Programming Principle (Proposition 2.3), for any $\epsilon > 0$, there exists $\alpha \in \mathcal{A}$ such that

$$w(x_0, y_0) \leq \mathbb{E} \left[w \left(X_{\theta_h^\alpha \wedge \tau^{y_0, \alpha}}^{x_0}, Y_{\theta_h^\alpha \wedge \tau^{y_0, \alpha}}^{y_0, \alpha} \right) \right] + \epsilon.$$

Then using the previous two inequalities, and applying Itô's change of variable formula (See Section 3.3 in Karatzas-Shreve [KS91]) to the smooth function ϕ , we compute

$$\begin{aligned} \phi(x_0, y_0) &= w(x_0, y_0) \\ &\leq \mathbb{E} \left[w \left(X_{\theta_h^\alpha \wedge \tau^{y_0, \alpha}}^{x_0}, Y_{\theta_h^\alpha \wedge \tau^{y_0, \alpha}}^{y_0, \alpha} \right) \right] + \epsilon \\ &\leq \mathbb{E} \left[\phi \left(X_{\theta_h^\alpha \wedge \tau^{y_0, \alpha}}^{x_0}, Y_{\theta_h^\alpha \wedge \tau^{y_0, \alpha}}^{y_0, \alpha} \right) \right] + \epsilon \\ &= \phi(x_0, y_0) + \mathbb{E} \left[\int_0^{\theta_h^\alpha \wedge \tau^{y_0, \alpha}} (\mathcal{L}^\alpha \phi)(s, X_s^{x_0}, Y_s^{y_0, \alpha}) ds \right] + \epsilon, \end{aligned}$$

where

$$\mathcal{L}^\alpha \phi(u, x, y) := \frac{1}{2} \phi_{xx}(x, y) + \alpha_s \phi_{xy}(x, y) + \frac{1}{2} \alpha_s^2 \phi_{yy}(x, y) - \phi_y(x, y).$$

Re-arranging the inequality above and examining convergence as $h \rightarrow 0^+$, we see

$$\begin{aligned} -\epsilon &\leq \limsup_{h \rightarrow 0^+} \mathbb{E} \left[\frac{1}{h} \int_0^{\theta_h^\alpha \wedge \tau^{y_0, \alpha}} (\mathcal{L}^\alpha \phi)(X_s^{x_0}, Y_s^{y_0, \alpha}) ds \right] \\ &\leq \mathbb{E} \left[\limsup_{h \rightarrow 0^+} \frac{1}{h} \int_0^{\theta_h^\alpha \wedge \tau^{y_0, \alpha}} (\mathcal{L}^\alpha \phi)(X_s^{x_0}, Y_s^{y_0, \alpha}) ds \right] \\ &= (\mathcal{L}^\alpha \phi)(0, x_0, y_0) \\ &= \frac{1}{2} \phi_{xx}(x_0, y_0) + \alpha_0 \phi_{xy}(x_0, y_0) + \frac{1}{2} \alpha_0^2 \phi_{yy}(x_0, y_0) - \phi_y(x_0, y_0). \end{aligned}$$

But then we conclude that w is a viscosity subsolution of (2.6). □

We now aim to show a uniqueness result for viscosity solutions in order to characterize the auxiliary value function w in terms of (2.6). We begin with a comparison principle.

Theorem 2.4 (Comparison Principle). *Let $\underline{u}, \bar{u} : \mathbb{R} \times [0, \infty) \rightarrow \mathbb{R}$ be a uniformly continuous viscosity subsolution and viscosity supersolution of (2.6), respectively. Suppose there exists $C > 0$ such that*

$$|\underline{u}(x, y)| + |\bar{u}(x, y)| \leq C(1 + |x| + \sqrt{y})$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$. Then

$$\underline{u}(x, y) \leq \bar{u}(x, y)$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$.

We remind the reader that the key difficulty in this proof is that the Hamiltonian in (2.6) is not continuous with respect to the Hessian D^2u because of the unbounded supremum. The key to the following proof will how we obtain a bound which is independent of the control $\alpha \in \mathbb{R}$.

Proof. 1. Assume to the contrary that

$$\sigma := (\underline{u} - \bar{u})(x_0, y_0) > 0$$

for some $(x_0, y_0) \in \mathbb{R} \times [0, \infty)$. Choose $\epsilon, \lambda > 0$ both to be small and define

$$M := \sup_{\mathbb{R}^2 \times [0, \infty)^2} \Phi(x_1, x_2, y_1, y_2),$$

where

$$\Phi(x_1, x_2, y_1, y_2) := \underline{u}(x_1, y_1) - \bar{u}(x_2, y_2) - \phi(x_1, x_2, y_1, y_2)$$

and

$$\phi(x_1, x_2, y_1, y_2) := \frac{1}{2}\epsilon^{-2} \left((x_1 - x_2)^2 + (y_1 - y_2)^2 \right) + \frac{1}{2}\epsilon (x_1^2 + x_2^2) + \lambda(y_1 + y_2).$$

Then because of the growth bounds on \underline{u} and \bar{u} (linear in x and sub-linear in y), there exists $(\bar{x}_1, \bar{x}_2, \bar{y}_1, \bar{y}_2) \in \mathbb{R}^2 \times [0, \infty)^2$ such that

$$M = \Phi(\bar{x}_1, \bar{x}_2, \bar{y}_1, \bar{y}_2).$$

2. We may take $\epsilon, \lambda > 0$ sufficiently small that

$$\Phi(\bar{x}_1, \bar{x}_2, \bar{y}_1, \bar{y}_2) \geq \Phi(x_0, x_0, y_0, y_0) \geq \frac{\sigma}{2}. \quad (2.12)$$

In addition, $\Phi(\bar{x}_1, \bar{x}_2, \bar{y}_1, \bar{y}_2) \geq \Phi(0, 0, 0, 0)$, so we see

$$\begin{aligned} \lambda(\bar{y}_1 + \bar{y}_2) + \frac{1}{2}\epsilon^{-2} \left((\bar{x}_1 - \bar{x}_2)^2 + (\bar{y}_1 - \bar{y}_2)^2 \right) + \frac{1}{2}\epsilon (\bar{x}_1^2 + \bar{x}_2^2) \\ \leq \underline{u}(\bar{x}_1, \bar{y}_1) - \bar{u}(\bar{x}_2, \bar{y}_2) - \underline{u}(0, 0) + \bar{u}(0, 0) \\ \leq C \left(4 + |\bar{x}_1| + |\bar{x}_2| + \bar{y}_1^{1/2} + \bar{y}_2^{1/2} \right). \end{aligned}$$

Noting that

$$C \left(4 + |\bar{x}_1| + |\bar{x}_2| + \bar{y}_1^{1/2} + \bar{y}_2^{1/2} \right) \leq C \left(4 + C\epsilon^{-1} + \frac{1}{2}C\lambda^{-1} \right) + \frac{1}{2}\epsilon (\bar{x}_1^2 + \bar{x}_2^2) + \lambda(\bar{y}_1 + \bar{y}_2),$$

we can put these two inequalities, we see

$$\frac{1}{2}\epsilon^{-2} \left((\bar{x}_1 - \bar{x}_2)^2 + (\bar{y}_1 - \bar{y}_2)^2 \right) \leq C \left(4 + C\epsilon^{-1} + \frac{1}{2}C\lambda^{-1} \right).$$

From this, we deduce that

$$|\bar{x}_1 - \bar{x}_2|, |\bar{y}_1 - \bar{y}_2| \in O(\epsilon^{1/2})$$

for fixed $\lambda > 0$ as $\epsilon \rightarrow 0$.

3. Because \underline{u} is uniformly continuous, we can write $\underline{\omega}(\cdot)$ to denote its modulus of continuity; this is,

$$|\underline{u}(x_1, y_1) - \underline{u}(x_2, y_2)| \leq \underline{\omega}(|x_1 - x_2| + |y_1 - y_2|)$$

for all $(x_1, y_1), (x_2, y_2) \in \mathbb{R} \times [0, \infty)$ and $\underline{\omega}(r) \rightarrow 0$ as $r \rightarrow 0$. Similarly, we denote the modulus of continuity of \bar{u} by $\bar{\omega}(\cdot)$. Then (2.12) implies

$$\begin{aligned} \frac{\sigma}{2} &\leq \underline{u}(\bar{x}_1, \bar{y}_1) - \bar{u}(\bar{x}_2, \bar{y}_2) \\ &\leq \underline{u}(\bar{x}_1, 0) - \bar{u}(\bar{x}_2, 0) + \underline{\omega}(\bar{y}_1) + \bar{\omega}(\bar{y}_2) \\ &\leq f(\bar{x}_1) - f(\bar{x}_2) + \underline{\omega}(\bar{y}_1) + \bar{\omega}(\bar{y}_2) \\ &\leq L|\bar{x}_1 - \bar{x}_2| + \underline{\omega}(\bar{y}_1) + \bar{\omega}(\bar{y}_2), \end{aligned}$$

where $L > 0$ is the Lipschitz constant of f .

Suppose $\bar{y}_1 = 0$. Then $|\bar{x}_1 - \bar{x}_2|, |\bar{y}_1 - \bar{y}_2| \in O(\epsilon^{1/2})$ implies there exists $C > 0$ such that

$$\frac{\sigma}{2} \leq CL\epsilon^{1/2} + \underline{\omega}(0) + \bar{\omega}(C\epsilon^{1/2}).$$

But taking $\epsilon \rightarrow 0$, this is a contradiction. Therefore, $\bar{y}_1 > 0$ for sufficiently small $\epsilon > 0$. A similar argument shows that $\bar{y}_2 > 0$ for sufficiently small $\epsilon > 0$.

4. Now $(\bar{x}_1, \bar{x}_2, \bar{y}_1, \bar{y}_2) \in \mathbb{R}^2 \times (0, \infty)^2$ for sufficiently small $\epsilon > 0$, so we can apply the Crandall-Ishii Lemma (See Crandall-Ishii-Lions [CIL92]). We state the result in terms of smooth test functions instead of sub- and super-jets and take $\rho := \epsilon^2$. There exists smooth $\underline{\phi}, \bar{\phi} : \mathbb{R} \times [0, \infty)$ such that $(\underline{u} - \underline{\phi})$ attains a local maximum at (\bar{x}_1, \bar{y}_1) , $(\bar{u} - \bar{\phi})$ attains a local minimum at (\bar{x}_2, \bar{y}_2) ,

$$D\underline{\phi}(\bar{x}_1, \bar{y}_1) = \begin{pmatrix} \epsilon^{-2}(\bar{x}_1 - \bar{x}_2) + \epsilon\bar{x}_1 \\ \epsilon^{-2}(\bar{y}_1 - \bar{y}_2) + \lambda \end{pmatrix},$$

$$D\bar{\phi}(\bar{x}_2, \bar{y}_2) = \begin{pmatrix} \epsilon^{-2}(\bar{x}_1 - \bar{x}_2) - \epsilon\bar{x}_2 \\ \epsilon^{-2}(\bar{y}_1 - \bar{y}_2) - \lambda \end{pmatrix},$$

and (for $\epsilon \ll 1$)

$$\begin{pmatrix} D^2\underline{\phi}(\bar{x}_1, \bar{y}_1) & 0 \\ 0 & -D^2\bar{\phi}(\bar{x}_2, \bar{y}_2) \end{pmatrix} \leq 5\epsilon^{-2}A_4 + 2\epsilon B_4,$$

where

$$A_4 := \begin{pmatrix} 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \\ -1 & 0 & 1 & 0 \\ 0 & -1 & 0 & 1 \end{pmatrix} \text{ and } B_4 := \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

By the matrix inequality above, we mean with respect to the partial order induced by the positive semi-definite cone. The key property is that, for any $\alpha \in \mathbb{R}$, if we conjugate the matrix inequality by $(1, \alpha, -1, -\alpha)^\top$, then we conclude

$$\begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top D^2 \underline{\phi}(\bar{x}_1, \bar{y}_1) \begin{pmatrix} 1 \\ \alpha \end{pmatrix} \leq \begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top D^2 \bar{\phi}(\bar{x}_2, \bar{y}_2) \begin{pmatrix} 1 \\ \alpha \end{pmatrix} + 2\epsilon. \quad (2.13)$$

5. By the viscosity subsolution property, there exists $\alpha \in \mathbb{R}$ such that

$$\begin{aligned} \underline{\phi}_y(\bar{x}_1, \bar{y}_1) &\leq \epsilon + \frac{1}{2} \underline{\phi}_{xx}(\bar{x}_1, \bar{y}_1) + \alpha \underline{\phi}_{xy}(\bar{x}_1, \bar{y}_1) + \frac{1}{2} \alpha^2 \underline{\phi}(\bar{x}_1, \bar{y}_1) \\ &= \epsilon + \frac{1}{2} \begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top D^2 \underline{\phi}(\bar{x}_1, \bar{y}_1) \begin{pmatrix} 1 \\ \alpha \end{pmatrix}. \end{aligned}$$

But by the properties of $\underline{\phi}$ from the previous step, this implies

$$\epsilon^{-2}(\bar{y}_1 - \bar{y}_2) + \lambda \leq \epsilon + \frac{1}{2} \begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top D^2 \underline{\phi}(\bar{x}_1, \bar{y}_1) \begin{pmatrix} 1 \\ \alpha \end{pmatrix}. \quad (2.14)$$

Similarly, by the viscosity supersolution property, we have

$$\begin{aligned} \bar{\phi}_y(\bar{x}_2, \bar{y}_2) &\geq \frac{1}{2} \bar{\phi}_{xx}(\bar{x}_2, \bar{y}_2) + \alpha \bar{\phi}_{xy}(\bar{x}_2, \bar{y}_2) + \frac{1}{2} \alpha^2 \bar{\phi}(\bar{x}_2, \bar{y}_2) \\ &= \frac{1}{2} \begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top D^2 \bar{\phi}(\bar{x}_2, \bar{y}_2) \begin{pmatrix} 1 \\ \alpha \end{pmatrix}. \end{aligned}$$

By the properties of $\bar{\phi}$ from the previous step, this implies

$$\epsilon^{-2}(\bar{y}_1 - \bar{y}_2) - \lambda \geq \frac{1}{2} \begin{pmatrix} 1 \\ \alpha \end{pmatrix}^\top D^2 \bar{\phi}(\bar{x}_2, \bar{y}_2) \begin{pmatrix} 1 \\ \alpha \end{pmatrix}. \quad (2.15)$$

Putting together (2.13), (2.14), and (2.15), we conclude

$$3\epsilon \geq \lambda > 0.$$

But then taking $\epsilon > 0$ sufficiently small, we obtain a contradiction. □

We can then immediately state a uniqueness result which characterizes the auxiliary value function in terms of the HJB PDE (2.6).

Corollary 2.3. *The auxiliary value function w is the unique uniformly continuous viscosity solution of (2.6) for which there exists $C > 0$ such that*

$$|u(x, y)| \leq C(1 + |x| + \sqrt{y})$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$.

Proof. We note that w is uniformly continuous and satisfies the stated growth condition by Corollary 2.2. Suppose that there $u : \mathbb{R} \times [0, \infty)$ is a different uniformly continuous viscosity solution of (2.6). Then by Theorem 2.4 we immediately see

$$w(x, y) = u(x, y)$$

for all $(x, y) \in \mathbb{R} \times [0, \infty)$. □

Chapter 3

A Time-Inconsistent Optimal Stochastic Control Problem

3.1 Introduction

The following chapter is based upon the joint work of Miller-Yang [MY15], in which we consider a class of continuous-time stochastic optimal control problems, including those with Conditional Value-at-Risk (CVaR) appearing in the objective function. The original paper generalizes to allow optimal stochastic control when the objective function includes several other time-inconsistent features, such as variance and median absolute deviation. In this chapter, we focus on the specific case where the objective function represents a trade-off between expectation and CVaR. The emphasis, for the purposes of this thesis, is on how to convert the time-inconsistent stochastic control problem into an optimization problem over the value function of a related time-consistent stochastic control problem. We then consider an application in portfolio selection.

3.1.1 Mathematical setup

In this chapter, we consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which supports a standard Brownian motion W . For the purposes of this thesis, we take W to be one-dimensional, but this is extended in Miller-Yang [MY15]. We let $\mathbb{F} := \{\mathcal{F}_t\}_{t \geq 0}$ denote a filtration, which is assumed to be right-continuous and have all \mathbb{P} -negligible sets contained in \mathcal{F}_0 . We let \mathbb{A} be a compact and finite-dimensional set of controls, and let

$$\mathcal{A} := \left\{ \alpha : \Omega \times [0, T] \rightarrow \mathbb{R} \mid \alpha \text{ is progressively-measurable and } \mathbb{E} \int_0^T \alpha_t^2 dt < +\infty \right\}.$$

For each control process $\alpha \in \mathcal{A}$, we consider the controlled stochastic differential equation,

$$dX_t^\alpha := \mu(X_t^\alpha, \alpha_t) dt + \sigma(X_t^\alpha, \alpha_t) dW_t. \quad (3.1)$$

We take $\mu, \sigma : \mathbb{R} \times \mathbb{A} \rightarrow \mathbb{R}$ to be continuous functions such that, for $C > 0$ large enough, we have

$$|\mu(x, a) - \mu(x', a)| + |\sigma(x, a) - \sigma(x', a)| \leq C|x - x'| \quad (3.2)$$

$$|\mu(x, a)| + |\sigma(x, a)| \leq C(1 + |x|) \quad (3.3)$$

for all $x, x' \in \mathbb{R}$ and $a \in \mathbb{A}$. For any choice of $(t, x) \in [0, T] \times \mathbb{R}$ and $\alpha \in \mathcal{A}$, we write $\{X_s^{t,x,\alpha} \mid t \leq s \leq T\}$ to denote the solution of (3.1) with initial condition $X_t^{t,x,\alpha} = x$. The conditions above suffice to guarantee this solution is unique (See Section 5.2 in Karatzas-Shreve [KS91]).

We next recall definitions of value-at-risk (VaR) and conditional value-at-risk (CVaR) which will be used in this chapter. We choose to view these as functions on a space of probability measures, rather than functions of random variables, to emphasize that they do not depend upon the choice of probability space itself and to set up for certain analytical observations later in the chapter.

Definition 3.1. Let $\mathcal{P}(\mathbb{R})$ be the collection of Radon probability measures on \mathbb{R} . For any $p \in (0, 1)$ and $\mu \in \mathcal{P}(\mathbb{R})$, we define value-at-risk (at probability p) as a function $\text{VaR}_p : \mathcal{P}(\mathbb{R}) \rightarrow \mathbb{R}$ given by

$$\text{VaR}_p(\mu) := \inf \{y \in \mathbb{R} \mid \mu((-\infty, y]) \geq p\}.$$

Definition 3.2. Let $\mathcal{P}_1(\mathbb{R}) \subset \mathcal{P}(\mathbb{R})$ be the subset of Radon probability measures on \mathbb{R} which have finite first moment. For any $p \in (0, 1)$ and $\mu \in \mathcal{P}_1(\mathbb{R})$, we define conditional value-at-risk (at probability p) as a function $\text{CVaR}_p : \mathcal{P}_1(\mathbb{R}) \rightarrow \mathbb{R}$ given by

$$\begin{aligned} \text{CVaR}_p(\mu) &:= y^* - p^{-1} \int_{\mathbb{R}} (y^* - x) \mu(dx) \\ &= p^{-1} \left[\int_{(-\infty, y^*)} x \mu(dx) + y^* (p - \mu((-\infty, y^*))) \right], \end{aligned} \tag{3.4}$$

where $y^* := \text{VaR}_p(\mu)$.

There are many competing definitions of CVaR in the literature, but they all satisfy the intuitive property that for any $X \in L^1(\Omega, \mathbb{P})$ such that the pull-back measure $\mathbb{P} \circ X^{-1}$ contains no atoms, we have

$$\text{CVaR}_p(\mathbb{P} \circ X^{-1}) = \mathbb{E}[X \mid X \leq \text{VaR}_p(\mathbb{P} \circ X^{-1})]. \tag{3.5}$$

The intuition is that VaR_p represents the p th percentile worst-case outcome of a distribution, while CVaR_p represents the expected outcome conditional upon being in one of the p th percentile worst-case outcomes. For this reason, both VaR and CVaR are popular measures of tail-risk.

In the case that the distribution of X contains atoms, the more general definition (3.4) can be re-written as

$$\text{CVaR}_p(\mathbb{P} \circ X^{-1}) = p^{-1} (\mathbb{E}[X \mid X < y^*] \mathbb{P}[X < y^*] + y^* (p - \mathbb{P}[X < y^*])).$$

The intuition behind this equality is that, if $\mathbb{P}[X = y^*] > 0$, then we include only a fraction of the atom which corresponds to probability up to the p th tail when computing CVaR_p .

In the remainder of this chapter, we will often abuse notation and write $\text{VaR}_p[X]$ or $\text{CVaR}_p[X]$, which denote applying these operations to the pull-back measure on X induced by \mathbb{P} . We use this more general definition partially to simplify analysis, but we also show that (3.4) is the only continuous function (with respect to a Wasserstein metric) which satisfies an analogue of the intuitive property (3.5) when the measure contains no atoms.

3.1.2 Overview of results

In the remainder of the chapter, we fix some $x_0 \in \mathbb{R}$, $\lambda \geq 0$, and $p \in (0, 1)$. We then define the main problem considered in this chapter.

Definition 3.3. *The time-inconsistent optimal stochastic control of Mean-CVaR is to compute*

$$p^* := \sup_{\alpha \in \mathcal{A}} [\mathbb{E} [X_T^{0,x_0,\alpha}] + \lambda CVaR_p [X_T^{0,x_0,\alpha}]] \quad (3.6)$$

and to find $\alpha^* \in \mathcal{A}$ for which the supremum is attained.

The intuition is that the Mean-CVaR optimal stochastic control problem represents a trade-off between maximizing expectation while minimizing tail-risk.

As usual, we define a corresponding value function.

Definition 3.4. *We define the value function for the time-inconsistent stochastic control of Mean-CVaR as*

$$v(t, x) := \sup_{\alpha \in \mathcal{A}} [\mathbb{E} [X_T^{t,x,\alpha}] + \lambda CVaR_p [X_T^{t,x,\alpha}]]$$

for each $(t, x) \in [0, T] \times \mathbb{R}$.

Notice that $v(0, x_0) = p^*$.

This value function is time-inconsistent due to the CVaR term appearing in the objective function. The main observation of this chapter is that we can re-write the value function v as a maximization over the value function for a time-consistent stochastic optimal control problem.

Definition 3.5. *We define the value function for an auxiliary time-consistent stochastic control problem to be*

$$w(t, x, y) := \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[X_T^{t,x,\alpha} - \lambda p^{-1} (y - X_T^{t,x,\alpha})^+ \right]$$

for each $(t, x, y) \in [0, T] \times \mathbb{R} \times \mathbb{R}$.

Then the main observation of this chapter is that

$$v(x) = \sup_{y \in \mathbb{R}} [w(0, x, y) + \lambda y] \quad (3.7)$$

for all $x \in \mathbb{R}$.

The upside of this approach is that w is the value function of a time-consistent stochastic control problem, so it can be characterized as the unique viscosity solution of the HJB PDE:

$$\begin{cases} u_t + \sup_{a \in \mathcal{A}} [\mu(x, a)u_x + \frac{1}{2}\sigma(x, a)^2u_{xx}] = 0 & \text{in } [0, T) \times \mathbb{R} \times \mathbb{R} \\ u = x - \lambda p^{-1}(y - x)^+ & \text{on } \{t = T\} \times \mathbb{R} \times \mathbb{R}. \end{cases} \quad (3.8)$$

We can then compute w via standard methods, solve the optimization problem (3.7), then obtain a solution to (3.6).

3.1.3 Overview of existing literature

Conditional value-at-risk (CVaR) has received significant attention over the past two decades as a tool for managing risk. CVaR measures the expected value conditional upon being within some percentage of the worst-case loss scenarios. While both value-at-risk (VaR) and CVaR are risk measures, only CVaR is *coherent* in the sense of Artzner et al. [ADEH99]. One common criticism of VaR stems from its inability to distinguish based on the magnitude of losses in the tails of a distribution. In contrast, CVaR takes into account the magnitude of losses when in values exceeding VaR.

Due to the superior mathematical properties and practical implications, CVaR has gained popularity in risk management.¹ In particular, *static* or single-stage optimization with CVaR functions can be efficiently performed via convex and linear programming methods (See Rockafellar-Uryasev [RU00] and Mansini et al. [MOS07]). With the advances in optimization algorithms for CVaR, this risk measure has shown to be useful in various finance and engineering applications.

Dynamic or sequential optimization of CVaR is often of interest when decisions can be made at multiple stages. In such an optimal control setting, we can optimize a control action at a certain time based on the information from observations up to that time. This dynamic control approach enjoys an effective usage of information gathered in the process of making decisions under uncertainty. The need for efficient optimal control tools with CVaR is also motivated by emerging dynamic risk management methods in engineering and finance (See Qin-Su-Rajagopal [QSR13] and Yang-Callaway-Tomlin [YCT a]).

The major challenge in optimal control involving CVaR arises from its *time-inconsistency* (See Artzner et al. [ADEHK07]). Mathematically, this time-inconsistency prevents us from directly applying dynamic programming, in contrast with problems involving Markov risk measures (See Ruszczyński [Rus10], Cavuş-Ruszczyński [CR14], and Ruszczyński-Yao [RY15]) or risk-sensitive criteria (See James-Baras-Elliot [JBE94] and Fleming-McEneaney [FM95]). To overcome this difficulty, several methods have been proposed. A state-space lifting approach for dynamic programming with a discrete-time and discrete-state Markov decision process (MDP) setting is first proposed in Bäuerle-Ott [BO11].

Another lifting method and relevant algorithms are developed in Pflug-Pichler [PP16b] and Chow et al. [CTMP15], relying on a so-called CVaR Decomposition Theorem of Pflug-Pichler [PP16b]. This approach uses a dual representation of CVaR and hence requires optimization over a space of probability densities when solving an associated Bellman equation. This optimization problem can be effectively solved in discrete-time and finite discrete-state MDPs. However, it becomes computationally intractable in (uncountable) continuous-state optimal control problems as the space of densities is infinite dimensional. In Haskell-Jain [HJ15], a different approach is developed for risk-aware discrete-time finite-state MDPs, which is based on occupation measures. Due to the nonconvexity of the resulting infinite-

¹More detailed comparisons between VaR and CVaR, in terms of stability of statistical estimation and simplicity of optimization procedures, can be found in Sarykalin et al. [SSU08].

dimensional optimization problem, this method uses a successive linear approximation procedure.

In this chapter, we demonstrate a solution of the continuous-time and continuous-space optimal control of Mean-CVaR using the so-called *extremal* representation of CVaR originally proposed in Rockafellar-Uryasev [RU00]. We reformulate the optimal control problem as a bi-level optimization problem in which the outer optimization problem is convex and the inner optimization problem is standard stochastic optimal control. We note that, while the auxiliary time-consistent stochastic control problem features an extra state variable, in practice we can perform gradient descent-based optimization over the value of the extra state variable rather than computing the auxiliary value function as a function in a higher-dimensional state space.

In the final section of this chapter, we demonstrate a practical implementation of our methodology in an optimal investment problem subject to CVaR constraints. To our knowledge, this is the first solution of a dynamic portfolio optimization problem subject to tail-risk constraints in continuous time. The closest comparisons to our results are given by approximate equilibrium solutions (See Dong-Sircar [DS14]), mean-field control approaches (See Pfeiffer [Pfe16]), or in mean-variance frameworks (See Pedersen-Peskir [PP16a]).

3.2 Equivalent Sequential Time-Consistent Problem

The goal of this section is to demonstrate how we can convert the time-inconsistent optimal stochastic control problem (3.6) into a sequential optimization problem involving a time-consistent control problem.

As in the previous chapter, we proceed by re-writing the time-inconsistent feature (CVaR) in terms of a family of stochastic optimal control problems in an enlarged state-space. We then show how to construct an optimal control for the time-inconsistent problem by picking an optimal member of the family of time-consistent problems and generating an optimal solution of the corresponding time-consistent stochastic control problem.

In the problem considered in this chapter, it will turn out that the extra state variable represents the value-at-risk (VaR) of the optimal control. The dynamics of this extra state variable are actually trivial, in contrast with the dynamics of the previous chapter.

3.2.1 Main equivalence result

In this section, we demonstrate the main equivalence between the time-inconsistent value function v and an optimization problem over the time-consistent auxiliary value function w . We begin by proving a lemma regarding a representation of CVaR as a maximization problem.

The following representation of CVaR dates back to Rockafellar-Uryasev [RU00] in the case of probability measures with no atoms. We provide a complete proof of the more general version for the sake of completeness.

Recall the definition of VaR and CVaR given in (3.4) and (3.4).

Lemma 3.1. *For any $\mu \in \mathcal{P}_1(\mathbb{R})$, we have*

$$CVaR_p(\mu) = \sup_{y \in \mathbb{R}} \left[y - p^{-1} \int_{\mathbb{R}} (y - x)^+ \mu(dx) \right].$$

Furthermore, the maximum is achieved at $y^* := VaR_p(\mu)$.

Proof. Define a function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ as

$$\phi(y) := y - p^{-1} \int_{\mathbb{R}} (y - X)^+ \mu(dx)$$

for all $y \in \mathbb{R}$. Let $y^* := VaR_p(\mu)$. It is obvious that

$$CVaR_p(\mu) = \phi(y^*)$$

by the definition in (3.4). We now aim to show that y^* maximizes ϕ .

1. Let $y \in \mathbb{R}$ satisfy $y < y^*$. Then we can compute

$$\begin{aligned} \phi(y^*) - \phi(y) &= y^* - y + p^{-1} \int_{\mathbb{R}} ((y - x)^+ - (y^* - x)^+) \mu(dx) \\ &= y^* - y + p^{-1} \int_{(-\infty, y)} (y - x) \mu(dx) - p^{-1} \int_{(-\infty, y^*)} (y^* - x) \mu(dx) \\ &= y^* - y + p^{-1} \int_{(-\infty, y^*)} (y - y^*) \mu(dx) - p^{-1} \int_{[y, y^*)} (y - x) \mu(dx) \\ &\geq (y^* - y) (1 - p^{-1} \mu((-\infty, y^*))). \end{aligned}$$

But for any $n \geq 1$, we have $\mu((-\infty, y^* - n^{-1})) < p$ by the definition of VaR. The union of these sets is $(-\infty, y^*)$, so by the continuity properties of measures, we have

$$\mu((-\infty, y^*)) = \lim_{n \rightarrow \infty} \mu((-\infty, y^* - n^{-1})) \leq p.$$

Then these two inequalities imply $\phi(y^*) \geq \phi(y)$.

2. Let $y \in \mathbb{R}$ satisfy $y > y^*$. By a slight variation on the previous computation, we see

$$\begin{aligned} \phi(y^*) - \phi(y) &= y^* - y + p^{-1} \int_{\mathbb{R}} ((y - x)^+ - (y^* - x)^+) \mu(dx) \\ &= y^* - y + p^{-1} \int_{(-\infty, y]} (y - x) \mu(dx) - p^{-1} \int_{(-\infty, y^*]} (y^* - x) \mu(dx) \\ &= y^* - y + p^{-1} \int_{(-\infty, y^*]} (y - y^*) \mu(dx) + p^{-1} \int_{(y^*, y]} (y - x) \mu(dx) \\ &\geq (y^* - y) (1 - p^{-1} \mu((-\infty, y^*])). \end{aligned}$$

But for any $n \geq 1$, we have $\mu((-\infty, y^* + n^{-1})) \geq p$ by the definition of VaR and sub-additivity of measures. The intersection of these sets is $(-\infty, y^*]$, so again by the continuity properties of measures, we have

$$\mu((-\infty, y^*)) = \lim_{n \rightarrow \infty} \mu((-\infty, y^* + n^{-1})) \geq p.$$

Combining these two inequalities then implies $\phi(y^*) \geq \phi(y)$.

□

We this lemma in hand, we can immediately re-write the time-inconsistent value function. Recall the definitions of v in (3.7) and the time-consistent auxiliary value function w in (3.7).

Theorem 3.1. *We have*

$$v(t, x) = \sup_{y \in \mathbb{R}} [w(t, x, y) + \lambda y].$$

for all $(t, x) \in [0, T] \times \mathbb{R}$.

Proof. Fix $(t, x) \in [0, T] \times \mathbb{R}$. Let $y \in \mathbb{R}$ and $\alpha \in \mathcal{A}$ be arbitrary. Then by Definition 3.4 and Lemma 3.1, we have

$$\begin{aligned} v(t, x) &\geq \mathbb{E} [X_T^{t,x,\alpha}] + \lambda \text{CVaR}_p [X_T^{t,x,\alpha}] \\ &\geq \mathbb{E} [X_T^{t,x,\alpha}] + \lambda \left(y - p^{-1} \mathbb{E} \left[(y - X_T^{t,x,\alpha})^+ \right] \right) \\ &= \mathbb{E} \left[X_T^{t,x,\alpha} - \lambda p^{-1} (y - X_T^{t,x,\alpha})^+ \right] + \lambda y. \end{aligned}$$

But because $\alpha \in \mathcal{A}$ was arbitrary, this implies

$$v(t, x) \geq w(t, x, y) + \lambda y.$$

Because y was arbitrary, the claimed result holds. □

3.2.2 Construction of optimal pre-commitment strategies

We have shown from Theorem 3.1 that we can recover p^* in the time-inconsistent optimal control problem (3.6) by maximizing over choice of $y \in \mathbb{R}$ and control $\alpha \in \mathcal{A}$. However, it remains to be shown that we can construct (approximate) optimal pre-commitment controls for the original time-inconsistent problem.

The next theorem relates nearly-optimal choices of $(y, \alpha) \in \mathbb{R} \times \mathcal{A}$ with nearly-optimal choices of pre-commitment control in the original problem.

Theorem 3.2. *For any $\epsilon \geq 0$, let $y \in \mathbb{R}$ satisfy*

$$v(0, x_0) \leq w(0, x_0, y) + \lambda y + \epsilon$$

and let $\alpha \in \mathcal{A}$ satisfy

$$w(0, x_0, y) \leq \mathbb{E} \left[X_T^{0, x_0, \alpha} - \lambda p^{-1} (y - X_T^{0, x_0, \alpha})^+ \right] + \epsilon.$$

Then $\alpha \in \mathcal{A}$ satisfies

$$p^* - 2\epsilon \leq \mathbb{E} [X_T^{0, x_0, \alpha}] + \lambda CVaR_p [X_T^{0, x_0, \alpha}] \leq p^*.$$

Proof. Recall from Definition 3.4 that $p^* = v(0, x_0)$. Then by combining this with the two assumed inequalities and Lemma 3.1, we have

$$\begin{aligned} p^* &= v(0, x_0) \\ &\leq w(0, x_0, y) + \lambda y + \epsilon \\ &\leq \mathbb{E} \left[X_T^{0, x_0, \alpha} - \lambda p^{-1} (y - X_T^{0, x_0, \alpha})^+ \right] + \lambda y + 2\epsilon \\ &= \mathbb{E} [X_T^{0, x_0, \alpha}] + \lambda \left(y - p^{-1} \mathbb{E} \left[(y - X_T^{0, x_0, \alpha})^+ \right] \right) + 2\epsilon \\ &\leq \mathbb{E} [X_T^{0, x_0, \alpha}] + \lambda CVaR_p [X_T^{0, x_0, \alpha}] + 2\epsilon. \end{aligned}$$

But, of course, by the definition of p^* , we also have

$$\mathbb{E} [X_T^{0, x_0, \alpha}] + \lambda CVaR_p [X_T^{0, x_0, \alpha}] \leq p^*.$$

□

Then we can record the following corollary regarding obtaining an optimal control.

Corollary 3.1. *Let $y^* \in \mathbb{R}$ satisfy*

$$w(0, x_0, y^*) + \lambda y^* = \max_{y \in \mathbb{R}} [w(0, x_0, y) + \lambda y]$$

and let $\alpha^* \in \mathcal{A}$ satisfy

$$\mathbb{E} \left[X_T^{0, x_0, \alpha^*} - \lambda p^{-1} (y^* - X_T^{0, x_0, \alpha^*})^+ \right] = w(0, x_0, y^*).$$

Then α^* is an optimal stopping time for (3.6). That is,

$$p^* = \mathbb{E} [X_T^{0, x_0, \alpha^*}] + \lambda CVaR_p [X_T^{0, x_0, \alpha^*}].$$

This is straightforward from Theorem 3.2 with $\epsilon = 0$.

3.3 Properties of the Auxiliary Value Function

In this section, we investigate various properties of the auxiliary value function w , which is defined in (3.7). In particular, we demonstrate that w is the unique locally Hölder continuous viscosity solution of (3.8) which satisfies certain growth conditions to be specified.

We also include in this section some additional properties of both practical and theoretical interest. First, we include some sufficient conditions for the auxiliary value function w to be concave in y . In these cases, we can optimize over the auxiliary value function via gradient descent. Furthermore, we include a characterization of the definition of CVaR used in this chapter in (3.4) as the unique continuous function satisfying certain intuitive properties.

3.3.1 Analytical properties of the auxiliary value function

We begin by observing a trivial boundary condition for w .

Proposition 3.1. *For each $(x, y) \in \mathbb{R} \times \mathbb{R}$, we have*

$$w(T, x, y) = x - \lambda p^{-1}(y - x)^+.$$

The proof of this is straightforward from the definition of w .

We next consider a much more subtle regularity result of the auxiliary value function. In particular, we show that w is Lipschitz in (x, y) but only locally 1/2-Hölder continuous in t .

Proposition 3.2. *There exists $C > 0$, which depends only on μ , σ , and T , such that for any $(t_1, x_1, y_1), (t_2, x_2, y_2) \in [0, T] \times \mathbb{R} \times \mathbb{R}$ we have*

$$|w(t_1, x_1, y_1) - w(t_2, x_2, y_2)| \leq C (|x_1 - x_2| + |y_1 - y_2| + (1 + |x_1| + |x_2|)|t_1 - t_2|^{1/2}).$$

That is, the auxiliary value function is Lipschitz continuous in (x, y) and locally 1/2-Hölder continuous in t .

Proof. 1. Fix any $y \in \mathbb{R}$ and $(t_1, x_1), (t_2, x_2) \in [0, T] \times \mathbb{R}$. Let $\alpha \in \mathcal{A}$ be arbitrary. Then we see

$$\begin{aligned} w(t_2, x_2, y) &\geq \mathbb{E} \left[X_T^{t_2, x_2, \alpha} - \lambda p^{-1} (y - X_T^{t_2, x_2, \alpha})^+ \right] \\ &\geq \mathbb{E} \left[X_T^{t_1, x_2, \alpha} - \lambda p^{-1} (y - X_T^{t_1, x_2, \alpha})^+ \right] - (1 + \lambda p^{-1}) \mathbb{E} [|X_T^{t_1, x_2, \alpha} - X_T^{t_2, x_2, \alpha}|]. \end{aligned}$$

Our goal is to bound the second term on the right-hand-side independently of α . By Hölder's Inequality it suffices to consider the squared-expectation. We begin by assuming that $t_1 \leq t_2$. Define $\phi_1 : [t_1, t_2] \rightarrow \mathbb{R}$ by

$$\phi_1(s) := \mathbb{E} (X_s^{t_1, x_1, \alpha} - x_2)^2$$

for each $s \in [t_1, t_2]$. Using Itô's Lemma and the growth bounds on μ and σ given in (3.3), we compute

$$\begin{aligned}
\phi_1(s) - (x_1 - x_2)^2 &= \mathbb{E} \int_{t_1}^s (2(X_u^{t_1, x_1, \alpha} - x_2) \mu(X_u^{t_1, x_1, \alpha}, \alpha_u) + \sigma(X_u^{t_1, x_1, \alpha}, \alpha_u)^2) du \\
&\leq \mathbb{E} \int_{t_1}^s (X_u^{t_1, x_1, \alpha} - x_2)^2 du + C^2 \mathbb{E} \int_{t_1}^s (1 + |X_u^{t_1, x_1, \alpha}|)^2 du \\
&\leq \mathbb{E} \int_{t_1}^s (X_u^{t_1, x_1, \alpha} - x_2)^2 du + C^2 \mathbb{E} \int_{t_1}^s (1 + |X_u^{t_1, x_1, \alpha} - x_2| + |x_2|)^2 du \\
&\leq (1 + 3C^2) \mathbb{E} \int_{t_1}^s (X_u^{t_1, x_1, \alpha} - x_2)^2 du + 3C^2(1 + |x_2|^2)(s - t_1) \\
&\leq (1 + 3C^2) \int_{t_1}^s \phi_1(u) du + 3C^2(1 + |x_2|)^2 |s - t_1|.
\end{aligned}$$

But then by Gronwall's Inequality, we see

$$\phi_1(t_2) \leq ((x_1 - x_2)^2 + 3C^2(1 + |x_2|)^2 |t_1 - t_2|) \exp((1 + 3C^2)T).$$

Similarly, we can define $\phi_2 : [t_2, T] \rightarrow \mathbb{R}$ by

$$\phi_2(s) := \mathbb{E} (X_s^{t_1, x_1, \alpha} - X_s^{t_2, x_2, \alpha})^2$$

for each $s \in [t_2, T]$. Using Itô's Lemma and the Lipschitz bounds on μ and σ given in (3.2), we compute

$$\begin{aligned}
\phi_2(s) - \phi_1(t_2) &= \mathbb{E} \int_t^s 2(X_u^{t, x_1, \alpha} - X_u^{t, x_2, \alpha}) (\mu(X_u^{t, x_1, \alpha}, \alpha_u) - \mu(X_u^{t, x_2, \alpha}, \alpha_u)) du \\
&\quad + \mathbb{E} \int_t^s (\sigma(X_u^{t, x_1, \alpha}, \alpha_u) - \sigma(X_u^{t, x_2, \alpha}, \alpha_u))^2 du \\
&\leq (2C + C^2) \mathbb{E} \int_t^s (X_u^{t, x_1, \alpha} - X_u^{t, x_2, \alpha})^2 du \\
&\leq 2(1 + C^2) \int_t^s \phi_2(u) du.
\end{aligned}$$

But then by Gronwall's Inequality, we see

$$\begin{aligned}
\mathbb{E} (X_T^{t, x_1, \alpha} - X_T^{t, x_2, \alpha})^2 &= \phi_2(T) \\
&\leq \phi_1(t_2) \exp(2(1 + C^2)T) \\
&\leq (1 + 3C^2) (|x_1 - x_2| + (1 + |x_2|)|t_1 - t_2|^{1/2})^2 \exp((3 + 5C^2)T).
\end{aligned}$$

Repeating this argument for the case $t_2 < t_1$, we see

$$\mathbb{E} \left[X_T^{t, x_1, \alpha} - \lambda p^{-1} (y_1 - X_T^{t, x_1, \alpha})^+ \right] \leq w(t, x_2, y_2) + \bar{C} (|x_1 - x_2| + (1 + |x_1| + |x_2|)|t_1 - t_2|^{1/2}),$$

if we take $\bar{C} := (1 + \lambda p^{-1})(1 + 2C) \exp(3(1 + C^2)T)$. Because $\alpha \in \mathcal{A}$ was arbitrary, we then conclude

$$w(t_1, x_1, y) \leq w(t_2, x_2, y) + \bar{C} (|x_1 - x_2| + (1 + |x_1| + |x_2|)|t_1 - t_2|^{1/2}).$$

By reversing the roles of (t_1, x_1) and (x_2, x_2) , we also see

$$|w(t_1, x_1, y) - w(t_2, x_2, y)| \leq \bar{C} (|x_1 - x_2| + (1 + |x_1| + |x_2|)|t_1 - t_2|^{1/2}).$$

2. Fix any $(t, x) \in [0, T] \times \mathbb{R}$ and $y_1, y_2 \in \mathbb{R}$. Let $\alpha \in \mathcal{A}$ be arbitrary. Then we see

$$\begin{aligned} w(t, x, y_2) &\geq \mathbb{E} \left[X_T^{t,x,\alpha} - \lambda p^{-1} (y_2 - X_T^{t,x,\alpha})^+ \right] \\ &\geq \mathbb{E} \left[X_T^{t,x,\alpha} - \lambda p^{-1} (y_1 - X_T^{t,x,\alpha})^+ \right] - \lambda p^{-1} |y_1 - y_2|. \end{aligned}$$

Because $\alpha \in \mathcal{A}$ was arbitrary, we then conclude

$$w(t, x, y_1) \leq w(t, x, y_2) + \lambda p^{-1} |y_1 - y_2|.$$

By reversing the roles of y_1 and y_2 , we also see

$$|w(t, x, y_1) - w(t, x, y_2)| \leq \lambda p^{-1} |y_1 - y_2|.$$

Of course, then for any $(t_1, x_1, y_1), (t_2, x_2, y_2) \in [0, T] \times \mathbb{R} \times \mathbb{R}$, we can put these separate results together to see

$$\begin{aligned} |w(t_1, x_1, y_1) - w(t_2, x_2, y_2)| &\leq |w(t_1, x_1, y_1) - w(t_1, x_1, y_2)| + |w(t_1, x_1, y_2) - w(t_2, x_2, y_2)| \\ &\leq \lambda p^{-1} |y_1 - y_2| + \bar{C} (|x_1 - x_2| + (1 + |x_1| + |x_2|)|t_1 - t_2|^{1/2}) \\ &\leq \bar{C} (|x_1 - x_2| + |y_1 - y_2| + (1 + |x_1| + |x_2|)|t_1 - t_2|^{1/2}), \end{aligned}$$

because $\bar{C} \geq \lambda p^{-1}$.

□

From this we immediately obtain a weaker (but more easily stated) continuity result and a linear asymptotic growth bound.

Corollary 3.2. *The auxiliary value function w is locally Hölder continuous satisfies*

$$|w(t, x, y)| \leq C (1 + |x| + |y|)$$

for all $(t, x, y) \in [0, T] \times \mathbb{R} \times \mathbb{R}$, for $C > 0$ which depends only upon μ , σ , and T .

We end this section by stating the corresponding Dynamic Programming Principle for the auxiliary value function w . As in the previous chapter, the purpose is to later pass this functional equality to smooth test functions to show that w is a viscosity solution of (3.8).

Proposition 3.3 (Dynamic Programming Principle). *Fix $(t, x, y) \in [0, T) \times \mathbb{R} \times \mathbb{R}$ and let $\{\theta^\alpha\}$ be any collection of stopping times valued in $[t, T]$ which are indexed by $\alpha \in \mathcal{A}$. Then for any $h > 0$ such that $t + h \leq T$, we have*

$$w(t, x, y) = \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[w \left((t + h) \wedge \theta^\alpha, X_{(t+h) \wedge \theta^\alpha}^{t, x, \alpha}, y \right) \right].$$

Because the controls are contained in a bounded set and we have shown the auxiliary value function w is continuous, this is a standard result which may be found, for instance, in Chapter 5 of Fleming-Soner [FS06].

3.3.2 Viscosity solution characterization

The goal of this section is to characterize the auxiliary value function w as the unique locally Hölder continuous viscosity solution of (3.8) which satisfies the growth condition from Corollary 3.2. Compared to the results of the previous chapter, this characterization is very standard.

For the sake of completeness, we recall a working definition of viscosity solution for (3.8).

Definition 3.6. *Let $u : \mathbb{R} \times [0, T] \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a continuous function.*

1. *We say that u is a viscosity supersolution of (3.8) if*

- a) $u(T, x, y) \geq x - \lambda p^{-1}(y - x)^+$ for all $(x, y) \in \mathbb{R} \times \mathbb{R}$, and
- b) For any $(t_0, x_0, y_0) \in [0, T) \times \mathbb{R} \times \mathbb{R}$, any smooth function $\phi : [0, T] \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ for which $(u - \phi)$ attains a local minimum at (t_0, x_0, y_0) , we have

$$\phi_t(t_0, x_0, y_0) + \sup_{a \in \mathbb{A}} \left[\mu(x_0, a) \phi_x(t_0, x_0, y_0) + \frac{1}{2} \sigma(x_0, a)^2 \phi_{xx}(t_0, x_0, y_0) \right] \geq 0.$$

2. *We say that u is a viscosity subsolution of (3.8) if*

- a) $u(T, x, y) \leq x - \lambda p^{-1}(y - x)^+$ for all $(x, y) \in \mathbb{R} \times \mathbb{R}$, and
- b) For any $(t_0, x_0, y_0) \in [0, T) \times \mathbb{R} \times \mathbb{R}$, any smooth function $\phi : [0, T] \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ for which $(u - \phi)$ attains a local maximum at (t_0, x_0, y_0) , we have

$$\phi_t(t_0, x_0, y_0) + \sup_{a \in \mathbb{A}} \left[\mu(x_0, a) \phi_x(t_0, x_0, y_0) + \frac{1}{2} \sigma(x_0, a)^2 \phi_{xx}(t_0, x_0, y_0) \right] \leq 0.$$

3. *We say that u is a viscosity solution of (3.8) if it is both a viscosity supersolution and a viscosity subsolution.*

Proposition 3.4. *The auxiliary value function w is a viscosity solution of (3.8).*

This follows by using the Dynamic Programming Principle functional equality from Proposition 3.3 to smooth test functions exactly as in Proposition 2.4 in the previous chapter.

Theorem 3.3. *The auxiliary value function w is the unique locally Hölder continuous viscosity solution of (3.8) for which there exists $C > 0$ such that*

$$|u(t, x, y)| \leq C(1 + |x| + |y|)$$

for all $(t, x, y) \in [0, T] \times \mathbb{R} \times \mathbb{R}$.

Again, this follows immediately from standard theory because the control set \mathbb{A} is compact and the viscosity solution w is assumed continuous and of linear asymptotic growth. For a proof of this statement, see Fleming-Soner [FS06] or Touzi [Tou13].

3.3.3 Sufficient conditions for concavity in the additional state variable

Recall that the main result of this chapter in Theorem 3.1 states that we can solve the original time-inconsistent stochastic control problem by maximizing the auxiliary value function over choice of the initial condition of an additional state variable y . Therefore, a natural condition to investigate is when the auxiliary value function w is concave in the additional state variable.

Recall that the auxiliary value function w is defined in (3.7) as a supremum over concave functions of y . Then it is not necessarily concave. However, we can demonstrate a sufficient condition for concavity.

Recall the following definition of a partial order on $L^2(\Omega, \mathbb{P})$.

Definition 3.7. *Let \mathcal{D} denote the collection of all functions $\phi : \mathbb{R} \rightarrow \mathbb{R}$ which are non-decreasing, concave, and for which there exists $C > 0$ such that $|\phi(x)| \leq C(1 + x^2)$ for all $x \in \mathbb{R}$. We define a partial ordering \preceq on $L^2(\Omega, \mathbb{P})$, called second-order stochastic dominance, as*

$$X \preceq Y$$

if and only if

$$\mathbb{E}[\phi(X)] \leq \mathbb{E}[\phi(Y)]$$

for all $\phi \in \mathcal{D}$.

We then state a general result regarding concavity of the auxiliary value function w in the additional state variable y .

Proposition 3.5. *Fix $(t, x) \in [0, T] \times \mathbb{R}$. If the map*

$$\alpha \mapsto X_T^{t,x,\alpha}$$

is concave with respect to second-order stochastic dominance, \preceq , then the map

$$y \mapsto w(t, x, y)$$

is concave.

Proof. Let $y, y' \in \mathbb{R}$ and $\theta \in [0, 1]$. Let $\alpha, \alpha' \in \mathcal{A}$ be arbitrary controls. Note that the function $\phi : \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$\phi(x) := x + \lambda p^{-1} (\theta y + (1 - \theta)y' - x)^+$$

is non-decreasing, concave, and bounded by a quadratic asymptotically. Then by the concavity with respect to \preceq , we see

$$\begin{aligned} w(t, x, \theta y + (1 - \theta)y') &\geq \mathbb{E} \left[X_T^{t, x, \theta\alpha + (1-\theta)\alpha'} - \lambda p^{-1} (\theta y + (1 - \theta)y' - X_T^{t, x, \theta\alpha + (1-\theta)\alpha'})^+ \right] \\ &= \mathbb{E} \left[\phi \left(X_T^{t, x, \theta\alpha + (1-\theta)\alpha'} \right) \right] \\ &\geq \mathbb{E} \left[\phi \left(\theta X_T^{t, x, \alpha} + (1 - \theta)X_T^{t, x, \alpha'} \right) \right] \\ &= \mathbb{E} \left[\theta X_T^{t, x, \alpha} + (1 - \theta)X_T^{t, x, \alpha'} \right] \\ &\quad - \lambda p^{-1} \mathbb{E} \left[(\theta y + (1 - \theta)y' - \theta X_T^{t, x, \alpha} - (1 - \theta)X_T^{t, x, \alpha'})^+ \right]. \end{aligned}$$

We also note that the map

$$(x, y) \mapsto (y - x)^+ = \max\{y - x, 0\}$$

is (jointly) convex as the maximum of two affine functions. Then we see

$$\mathbb{E} \left[(\theta y + (1 - \theta)y' - \theta X_T^{t, x, \alpha} - (1 - \theta)X_T^{t, x, \alpha'})^+ \right] \leq \mathbb{E} \left[\theta (y - X_T^{t, x, \alpha})^+ + (1 - \theta) (y' - X_T^{t, x, \alpha'})^+ \right].$$

Combining these two inequalities, we see

$$w(t, x, \theta y + (1 - \theta)y') \geq \theta \mathbb{E} \left[X_T^{t, x, \alpha} - \lambda p^{-1} (y - X_T^{t, x, \alpha})^+ \right] + (1 - \theta) \mathbb{E} \left[X_T^{t, x, \alpha'} - \lambda p^{-1} (y' - X_T^{t, x, \alpha'})^+ \right].$$

Because α, α' were taken to be arbitrary, this implies

$$w(t, x, \theta y + (1 - \theta)y') \geq \theta w(t, x, y) + (1 - \theta)w(t, x, y').$$

□

In general, it is difficult to verify the condition that $\alpha \mapsto X_T^{t, x, \alpha}$ is concave with respect to second order stochastic dominance. However, in the following corollary, we make note of a special case which shows up in the application at the end of this chapter.

Corollary 3.3. *Suppose that $\mu : \mathbb{R} \times \mathbb{A} \rightarrow \mathbb{R}$ is jointly concave in (x, a) and non-decreasing in x . Suppose also that $\sigma : \mathbb{R} \times \mathbb{A} \rightarrow \mathbb{R}$ is affine in (x, a) and independent of x . Then for any $(t, x) \in [0, T] \times \mathbb{R}$, the map*

$$y \mapsto w(t, x, y)$$

is concave.

Proof. By Proposition 3.5 it suffices to show that the map

$$\alpha \mapsto X_T^{t,x,\alpha}$$

is concave with respect to second order stochastic dominance. We aim to demonstrate a stronger statement – that is it concave almost-surely.

Fix $\alpha, \alpha' \in \mathcal{A}$ and $\theta \in [0, 1]$. For notational convenience, we define

$$\begin{cases} \underline{X}_s & := \theta X_s^{t,x,\alpha} + (1-\theta)X_s^{t,x,\alpha'} \\ \overline{X}_s & := X_s^{t,x,\theta\alpha+(1-\theta)\alpha'} \end{cases}$$

for each $s \in [t, T]$. Then for any $s \in [t, T]$ we can compute

$$\begin{aligned} \underline{X}_s &= x + \theta \left(\int_t^s \mu(X_u^{t,x,\alpha}, \alpha_u) du + \int_t^s \sigma(X_u^{t,x,\alpha}, \alpha_u) dW_u \right) \\ &\quad + (1-\theta) \left(\int_t^s \mu(X_u^{t,x,\alpha'}, \alpha'_u) ds + \int_t^s \sigma(X_u^{t,x,\alpha'}, \alpha'_u) dW_u \right) \\ &\leq x + \int_t^s \mu(\underline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) du + \int_t^s \sigma(\underline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) dW_u, \end{aligned}$$

where we used the fact that μ is concave and σ is affine. We can also compute

$$\begin{aligned} \overline{X}_s &= x + \int_t^s \mu(\overline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) du + \int_t^s \sigma(\overline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) dW_u \\ &= x + \int_t^s \mu(\overline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) du + \int_t^s \sigma(\underline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) dW_u, \end{aligned}$$

where we used the fact that σ is independent of x .

But then subtracting these two results, we see

$$\overline{X}_s - \underline{X}_s \geq \int_t^s \mu(\overline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) - \mu(\underline{X}_u, \theta\alpha_u + (1-\theta)\alpha'_u) du$$

for each $s \in [t, T]$. Because μ is non-decreasing in x and $\overline{X}_t = \underline{X}_t = x$, this implies

$$X_T^{t,x,\theta\alpha+(1-\theta)\alpha'} = \overline{X}_T \geq \underline{X}_T = \theta X_T^{t,x,\alpha} + (1-\theta)X_T^{t,x,\alpha'},$$

almost-surely. □

3.3.4 Some additional properties of conditional value-at-risk

The goal of this section is to examine some analytical properties of CVaR as defined in (3.4). The purpose of this is to justify our choice of definition for CVaR in terms of natural properties because there are several competing definitions in the literature.

Recall the following common choice of metric for \mathcal{P}_1 :

Definition 3.8. We define the 1-Wasserstein metric on \mathcal{P}_1 as follows: For any $\mu_1, \mu_2 \in \mathcal{P}_1(\mathbb{R})$, let

$$W_1(\mu_1, \mu_2) := \inf_{\gamma \in \Gamma(\mu_1, \mu_2)} \int_{\mathbb{R} \times \mathbb{R}} |x_1 - x_2| \gamma(dx_1, dx_2), \quad (3.9)$$

where $\Gamma(\mu_1, \mu_2)$ denotes the collection of all probability measures on $\mathbb{R} \times \mathbb{R}$ with marginals μ_1 and μ_2 on the first and second coordinates respectively.

We then immediately have a continuity result about CVaR as defined in this chapter.

Proposition 3.6. For any $\mu_1, \mu_2 \in \mathcal{P}_1(\mathbb{R})$, we have

$$|CVaR_p(\mu_1) - CVaR_p(\mu_2)| \leq p^{-1}W_1(\mu_1, \mu_2).$$

Then $CVaR_p$ is Lipschitz continuous with respect to the 1-Wasserstein metric.

Proof. Fix any $y \in \mathbb{R}$ and $\epsilon > 0$. By (3.9), there exists a probability measure γ on $\mathbb{R} \times \mathbb{R}$, with marginals μ_1 and μ_2 on the first and second coordinates respectively, such that

$$W(\mu_1, \mu_2) + \epsilon \geq \int_{\mathbb{R} \times \mathbb{R}} |x_1 - x_2| \gamma(dx_1, dx_2).$$

But then by Lemma 3.1, we have

$$\begin{aligned} y + p^{-1} \int_{\mathbb{R}} (y - x)^+ \mu_1(dx) &= y + p^{-1} \int_{\mathbb{R} \times \mathbb{R}} (y - x_1)^+ \gamma(dx_1, dx_2) \\ &\leq y + p^{-1} \int_{\mathbb{R} \times \mathbb{R}} (y - x_2)^+ \gamma(dx_1, dx_2) + p^{-1} \int_{\mathbb{R} \times \mathbb{R}} |x_1 - x_2| \gamma(dx_1, dx_2) \\ &\leq y + p^{-1} \int_{\mathbb{R}} (y - x)^+ \mu_2(dx) + p^{-1} (W_1(\mu_1, \mu_2) + \epsilon) \\ &\leq CVaR_p(\mu_2) + p^{-1}W_1(\mu_1, \mu_2) + p^{-1}\epsilon. \end{aligned}$$

Recalling that $y \in \mathbb{R}$ and $\epsilon > 0$ were both arbitrary, we conclude

$$CVaR_p(\mu_1) \leq CVaR_p(\mu_2) + p^{-1}W_1(\mu_1, \mu_2).$$

Reversing the roles of μ_1 and μ_2 , we obtain the claimed Lipschitz bound. \square

Corollary 3.4. Suppose that $F : \mathcal{P}_1(\mathbb{R}) \rightarrow \mathbb{R}$ is continuous with respect to the 1-Wasserstein metric and satisfies

$$F(\mu) = \mu((-\infty, VaR_p(\mu)])^{-1} \int_{(-\infty, VaR_p(\mu)]} x \mu(dx)$$

for any $\mu \in \mathcal{P}_1(\mathbb{R})$ which has no atoms. Then $F = CVaR_p$.

Proof. 1. We start by verifying that CVaR_p satisfies these properties. It is continuous with respect to the 1-Wasserstein metric by Proposition 3.6. Let $\mu \in \mathcal{P}_1(\mathbb{R})$ be a measure with no atoms. Then in particular, $\mu(\{y^*\}) = 0$, where we let $y^* := \text{VaR}_p(\mu)$. Then we can compute

$$\begin{aligned} \text{CVaR}_p(\mu) &= y^* - p^{-1} \int_{\mathbb{R}} (y^* - x)^+ \mu(dx) \\ &= p^{-1} \left[\int_{(-\infty, y^*)} x \mu(dx) + y^* (p - \mu((-\infty, y^*))) \right] \\ &= p^{-1} \int_{(-\infty, y^*]} x \mu(dx) + p^{-1} y^* (p - \mu((-\infty, y^*])). \end{aligned}$$

By the same arguments as in the proof of Lemma 3.1, we see that $\mu((-\infty, y^*)) \leq p$ and $\mu((-\infty, y^*]) \geq p$. But because $\mu(\{y^*\}) = 0$, we can check

$$p \leq \mu((-\infty, y^*]) = \mu((-\infty, y^*)) \leq p.$$

Putting this together with the computation above, we see

$$\text{CVaR}_p(\mu) = \mu((-\infty, y^*])^{-1} \int_{(-\infty, y^*]} x \mu(dx).$$

2. Suppose that $F : \mathcal{P}_1 \rightarrow \mathbb{R}$ is any other function which is continuous with respect to the 1-Wasserstein metric and satisfies the property in the statement. Let $\mu \in \mathcal{P}_1(\mathbb{R})$ be any probability measure for which $F(\mu) \neq \text{CVaR}_p(\mu)$.

For any $\epsilon > 0$, define $\mu_\epsilon \in \mathcal{P}_1$ to be distribution of the sum of a draw from μ and an independent draw from a mean-zero normal distribution with variance ϵ . We can write μ_ϵ explicitly in terms of a convolution as

$$\mu_\epsilon(A) := \int_{\mathbb{R}} \int_{\mathbb{R}} 1_A(x+y) \epsilon^{-1} \phi(\epsilon^{-1}y) \mu(dx) dy$$

for any $A \in \mathcal{B}(\mathbb{R})$. It is simple to verify that has finite first moment, so $\mu_\epsilon \in \mathcal{P}_1(\mathbb{R})$. It is also clear from the convolution formula that μ_ϵ has no atoms.

Taking $\gamma \in \Gamma(\mu, \mu_\epsilon)$ to the joint distribution of $(X, X+Y)$, where X is a draw from μ and ξ is an independent draw from a mean-zero normal distribution with variance ϵ , it is clear that the marginals of γ are μ and μ_ϵ respectively. Then we can compute

$$W_1(\mu, \mu_\epsilon) \leq \int_{\mathbb{R} \times \mathbb{R}} |x_1 - x_2| \gamma(dx_1, dx_2) \leq \sqrt{\mathbb{E}[Y^2]} = \epsilon^{1/2}.$$

Then $\mu_\epsilon \rightarrow \mu$ in the 1-Wasserstein metric as $\epsilon \rightarrow 0$.

But then $F(\mu_\epsilon) = \text{CVaR}_p(\mu_\epsilon)$ for all $\epsilon > 0$ because $\mu_\epsilon \in \mathcal{P}_1(\mathbb{R})$, but by the continuous of each function, $F(\mu_\epsilon) \rightarrow F(\mu)$ and $\text{CVaR}_p(\mu_\epsilon) \rightarrow \text{CVaR}_p(\mu)$ as $\epsilon \rightarrow 0$. The contradicts the assumption that $F(\mu) \neq \text{CVaR}_p(\mu)$.

□

3.4 Application to Mean-CVaR Portfolio Optimization

In this section, we illustrate a practical use of our main results in an application to portfolio optimization under a Mean-CVaR objective. Our goal is to ultimately use this methodology to compute the efficient frontier representing the trade-off between maximizing expected log-return and minimizing the CVaR of losses. We emphasize that dynamic optimization can significantly reduce CVaR while maintaining the same expected return as compared to optimal static investment strategies.

3.4.1 Problem formulation

Consider a market consisting of n risky assets evolving via the SDE

$$\frac{dS_t^{(i)}}{S_t^{(i)}} = \mu_i dt + \Sigma_{i,j}^{1/2} dW_t^{(j)}$$

for each $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, d\}$. Here $\mu \in \mathbb{R}^n$ is a vector of drifts and Σ is the covariance matrix of returns. The covariance matrix is assumed positive semi-definite so we take $\Sigma^{1/2}$ to denote its Cholesky Decomposition. We also assume that there exists a risk-free asset with drift r .

We assume that we choose a control α , which is a progressively-measurable process lying in some compact set \mathbb{A} , representing the percent of the portfolio exposed to each of the n risky assets. For example, we might choose $\mathbb{A} := \{a \in \mathbb{R}^n \mid a^\top \Sigma a \leq l\}$ for a constant l corresponding to a hard portfolio risk cap.

With this setup, our portfolio value Z evolves via the SDE

$$\frac{dZ_t^\alpha}{Z_t^\alpha} = [r + \alpha_t^\top (\mu - r \mathbf{1})] dt + \alpha_t^\top \Sigma^{1/2} dW_t.$$

For simplicity, we consider the log value of the portfolio, $X_t^\alpha := \log Z_t^\alpha$, which can be seen to solve the SDE

$$dX_t^\alpha = \left[r + \alpha_t^\top (\mu - r \mathbf{1}) - \frac{1}{2} \alpha_t^\top \Sigma \alpha_t \right] dt + \alpha_t^\top \Sigma^{1/2} dW_t. \quad (3.10)$$

Without loss of generality, we assume $Z_0^\alpha = S_0 = 1$. Then, $X_0^\alpha = 0$ and we can interpret X_t^α as the log-returns of the portfolio up to time t .

In this section, we consider the problem of maximizing a Mean-CVaR objective,

$$p^* := \sup_{\alpha \in \mathcal{A}} [\mathbb{E} [X_T^{0,0,\alpha}] + \lambda \text{CVaR}_p [X_T^{0,0,\alpha}]] \quad (3.11)$$

for fixed $\lambda \geq 0$ and $p \in (0, 1)$. By varying λ , we can compute a subset of the efficient frontier between expected log-return and the CVaR of returns.

3.4.2 Solution via gradient descent

By Theorem 3.1, the problem (3.11) is equivalent to the bi-level optimization

$$p^* = \sup_{y \in \mathbb{R}} (w(0, 0, y) + \lambda y), \quad (3.12)$$

where

$$w(t, x, y) := \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[X_T^{t,x,\alpha} - \lambda p^{-1} (y - X_T^{t,x,\alpha})^+ \right].$$

By Theorem 3.3, the auxiliary value function w is the unique locally Hölder continuous viscosity solution of the HJB PDE (3.8) satisfying certain growth bounds. In particular, the PDE is independent of y except in the terminal condition, so in practice we can compute $w(0, 0, y)$ using only a grid in (t, x) and a fixed value of $y \in \mathbb{R}$.

It is simple to check that the dynamics in (3.10) satisfy the conditions of Corollary 3.3. Then the auxiliary value function w is also concave in y . Therefore, we conclude that it suffices to obtain a maximum in (3.12) by gradient descent along y because the objective function is concave in y .

3.4.3 Numerical results

In this section we consider a concrete example involving selection between a single risky asset, representing a US stock index, and a risk-free asset. We compute an efficient frontier representing the trade-off between expected log-return and CVaR when using optimal dynamic strategies. For comparison, we compare to an efficient frontier when restricting to static strategies, i.e. strategies where A is constant over time, representing a fixed leverage ratio.

For our example, we choose $\mu = 11\%$, $\sigma = 20\%^2$, and $r = 1\%$ as market parameters. We take our time horizon as $T = 1$ and constrain our leverage ratio to lie within the range $\mathbb{A} := [-6, +6]$.³ Finally, we consider CVaR at the $\alpha = 95\%$ threshold.

For each fixed $\lambda > 0$, we solve the corresponding dynamic mean-CVaR optimization problem using the techniques outlined in the previous section. To obtain numerical solutions of (3.8), we employ finite-difference solvers with upwinding to guarantee a monotone scheme (See Courant-Isaacson-Rees [CIR52] and Barles-Souganidis [BS91]). For the purposes of this chapter, we obtain numerical supergradients in y through a finite-difference approximation. In Miller-Yang [MY15], there is a nuanced discussion of how to obtain a supergradient by solving a PDE corresponding to the formal linearization of (3.8).

²This choice corresponds, roughly, to the historical arithmetic mean and standard deviation of annual returns on the S&P 500, including dividend reinvestment, over the period 1928–2014. However, we emphasize that the exact choice of parameters should not be taken too seriously in this example.

³We choose this range to correspond, roughly, to the maximum leverage a qualifying US investor can achieve with a portfolio margin policy, as described at <http://www.finra.org/industry/portfolio-margin-faq>. In practice, the exact constraints depend upon the type of investor and financial instruments used for investment. We emphasize that this choice is meant for illustration only.




Figure 3.1: The efficient frontier of Mean-CVaR portfolio optimization, representing the possible trade-off between maximizing expected log-returns and minimizing CVaR, computed by varying $\lambda \in (0, 1]$.

We compute points on the efficient frontier between expected log-return and CVaR by varying λ over the interval $(0, 1]$. For the purposes of this chapter, we compute expected log-return and corresponding CVaR using Monte Carlo simulation of optimal trajectories for each fixed value of λ . The resulting frontier is shown in Figure 3.1 (solid).

For comparison, we consider the same optimization problem when restricted to a subcollection of static controls, defined as

$$\mathcal{A}_{\text{static}} := \{A \in \mathcal{A} \mid \text{there exists } a \in \mathbb{A} \text{ such that } A(t) = a \text{ for all } t \in [0, T] \text{ a.s.}\}.$$

These strategies represent constant leverage portfolios. An important example of these is the “buy and hold” strategy, e.g. $A(t) \equiv 1$. Under this class of controls, X_T^A is normally-distributed. Therefore, we can directly compute optimal strategies and construct the efficient frontier.

In Figure 3.1, we illustrate a comparison between the efficient frontier under our dynamic strategies and under static strategies. We see that by employing strategies with dynamic leverage, we can significantly reduce CVaR at the 95% quantile while maintaining the same expected log-return, as compared to a static leverage strategy. Similarly, we can increase expected log-return while maintaining the same CVaR using a dynamic strategy. For example, the static buy-and-hold strategy, $A(t) \equiv 1$, has an expected log-return of 9% and CVaR of approximately 32%. By employing strategies with dynamic leverage, we can reduce CVaR by approximately 50% while maintaining the same expected log-return, or alternatively increase expected log-return by approximately 30% while maintaining the same CVaR.

We next turn our attention to an examination of statistical and qualitative properties of the optimal dynamic control and resulting returns. In Figure 3.2, we illustrate the cumulative distribution function (CDF) of X_T^A under the optimal dynamic control corresponding an expected log-return of 9%. We compare this to the CDF of X_T^A under the buy-and-hold strategy, which follows a normal distribution. While both of these distributions have the




Figure 3.2: The cumulative distribution function of X_T^A when following the static buy-and-hold strategy and the optimal dynamic strategy which achieves the same expected log-return.




Figure 3.3: (a) A sample path of stock prices and the corresponding portfolio log-return process (X^{A^*}), and (b) the corresponding optimal leverage process (A^*).

same expected value, the one corresponding to the optimal dynamic strategy has significantly fatter (right) tails on the upside and an effective (left) floor on losses on the downside. We attribute this to a (de-)leveraging effect of the dynamic strategy whereby it increases leverage significantly once it has “locked in” gains and will de-leverage only as needed to discourage losses exceeding a certain threshold.

This qualitative tendency of the optimal strategy to increase in leverage once it has locked in gains is emphasized further by sample paths illustrated in Figure 3.3. Here, we illustrate a particular sample path of stock prices (quoted as log-return), as well as the




Figure 3.4: (a) A sample path of stock prices and the corresponding portfolio log-return process (X^{A^*}), and (b) the corresponding optimal leverage process (A^*).

corresponding optimal dynamic leverage process, A^* , and the resulting portfolio log-return process, X^{A^*} . Note that the stock price corresponds to the log-return under the static buy-and-hold strategy, $A(t) \equiv 1$. We observe that early on in the period, the leverage process increases or decreases in sync with overall portfolio returns. However, as it becomes later in the period and the portfolio return is positive, the optimal leverage increases significantly before being capped at a fixed value. The optimal strategy generally does not appear to decrease leverage late in the period, even with stock price declines, unless it is risking falling below the loss threshold seen in the jump in Figure 3.2.

In Figure 3.4, we illustrate an alternative sample path which emphasizes how the increasing leverage can lead to large returns on the upside. In this path, the leverage process, A^* , initially decreases to lower risk as the portfolio takes initial losses. However, in the latter half of the period, as stock prices rise, the increasing leverage leads to a return on the portfolio which significantly exceeds that of the buy-and-hold strategy. It is this transition from low leverage when avoiding tail losses to high leverage when locking in gains which allows the strategy to maintain a low CVaR while maximizing expected log-return.

The tendency of the optimal dynamic strategy to keep leverage higher than the static strategy unless it is facing losses also helps explain the skew seen in Figure 3.2. Because the dynamic strategy has the option to decrease its leverage to stop losses, it can achieve a significantly lower CVaR while maintaining a preference for high leverage, which contributes to large returns in positive outcomes. However, there is no such thing as a free lunch; in neutral outcomes, the positive correlation between log-returns and leverage leads to decay in portfolio value from convexity (See Perold-Sharpe [PS88]). In this sense, the optimal

dynamic strategy shares many qualitative features with Constant Proportion Portfolio Insurance (CPPI) strategies (See Black-Perold [BP92]). This makes sense as CPPI strategies are generally employed to limit downside losses, while maintaining upside gains, using dynamic trading.

Chapter 4

Distribution-Constrained Optimal Stopping

4.1 Introduction

The following chapter is based upon the joint work of Bayraktar-Miller [BM16], in which we consider the problem of choosing an optimal stopping time for a Brownian motion when constrained in the choice of distribution for the stopping time. We demonstrate that if the stopping time is constrained to have a distribution consisting of finitely-many atoms then this problem can be re-written as a sequence of time-consistent state-constrained optimal stochastic control problems.

4.1.1 Problem formulation

In this chapter, we consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which supports a standard Brownian motion W . We let $\mathbb{F} := \{\mathcal{F}_t\}_{t \geq 0}$ denote filtration, which is assumed to be right-continuous and have all \mathbb{P} -negligible sets contained in \mathcal{F}_0 . We consider a given pay-off function $f : \mathbb{R} \rightarrow \mathbb{R}$ which is assumed to be Lipschitz continuous. We also use the notation

$$X_t^x := x + W_t$$

for any $(x, t) \in \mathbb{R} \times [0, \infty)$.

In this chapter, we are also given a target distribution μ , which is supported on $(0, \infty)$ and assumed to consist of finitely-many atoms. Without loss of generality, we assume the following representation:

$$\mu = \sum_{k=1}^r p_k \delta_{t_k}, \quad (4.1)$$

where $r \in \mathbb{N}$, $0 = t_0 < t_1 < \dots < t_r$, $p_1 + \dots + p_r = 1$, and $p_1, \dots, p_r > 0$. We also introduce the convenient notation $\Delta t_k := t_k - t_{k-1}$ for each $k \in \{1, \dots, r\}$. For some fixed $x_0 \in \mathbb{R}$, we define the main problem considered in this chapter.

Definition 4.1. *The distribution-constrained optimal stopping problem is to compute*

$$p^* := \sup_{\tau \in \mathcal{T}(\mu)} \mathbb{E}[f(X_\tau^{x_0})],$$

where we take $\mathcal{T}(\mu)$ to be the collection of all finite-valued \mathbb{F} -stopping times whose distribution is equal to μ , and to find $\tau^* \in \mathcal{T}(\mu)$ which attains the supremum.

That is, we choose a stopping time τ whose distribution is equal to μ in order to maximize the expected pay-off of a stopped Brownian motion starting at x_0 .

We note that, for simplicity of notation, we often choose to write the distribution-constrained optimal stopping problem interchangeably as

$$\begin{aligned} p^* &= \sup_{\tau \in \mathcal{T}} \mathbb{E}[f(X_\tau^{x_0})] \\ \text{s.t.} \quad &\tau \sim \sum_{k=1}^r p_k \delta_{t_k}, \end{aligned}$$

where \mathcal{T} is the collection of all finite-valued \mathbb{F} -stopping times.

4.1.2 Overview of previous literature

While standard optimal stopping theory has focused primarily on unconstrained finite- and infinite-horizon stopping times (e.g., Peskir-Shiryaev [PS06] and Shiryaev [Shi08]) and very recently on constraints on the first moment of the stopping time (e.g. Miller [Mil16], Pedersen-Peskir [PP13], and Ankirchner-Klein-Kruse [AKK15]), the paper on which this chapter is based was the first on the problem of optimal stopping under distribution constraints on the stopping time.

It turns out that distribution-constrained optimal stopping is a difficult problem, with stopping strategies depending path-wise on the Brownian motion in general. This is to be expected because a constraint on the stopping time's distribution forces the stopper to consider what he would have done along all other paths of the Brownian motion when deciding whether to stop. The main task at hand is to identify sufficient statistics and then transform the problem so that it can be analyzed by standard methods.

In this chapter we illustrate a solution in the special case that the target distribution consists of finitely-many atoms. Our approach consists of an iterative stochastic control-based solution wherein we introduce controlled processes representing the conditional distribution of the stopping time. We then characterize the value function of the distribution-constrained optimal stopping problem in terms of the value functions of a finite number of state-constrained optimal control problems.

The key mathematical contributions of this chapter lie in our proof of a dynamic programming principle relating each of the sequential optimal control problems. We provide an argument which avoids the use of measurable selections, similar to the proofs of weak dynamic programming principles in Bouchard-Touzi [BT11], Bouchard-Nutz [BN12], and Bayraktar-Yao [BY13]. However, we deal with state-constraints in a novel way which relies on some a priori regularity of the value functions (e.g. continuity and concavity in particular directions).

While the problem of distribution-constrained optimal stopping is of mathematical interest in its own right, we emphasize that there is room for applications in mathematical finance and optimal control theory. For instance, we demonstrate an application to model-free superhedging of financial derivatives when one has an outlook on the quadratic variation of an asset price. Here, the distribution on the quadratic variation corresponds to that of a stopping time by the martingale time-change methods utilized recently in Bonnans-Tan [BT13] and Galichon-Henry-Labordère-Touzi [GHLT14]. Furthermore, the problem of optimal stopping under moment constraints on the stopping time reduces to the distribution-constrained optimal stopping problem in cases where there exists a unique atomic representing measure in the truncated moment problem (e.g., Curto-Fialkow [CF91] and Lasserre [Las10]).

4.2 Main Results

In the following section we give an outline of the main results of this chapter. Given the technical nature of the proofs of Lemmas 4.3–4.5, we relegate the full details to a later section.

4.2.1 Construction of distribution-constrained stopping times

There are multiple ways to naturally represent a stopping time satisfying a distribution constraint. In this section, we outline two particular such representations and illustrate how they immediately lead to constructions of such stopping times.

We first provide a characterization of distribution-constrained stopping times in terms of a partitioning of path space into regions with specified measure. Later, we make a connection with controlled processes.

Lemma 4.1. *A stopping time τ has the distribution μ if and only if it is of the following form:*

$$\tau = \sum_{k=1}^r t_k \mathbb{1}_{A_k},$$

almost-surely, where $\{A_1, \dots, A_r\}$ partition Ω and, for each $k \in \{1, \dots, r\}$, A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$.

Proof. It is clear from the construction that such a τ is a \mathbb{F} -stopping time and $\tau \sim \mu$. The converse follows by taking a stopping time τ such that $\tau \sim \mu$ and defining the sets $A_k := \{\tau = t_k\}$ for each $k \in \{1, \dots, r\}$. □

With this in mind, we can immediately explicitly construct a stopping time with given distribution.

Corollary 4.1. *There exists a stopping time τ such that $\tau \sim \mu$.*

Proof. Define a partition $\{A_1, \dots, A_r\}$ of Ω as

$$\begin{aligned} A_1 &:= \{W_{t_1} - W_0 \leq \sqrt{t_1} \Phi^{-1}(p_1)\} \\ A_2 &:= \left\{ W_{t_2} - W_{t_1} \leq \sqrt{t_2 - t_1} \Phi^{-1} \left(\frac{p_2}{p_2 + \dots + p_r} \right) \right\} \setminus A_1 \\ &\vdots \\ A_k &:= \left\{ W_{t_k} - W_{t_{k-1}} \leq \sqrt{t_k - t_{k-1}} \Phi^{-1} \left(\frac{p_k}{p_k + \dots + p_r} \right) \right\} \setminus (A_1 \cup \dots \cup A_{k-1}) \\ &\vdots \\ A_r &:= \Omega \setminus (A_1 \cup \dots \cup A_{r-1}), \end{aligned}$$

where Φ is the cumulative distribution function of the standard normal distribution. It is clear that A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$ for each $k \in \{1, \dots, r\}$. Then, by Lemma 4.1, $\tau := \sum_{k=1}^r t_k 1_{A_k}$ defines a stopping time with $\tau \sim \mu$. \square

The proof above constructs a stopping time which roughly stops when there are events in the left-tail of a distribution. However, one could easily modify the construction to stop in right-tail events, events near the median, or on the image of any Borel set of appropriate measure under Φ .

While this construction may suggest converting the distribution-constrained optimal stopping problem into optimization over Borel sets of specified measure, we emphasize next that there is no reason to expect the stopping times to be measurable with respect to $\sigma(W_{t_1}, \dots, W_{t_r})$. In particular, in the next example, we show a construction of a distribution-constrained stopping time which is entirely path-dependent.

Corollary 4.2. *There exists a stopping time τ , independent of $(W_{t_1}, \dots, W_{t_r})$, satisfying $\tau \sim \mu$.*

Proof. Define a sequence of random variables (M_1, \dots, M_r) as

$$M_k := (t_k - t_{k-1})^{-1/2} \max_{t_{k-1} \leq s \leq t_k} \left| W_s - W_{t_{k-1}} - (s - t_{k-1}) \frac{W_{t_k} - W_{t_{k-1}}}{t_k - t_{k-1}} \right|$$

for each $k \in \{1, \dots, r\}$. Then each M_k is the absolute maximum of a Brownian bridge over $[t_{k-1}, t_k]$, scaled by the length of the time interval. In particular, each M_k is \mathcal{F}_{t_k} -measurable, independent of $(W_{t_1}, \dots, W_{t_r})$, and equal in distribution to the absolute maximum of a standard Brownian bridge on $[0, 1]$, the cumulative distribution function of which we denote by Φ_{BB} .

Define a partition $\{A_1, \dots, A_r\}$ of Ω as

$$\begin{aligned} A_1 &:= \{M_1 \leq \Phi_{BB}^{-1}(p_1)\} \\ A_2 &:= \left\{ M_2 \leq \Phi_{BB}^{-1} \left(\frac{p_2}{p_2 + \dots + p_r} \right) \right\} \setminus A_1 \\ &\vdots \\ A_k &:= \left\{ M_k \leq \Phi_{BB}^{-1} \left(\frac{p_k}{p_k + \dots + p_r} \right) \right\} \setminus (A_1 \cup \dots \cup A_{k-1}) \\ &\vdots \\ A_r &:= \Omega \setminus (A_1 \cup \dots \cup A_{r-1}). \end{aligned}$$

It is clear that A_k is \mathcal{F}_{t_k} -measurable with $\mathbb{P}[A_k] = p_k$ for each $k \in \{1, \dots, r\}$. Then, by Lemma 4.1, $\tau := \sum_{k=1}^r t_k 1_{A_k}$ defines a stopping time with $\tau \sim \mu$ which is independent of $(W_{t_1}, \dots, W_{t_r})$. \square

Clearly, the stopping time constructed above is an admissible stopping time in the distribution-constrained optimal stopping problem, but there is no hope to express it in terms of the value of the Brownian motion at each potential time to stop. While stopping times involving the Brownian bridge may seem unnatural at first, their use is a key idea in the proofs of Lemma 4.3 and Lemma 4.4.

It turns out that we can obtain a more manageable representation if we introduce an extra controlled processes which represent the conditional probability of the stopping time taking on each possible value. This vector-valued stochastic process is a martingale in a probability simplex. In the next result, we make clear the connection between this process and a distribution-constrained stopping time.

It turns out that we can obtain a more manageable representation if we introduce extra controlled processes which represent the conditional probability of the stopping time taking on each possible value. This vector-valued stochastic process is a martingale in a probability simplex. In the next result, we make clear the connection between this process and a distribution-constrained stopping time.

In the remainder of the chapter, we define

$$\mathcal{A} := \left\{ \alpha : \Omega \times [0, \infty) \rightarrow \mathbb{R}^r \mid \alpha \text{ is progressively-measurable and } \mathbb{E} \int_0^\infty \|\alpha_t\|^2 dt < +\infty \right\}.$$

For any choice of $y \in \mathbb{R}^r$ and $\alpha \in \mathcal{A}$, we denote

$$Y_t^{y,\alpha} := y + \int_0^t \alpha_s dW_s,$$

for all $t \in [0, \infty)$. When needed, we will denote the k th coordinate of this vector-valued process by $Y^{(k),y,\alpha}$. We will occasionally abuse notation and leave out subscripts when they are clearly implied by the context.

We also denote by Δ the following closed and convex set:

$$\Delta := \{y = (y_1, \dots, y_r) \in [0, 1]^r \mid y_1 + \dots + y_r = 1\} \subset \mathbb{R}^r.$$

We then can state a lemma regarding a characterization of distribution-constrained stopping times in terms of a state-constrained controlled martingale.

Lemma 4.2. *A stopping time $\tau \in \mathcal{T}$ has the distribution μ if and only if it is of the form*

$$\tau = \min_{k \in \{1, \dots, r\}} \left\{ t_k \mid Y_{t_k}^{(k),p,\alpha} = 1 \right\},$$

almost-surely, for some $\alpha \in \mathcal{A}$ such that

$$Y_t^{p,\alpha} \in \Delta,$$

almost-surely, for all $t \geq 0$, and

$$Y_{t_k}^{(k),p,\alpha} \in \{0, 1\},$$

almost-surely, for each $k \in \{1, \dots, r\}$.

Proof. 1. Let $\alpha \in \mathcal{A}$ be a control for which $Y_t^{p,\alpha} \in \Delta$, almost-surely, for all $t \geq 0$ and $Y_{t_k}^{(k),p,\alpha} \in \{0, 1\}$, almost-surely, for each $k \in \{1, \dots, r\}$. Define τ as

$$\tau := \min_{k \in \{1, \dots, r\}} \left\{ t_k \mid Y_{t_k}^{(k),p,\alpha} = 1 \right\}.$$

It is clear from the properties above that $Y_{t_r}^{(k),p,\alpha} \in \{0, 1\}$ for every $k \in \{1, \dots, r\}$ and $Y_{t_r}^{p,\alpha} \in \Delta$, which implies that $\tau \leq t_r$, almost-surely. Then $\tau \in \mathcal{T}$, but we must check that it has μ as its distribution.

Fix $k \in \{1, \dots, r\}$ and note that

$$\mathbb{P}[\tau = t_k] = \mathbb{P} \left[\underbrace{\{Y_{t_1}^{(1),p,\alpha} = 0\} \cap \dots \cap \{Y_{t_{k-1}}^{(k-1),p,\alpha} = 0\}}_A \cap \underbrace{\{Y_{t_k}^{(k),p,\alpha} = 1\}}_B \right].$$

Note that $B \subset A$ up to a set of measure zero because in the set $B \setminus A$, we have $Y_{t_k}^{(k),p,\alpha} = 1$ as well as $Y_{t_\ell}^{(\ell),p,\alpha} = 1$ for some $\ell < k$. Because $Y^{p,\alpha}$ is a martingale constrained to Δ , this implies $Y_{t_k}^{(\ell),p,\alpha} = 1$, almost-surely, which contradicts $Y_{t_k} \in \Delta$. Then we can conclude

$$\mathbb{P}[\tau = t_k] = \mathbb{P} \left[Y_{t_k}^{(k),p,\alpha} = 1 \right] = p_k$$

because $Y_0^{(k),p,\alpha} = p_k$ and $Y_t^{(k),p,\alpha}$ is a martingale taking values zero and one at t_k .

2. Let τ be a stopping time such that $\tau \sim \mu$. Then define the $[0, 1]^r$ -valued process \bar{Y} as

$$\bar{Y}_t^{(k)} := \mathbb{E} \left[1_{\{\tau = t_k\}} \mid \mathcal{F}_t \right].$$

Note that $\bar{Y}_0 = p$. By the Martingale Representation Theorem, there exists a control $\alpha \in \mathcal{A}$ for which $Y_t^{p,\alpha} = \bar{Y}_t$, almost-surely, for all $t \geq 0$. We can then check that,

$$Y_t^{(1),p,\alpha} + \dots + Y_t^{(r),p,\alpha} = \mathbb{E} \left[1_{\{\tau = t_1\}} + \dots + 1_{\{\tau = t_r\}} \mid \mathcal{F}_t \right] = 1,$$

so $Y_t^{p,\alpha} \in \Delta$ for all $t \geq 0$, almost-surely. Finally, for any $k \in \{1, \dots, r\}$, we have $Y_{t_k}^{(k),p,\alpha} = 1_{\{\tau = t_k\}} \in \{0, 1\}$ because $\{\tau = t_k\}$ is \mathcal{F}_{t_k} -measurable.

Define a stopping time σ as

$$\sigma := \min_{k \in \{1, \dots, r\}} \left\{ t_k \mid Y_{t_k}^{(k),p,\alpha} = 1 \right\}$$

and suppose that there exists a set A of non-zero probability on which $\tau \neq \sigma$. Then for some $k, \ell \in \{1, \dots, r\}$ such that $k \neq \ell$, the set $B := A \cap \{\tau = t_k\} \cap \{\sigma = t_\ell\}$ has non-zero probability.

Suppose that $\ell < k$. Then $Y_{t_\ell}^{(\ell),p,\alpha} = 1$ on B and because $Y^{p,\alpha}$ is a martingale constrained to Δ , it follows that $Y_{t_k}^{(\ell),p,\alpha} = 1$ on B , and consequently, $Y_{t_k}^{(k),p,\alpha} = 1_{\{\tau=t_k\}} = 0$, which contradicts $\tau = t_k$ on B . On the other hand, suppose that $\ell > k$. Then $Y_{t_k}^{(k),p,\alpha} \neq 1$ on B , but because $Y_{t_k}^{(k),p,\alpha} = 1_{\{\tau=t_k\}}$ this also contradicts $\tau = t_k$ on B . Then we conclude $\tau = \sigma$, almost-surely. \square

4.2.2 Solution via iterated stochastic control

We begin this section by defining a sequence of iterated distribution-constrained optimal stopping problems.

It is convenient to define a sequence of sets which will be important in the remainder of the chapter. For each $k \in \{1, \dots, r\}$, define

$$\Delta_k := \{(y_1, \dots, y_r) \in \Delta \mid y_\ell = 0 \text{ for each } \ell \in \{1, \dots, k-1\}\} \subseteq \Delta.$$

Note that each set is closed and convex and $\Delta_{k+1} \subset \Delta_k$ for each $k \in \{1, \dots, r-1\}$.

We then define a sequence of iterated distribution-constrained optimal stopping problems.

Definition 4.2. For each $k \in \{1, \dots, r\}$, define a function $v_k : \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ as

$$v_k(x, y) := \sup_{\tau \in \mathcal{T}} \mathbb{E}[f(X_\tau^x)] \quad (4.2)$$

s.t. $\tau \sim \sum_{\ell=k}^r y_\ell \delta_{t_\ell - t_{k-1}}$.

Note that $p^* = v_1(x_0, p)$. Also, we emphasize that while each v_k is written as a function depending on an entire tuple $y = (y_1, \dots, y_r) \in \Delta_k$, we have $y_1 = \dots = y_{k-1} = 0$ by the definition of Δ_k .

Our goal is to convert these iterated distribution-constrained optimal stopping problems into iterated state-constrained stochastic control problems.

First, we record a growth and continuity estimate for each v_k .

Proposition 4.1. *There exists $C > 0$, which depends only on f and μ , for which*

$$\begin{aligned} |v_k(x, y)| &\leq C(1 + |x|) \\ |v_k(x, y) - v_k(x', y)| &\leq C|x - x'| \end{aligned}$$

for each $k \in \{1, \dots, r\}$ and all $(x, y) \in \mathbb{R} \times \Delta_k$ and $x' \in \mathbb{R}$.

We emphasize that we do not at this point have any guaranteed continuity in y . Therefore, the first inequality does not follow from the second. It is important to note that the second inequality is a Lipschitz continuity estimate which is *uniform* for all values of $y \in \Delta_k$.

Proof. Recall that f is assumed to be Lipschitz continuity. Fix $k \in \{1, \dots, r\}$ and $(x, y) \in \mathbb{R} \times \Delta_k$. Let $\tau \in \mathcal{T}$ be an arbitrary stopping time such that $\tau \sim \sum_{\ell=k}^r y_\ell \delta_{t_\ell - t_{k-1}}$ (such a stopping time exists by Corollary 4.1). Then we have

$$\begin{aligned} |\mathbb{E}[f(X_\tau^x)]| &\leq \mathbb{E}[|f(X_\tau^x)|] \\ &\leq |f(0)| + L(|x| + \mathbb{E}[|W_\tau|]) \\ &\leq |f(0)| + L(|x| + \mathbb{E}[|W_{t_r}|]) \\ &\leq |f(0)| + L\left(|x| + \sqrt{\frac{2}{\pi}t_r}\right). \end{aligned}$$

Because τ was arbitrary, we conclude

$$|v_k(x, y)| \leq \left(|f(0)| + L + \sqrt{\frac{2}{\pi}t_r}\right) (1 + |x|).$$

Similarly, for any $x' \in \mathbb{R}$, we have

$$\begin{aligned} v_k(x', y) &\geq \mathbb{E}[f(X_\tau^{x'})] \\ &\geq \mathbb{E}[f(X_\tau^x)] - L|x - x'|. \end{aligned}$$

Because τ was arbitrary, we conclude

$$v_k(x, y) \leq v_k(x', y) + L|x - x'|.$$

Reversing the roles of x and x' , we see

$$|v_k(x, y) - v_k(x', y)| \leq L|x - x'|.$$

Then the result holds for $C > 0$ sufficiently large. \square

In the remainder of the chapter, it will prove useful to consider a type of perspective map on the sets Δ_k . For each $k \in \{1, \dots, r\}$, define $P_k : \Delta_k \rightarrow \Delta_k$ as

$$P_k(y_1, \dots, y_r) := \begin{cases} (y_1, \dots, y_r) & \text{if } y_k = 1 \\ (y_{k+1} + \dots + y_r)^{-1}(0, \dots, 0, y_{k+1}, \dots, y_r) & \text{if } y_k < 1. \end{cases} \quad (4.3)$$

We note three key properties of this map.

1. For any $y \in \Delta_k \setminus \{e_k\}$, we have $P_k(y) \in \Delta_{k+1}$,
2. For any $y \in \Delta_k$, the k th coordinate of $P_k(y)$ is either zero or one, and
3. The map P_k is continuous on $\Delta_k \setminus \{e_k\}$.

We now provide a dynamic programming lemma whose proof has the same flavor of the weak dynamic programming results in Bouchard-Touzi [BT11], Bouchard-Nutz [BN12], and Bayraktar-Yao [BY13]. Compared to these previous results, we have a priori continuity of the value functions on the right-hand-side, so we do not need to consider upper- and lower-semicontinuous envelopes. However, we still need to avoid measurable selection, which is a non-trivial task in state-constrained problems. We extend the ideas of a countable covering of the state-space by balls, each associated with a nearly optimal stopping time. To deal with the state-constraints, we employ an argument that utilizes the compactness and convexity of Δ_k along with the continuity of v_{k+1} . The proof of this lemma is largely the heart of the chapter, but is quite involved, so it is relegated to a later section.

Lemma 4.3 (Dynamic Programming). *Suppose that for some $k \in \{1, \dots, r-1\}$, the value function $v_{k+1} : \mathbb{R} \times \Delta_{k+1} \rightarrow \mathbb{R}$ is continuous. Then for every $(x, y) \in \mathbb{R} \times \Delta_k$, we have*

$$v_k(x, y) = \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[Y_{\Delta t_k}^{(k), y, \alpha} f(X_{\Delta t_k}^x) + (1 - Y_{\Delta t_k}^{(k), y, \alpha}) v_{k+1}(X_{\Delta t_k}^x, Y_{\Delta t_k}^{y, \alpha}) \right] \quad (4.4)$$

*s.t. $Y_t^{y, \alpha} \in \Delta_k$ for all $t \geq 0$
 $Y_{\Delta t_k}^{(k), y, \alpha} \in \{0, 1\}$, almost-surely.*

Proof. See Section 4.4.1. □

Next, we provide an inductive lemma which shows that we may relax the terminal constraint. The proof of this idea relies on a careful construction of a perturbed martingale which satisfies the terminal constraints of the previous problem, but does not significantly change the expected pay-off. The proof of this result shares many of the key ideas as used in that of the previous lemma. For the sake of exposition, we provide this proof in the appendix as well.

Lemma 4.4 (Constraint Relaxation). *Suppose that for some $k \in \{1, \dots, r-1\}$, the value function $v_{k+1} : \mathbb{R} \times \Delta_{k+1} \rightarrow \mathbb{R}$ is continuous. Then for every $(x, y) \in \mathbb{R} \times \Delta_k$, we have*

$$v_k(x, y) = \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[Y_{\Delta t_k}^{(k), y, \alpha} f(X_{\Delta t_k}^x) + (1 - Y_{\Delta t_k}^{(k), y, \alpha}) v_{k+1}(X_{\Delta t_k}^x, P_k(Y_{\Delta t_k}^{y, \alpha})) \right] \quad (4.5)$$

s.t. $Y_t^{y, \alpha} \in \Delta_k$ for all $t \geq 0$, almost-surely,

where $P_k : \Delta_k \rightarrow \Delta_k$ is the perspective map defined in (4.3).

Note, even though $P_k(e_k) \notin \Delta_{k+1}$, the right-hand-side of (4.5) is well-defined because v_{k+1} is known to be bounded and continuous. Then there is a unique continuous extension of the map $(x, y) \mapsto (1 - y_k)v_{k+1}(x, y)$ to from $\Delta_k \setminus \{e_k\}$ to Δ_k . That is, taking the right-hand-side to be zero when $y = e_k$.

Proof. See Section 4.4.2. □

Lastly, we record an inductive lemma which provides basic regularity of the form of continuity of each value function and concavity with respect to the extra state-variables. We note that concavity is mainly used as a tool to obtain continuity in the extra state-variables, which is the key property used in the proof of Lemma 4.3. We provide this proof in the final appendix of the chapter.

Lemma 4.5 (Regularity). *Suppose that for some $k \in \{1, \dots, r-1\}$, we have $v_{k+1} \in C^0(\mathbb{R} \times \Delta_{k+1})$ and the map*

$$y \mapsto v_{k+1}(x, y)$$

is concave for each $x \in \mathbb{R}$. Then $v_k \in C^0(\mathbb{R} \times \Delta_k)$ and the map

$$y \mapsto v_k(x, y)$$

is concave for each $x \in \mathbb{R}$.

Proof. See Section 4.4.3. □

With these three lemmas in hand, we can now state the main result of this chapter.

Theorem 4.1. *The function $v_r : \mathbb{R} \times \Delta_r \rightarrow \mathbb{R}$ satisfies*

$$v_r(x, y) = \mathbb{E} [f(X_{\Delta_r}^x)]$$

for every $(x, y) \in \mathbb{R} \times \Delta_r$.

For each $k \in \{1, \dots, r-1\}$, the function $v_k : \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ is the value function of the following state-constrained stochastic control problem:

$$v_k(x, y) = \sup_{\alpha \in \mathcal{A}} \mathbb{E}^{x, y} \left[Y_{\Delta_k}^{(k), y, \alpha} f(X_{\Delta_k}^x) + (1 - Y_{\Delta_k}^{(k), y, \alpha}) v_{k+1}(X_{\Delta_k}^x, P_k(Y_{\Delta_k}^{y, \alpha})) \right]$$

s.t. $Y_t^{y, \alpha} \in \Delta_k$ for all $t \geq 0$, almost-surely,

where $P_k : \Delta_k \rightarrow \Delta_k$ is defined as in (4.3).

Of course, we then have

$$v^* = v_1(x_0, p_1, \dots, p_r).$$

Proof. It is clear that v_r has the representation above because there is only one admissible stopping time. The value function v_r is continuous by the smoothing properties of the heat equation (See Evans [Eva10]). For each fixed $x \in \mathbb{R}$, the map $y \mapsto v_r(x, y)$ is trivially concave because Δ_r is a singleton set. The result follows by iteratively applying Lemmas 4.3–4.5. □

4.2.3 Time-dependent value functions and an associated HJB equation

For the purposes of this chapter, we consider the results of Theorem 4.1 as a solution to the distribution-constrained optimal stopping problem. However, we can perform one more transformation which will put the problem in a form more amenable to practical solution via numerical methods.

In particular, we convert to a time-dependent version of the state-constrained problems, which will have a corresponding parabolic Hamilton-Jacobi-Bellman (HJB) equation.

We need to first introduce some extra notation which is specific to the time-dependent problem. In the remainder of the chapter, we will denote

$$\begin{aligned} X_u^{t,x} &:= x + W_u - W_t \\ Y_u^{t,y,\alpha} &:= y + \int_t^u \alpha_s dW_s \end{aligned}$$

for any $(t, x, y) \in [0, \infty) \times \mathbb{R} \times \mathbb{R}^r$, $u \in [t, \infty)$, and $\alpha \in \mathcal{A}$. As before, we will occasionally denote the k th coordinate of $Y^{t,y,\alpha}$ by $Y^{(k),t,y,\alpha}$.

Definition 4.3. Define a function $w_r : [0, \Delta t_r] \times \mathbb{R} \times \Delta_r \rightarrow \mathbb{R}$ as

$$w_r(t, x, y) := \mathbb{E} [f(X_{\Delta t_r}^{t,x})].$$

For each $k \in \{1, \dots, r-1\}$, define a function $w_k : [0, \Delta t_k] \times \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ as

$$\begin{aligned} w_k(t, x, y) &:= \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[Y_{\Delta t_k}^{(k),t,y,\alpha} f(X_{\Delta t_k}^{t,x}) + (1 - Y_{\Delta t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{\Delta t_k}^{t,x}, P_k(Y_{\Delta t_k}^{t,y,\alpha})) \right] \\ &\text{s.t. } Y_u^{t,y,\alpha} \in \Delta_k \text{ for all } u \geq t, \text{ almost-surely,} \end{aligned}$$

where $P_k : \Delta_k \rightarrow \Delta_k$ is defined as in (4.3).

We note an immediate relationship with the value functions of Section 4.2.2.

Proposition 4.2. For each $k \in \{1, \dots, r\}$ we have

$$v_k(x, y) = w_k(0, x, y)$$

for all $(x, y) \in \mathbb{R} \times \Delta_k$.

Proof. This result is obvious from the definition of w_k and Theorem 4.1. \square

Before stating a Dynamic Programming Principle for the time-dependent value functions, we first investigate their regularity. In particular, we aim to demonstrate that w_k is continuous on $[0, \Delta t_k] \times \mathbb{R} \times \Delta_k$ and lower semi-continuous on the boundary.

Proposition 4.3. *There exists $C > 0$, which depends only on f and μ , such that for each $k \in \{1, \dots, r-1\}$, we have*

$$w_k(t, x, y) - w_k(t', x', y) \leq C \left(|t - t'|^{1/2} + |x - x'| \right)$$

for all $(t, x, y) \in [0, \Delta t_k] \times \mathbb{R} \times \Delta_k$ and $(t', x') \in [0, \Delta t_k] \times \mathbb{R}$ such that $t' \leq t$. Furthermore,

$$w_k(t', x', y) - w_k(t, x, y) \leq C \left(\frac{|t - t'|}{\Delta t_k - t} + |x - x'| \right)$$

for all $(t, x, y) \in [0, \Delta t_k] \times \mathbb{R} \times \Delta_k$, and $(t', x') \in [0, \Delta t_k] \times \mathbb{R}$ such that $t' \leq t$.

Proof. 1. Fix $k \in \{1, \dots, r-1\}$ as well as $(t, x, y) \in [0, \Delta t_k] \times \mathbb{R} \times \Delta_k$ and $(t', x') \in [0, \Delta t_k] \times \mathbb{R}$ such that $t' \leq t$. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \geq t$, almost-surely. Define a new control $\alpha' \in \mathcal{A}$ as

$$\alpha'_u := 1_{\{u \geq t\}} \alpha_u,$$

for all $u \geq t'$. In particular, we see that $Y_u^{t',y,\alpha'} \in \Delta_k$ for all $u \in [t', \Delta t_k]$ and $Y_{\Delta t_k}^{t',y,\alpha'} = Y_{\Delta t_k}^{t,y,\alpha}$, almost-surely. Then

$$\begin{aligned} w_k(t', x', y) &\geq \mathbb{E} \left[Y_{\Delta t_k}^{(k),t',y,\alpha'} f(X_{\Delta t_k}^{t',x'}) + (1 - Y_{\Delta t_k}^{(k),t',y,\alpha'}) v_{k+1}(X_{\Delta t_k}^{t',x'}, P_k(Y_{\Delta t_k}^{t',y,\alpha'})) \right] \\ &= \mathbb{E} \left[Y_{\Delta t_k}^{(k),t,y,\alpha} f(X_{\Delta t_k}^{t',x'}) + (1 - Y_{\Delta t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{\Delta t_k}^{t',x'}, P_k(Y_{\Delta t_k}^{t,y,\alpha})) \right] \\ &\geq \mathbb{E} \left[Y_{\Delta t_k}^{(k),t,y,\alpha} f(X_{\Delta t_k}^{t,x}) + (1 - Y_{\Delta t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{\Delta t_k}^{t,x}, P_k(Y_{\Delta t_k}^{t,y,\alpha})) \right] \\ &\quad - 2C \left(\mathbb{E} \left| X_{\Delta t_k}^{t,x} - X_{\Delta t_k}^{t',x} \right| + |x - x'| \right), \end{aligned}$$

where $C > 0$ is at least as large as the Lipschitz constants for f and v_{k+1} . But recall that for Brownian motion we can find $C > 0$ such that

$$\mathbb{E} \left| X_{\Delta t_k}^{t,x} - X_{\Delta t_k}^{t',x} \right| = \mathbb{E} |W_{t'} - W_t| \leq C |t - t'|^{1/2}.$$

Using this and the fact that α was arbitrary, we then conclude

$$w_k(t', x', y) \geq w_k(t, x, y) - 2C(C+1) \left(|t - t'|^{1/2} + |x - x'| \right).$$

2. Fix $k \in \{1, \dots, r-1\}$ as well as $(t, x, y) \in [0, \Delta t_k] \times \mathbb{R} \times \Delta_k$ and $(t', x') \in [0, \Delta t_k] \times \mathbb{R}$ such that $t' \leq t$. Define $\eta := \sqrt{\frac{\Delta t_k - t'}{\Delta t_k - t}} \geq 1$. Let $\alpha' \in \mathcal{A}$ be an arbitrary control for which $Y_u^{t',y,\alpha'} \in \Delta_k$ for all $u \geq t'$, almost-surely. Define new control $\alpha \in \mathcal{A}$ as

$$\alpha_u := \eta \alpha'_{\tau_u},$$

where

$$\tau_u := \eta^2(u - t) + t'$$

for all $u \in [t, \Delta t_k]$. Note that $\alpha_u \in \mathcal{F}_{\tau_u}$ by definition. Because $\tau_u \leq u$, we then have $\alpha_u \in \mathcal{F}_u$ so it is an adapted control. We can also check by the time-change properties of the Itô Integral that

$$(W_{\Delta t_k} - W_t, Y_{\Delta t_k}^{t,y,\alpha}) \stackrel{(d)}{=} \left(\eta^{-1}(W_{\Delta t_k} - W_{t'}), Y_{\Delta t_k}^{t',y,\alpha'} \right).$$

Then $Y_u^{t,y,\alpha} \in \Delta_k$ for all $u \in [t, \Delta t_k]$, almost-surely, by the convexity of Δ_k and the martingale property of Y . Then α is an admissible control.

Then we can compute

$$\begin{aligned} w_k(t, x, y) &\geq \mathbb{E} \left[Y_{\Delta t_k}^{(k),t,y,\alpha} f(X_{\Delta t_k}^{t,x}) + (1 - Y_{\Delta t_k}^{(k),t,y,\alpha}) v_{k+1}(X_{\Delta t_k}^{t,x}, P_k(Y_{\Delta t_k}^{t,y,\alpha})) \right] \\ &\geq \mathbb{E} \left[Y_{\Delta t_k}^{(k),t',y,\alpha'} f(X_{\Delta t_k}^{t',x'}) + (1 - Y_{\Delta t_k}^{(k),t',y,\alpha'}) v_{k+1}(X_{\Delta t_k}^{t',x'}, P_k(Y_{\Delta t_k}^{t',y,\alpha'})) \right] \\ &\quad - 2C (|x - x'| + (1 - \eta^{-1}) \mathbb{E} |W_{\Delta t_k} - W_{t'}|). \end{aligned}$$

Now we proceed to bound the final term in this inequality. First, note that by the convexity of $x \mapsto x^{-1/2}$, we can bound

$$\eta^{-1} = \left(1 + \frac{t - t'}{\Delta t_k - t} \right)^{-1/2} \geq 1 - \frac{t - t'}{2(\Delta t_k - t)}.$$

Furthermore, for large enough $C > 0$, depending only upon t_r , we have $\mathbb{E} |W_u| \leq C$ for all $u \in [0, t_r]$. Then we can estimate

$$(1 - \eta^{-1}) \mathbb{E} |W_{\Delta t_k} - W_{t'}| \leq 2C \frac{t - t'}{2(\Delta t_k - t)}.$$

Putting these together and recalling that α was arbitrary, we conclude

$$w_k(t, x, y) \geq w(t', x', y) - 2C(1 + C) \left(|x - x'| + \frac{|t - t'|}{\Delta t_k - t} \right).$$

□

Then we can immediately make the following claim:

Corollary 4.3. *The function $w_r : [0, \Delta t_r] \times \mathbb{R} \times \Delta_r \rightarrow \mathbb{R}$ is continuous. For each $k \in \{1, \dots, r - 1\}$, the function $w_k : [0, \Delta t_k] \times \mathbb{R} \times \Delta_k \rightarrow \mathbb{R}$ is lower semi-continuous, concave in y , and continuous when restricted to $[0, \Delta t_k) \times \mathbb{R} \times \Delta_k$.*

Proof. The continuity of w_r is a standard result because there are no controls involved in the definition and the terminal pay-off is assumed Lipschitz. The remaining claims follow from the same argument as in the proof of Lemma 4.5 when using the estimates from Proposition 4.3. \square

The upside of this representation is that we can characterize each time-dependent value function w_k as a viscosity solution of a corresponding HJB equation. At this point, we can prove a Dynamic Programming Principle for the time-dependent value functions. While these are state-constrained stochastic control problems, we can directly use the a priori continuity of w_k in y and convexity of Δ_k as in the proof of Lemma 4.3.

For every $t \geq 0$, define \mathcal{A}_t as the sub-collection of controls in \mathcal{A} which are independent of \mathcal{F}_t . Then we have the following result.

Theorem 4.2. *Fix $k \in \{1, \dots, r-1\}$, $(t, x, y) \in [0, \Delta t_k) \times \mathbb{R} \times \Delta_k$, and any $h > 0$ such that $t + h < \Delta t_k$. Let $\{\tau^\alpha\}_{\alpha \in \mathcal{A}_t}$ be a family of stopping times independent of \mathcal{F}_t and valued in $[t, t + h]$. Then*

$$w_k(t, x, y) = \sup_{\alpha \in \mathcal{A}_t} \mathbb{E} [w(\tau^\alpha, X_{\tau^\alpha}^{t,x}, Y_{\tau^\alpha}^{t,y,\alpha})] \\ \text{s.t. } Y_u^{t,y,\alpha} \in \Delta_k \text{ for all } u \geq t, \text{ almost-surely.}$$

Proof. See the Appendix of Bayraktar-Miller [BM16]. \square

From this result, we immediately can verify that each time-dependent value function is a viscosity solution of an HJB. Once we have the Dynamic Programming Principle in hand, this result becomes reasonably standard, so we direct the interested reader to Katsoulakis [Kat94], Bouchard-Nutz [BN12], and Rokhlin [Rok14].

Proposition 4.4. *The function $w_r : [0, \Delta t_r] \times \mathbb{R} \times \Delta_r \rightarrow \mathbb{R}$ is the unique solution of the following heat equation (in reversed time):*

$$\begin{cases} u_t + \frac{1}{2}u_{xx} = 0 & \text{in } [0, \Delta t_r) \times \mathbb{R} \times \Delta_r \\ u = f & \text{on } \{t = \Delta t_r\} \times \mathbb{R} \times \Delta_r. \end{cases}$$

For each $k \in \{1, \dots, r-1\}$, if $w_k : [0, \Delta t_k] \times \mathbb{R} \times \Delta_k \times \mathbb{R}$ is a lower semi-continuous viscosity solution of the following HJB equation:

$$\begin{cases} u_t + \sup_{a \in \mathbb{A}_k(y)} [\frac{1}{2}u_{xx} + a \cdot D_y u_x + \frac{1}{2}a^\top D_y^2 u a] = 0 & \text{in } [0, \Delta t_k) \times \mathbb{R} \times \Delta_k \\ u = y_k f(x) + (1 - y_k) w_{k+1}(0, x, P_k(y)) & \text{on } \{t = \Delta t_k\} \times \mathbb{R} \times \Delta_k, \end{cases}$$

where $\mathbb{A}_k(y) := \{a \in \mathbb{R}^r \mid \exists \epsilon > 0 \text{ s.t. } y + a(-\epsilon, \epsilon) \subset \Delta_k\}$.

One would then expect to be able to prove a comparison principle for these HJB equation. Because the controls are unbounded, the Hamiltonian is potentially discontinuous, so this is not an immediate result. However, one would expect to be able to show that each w_k is the unique lower semi-continuous viscosity solution which has at most linear asymptotic growth in x . We leave the details of this procedure to future work.

4.3 Application to Superhedging with a Volatility Outlook

In this section, we consider a particular example of an application of distribution-constrained optimal stopping in mathematical finance. In particular, we consider the problem of model-free superhedging a contingent claim with payoff $f(X_T)$ using only dynamic trading in an underlying asset X .

We assume that the price process X_t is a martingale under some unknown martingale measure \mathbb{Q} , but do not specify the exact volatility dynamics. However, in this problem we assume that we have an outlook on the volatility in the form of the distribution of the quadratic variation, $\langle X \rangle_T$.¹

4.3.1 Model-free super-hedging setup

We follow the model-free setting of Galichon-Henry-Labordère-Touzi [GHLT14] and Bonnans-Tan [BT13]. Let $\Omega := \{\omega \in C([0, T], \mathbb{R}) \mid \omega_0 = 0\}$ be the canonical space equipped with uniform norm $\|\omega\|_\infty := \sup_{0 \leq t \leq T} |\omega_t|$, B the canonical process, \mathbb{Q}_0 the Wiener measure, $\mathbb{F} := \{\mathcal{F}_t\}_{0 \leq t \leq T}$ the filtration generated by B , and $\mathbb{F}^+ := \{\mathcal{F}_t^+\}_{0 \leq t \leq T}$ the right-limit of \mathbb{F} .

Fix some initial value $x_0 \in \mathbb{R}$. Then we denote

$$X_t := x_0 + B_t.$$

For any real-valued, \mathbb{F} -progressively measurable process α satisfying $\int_0^T \alpha_s^2 ds < \infty$, \mathbb{Q}_0 -a.s., we define the probability measure on (Ω, \mathcal{F}) ,

$$\mathbb{Q}^\alpha := \mathbb{Q}_0 \circ (X^\alpha)^{-1},$$

where

$$X_t^\alpha := x_0 + \int_0^t \alpha_r dB_r.$$

Then X^α is a \mathbb{Q}^α -local martingale. We denote by \mathcal{Q} the collection of all such probability measures \mathbb{Q} on (Ω, \mathcal{F}) under which X is a \mathbb{Q} -uniformly integrable martingale. The quadratic variation process $\langle X \rangle = \langle B \rangle$ is universally defined under any $\mathbb{Q} \in \mathcal{Q}$, and takes values in the set of all non-decreasing continuous functions from \mathbb{R}_+ to \mathbb{R}_+ .

Let μ be a given probability distribution of the form (4.1). Then we consider the problem:

$$\begin{aligned} \bar{U} := & \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}^{\mathbb{Q}} [f(X_T)] \\ & \text{s.t. } \langle X \rangle_T \sim \mu, \end{aligned}$$

¹We note that, while it may seem unlikely that we have an atomic measure representing our volatility outlook, this is a reasonable starting place for two reasons. It is possible to approximate more general measures by atomic measures since it is possible to prove continuity of the value function in the Wasserstein topology (See e.g. Lemma 3.1 in Cox-Källblad [CK15]). Second, pricing by allowing only a finite number of scenarios, as opposed to specifying a full continuous-valued model, is sometimes the standard in industry (e.g. the specification of rates, default, and prepayment scenarios in standard models for securitized products).

where \mathcal{Q} is a collection of admissible martingale measures. This corresponds to a model-free superhedging price in a sense made clear by the duality results in, for example, Bonnans-Tan [BT13].

4.3.2 Equivalence with distribution-constrained optimal stopping

We show that this problem is equivalent to distribution-constrained optimal stopping of Brownian motion.

Proposition 4.5. *We have*

$$\begin{aligned} \bar{U} &:= \sup_{\mathbb{Q} \in \mathcal{Q}} \mathbb{E}^{\mathbb{Q}} [f(X_T)] = \sup_{\tau \in \mathcal{T}(\mu)} \mathbb{E}^{\mathbb{Q}_0} [f(X_\tau)] \\ &\text{s.t. } \langle X \rangle_T \sim \mu, \end{aligned}$$

where \mathbb{Q}_0 is the measure under which X_t is a Brownian motion.

Proof. This argument can be found in Theorem 2.4 of Bonnans-Tan [BT13]. For completeness, we reproduce it below.

Let $\mathbb{Q} \in \mathcal{Q}$ such that the \mathbb{Q} -distribution of $\langle X \rangle_T$ is μ . It follows by the time-change martingale theorem that $X_T = x + \tilde{W}_{\langle X \rangle_T}$ where \tilde{W} is a standard Brownian motion and $\tau := \langle X \rangle_T$ is a stopping time with respect to the time-changed filtration with distribution μ . Then $\bar{U} \leq \sup_{\tau \sim \mu} \mathbb{E}^{\mathbb{Q}_0} [f(X_\tau)]$.

Let τ be a stopping time such that $\tau \sim \mu$. Define a process X^τ as

$$X_t^\tau := x + B_{\tau \wedge \frac{t}{T-t}}.$$

Then X^τ is a continuous martingale on $[0, T]$ with $\langle X^\tau \rangle_T = \tau$. Then X^τ induces a probability measure $\mathbb{Q} \in \mathcal{Q}$ such that $\langle X^\tau \rangle_T = \tau \sim \mu$. Then the opposite inequality holds. \square

Then one can obtain a model-free super-hedging price with a volatility outlook by solving the iterated stochastic control problem in Section 4.2.2.

4.3.3 Numerical example

In this section we obtain approximate numerical solutions of the distribution-constrained optimal stopping problem using finite-difference schemes.

In particular, we consider two potential outlooks on volatility. In the first, the binary outlook, we assume equal probability between a high- and low-volatility scenario:

$$\mu_2 := \frac{1}{2} \delta_{10} + \frac{1}{2} \delta_{20}.$$

In the second, we augment the binary outlook with a third extreme volatility scenario which occurs with small probability:

$$\mu_3 := \frac{9}{20}\delta_{10} + \frac{9}{20}\delta_{20} + \frac{1}{10}\delta_{100}.$$

Our goal is to compute the model-free superhedging price of a European call option under each volatility outlook. Because we do not restrict to models where the price process is non-negative, we can take the pay-off to be $f(x) := x^+$ without loss of generality.

Then, as before, we define value functions for each outlook as

$$v_2(x) := \sup_{\tau \in \mathcal{T}(\mu_2)} \mathbb{E}^x [f(W_\tau)] \text{ and } v_3(x) := \sup_{\tau \in \mathcal{T}(\mu_3)} \mathbb{E}^x [f(W_\tau)].$$

We solve the problem using the iterated stochastic control approach from Section 4.2.2. In particular, obtaining a viscosity solution of the corresponding Hamilton-Jacobi-Bellman equation in Section 4.2.3 using a finite-difference scheme. It is important to emphasize that, because of potential degeneracy due to the extra state-variables in w_2 and w_3 , it is critical to use a monotone numerical scheme.

In these results, we apply a version of the wide-stencil scheme introduced in Oberman [Obe07]. In particular, we approximate the non-linear terms in each equation by monotone finite-difference approximations of the following form:

$$\sup_{a \in \mathbb{R}} \begin{pmatrix} 1 \\ a \end{pmatrix}^\top \begin{pmatrix} u_{xx} & u_{xy} \\ u_{xy} & u_{yy} \end{pmatrix} \begin{pmatrix} 1 \\ a \end{pmatrix} \approx \max_{k \in \mathcal{K}(t,x,y)} \frac{u(x+h, t, y+k) - 2u(x, t, y) + u(x-h, t, y-k)}{h^2},$$

where the set $\mathcal{K}(t, x, y)$ is a collection such that $y \pm k$ lies on nearby grid-points. For a rigorous analysis of wide-stencil schemes for degenerate elliptic equations, we refer the reader to Oberman [Obe08; Obe07] and Froese-Oberman [FO11].

For comparison, we consider two main special cases, which we refer to as the “mean volatility” value and the “support-constrained” value. We define the mean volatility value as the model-free superhedging price obtained by assuming the quadratic variation will be equal to the mean of the distribution in the corresponding distribution-constrained problem. We define their corresponding value functions as \underline{v}_2 and \underline{v}_3 , respectively. On the other hand, we define the support-constrained value as the model-free superhedging price obtained when only restricting the quadratic variation to have the same support as that of the distribution in the corresponding distribution-constrained problem. We define their corresponding value functions as \bar{v}_2 and \bar{v}_3 , respectively.

We expect the following ordering:

$$f(x) \leq \underline{v}_2(x) \leq v_2(x) \leq \bar{v}_2(x)$$

and

$$f(x) \leq \underline{v}_3(x) \leq v_3(x) \leq \bar{v}_3(x).$$




Figure 4.1: Comparison of the model-free superhedging values for with distribution constraints on quadratic variation, support constraints on quadratic variation, and under averaged quadratic variation. Each of these is in the two-atom (binary) volatility outlook. The distribution-constrained value corresponds with the value function of an optimal stopping problem under a two-atom distribution constraint.

Furthermore, we note that we can compute \underline{v}_2 , \bar{v}_2 , \underline{v}_3 , and \bar{v}_3 explicitly in terms of heat kernels.

We illustrate the value function for the two- and three-atom problem in Figure 4.1 and Figure 4.2, respectively. As expected, we see a superhedging value which is increasing in the underlying asset price (or, equivalently, decreasing in the strike price) and respects the bounds implied by the support-constrained and average-volatility models. As expected, the bound provided by the support-constrained superhedging problem is particularly poor in the three-model volatility outlook, where we stipulate that the high volatility (high value) case is rare.

It is interesting to note that careful comparison of the two figures illustrates an increase in superhedging value between the two volatility outlooks which is roughly proportional to the increase in square-root of expected quadratic variation. For example, there is approximately a 25% increase in value at $x = 0$, which is essentially exactly in-line with the 25.2% increase in square-root of expected quadratic variation between the two outlooks. This matches our intuition that call option superhedging prices should be proportional to expected volatility to first order.

In Figure 4.3, we provide a probability density estimate of W_{10} conditional on $\tau = 10$ and $\tau = 20$ for an approximate optimal stopping time for the two-atom volatility outlook model




Figure 4.2: Comparison of the model-free superhedging values for with distribution constraints on quadratic variation, support constraints on quadratic variation, and under averaged quadratic variation. Each of these is in the three-atom (trinary) volatility outlook. The distribution-constrained value corresponds with the value function of an optimal stopping problem under a three-atom distribution constraint.

starting from $W_0 = 0$. We obtain these estimates by performing Monte Carlo simulations with controls estimated from a numerical solution of the associated HJB equations. We use grid spacings $dx = 0.1$, $dy = 0.005$, and $dt = 0.01$. We perform 10^7 simulations and verify that relevant statistics from the Monte Carlo simulation match those from the finite-difference solutions (e.g. expected pay-off, distribution and moments of the stopping time and stopped process) to within a reasonable margin of error.

The density estimates provide insight into form of an optimal strategy. Recall, the payoff is locally-affine at all points except $x = 0$, where it is strictly convex instead. Then we expect an optimal stopping strategy to be one which maximizes local time accumulated at the origin. As expected, we find that the density of W_{10} conditional upon $\tau = 10$ is largely concentrated on points away from $x = 0$, at which the pay-off process is unlikely to spend significant time as a sub-martingale if we were to choose not to stop.

It is interesting to note the lack of sharp cut-off between the two density estimates. One might expect the optimal strategy is of a form where there exists a “stopping region” and a “continuation region.” On the contrary, the smooth overlap of the two density estimates is persistent even as we vary the resolution of the finite-difference solver, which suggests that the true optimal stopping strategy is not of the form $\{\tau = 10\} \subset \sigma(W_{10})$. The numerics suggest that optimal stopping strategies may be path-dependent even in simple examples.




Figure 4.3: Probability density estimates of W_{10} conditional on $\tau = 10$ and $\tau = 20$ for an optimal stopping time for the two-atom volatility outlook model starting from $W_0 = 0$. Density estimates were made by Monte Carlo simulations on high-resolution solutions to the associated HJB equations. Sample size, $N = 10^7$.

4.4 Proof of Lemmas 4.3–4.5

In this section, we provide a full proof of the three main lemmas contained in this chapter.

4.4.1 Proof of Lemma 4.3

This first argument is in the spirit of proofs of the weak dynamic programming which avoid measurable selection, as in Bouchard-Touzi [BT11], Bouchard-Nutz [BN12], and Bayraktar-Yao [BY13]. In these arguments, the authors typically use a covering argument to find a countable selection of ϵ -optimal controls on small balls of the state-space. The main difficulty here is that, while a control may be admissible for the state-constrained problem at one point in state-space, there is no reason to expect it to satisfy the state constraints starting from nearby states.

The new idea in our approach is to cover Δ_{k+1} with a finite mesh. We show that we can replace the process Y by a modified process Y^ϵ , which lies on the mesh points almost-surely at the terminal time. We construct the new process in a measurable way using the Martingale Representation Theorem on a carefully constructed random variable. Then we show that, using the continuity of v_{k+1} , that the objective function along Y is close to that along Y^ϵ for a fine enough grid.

Once we know we can consider a perturbed process Y^ϵ which lies on a finite number of points in Δ_{k+1} at the terminal time almost-surely, we can construct ϵ -optimal stopping times using a standard Lindelöf covering argument in \mathbb{R} .

Proof. Fix $(x, y) \in \mathbb{R} \times \Delta_k$. For convenience of notation, define $\theta := \Delta t_k$ and

$$A := \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha}) \right]$$

s.t. $Y_t^{y,\alpha} \in \Delta_k$ for all $t \geq 0$
 $Y_\theta^{(k),y,\alpha} \in \{0, 1\}$, almost-surely.

Let $\epsilon > 0$, $R > 0$, $\delta > 0$, and $h > 0$ be constants to be fixed later.

1. We start by constructing a finite mesh on Δ_{k+1} . By the continuity of v_{k+1} , we can take δ small enough such that

$$|v_{k+1}(x', y') - v_{k+1}(x', y'')| \leq \epsilon$$

for every $x' \in \mathbb{R}$ and $y', y'' \in \Delta_{k+1}$ such that $|x - x'| \leq R$ and $|y' - y''| \leq \delta$. Let $\mathcal{P} := \{y_j\}_{j=1}^N$ be a finite subset of Δ_{k+1} with the property that

- The convex hull of \mathcal{P} is Δ_{k+1} , and
- Any point $y \in \Delta_{k+1}$ can be written as a convex combination of finitely-many points in \mathcal{P} , each contained in a δ -neighborhood of y .

This is possible by compactness and convexity of Δ_{k+1} . In particular, we can define a continuous function $T : \Delta_{k+1} \rightarrow [0, 1]^N$ with the properties that

- $T_j(y) = 0$ for all $y \in \Delta_{k+1}$ such that $|y - y_j| > \delta$
- $\sum_{j=1}^N T_j(y) = 1$ for all $y \in \Delta_{k+1}$, and
- $\sum_{j=1}^N y_j T_j(y) = y$ for all $y \in \Delta_{k+1}$.

This corresponds to a continuous map from a point $y \in \Delta_{k+1}$ to a probability weighting of points in \mathcal{P} such that y is a convex combination of nearby points in \mathcal{P} . Such a map can be obtained by an ℓ^2 -minimization problem, for instance.

2. Let $L > 0$ denote the Lipschitz constant of f . Recall by Proposition 4.1 that L is also a Lipschitz constant for v_{k+1} in x . Let $\{A_i\}_{i \geq 1}$ be a countable and disjoint covering of \mathbb{R} with an associated set of points $\{x_i\}$ such that the ball of diameter ϵL^{-1} centered at x_i contains the set A_i .

For each $i \geq 1$ and $j \in \{1, \dots, N\}$, let $\tau_{i,j}$ be a stopping time satisfying $\tau_{i,j} \sim \sum_{\ell=k+1}^r y_j^{(\ell)} \delta_{t_\ell - t_k}$ such that

$$\mathbb{E} \left[f(X_{\tau_{i,j}}^{x_i}) \right] \geq v_{k+1}(x_i, y_j) - \epsilon.$$

By the Lipschitz continuity of f and v_{k+1} and the definition of the sets A_i , we have

$$\begin{aligned} v_{k+1}(x_i, y_j) &\geq v_{k+1}(x, y_j) - \epsilon \\ \mathbb{E} \left[f(X_{\tau_{i,j}}^x) \right] &\geq \mathbb{E} \left[f(X_{\tau_{i,j}}^{x_i}) \right] - \epsilon \end{aligned}$$

for all $x \in A_i$.

Putting these inequalities together, we conclude that

$$\mathbb{E} \left[f(X_{\tau_{i,j}}^x) \right] \geq v_{k+1}(x, y_j) - 3\epsilon$$

for all $i \geq 1$, $j \in \{1, \dots, N\}$, and $x \in A_i$.

3. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_t^{y,\alpha} \in \Delta_k$ for $t \geq 0$ and $Y_\theta^{(k),y,\alpha} \in \{0, 1\}$ almost-surely. For any $0 < h \ll \theta$, define two random variables, M_1 and M_2 , as

$$\begin{aligned} M_1 &:= h^{-1/2} (W_\theta - W_{\theta-h}) \\ M_2 &:= h^{-1/2} \max_{\theta-h \leq s \leq \theta} |W_s - W_{\theta-h} - \delta^{-1} (s - \theta + h) (W_\theta - W_{\theta-h})|. \end{aligned} \tag{4.6}$$

Then M_1 and M_2 are \mathcal{F}_θ -measurable and independent of each other. M_1 is equal in distribution to a standard normal distribution, the cumulative distribution function of which we denote by Φ . Similarly, M_2 is equal in distribution to the absolute maximum

of a standard Brownian bridge on $[0, 1]$, the cumulative distribution function of which we denote by Φ_{BB} . Furthermore, if we define $\mathcal{G} := \sigma(\mathcal{F}_{\theta-h} \cup \sigma(W_\theta))$, then M_1 is \mathcal{G} -measurable, while M_2 is independent of \mathcal{G} .

Define a random vector \bar{Y}_θ as

$$\bar{Y}_\theta^{(k)} := 1_{\{M_2 \leq \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}}$$

and

$$\bar{Y}_\theta^{(k+1):r} := 1_{\{M_2 > \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}} \sum_{j=1}^N y_j 1_{\{\Phi^{-1}(\sum_{i=1}^{j-1} T_i(P_k(Y_{\theta-h}^{y,\alpha}))) < M_1 \leq \Phi^{-1}(\sum_{i=1}^j T_i(P_k(Y_{\theta-h}^{y,\alpha})))\}},$$

where we follow the conventions that $\Phi^{-1}(0) = -\infty$, $\Phi^{-1}(1) = +\infty$, and that sums over an empty set are zero. Then $\bar{Y}_\theta \in \Delta_k$ is \mathcal{F}_θ -measurable and is constructed to have the key property that $\mathbb{E}[\bar{Y}_\theta \mid \mathcal{F}_{\theta-h}] = Y_{\theta-h}^{y,\alpha}$, almost-surely.

By the Martingale Representation Theorem, there exists $\alpha_\epsilon \in \mathcal{A}$ for which $Y_\theta^{y,\alpha_\epsilon} = \bar{Y}_\theta$ almost-surely. Then, by construction, $Y_t^{y,\alpha_\epsilon} \in \Delta_k$ for all $t \geq 0$, $Y_\theta^{(k),y,\alpha_\epsilon} \in \{0, 1\}$, and $Y_\theta^{y,\alpha_\epsilon} \in \mathcal{P}$ when $Y_\theta^{(k),y,\alpha_\epsilon} = 0$, almost-surely.

We now perform a key computation. First note that

$$\begin{aligned} & \mathbb{E} \left[Y_\theta^{(k),y,\alpha_\epsilon} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha_\epsilon}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] \\ &= \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=1\}} f(X_\theta^x) \right] + \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} v_{k+1}(X_\theta^x, \bar{Y}_\theta) \right]. \end{aligned}$$

For the first term on the right-hand-side, we simply compute

$$\begin{aligned} \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=1\}} f(X_\theta^x) \right] &= \mathbb{E} \left[1_{\{M_2 \leq \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}} f(X_\theta^x) \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[1_{\{M_2 \leq \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}} \mid \mathcal{G} \right] f(X_\theta^x) \right] \\ &= \mathbb{E} \left[Y_{\theta-h}^{(k),y,\alpha} f(X_\theta^x) \right]. \end{aligned}$$

We deal with the second term in a similar way, but the computation is more involved. Note that by construction we have

$$\|\bar{Y}_\theta - P_k(Y_{\theta-h}^{y,\alpha})\|_{\ell^\infty} \leq \delta$$

almost-surely in the set $\{\bar{Y}_\theta^{(k)} = 0\}$. Recall we also took δ small enough such that

$$|v_{k+1}(x', y') - v_{k+1}(x', y'')| \leq \epsilon$$

for all $x' \in \mathbb{R}$ and $y', y'' \in \Delta_{k+1}$ such that $|x - x'| \leq R$ and $|y' - y''| \leq \delta$. But then we can compute

$$\begin{aligned}
 & \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} v_{k+1}(X_\theta^x, \bar{Y}_\theta) \right] \\
 &= \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} 1_{\{|W_\theta| \leq R\}} v_{k+1}(X_\theta^x, \bar{Y}_\theta) \right] + \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} 1_{\{|W_\theta| \geq R\}} v_{k+1}(X_\theta^x, \bar{Y}_\theta) \right] \\
 &\geq \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} 1_{\{|W_\theta| \leq R\}} v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\
 &\quad + \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} 1_{\{|W_\theta| \geq R\}} v_{k+1}(X_\theta^x, \bar{Y}_\theta) \right] - \epsilon \\
 &\geq \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] - \\
 &\quad \mathbb{E} \left[1_{\{|W_\theta| \geq R\}} (|v_{k+1}(X_\theta^x, \bar{Y}_\theta)| + |v_{k+1}(X_\theta^x, Y_{\theta-h}^{y,\alpha})|) \right] - \epsilon \\
 &\geq \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] - \sqrt{\mathbb{P}[|W_\theta| \geq R]} \sqrt{2C(1 + |x|)} - \epsilon \\
 &\geq \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] - R^{-1} \sqrt{C\theta(1 + |x|)} - \epsilon.
 \end{aligned}$$

With this in hand, we now complete the analysis of the second term:

$$\begin{aligned}
 \mathbb{E} \left[1_{\{\bar{Y}_\theta^{(k)}=0\}} v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] &= \mathbb{E} \left[1_{\{M_2 > \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}} v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\
 &= \mathbb{E} \left[\mathbb{E} \left[1_{\{M_2 > \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}} \mid \mathcal{G} \right] v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\
 &= \mathbb{E} \left[(1 - Y_{\theta-h}^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right].
 \end{aligned}$$

Using the continuity of f , v_{k+1} , and P_k , along with the Dominated Convergence Theorem, we note

$$\begin{aligned}
 \lim_{h \rightarrow 0} \mathbb{E} \left[Y_{\theta-h}^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_{\theta-h}^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\
 &= \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y,\alpha})) \right] \\
 &= \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha}) \right].
 \end{aligned}$$

Then putting these results together, we see that for large enough R and small enough h we have

$$\begin{aligned}
 & \mathbb{E} \left[Y_\theta^{(k),y,\alpha_\epsilon} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha_\epsilon}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] \\
 &\geq \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha}) \right] - 3\epsilon.
 \end{aligned}$$

4. Lastly, we intend to construct an ϵ -optimal stopping time using the covering from the second step. Define a stopping time τ_ϵ as

$$\tau_\epsilon := \theta + 1_{\{Y_\theta^{(k),y,\alpha_\epsilon}=0\}} \sum_{i=1}^{\infty} \sum_{j=1}^N \tau_{i,j} 1_{\{X_\theta^x \in A_i\}} 1_{\{Y_\theta^{y,\alpha_\epsilon}=y_j\}}.$$

By construction, we have $\tau_\epsilon \sim \sum_{\ell=k}^r y_\ell \delta_{t_\ell - t_{k-1}}$. We proceed to make a careful computation. First, note that

$$\mathbb{E} [f(X_{\tau_\epsilon}^x)] = \mathbb{E} [1_{\{\tau_\epsilon = \theta\}} f(X_\theta^x)] + \mathbb{E} [1_{\{\tau_\epsilon > \theta\}} f(X_{\tau_\epsilon}^x)].$$

We focus on the second term. In particular, we have

$$\begin{aligned} \mathbb{E} [1_{\{\tau_\epsilon > \theta\}} f(X_{\tau_\epsilon}^x)] &= \sum_{i=1}^{\infty} \sum_{j=1}^N \mathbb{E} \left[1_{\{Y_\theta^{(k),y,\alpha_\epsilon} = 0\}} 1_{\{X_\theta^x \in A_i\}} 1_{\{Y_\theta^{y,\alpha_\epsilon} = y_j\}} f(X_{\theta + \tau_{i,j}}^x) \right] \\ &= \sum_{i=1}^{\infty} \sum_{j=1}^N \mathbb{E} \left[1_{\{Y_\theta^{(k),y,\alpha_\epsilon} = 0\}} 1_{\{X_\theta^x \in A_i\}} 1_{\{Y_\theta^{y,\alpha_\epsilon} = y_j\}} \mathbb{E} \left[f(X_{\theta + \tau_{i,j}}^x) \mid \mathcal{F}_\theta \right] \right] \\ &\geq \sum_{i=1}^{\infty} \sum_{j=1}^N \mathbb{E} \left[1_{\{Y_\theta^{(k),y,\alpha_\epsilon} = 0\}} 1_{\{X_\theta^x \in A_i\}} 1_{\{Y_\theta^{y,\alpha_\epsilon} = y_j\}} v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] - 3\epsilon \\ &= \mathbb{E} \left[(1 - Y_\theta^{(k),y,\alpha_\epsilon}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] - 3\epsilon. \end{aligned}$$

Then we conclude

$$\mathbb{E} [f(X_{\tau_\epsilon}^x)] \geq \mathbb{E} \left[Y_\theta^{(k),y,\alpha_\epsilon} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha_\epsilon}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] - 3\epsilon.$$

Combining this with the main inequality from the previous step, we obtain

$$\begin{aligned} v_k(x, y) &\geq \mathbb{E} [f(X_{\tau_\epsilon}^x)] \\ &\geq \mathbb{E} \left[Y_\theta^{(k),y,\alpha_\epsilon} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha_\epsilon}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] - 3\epsilon \\ &\geq \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha}) \right] - 6\epsilon. \end{aligned}$$

Because ϵ and α were arbitrary, then we conclude $A \leq v_k(x, y)$.

5. Let $\tau \in \mathcal{T}$ be an arbitrary stopping time such that $\tau \sim \sum_{i=k}^r y_i \delta_{t_i - t_{k-1}}$. Define a martingale as

$$Y_t^{(i)} := \mathbb{E} [1_{\{\tau = t_i - t_{k-1}\}} \mid \mathcal{F}_t]$$

for all $t \geq 0$ and each $i \in \{k, \dots, r\}$. We can easily check that $Y_0^{(i)} = y_i$ for each $i \in \{k, \dots, r\}$ and

$$Y_t^{(k)} + \dots + Y_t^{(r)} = \mathbb{E} [1_{\{\tau = t_k - t_{k-1}\}} + \dots + 1_{\{\tau = t_r - t_{k-1}\}} \mid \mathcal{F}_t] = 1.$$

Then if we consider Y as an \mathbb{R}^r -valued martingale with $Y_t^{(i)} \equiv 0$ for all $i \in \{1, \dots, k-1\}$, then we see $Y_t \in \Delta_k$ for each $t \geq 0$. Finally, we have

$$Y_\theta^{(k)} = \mathbb{E} [1_{\{\tau = \theta\}} \mid \mathcal{F}_\theta] = 1_{\{\tau = \theta\}} \in \{0, 1\}.$$

Then by the Martingale Representation Theorem, there exists $\alpha \in \mathcal{A}$ for which $Y_t^{y,\alpha} = Y_t$ for all $t \geq 0$, almost-surely. We can compute

$$\begin{aligned} \mathbb{E}[f(X_\tau^x)] &= \mathbb{E}[1_{\{\tau=\theta\}}f(X_\theta^x) + 1_{\{\tau>\theta\}}f(X_\tau^x)] \\ &= \mathbb{E}\left[Y_\theta^{(k),y,\alpha}f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha})\mathbb{E}[f(X_\tau^x) \mid \mathcal{F}_\theta]\right]. \end{aligned}$$

On the set $\{\tau > \theta\}$, we have

$$\mathbb{P}[\tau - \theta = t_i - t_k \mid \mathcal{F}_\theta] = \mathbb{E}[1_{\{\tau=t_i-t_k\}} \mid \mathcal{F}_\theta] = Y_\theta^{(i)}$$

for each $i \in \{k+1, \dots, r\}$. For almost every $\omega \in \{\tau > \theta\}$, we have

$$\mathbb{E}[f(X_\tau^x) \mid \mathcal{F}_\theta] \leq v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha})$$

by the Strong Markov Property of Brownian motion. Then we conclude

$$\begin{aligned} \mathbb{E}[f(X_\tau^x)] &= \mathbb{E}\left[Y_\theta^{(k),y,\alpha}f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha})\mathbb{E}[f(X_\tau^x) \mid \mathcal{F}_\theta]\right] \\ &\leq \mathbb{E}\left[Y_\theta^{(k),y,\alpha}f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha})v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha})\right] \\ &\leq A. \end{aligned}$$

Because τ was an arbitrary stopping time, this implies

$$v_k(x, y) \leq A$$

□

4.4.2 Proof of Lemma 4.4

The main idea of this argument is that we can take a controlled process Y , which does not satisfy $Y_{\Delta t_k}^{(k)} \in \{0, 1\}$, and modify it on an interval $[\Delta t_k - h, \Delta t_k]$ to a perturbed process Y^ϵ with the properties that $Y_{\Delta t_k - h} = Y_{\Delta t_k - h}^\epsilon$ and $Y_{\Delta t_k}^{\epsilon, (k)} \in \{0, 1\}$. In particular, we may do this in a way that does not appreciably change the expected pay-off.

One key idea which we draw the reader's attention toward is the use of the Brownian bridge over $[\Delta t_k - h, \Delta t_k]$ in the construction. This construction is in the spirit of Corollary 4.2. While one might initially attempt a construction similar to Corollary 4.1, using a Brownian bridge instead of Brownian increments allows us to condition on $W_{\Delta t_k}$ at a key point in the argument.

Proof. Fix $(x, y) \in \mathbb{R} \times \Delta_k$. For convenience of notation, define $\theta := \Delta t_k$,

$$\begin{aligned} A &:= \sup_{\alpha \in \mathcal{A}} \mathbb{E}\left[Y_\theta^{(k),y,\alpha}f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha})v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha})\right] \\ &\text{s.t. } Y_t^{y,\alpha} \in \Delta_k \text{ for } t \geq 0 \\ &\quad Y_\theta^{(k),y,\alpha} \in \{0, 1\} \text{ almost-surely,} \end{aligned}$$

and

$$B := \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[Y_{\theta}^{(k),y,\alpha} f(X_{\theta}^x) + (1 - Y_{\theta}^{(k),y,\alpha}) v_{k+1}(X_{\theta}^x, P_k(Y_{\theta}^{y,\alpha})) \right]$$

s.t. $Y_t^{y,\alpha} \in \Delta_k$ for $t \geq 0$.

By Lemma 4.3, we have $v_k(x, y) = A$.

1. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_t^{y,\alpha} \in \Delta_k$ for $t \geq 0$ and $Y_{\theta}^{(k),y,\alpha} \in \{0, 1\}$ almost-surely. Note that $Y_{\theta}^{y,\alpha} = P_k(Y_{\theta}^{y,\alpha})$ on the set $\{Y_{\theta}^{(k),y,\alpha} = 0\}$, almost-surely. Then

$$\begin{aligned} \mathbb{E} \left[Y_{\theta}^{(k),y,\alpha} f(X_{\theta}^x) + (1 - Y_{\theta}^{(k),y,\alpha}) v_{k+1}(X_{\theta}^x, Y_{\theta}^{y,\alpha}) \right] \\ = \mathbb{E} \left[Y_{\theta}^{(k),y,\alpha} f(X_{\theta}^x) + (1 - Y_{\theta}^{(k),y,\alpha}) v_{k+1}(X_{\theta}^x, P_k(Y_{\theta}^{y,\alpha})) \right] \\ \leq B. \end{aligned}$$

Because α was arbitrary, we conclude $A \leq B$.

2. Let $\alpha \in \mathcal{A}$ be an arbitrary control for which $Y_t^{y,\alpha} \in \Delta_k$ for $t \geq 0$, almost-surely. For any $0 < h \ll \theta$, define a random variable M as

$$M := h^{-1/2} \max_{\theta-h \leq s \leq \theta} |W_s - W_{\theta-h} - \delta^{-1}(s - \theta + h)(W_{\theta} - W_{\theta-h})|.$$

Then M is \mathcal{F}_{θ} -measurable and is equal in distribution to the absolute maximum of a standard Brownian bridge on $[0, 1]$, the cumulative distribution function of which we denote by Φ_{BB} . If we define $\mathcal{G} := \sigma(\mathcal{F}_{\theta-h} \cup \sigma(W_{\theta}))$, then M_2 is independent of \mathcal{G} .

Define a random vector \bar{Y}_{θ} as

$$\bar{Y}_{\theta}^{(k)} := 1_{\{M \leq \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}}$$

and

$$\bar{Y}_{\theta}^{(k+1):r} := P_k(Y_{\theta-h}^{y,\alpha}) 1_{\{M > \Phi_{BB}^{-1}(Y_{\theta-h}^{(k),y,\alpha})\}}.$$

Let $\bar{Y}_{\theta}^{(i)} \equiv 0$ for any $i \in \{1, \dots, k-1\}$. Then \bar{Y}_{θ} is \mathcal{F}_{θ} -measurable and has the key property that $\mathbb{E}[\bar{Y}_{\theta} | \mathcal{F}_{\theta-h}] = Y_{\theta-h}^{y,\alpha}$. We also note that $\mathbb{E}[1_{\{\bar{Y}_{\theta}^{(k)}=1\}} | \mathcal{G}] = Y_{\theta-h}^{(k),y,\alpha}$.

By the Martingale Representation Theorem, there exists $\alpha_{\epsilon} \in \mathcal{A}$ such that $Y_t^{y,\alpha_{\epsilon}} \in \Delta_k$ for $t \geq 0$, $Y_{\theta}^{(k),y,\alpha_{\epsilon}} \in \{0, 1\}$, and $Y_{\theta}^{y,\alpha_{\epsilon}} = \bar{Y}_{\theta}$ almost-surely. We can then compute

$$\begin{aligned} \mathbb{E} \left[Y_{\theta}^{(k),y,\alpha_{\epsilon}} f(X_{\theta}^x) + (1 - Y_{\theta}^{(k),y,\alpha_{\epsilon}}) v_{k+1}(X_{\theta}^x, Y_{\theta}^{y,\alpha_{\epsilon}}) \right] \\ = \mathbb{E} \left[1_{\{\bar{Y}_{\theta}^{(k)}=1\}} f(X_{\theta}^x) + 1_{\{\bar{Y}_{\theta}^{(k)}=0\}} v_{k+1}(X_{\theta}^x, \bar{Y}_{\theta}^{(k)}) \right] \\ = \mathbb{E} \left[1_{\{\bar{Y}_{\theta}^{(k)}=1\}} f(X_{\theta}^x) + 1_{\{\bar{Y}_{\theta}^{(k)}=0\}} v_{k+1}(X_{\theta}^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\ = \mathbb{E} \left[\mathbb{E} \left[1_{\{\bar{Y}_{\theta}^{(k)}=1\}} f(X_{\theta}^x) + 1_{\{\bar{Y}_{\theta}^{(k)}=0\}} v_{k+1}(X_{\theta}^x, P_k(Y_{\theta-h}^{y,\alpha})) \mid \mathcal{G} \right] \right] \\ = \mathbb{E} \left[\mathbb{E} \left[1_{\{\bar{Y}_{\theta}^{(k)}=1\}} \mid \mathcal{G} \right] f(X_{\theta}^x) + \mathbb{E} \left[1_{\{\bar{Y}_{\theta}^{(k)}=0\}} \mid \mathcal{G} \right] v_{k+1}(X_{\theta}^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\ = \mathbb{E} \left[Y_{\theta-h}^{(k),y,\alpha} f(X_{\theta}^x) + (1 - Y_{\theta-h}^{(k),y,\alpha}) v_{k+1}(X_{\theta}^x, P_k(Y_{\theta-h}^{y,\alpha})) \right]. \end{aligned}$$

But by the continuity and growth bounds of f and v_{k+1} , we can apply the Dominated Convergence Theorem to see

$$\begin{aligned} \lim_{\delta \rightarrow 0^+} \mathbb{E} \left[Y_{\theta-h}^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_{\theta-h}^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] \\ = \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y,\alpha})) \right]. \end{aligned}$$

So then for any $\epsilon > 0$, we may take $\delta > 0$ small enough that

$$\begin{aligned} \mathbb{E} \left[Y_\theta^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y,\alpha})) \right] \\ \leq \mathbb{E} \left[Y_{\theta-h}^{(k),y,\alpha} f(X_\theta^x) + (1 - Y_{\theta-h}^{(k),y,\alpha}) v_{k+1}(X_\theta^x, P_k(Y_{\theta-h}^{y,\alpha})) \right] + \epsilon \\ = \mathbb{E} \left[Y_\theta^{(k),y,\alpha_\epsilon} f(X_\theta^x) + (1 - Y_\theta^{(k),y,\alpha_\epsilon}) v_{k+1}(X_\theta^x, Y_\theta^{y,\alpha_\epsilon}) \right] + \epsilon \\ \leq A + \epsilon. \end{aligned}$$

Because ϵ and α were arbitrary, we conclude $B \leq A$.

□

4.4.3 Proof of Lemma 4.5

Proof. By Lemma 4.4, we can use either representation (4.2) or (4.5) of v_k , as is convenient, in this proof. Recall that there exists $C > 0$ large enough such that

$$|v_k(x, y) - v_k(x', y)| \leq C |x - x'|$$

for all $(x, y) \in \mathbb{R} \times \Delta_k$ and $x' \in \mathbb{R}$.

1. We first aim to demonstrate that the map $y \mapsto v_k(x, y)$ is concave for any $x \in \mathbb{R}$. The key observation is that the map

$$\begin{aligned} \Delta_k \setminus \{e_k\} \ni y &\mapsto (1 - y_k) v_{k+1}(x, P_k(y)) \\ &= (y_{k+1} + \dots + y_r) v_{k+1} \left(x, \frac{(0, \dots, 0, y_{k+1}, \dots, y_r)}{y_{k+1} + \dots + y_r} \right) \end{aligned}$$

is concave for every $x \in \mathbb{R}$ because it is the perspective transformation of the concave map $\Delta_{k+1} \ni y \mapsto v_{k+1}(x, y)$ (See Section 3.2.6 in Boyd-Vandenberghe [BV04]).

With this in mind, fix $x \in \mathbb{R}$, $y_1, y_2 \in \Delta_k$, and $\lambda \in [0, 1]$. Let $\alpha_1, \alpha_2 \in \mathcal{A}$ be arbitrary controls for which

$$Y_t^{y_1, \alpha_1}, Y_t^{y_2, \alpha_2} \in \Delta_k,$$

almost-surely, for all $t \geq 0$. Define $\bar{y} := \lambda y_1 + (1 - \lambda) y_2$ and $\bar{\alpha}_t := \lambda \alpha_{1,t} + (1 - \lambda) \alpha_{2,t}$. Then $\bar{\alpha} \in \mathcal{A}$ and

$$Y_t^{\bar{y}, \bar{\alpha}} \in \Delta_k,$$

almost-surely, for all $t \geq 0$ by the convexity of the set Δ_k .

Then using the concavity of the perspective map, we can compute

$$\begin{aligned}
 v_k(x, \bar{y}) &\geq \mathbb{E} \left[Y_\theta^{(k), \bar{y}, \bar{\alpha}} f(X_\theta^x) + (1 - Y_\theta^{(k), \bar{y}, \bar{\alpha}}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{\bar{y}, \bar{\alpha}})) \right] \\
 &\geq \mathbb{E}^x \left[Y_\theta^{(k), \bar{y}, \bar{\alpha}} f(X_\theta^x) + \lambda(1 - Y_\theta^{(k), y_1, \alpha_1}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y_1, \alpha_1})) \right. \\
 &\quad \left. + (1 - \lambda)(1 - Y_\theta^{(k), y_2, \alpha_2}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y_2, \alpha_2})) \right] \\
 &= \lambda \mathbb{E} \left[Y_\theta^{(k), y_1, \alpha_1} f(X_\theta^x) + (1 - Y_\theta^{(k), y_1, \alpha_1}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y_1, \alpha_1})) \right] \\
 &\quad + (1 - \lambda) \mathbb{E} \left[Y_\theta^{(k), y_2, \alpha_2} f(X_\theta^x) + (1 - Y_\theta^{(k), y_2, \alpha_2}) v_{k+1}(X_\theta^x, P_k(Y_\theta^{y_2, \alpha_2})) \right].
 \end{aligned}$$

But because α_1, α_2 were arbitrary, we conclude

$$v_k(x, \bar{y}) \geq \lambda v_k(x, y_1) + (1 - \lambda) v_k(x, y_2).$$

2. In particular, the concavity result implies that for any $x \in \mathbb{R}$, the map $\Delta_k \ni y \mapsto v_k(x, y)$ is continuous on the relative interior of Δ_k , and always lower semi-continuous. We claim these properties carry over to the function v_k by the uniform Lipschitz estimate of Proposition 4.1.

Let $\{(x_n, y_n)\}_{n \geq 1} \subset \mathbb{R} \times \Delta_k$ be a sequence converging to $(x_0, y_0) \in \mathbb{R} \times \Delta_k$ as $n \rightarrow \infty$. Then we can compute

$$\begin{aligned}
 v_k(x_n, y_n) &\leq |v_k(x_n, y_n) - v_k(x_0, y_n)| + v_k(x_0, y_n) \\
 &\leq C |x_n - x_0| + v_k(x_0, y_n).
 \end{aligned}$$

But this implies

$$\liminf_{n \rightarrow \infty} v_k(x_n, y_n) \leq v_k(x_0, y_0)$$

by the lower semi-continuity of the map $y \mapsto v_k(x_0, y)$. Then v_k is lower semi-continuous.

If $y_0 \in \text{rel int}(\Delta_k)$, then we similarly compute

$$\begin{aligned}
 |v_k(x_n, y_n) - v_k(x_0, y_0)| &\leq |v_k(x_n, y_n) - v_k(x_0, y_n)| + |v_k(x_0, y_n) - v_k(x_0, y_0)| \\
 &\leq C |x_n - x_0| + |v_k(x_0, y_n) - v_k(x_0, y_0)| \rightarrow 0
 \end{aligned}$$

as $n \rightarrow \infty$ by the continuity of $y \mapsto v_k(x_0, y)$ at y_0 . Then v_k is continuous for any point in the relative interior of $\mathbb{R} \times \Delta_k$.

3. Next we show that v_k is continuous near vertices of the simplex Δ_k . Fix $(x_0, y_0) \in \mathbb{R} \times \Delta_k$ where $y_{0, \ell} = 1$ for some $\ell \in \{k, \dots, r\}$. Denote by \mathcal{I} the subset of indices

$i \in \{k, \dots, r\}$ for which $i \neq \ell$. Note that there is only a single admissible stopping time at the point (x_0, y_0) , so

$$v_k(x_0, y_0) = \mathbb{E} [f(X_{t_\ell - t_{k-1}}^{x_0})].$$

Let $\{(x_n, y_n)\}_{n \geq 1} \subset \mathbb{R} \times \Delta_k$ be a sequence converging to (x_0, y_0) . For any $\epsilon > 0$, take $n \geq 1$ large enough that $\|y_n - y_0\|_{\ell^\infty} \leq \epsilon$ and $|x_n - x_0| \leq \epsilon$. Let τ_n be an arbitrary stopping time such that $\tau_n \sim \sum_{i=k}^r y_{n,i} \delta_{t_i - t_{k-1}}$. Then we can compute

$$\begin{aligned} \mathbb{E} [f(X_{\tau_n}^{x_n})] &\leq L|x_n - x_0| + \sum_{i=k}^r \mathbb{E} [1_{\{\tau_n = t_i - t_{k-1}\}} f(X_{t_i - t_{k-1}}^{x_0})] \\ &\leq v_k(x_0, y_0) + L\epsilon + \sum_{i \in \mathcal{I}} \mathbb{E} [1_{\{\tau_n = t_i - t_{k-1}\}} f(X_{t_i - t_{k-1}}^{x_0})] \\ &\quad + \mathbb{E} [(1 - 1_{\{\tau_n = t_\ell - t_{k-1}\}}) f(X_{t_\ell - t_{k-1}}^{x_0})]. \end{aligned}$$

Applying the Cauchy-Schwarz Inequality to the last two terms and using the Lipschitz assumption on f , we see that

$$\mathbb{E} [f(X_{\tau_n}^{x_n})] \leq v_k(x_0, y_0) + L\epsilon + C\epsilon^{1/2},$$

for $C > 0$ sufficiently large. This implies that v_k is also upper semi-continuous (hence continuous) at the point (x_0, y_0) .

4. Lastly, consider a point $(x_0, y_0) \in \mathbb{R} \times \Delta_k$ which is not in the relative interior of Δ_k and is also not a vertex, as considered in the previous step. Then there exists some subset of indices $\mathcal{I} \subset \{k, \dots, r\}$ for which $y_{0,i} = 0$ for $i \in \mathcal{I}$ and $y_{0,i} \in (0, 1)$ otherwise.

Let $\{(x_n, y_n)\}_{n \geq 1} \subset \mathbb{R} \times \Delta_k$ be a sequence converging to (x_0, y_0) . For any $\epsilon > 0$, take $n \geq 1$ large enough that $\|y_n - y_0\|_{\ell^\infty} \leq \epsilon$ and $|x_n - x_0| \leq \epsilon$. Let τ_n be an arbitrary stopping time such that $\tau_n \sim \sum_{i=k}^r y_{n,i} \delta_{t_i - t_{k-1}}$. Define $\ell := \max \{i \in \{k, \dots, r\} \mid i \notin \mathcal{I}\}$ and

$$\bar{y}_n := \begin{cases} 0 & \text{if } i \in \mathcal{I} \cup \{1, \dots, k-1\} \\ y_{n,i} & \text{if } i \in \{k, \dots, r\} \setminus (\mathcal{I} \cup \{\ell\}) \\ y_{n,\ell} + \sum_{i \in \mathcal{I}} y_{n,i} & \text{otherwise.} \end{cases}$$

The point is that \bar{y}_n is a nearby point that is on the ‘‘boundary.’’ Similarly, define a stopping time

$$\bar{\tau}_n := \sum_{i \in \{k, \dots, r\} \setminus \mathcal{I}} (t_i - t_{k-1}) 1_{\{\tau_n = t_i - t_{k-1}\}} + (t_\ell - t_{k-1}) \sum_{i \in \mathcal{I}} 1_{\{\tau_n = t_i - t_{k-1}\}}.$$

That is, when τ_n stops at $t_i - t_{k-1}$ for $i \in \mathcal{I}$, the stopping time $\bar{\tau}_n$ instead waits until $t_\ell - t_{k-1}$ to stop. Then we can check that $\bar{\tau}_n \sim \sum_{i=k}^r \bar{y}_{n,i} \delta_{t_i - t_{k-1}}$ and by the

same computation with the Cauchy-Schwarz Inequality as in the previous step, we can compute

$$\begin{aligned}\mathbb{E} [f(X_{\tau_n}^{x_n})] &\leq \mathbb{E} [f(X_{\bar{\tau}_n}^{x_0})] + L\epsilon + C\epsilon^{1/2} \\ &\leq v_k(x_0, \bar{y}_n) + L\epsilon + C\epsilon^{1/2},\end{aligned}$$

for $C > 0$ sufficiently large. Because τ_n was arbitrary, we conclude

$$v_k(x_n, y_n) \leq v_k(x_0, \bar{y}_n) + L\epsilon + C\epsilon^{1/2}.$$

However, note that by construction each \bar{y}_n is contained the convex hull

$$K := \text{Conv}(\{e_i\}_{i \in \{k, \dots, r\} \setminus \mathcal{I}}),$$

and y_0 is in the relative interior of K . The map $y \mapsto v_k(x_0, y)$ restricted to the convex hull K is concave and thus continuous at y_0 . Therefore, we conclude

$$\liminf_{n \rightarrow \infty} v_k(x_n, y_n) \leq \limsup_{n \rightarrow \infty} v_k(x_0, \bar{y}_n) = v_k(x_0, y_0),$$

so v_k is upper semi-continuous (hence continuous) at the point (x_0, y_0) .

□

Bibliography

- [ABPS13] B. Acciaio, M. Beiglböck, F. Penkner, and W. Schachermayer, “A model-free version of the fundamental theorem of asset pricing and the super-replication theorem”, *Mathematical Finance*, 2013.
- [ADEH99] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, “Coherent measures of risk”, *Math. Financ.*, vol. 9, no. 3, pp. 203–228, 1999.
- [ADEHK07] P. Artzner, F. Delbaen, J.-M. Eber, D. Heath, and H. Ku, “Coherent multi-period risk adjusted values and Bellman’s principle”, *Ann. Oper. Res.*, vol. 152, pp. 5–22, 2007.
- [AKK15] S. Ankirchner, M. Klein, and T. Kruse, “A verification theorem for optimal stopping problems with expectation constraints”, working paper or preprint, Nov. 2015, [Online]. Available: <https://hal.archives-ouvertes.fr/hal-01229024>.
- [ALP95] M. Avellaneda, A. Levy, and A. Parás, “Pricing and hedging derivative securities in markets with uncertain volatilities”, *Applied Mathematical Finance*, vol. 2, no. 2, pp. 73–88, 1995.
- [BEBD06] C. Barrera-Esteve, F. Bergeret, and C. Dossal, “Numerical methods for the pricing of swing options: a stochastic control approach”, *Methodology and computing in applied probability*, vol. 8, no. 4, pp. 517–540, 2006.
- [BL14] E. Bayraktar and M. Ludkovski, “Liquidation in limit order books with controlled intensity”, *Mathematical Finance*, vol. 24, no. 4, pp. 627–650, 2014.
- [BM14] T. Björk and A. Murgoci, “A theory of markovian time-inconsistent stochastic control in discrete time”, *Finance and Stochastics*, vol. 18, no. 3, pp. 545–592, 2014.
- [BM16] E. Bayraktar and C. W. Miller, “Distribution-Constrained Optimal Stopping”, *ArXiv e-prints*, Apr. 2016. arXiv: 1604.03042 [math.OC].
- [BN12] B. Bouchard and M. Nutz, “Weak dynamic programming for generalized state constraints”, *SIAM J. Control Optim.*, vol. 50, no. 6, pp. 3344–3373, 2012, ISSN: 0363-0129. DOI: 10.1137/110852942. [Online]. Available: <http://dx.doi.org/10.1137/110852942>.

- [BO11] N. Bäuerle and J. Ott, “Markov decision processes with average-value-at-risk criteria”, *Math. Meth. Oper. Res.*, vol. 74, pp. 361–379, 2011.
- [BP92] F. Black and A. Perold, “Theory of constant proportion portfolio insurance”, *Journal of Economic Dynamics and Control*, vol. 16, no. 3-4, pp. 403–426, 1992.
- [BS73] F. Black and M. Scholes, “The pricing of options and corporate liabilities”, *The journal of political economy*, pp. 637–654, 1973.
- [BS91] G. Barles and P. E. Souganidis, “Convergence of approximation schemes for fully nonlinear second order equations”, *Asymptot. Anal.*, vol. 4, no. 3, pp. 271–283, 1991, ISSN: 0921-7134.
- [BT04] B. Bouchard and N. Touzi, “Discrete-time approximation and monte-carlo simulation of backward stochastic differential equations”, *Stochastic Processes and their applications*, vol. 111, no. 2, pp. 175–206, 2004.
- [BT11] ———, “Weak dynamic programming principle for viscosity solutions”, *SIAM J. Control Optim.*, vol. 49, no. 3, pp. 948–962, 2011, ISSN: 0363-0129. DOI: 10.1137/090752328. [Online]. Available: <http://dx.doi.org/10.1137/090752328>.
- [BT13] J. F. Bonnans and X. Tan, “A model-free no-arbitrage price bound for variance options”, *Applied Mathematics & Optimization*, vol. 68, no. 1, pp. 43–73, 2013.
- [BV04] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004, ISBN: 0521833787.
- [BY13] E. Bayraktar and S. Yao, “A weak dynamic programming principle for zero-sum stochastic differential games with unbounded controls”, *SIAM Journal on Control and Optimization*, vol. 51, no. 3, pp. 2036–2080, 2013.
- [BZ03] J. F. Bonnans and H. Zidani, “Consistency of generalized finite difference schemes for the stochastic hjb equation”, *SIAM Journal on Numerical Analysis*, vol. 41, no. 3, pp. 1008–1021, 2003.
- [Bac00] L. Bachelier, *Théorie de la spéculation*. Gauthier-Villars, 1900.
- [Bel52] R. Bellman, “On the theory of dynamic programming”, *Proc. Nat. Acad. Sci. U. S. A.*, vol. 38, pp. 716–719, 1952, ISSN: 0027-8424.
- [Bis73] J.-M. Bismut, “Conjugate convex functions in optimal stochastic control”, *J. Math. Anal. Appl.*, vol. 44, pp. 384–404, 1973, ISSN: 0022-247x.
- [CF91] R. E. Curto and L. A. Fialkow, “Recursiveness, positivity, and truncated moment problems”, *Houston J. Math.*, vol. 17, no. 4, pp. 603–635, 1991, ISSN: 0362-1588.

- [CI90] M. G. Crandall and H. Ishii, “The maximum principle for semicontinuous functions”, *Differential Integral Equations*, vol. 3, no. 6, pp. 1001–1014, 1990, ISSN: 0893-4983.
- [CIL92] M. G. Crandall, H. Ishii, and P.-L. Lions, “Users guide to viscosity solutions of second order partial differential equations”, *Bulletin of the American Mathematical Society*, vol. 27, no. 1, pp. 1–67, 1992.
- [CIR52] R. Courant, E. Isaacson, and M. Rees, “On the solution of nonlinear hyperbolic differential equations by finite differences”, *Comm. Pure. Appl. Math.*, vol. 5, pp. 243–255, 1952, ISSN: 0010-3640.
- [CK15] A. M. G. Cox and S. Källblad, “Model-independent bounds for Asian options: a dynamic programming approach”, *ArXiv e-prints*, Jul. 2015. arXiv: 1507.02651 [q-fin.PR].
- [CL83] M. G. Crandall and P.-L. Lions, “Viscosity solutions of Hamilton-Jacobi equations”, *Trans. Amer. Math. Soc.*, vol. 277, no. 1, pp. 1–42, 1983, ISSN: 0002-9947. DOI: 10.2307/1999343. [Online]. Available: <http://dx.doi.org/10.2307/1999343>.
- [CPT14] J. Cvitanić, D. Possamaï, and N. Touzi, “Moral hazard in dynamic risk management”, *arXiv preprint arXiv:1406.5852*, 2014.
- [CPT15] —, “Dynamic programming approach to principal-agent problems”, *arXiv preprint arXiv:1510.07111*, 2015.
- [CR14] O. Cavus and A. Ruszczyński, “Risk-averse control of undiscounted transient markov models”, *SIAM Journal on Control and Optimization*, vol. 52, no. 6, pp. 3935–3966, 2014.
- [CST10] U. Cetin, H. M. Soner, and N. Touzi, “Option hedging for small investors under liquidity costs”, *Finance and Stochastics*, vol. 14, no. 3, pp. 317–341, 2010.
- [CSTV07] P. Cheridito, H. M. Soner, N. Touzi, and N. Victoir, “Second-order backward stochastic differential equations and fully nonlinear parabolic PDEs”, *Comm. Pure Appl. Math.*, vol. 60, no. 7, pp. 1081–1110, 2007, ISSN: 0010-3640. DOI: 10.1002/cpa.20168. [Online]. Available: <http://dx.doi.org/10.1002/cpa.20168>.
- [CTMP15] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, “Risk-sensitive and robust decision-making: a CVaR optimization approach”, in *NIPS*, 2015.
- [DL91] B. Dumas and E. Luciano, “An exact solution to a dynamic portfolio choice problem under transactions costs”, *The Journal of Finance*, vol. 46, no. 2, pp. 577–595, 1991.

- [DLL11] F. Da Lio and O. Ley, “Convex Hamilton-Jacobi equations under superlinear growth conditions on data”, *Appl. Math. Optim.*, vol. 63, no. 3, pp. 309–339, 2011, ISSN: 0095-4616. DOI: 10.1007/s00245-010-9122-9. [Online]. Available: <http://dx.doi.org/10.1007/s00245-010-9122-9>.
- [DS14] Y. Dong and R. Sircar, “Time-inconsistent portfolio investment problems”, in *Stochastic Analysis and Applications 2014*, Springer, 2014, pp. 239–281.
- [Dur10] R. Durrett, *Probability: theory and examples*. Cambridge university press, 2010.
- [EL10] I. Ekeland and A. Lazrak, “The golden rule when preferences are time inconsistent”, *Math. Financ. Econ.*, vol. 4, no. 1, pp. 29–55, 2010, ISSN: 1862-9679. DOI: 10.1007/s11579-010-0034-x. [Online]. Available: <http://dx.doi.org/10.1007/s11579-010-0034-x>.
- [Eva10] L. C. Evans, *Partial differential equations*, Second. American Mathematical Society, 2010, pp. xxii+749, ISBN: 978-0-8218-4974-3. DOI: 10.1090/gsm/019. [Online]. Available: <http://dx.doi.org/10.1090/gsm/019>.
- [FF14] M. Falcone and R. Ferretti, *Semi-Lagrangian approximation schemes for linear and Hamilton-Jacobi equations*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2014, pp. xii+319, ISBN: 978-1-611973-04-4.
- [FLO02] S. Frederick, G. Loewenstein, and T. O’Donoghue, “Time discounting and time preference: a critical review”, *Journal of economic literature*, vol. 40, no. 2, pp. 351–401, 2002.
- [FM95] W. H. Fleming and W. M. McEneaney, “Risk-sensitive control on an infinite time horizon”, *SIAM J. Control Optim.*, vol. 33, no. 6, pp. 1881–1915, 1995.
- [FO11] B. D. Froese and A. M. Oberman, “Convergent finite difference solvers for viscosity solutions of the elliptic Monge-Ampère equation in dimensions two and higher”, *SIAM J. Numer. Anal.*, vol. 49, no. 4, pp. 1692–1714, 2011, ISSN: 0036-1429. DOI: 10.1137/100803092. [Online]. Available: <http://dx.doi.org/10.1137/100803092>.
- [FPSS11] J.-P. Fouque, G. Papanicolaou, R. Sircar, and K. Sølna, *Multiscale stochastic volatility for equity, interest rate, and credit derivatives*. Cambridge University Press, 2011.
- [FS06] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*. Springer, 2006.
- [FTW11] A. Fahim, N. Touzi, and X. Warin, “A probabilistic numerical method for fully nonlinear parabolic PDEs”, *Ann. Appl. Probab.*, vol. 21, no. 4, pp. 1322–1364, 2011, ISSN: 1050-5164. DOI: 10.1214/10-AAP723. [Online]. Available: <http://dx.doi.org/10.1214/10-AAP723>.

- [GHLT14] A. Galichon, P. Henry-Labordère, and N. Touzi, “A stochastic control approach to no-arbitrage bounds given marginals, with an application to look-back options”, *Ann. Appl. Probab.*, vol. 24, no. 1, pp. 312–336, 2014, ISSN: 1050-5164. DOI: 10.1214/13-AAP925. [Online]. Available: <http://dx.doi.org/10.1214/13-AAP925>.
- [GRS11] S. Gökay, A. F. Roch, and H. M. Soner, “Liquidity models in continuous and discrete time”, in *Advanced mathematical methods for finance*, Springer, 2011, pp. 333–365.
- [HJ15] W. B. Haskell and R. Jain, “A convex analytic approach to risk-aware markov decision processes”, *SIAM Journal on Control and Optimization*, vol. 53, no. 3, pp. 1569–1598, 2015.
- [HJZ12] Y. Hu, H. Jin, and X. Zhou, “Time-inconsistent stochastic linear–quadratic control”, *SIAM J. Control Optim.*, vol. 50, pp. 1548–1572, 2012.
- [HKLW02] P. S. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward, “Managing smile risk”, *The Best of Wilmott*, p. 249, 2002.
- [HL91] S. J. Hoch and G. F. Loewenstein, “Time-inconsistent preferences and consumer self-control”, *Journal of consumer research*, vol. 17, no. 4, pp. 492–507, 1991.
- [HLOST+16] P. Henry-Labordere, J. Oblój, P. Spoida, N. Touzi, *et al.*, “The maximum maximum of a martingale with given n marginals”, *The Annals of Applied Probability*, vol. 26, no. 1, pp. 1–44, 2016.
- [HLOTTW16] P. Henry-Labordere, N. Oudjane, X. Tan, N. Touzi, and X. Warin, “Branching diffusion representation of semilinear pdes and monte carlo approximation”, *arXiv preprint arXiv:1603.01727*, 2016.
- [HM98] H. B. Humphreys and K. T. McClain, “Reducing the impacts of energy price volatility through dynamic portfolio selection”, *The Energy Journal*, pp. 107–131, 1998.
- [Hes93] S. L. Heston, “A closed-form solution for options with stochastic volatility with applications to bond and currency options”, *Review of financial studies*, vol. 6, no. 2, pp. 327–343, 1993.
- [Hor01] M. Horiguchi, “Markov decision processes with a stopping time constraint”, *Mathematical methods of operations research*, vol. 53, no. 2, pp. 279–295, 2001.
- [Ish89] H. Ishii, “On uniqueness and existence of viscosity solutions of fully nonlinear second-order elliptic PDEs”, *Comm. Pure Appl. Math.*, vol. 42, no. 1, pp. 15–45, 1989, ISSN: 0010-3640. DOI: 10.1002/cpa.3160420103. [Online]. Available: <http://dx.doi.org/10.1002/cpa.3160420103>.

- [JBE94] M. R. James, J. S. Baras, and R. J. Elliott, “Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems”, *IEEE Trans. Automat. Control*, vol. 39, no. 4, pp. 780–792, 1994.
- [Jen89] R. Jensen, “Uniqueness criteria for viscosity solutions of fully nonlinear elliptic partial differential equations”, *Indiana Univ. Math. J.*, vol. 38, no. 3, pp. 629–667, 1989, ISSN: 0022-2518. DOI: 10.1512/iumj.1989.38.38030. [Online]. Available: <http://dx.doi.org/10.1512/iumj.1989.38.38030>.
- [KMZ16] C. Karnam, J. Ma, and J. Zhang, “Dynamic Approaches for Some Time Inconsistent Problems”, *ArXiv e-prints*, Apr. 2016. arXiv: 1604.03913 [math.OA].
- [KS91] I. Karatzas and S. E. Shreve, *Brownian motion and stochastic calculus*, Second, ser. Graduate Texts in Mathematics. Springer-Verlag, New York, 1991, vol. 113, pp. xxiv+470, ISBN: 0-387-97655-8. DOI: 10.1007/978-1-4612-0949-2. [Online]. Available: <http://dx.doi.org/10.1007/978-1-4612-0949-2>.
- [Kat94] M. A. Katsoulakis, “Viscosity solutions of second order fully nonlinear elliptic equations with state constraints”, *Indiana Univ. Math. J.*, vol. 43, no. 2, pp. 493–519, 1994, ISSN: 0022-2518. DOI: 10.1512/iumj.1994.43.43020. [Online]. Available: <http://dx.doi.org/10.1512/iumj.1994.43.43020>.
- [Ken82] D. P. Kennedy, “On a constrained optimal stopping problem”, *Journal of Applied Probability*, vol. 19, no. 3, pp. 631–641, 1982, ISSN: 00219002. [Online]. Available: <http://www.jstor.org/stable/3213520>.
- [LP92] G. Loewenstein and D. Prelec, “Anomalies in intertemporal choice: evidence and an interpretation”, *The Quarterly Journal of Economics*, pp. 573–597, 1992.
- [LSMS95] F. López, M San Miguel, and G Sanz, “Lagrangean methods and optimal stopping”, *Optimization*, vol. 34, no. 4, pp. 317–327, 1995.
- [LZL02] X. Li, X. Y. Zhou, and A. E. Lim, “Dynamic mean-variance portfolio selection with no-shorting constraints”, *SIAM Journal on Control and Optimization*, vol. 40, no. 5, pp. 1540–1555, 2002.
- [Lai97] D. Laibson, “Golden eggs and hyperbolic discounting”, *The Quarterly Journal of Economics*, pp. 443–477, 1997.
- [Las10] J. B. Lasserre, *Moments, positive polynomials and their applications*, ser. Imperial College Press Optimization Series. Imperial College Press, London, 2010, vol. 1, pp. xxii+361, ISBN: 978-1-84816-445-1; 1-84816-445-9.

- [Lio83a] P.-L. Lions, “Optimal control of diffusion processes and Hamilton-Jacobi-Bellman equations. I. The dynamic programming principle and applications”, *Comm. Partial Differential Equations*, vol. 8, no. 10, pp. 1101–1174, 1983, ISSN: 0360-5302. DOI: 10.1080/03605308308820297. [Online]. Available: <http://dx.doi.org/10.1080/03605308308820297>.
- [Lio83b] —, “Optimal control of diffusion processes and Hamilton-Jacobi-Bellman equations. II. Viscosity solutions and uniqueness”, *Comm. Partial Differential Equations*, vol. 8, no. 11, pp. 1229–1276, 1983, ISSN: 0360-5302. DOI: 10.1080/03605308308820301. [Online]. Available: <http://dx.doi.org/10.1080/03605308308820301>.
- [Lio83c] —, “Optimal control of diffusion processes and Hamilton-Jacobi-Bellman equations. III. Regularity of the optimal cost function”, *Res. Notes in Math.* Vol. 93, pp. 95–205, 1983.
- [Lyo95] T. J. Lyons, “Uncertain volatility and the risk-free synthesis of derivatives”, *Applied mathematical finance*, vol. 2, no. 2, pp. 117–133, 1995.
- [MOS07] R. Mansini, W. Ogryczak, and M. G. Speranza, “Conditional value at risk and related linear programming models for portfolio optimization”, *Ann. Oper. Res.*, vol. 152, pp. 227–256, 2007.
- [MPY94] J. Ma, P. Protter, and J. Yong, “Solving forward-backward stochastic differential equations explicitly a four step scheme”, *Probability Theory and Related Fields*, vol. 98, no. 3, pp. 339–359, 1994.
- [MPZ15] A. Matoussi, D. Possamai, and C. Zhou, “Robust utility maximization in nondominated models with 2bsde: the uncertain volatility model”, *Mathematical Finance*, vol. 25, no. 2, pp. 258–287, 2015.
- [MY15] C. W. Miller and I. Yang, “Optimal control of conditional value-at-risk in continuous time”, *arXiv preprint arXiv:1512.05015*, 2015.
- [Mak09] C. Makasu, “Bounds for a constrained optimal stopping problem”, *Optimization Letters*, vol. 3, no. 4, pp. 499–505, 2009.
- [McK75] H. P. McKean, “Application of Brownian motion to the equation of Kolmogorov-Petrovskii-Piskunov”, *Comm. Pure Appl. Math.*, vol. 28, no. 3, pp. 323–331, 1975, ISSN: 0010-3640.
- [Mer69] R. C. Merton, “Lifetime portfolio selection under uncertainty: the continuous-time case”, *The review of Economics and Statistics*, pp. 247–257, 1969.
- [Mil16] C. W. Miller, “Non-linear PDE approach to time-inconsistent optimal stopping”, *SIAM Journal on Control and Optimization*, to appear, 2016.
- [NS12] M. Nutz and H. M. Soner, “Superhedging and dynamic risk measures under volatility uncertainty”, *SIAM Journal on Control and Optimization*, vol. 50, no. 4, pp. 2065–2089, 2012.

- [OR99] T. O’Donoghue and M. Rabin, “Doing it now or later”, *American Economic Review*, pp. 103–124, 1999.
- [OY03] H. Ou-Yang, “Optimal contracts in a continuous-time delegated portfolio management problem”, *Rev. Financ. Stud.*, vol. 16, no. 1, pp. 173–208, 2003.
- [Obe07] A. M. Oberman, “The convex envelope is the solution of a nonlinear obstacle problem”, *Proc. Amer. Math. Soc.*, vol. 135, no. 6, 1689–1694 (electronic), 2007, ISSN: 0002-9939. DOI: 10.1090/S0002-9939-07-08887-9. [Online]. Available: <http://dx.doi.org/10.1090/S0002-9939-07-08887-9>.
- [Obe08] A. Oberman, “Wide stencil finite difference schemes for the elliptic monge-ampre equation and functions of the eigenvalues of the hessian”, *Discrete and Continuous Dynamical Systems Series B*, vol. 10, pp. 221–238, 2008.
- [PBG64] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko, *The mathematical theory of optimal processes*, ser. Translated by D. E. Brown. A Pergamon Press Book. The Macmillan Co., New York, 1964, pp. vii+338.
- [PFV03] D. M. Pooley, P. A. Forsyth, and K. R. Vetzal, “Numerical convergence properties of option pricing PDEs with uncertain volatility”, *IMA J. Numer. Anal.*, vol. 23, no. 2, pp. 241–267, 2003, ISSN: 0272-4979. DOI: 10.1093/imanum/23.2.241. [Online]. Available: <http://dx.doi.org/10.1093/imanum/23.2.241>.
- [PP13] J. L. Pedersen and G. Peskir, “Optimal mean-variance portfolio selection”, Research Report, Probability and Statistics Group, School of Mathematics, The University of Manchester, Tech. Rep., 2013.
- [PP16a] —, “Optimal mean-variance selling strategies”, *Mathematics and Financial Economics*, vol. 10, no. 2, pp. 203–220, 2016.
- [PP16b] G. C. Pflug and A. Pichler, “Time-inconsistent multistage stochastic programs: Martingale bounds”, *European J. Oper. Res.*, vol. 249, no. 1, pp. 155–163, 2016.
- [PP90] É. Pardoux and S. G. Peng, “Adapted solution of a backward stochastic differential equation”, *Systems Control Lett.*, vol. 14, no. 1, pp. 55–61, 1990, ISSN: 0167-6911. DOI: 10.1016/0167-6911(90)90082-6. [Online]. Available: [http://dx.doi.org/10.1016/0167-6911\(90\)90082-6](http://dx.doi.org/10.1016/0167-6911(90)90082-6).
- [PRT+13] D. Possamai, G. Royer, N. Touzi, *et al.*, “On the robust superhedging of measurable claims”, *Preprint*, 2013.
- [PS06] G. Peskir and A. Shiryaev, *Optimal stopping and free-boundary problems*, ser. Lectures in Mathematics ETH Zürich. Birkhäuser Verlag, Basel, 2006, pp. xxii+500, ISBN: 978-3-7643-2419-3; 3-7643-2419-8.

- [PS88] A. F. Perold and W. F. Sharpe, “Dynamic strategies for asset allocation”, *Financ. Anal. J.*, vol. 44, no. 1, pp. 16–27, 1988.
- [PTZ15] D. Possamai, X. Tan, and C. Zhou, “Stochastic control for a class of non-linear kernels and applications”, *arXiv preprint arXiv:1510.08439*, 2015.
- [Pen90] S. Peng, “A general stochastic maximum principle for optimal control problems”, *SIAM Journal on control and optimization*, vol. 28, no. 4, pp. 966–979, 1990.
- [Pfe16] L. Pfeiffer, “Risk-averse merton’s portfolio problem”, in *P. IFAC CPDE*, 2016, pp. 266–271.
- [QSR13] J. Qin, H.-I. Su, and R. Rajagopal, “Storage in risk limiting dispatch: control and approximation”, in *P. Amer. Contr. Conf.*, 2013, pp. 4202–4208.
- [RBWV12] R. Rajagopal, E. Bitar, W. Wu, and P. Varaiya, “Risk limiting dispatch of wind power”, in *P. Amer. Contr. Conf.*, 2012, pp. 4417–4422.
- [RU00] R. T. Rockafellar and S. Uryasev, “Optimization of conditional value-at-risk”, *Journal of Risk*, vol. 2, pp. 21–42, 2000.
- [RY15] A. Ruszczyński and J. Yao, “Risk-averse control of diffusion processes”, *arXiv:1508.05316 [math.OC]*, 2015.
- [Rok14] D. B. Rokhlin, “Stochastic Perron’s method for optimal control problems with state constraints”, *Electron. Commun. Probab.*, vol. 19, no. 73, 15, 2014, ISSN: 1083-589X. DOI: 10.1214/ECP.v19-3616. [Online]. Available: <http://dx.doi.org/10.1214/ECP.v19-3616>.
- [Rus10] A. Ruszczyński, “Risk-averse dynamic programming for markov decision processes”, *Math. Program.*, vol. 125, pp. 235–261, 2010.
- [SSU08] S. Sarykalin, G. Serraino, and S. Uryasev, “Value-at-risk vs. conditional value-at-risk in risk management and optimization”, *Tutor. Oper. Res.*, pp. 270–294, 2008.
- [ST13] H. M. Soner and N. Touzi, “Homogenization and asymptotics for small transaction costs”, *SIAM Journal on Control and Optimization*, vol. 51, no. 4, pp. 2893–2921, 2013.
- [STZ+13] H. M. Soner, N. Touzi, J. Zhang, *et al.*, “Dual formulation of second order target problems”, *The Annals of Applied Probability*, vol. 23, no. 1, pp. 308–347, 2013.
- [STZ12] H. M. Soner, N. Touzi, and J. Zhang, “Wellposedness of second order backward sdes”, *Probability Theory and Related Fields*, vol. 153, no. 1-2, pp. 149–190, 2012.

- [Shi08] A. N. Shiryaev, *Optimal stopping rules*, ser. Stochastic Modelling and Applied Probability. Springer-Verlag, Berlin, 2008, vol. 8, pp. xii+217, Translated from the 1976 Russian second edition by A. B. Aries, Reprint of the 1978 translation, ISBN: 978-3-540-74010-0.
- [Str55] R. H. Strotz, “Myopia and inconsistency in dynamic utility maximization”, *The Review of Economic Studies*, vol. 23, no. 3, pp. 165–180, 1955.
- [TDR09] M. Thompson, M. Davison, and H. Rasmussen, “Natural gas storage valuation and optimization: a real options application”, *Naval Research Logistics (NRL)*, vol. 56, no. 3, pp. 226–238, 2009.
- [TK85] A. Tversky and D. Kahneman, “The framing of decisions and the psychology of choice”, in *Environmental Impact Assessment, Technology Assessment, and Risk Analysis*, Springer, 1985, pp. 107–129.
- [Tou10] N. Touzi, “Second order backward sdes, fully nonlinear pdes, and applications in finance”, in *Proceedings of the International Congress of Mathematicians*, vol. 4, 2010, pp. 3132–3150.
- [Tou13] ———, *Optimal stochastic control, stochastic target problems, and backward SDE*, ser. Fields Institute Monographs. Springer, New York; Fields Institute for Research in Mathematical Sciences, Toronto, ON, 2013, vol. 29, pp. x+214, With Chapter 13 by Angès Tourin, ISBN: 978-1-4614-4285-1; 978-1-4614-4286-8. DOI: 10.1007/978-1-4614-4286-8. [Online]. Available: <http://dx.doi.org/10.1007/978-1-4614-4286-8>.
- [Tri95] L. Trigeorgis, *Real options in capital investment: Models, strategies, and applications*. Greenwood Publishing Group, 1995.
- [XZ13] Z. Q. Xu and X. Y. Zhou, “Optimal stopping under probability distortion”, *Ann. Appl. Probab.*, vol. 23, no. 1, pp. 251–282, 2013, ISSN: 1050-5164. DOI: 10.1214/11-AAP838. [Online]. Available: <http://dx.doi.org/10.1214/11-AAP838>.
- [YCT a] I. Yang, D. S. Callaway, and C. J. Tomlin, “Variance-constrained risk sharing in stochastic systems”, *IEEE Trans. Automat. Control*, to appear.
- [YCT14] ———, “Dynamic contracts with partial observations: application to indirect load control”, in *P. Amer. Contr. Conf.*, 2014, pp. 1224–1230.
- [YCT15] ———, “Indirect load control for electricity market risk management via risk-limiting dynamic contracts”, in *P. Amer. Contr. Conf.*, 2015, pp. 3025–3031.
- [Yon12] J. Yong, “Time-inconsistent optimal control problems and equilibrium hjb equation”, *Math. Control Relat. Fields*, vol. 2, pp. 271–329, 2012.

- [ZL00] X. Y. Zhou and D. Li, “Continuous-time mean-variance portfolio selection: a stochastic lq framework”, *Applied Mathematics and Optimization*, vol. 42, no. 1, pp. 19–33, 2000.
- [ZY03] X. Y. Zhou and G. Yin, “Markowitz’s mean-variance portfolio selection with regime switching: a continuous-time model”, *SIAM Journal on Control and Optimization*, vol. 42, no. 4, pp. 1466–1482, 2003.
- [Zhu10] Y. Zhu, “Uncertain optimal control with application to a portfolio selection model”, *Cybernetics and Systems: An International Journal*, vol. 41, no. 7, pp. 535–547, 2010.