

Dynamic Portfolio Optimization

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE UNIVERSITY OF
MINNESOTA BY

Lei Yuan

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF DOCTOR OF PHILOSOPHY

Fadil Santosa, Daniel Mitchell

June, 2021

© Lei Yuan 2021

Acknowledgements

I would like to thank my advisor Fadil Santosa for accepting me as his student, guiding me through the graduate program, and helping me start my professional career. Working with him has always been an encouraging experience. I would also like to thank my co-advisor Daniel Mitchell, who has provided me with valuable research insights and advice. I am grateful to Gary Hatfield and Securian Financial for the support to my research and the opportunity to gain industry experience. Finally, I would like to thank my family and friends, without whom I could not have made it this far.

Dedication

To my wife and parents.

Abstract

The problem of portfolio optimization has attracted broad attention from the academia and financial industry. In this thesis, we will explore a variety of mathematical models related to portfolio optimization. We will derive equations that describe price dynamics of risky assets, and study portfolio optimization problems based on these equations.

Contents

Acknowledgements	i
Dedication	ii
Abstract	iii
List of Figures	vi
1 Introduction	1
2 Robust Portfolio Optimization	4
2.1 Robust Portfolio Selection Model	4
2.2 Linear Regression and Uncertainty Sets	6
2.3 Joint Ellipsoidal Uncertainty Set	8
2.4 Numerical Study	9
2.5 Discussion	11
3 Bayesian Methods	13
3.1 The Black-Litterman Model	14
3.2 Extension of Black-Litterman Model	16
3.3 Robust Bayesian Allocation	17
3.4 Numerical Study	20
4 Dynamic Portfolio Optimization	22
4.1 Robust Dynamic Programming	23
4.2 Merton's portfolio problem	26

5	Dynamic Portfolio Optimization with Bayesian Update - Discrete Model	31
5.1	A Discrete Model with an Uncertain Parameter	31
5.2	Transformation to a Dynamic Programming	34
5.3	Numerical Study	36
6	Dynamic Portfolio Optimization with Bayesian Update - Continuous Model	39
6.1	Conjugate Distributions	40
6.2	Stochastic Market Model	41
6.3	Stochastic Control Problem	42
6.4	Stochastic Control Problem - Unconstrained Action	45
6.5	Stochastic Control Problem - Single Action	53
6.6	Stochastic Control Problem - Finite Action Set	57
6.7	Stochastic Control Problem - Constrained Action	74
6.8	Discussion	85
7	Conclusion	87
	References	88
	Appendix A.	92
A.1	Derivation of HJB Equation	92

List of Figures

1	Summary statistics of asset returns	10
2	Portfolio allocations using GI and Lu models	11
3	Summary statistics and portfolio allocation using RBA model	21
4	Optimal action tree (no short sell allowed)	37
5	Optimal action tree (short sell allowed)	38
6	Numerical solution of A at some specific times	46
7	Numerical solution of $B(t, x)$ at some specific times	47
8	Numerical derivative of $B(t, x)$ w.r.t x at some specific times	48
9	Numerical solution of $B(t, x)$ at some specific times ($\alpha = 0.5$)	54
10	Numerical solution of $B(t, x)$ and optimal control	61
11	Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{0\}$ and $\mathcal{A} = \{1\}$)	62
12	Numerical derivatives of $B(t, x)$	64
13	Numerical solution of $B(t, x)$ and optimal control	66
14	Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$)	67
15	Numerical derivatives of $B(t, x)$	69
16	Numerical solution of $B(t, x)$ and optimal control	71
17	Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$)	72
18	Numerical derivatives of $B(t, x)$	73
19	Numerical solution of $B(t, x)$ and optimal control	76
20	Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{0\}$ and $\mathcal{A} = \{1\}$)	77

21	Numerical derivatives of $B(t, x)$	78
22	Numerical solution of $B(t, x)$ and optimal control	80
23	Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$)	81
24	Numerical derivatives of $B(t, x)$	82
25	Evolution of optimal control in time ($\mathcal{A} = [-1, 1]$)	84
26	Evolution of optimal control in time ($\mathcal{A} = [-30, 30]$)	85

Chapter 1

Introduction

Portfolio optimization is the problem of finding the optimal allocation of funds over a number of available assets in order to maximize the return on the investment while minimizing the risk. Solutions to this problem is important for both institutions and individuals. Institutions such as mutual funds and hedge funds make portfolio decision on a regular basis. Individuals planning for retirement are also concerned about the implications of their choices. Mathematical models have been developed to address various issues of relevance for the parties involved.

Makowitz [27] introduced the first mathematical model for portfolio optimization in the 1950s. Let the random vector $\mathbf{r} = (r_1, \dots, r_n)^T \in \mathbb{R}^n$ denote returns of n assets in the market, and $\boldsymbol{\phi} = (\phi_1, \dots, \phi_n)^T \in \mathbb{R}^n$ denote the proportions of wealth to be invested in the n risky assets such that $\mathbf{1}^T \boldsymbol{\phi} = 1$. Suppose \mathbf{r} has mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. According to Markowitz's theory, portfolio optimization is to find the optimal trade-off between "return" and "risk", which are represented by the expected return $\boldsymbol{\phi}^T \boldsymbol{\mu}$ and the variance $\boldsymbol{\phi}^T \boldsymbol{\Sigma} \boldsymbol{\phi}$, respectively. Mathematically speaking, given a minimum expected return $\mu_0 \in \mathbb{R}$, one has the optimization problem

$$\begin{aligned} & \underset{\boldsymbol{\phi} \in \mathbb{R}^n}{\text{minimize}} && \boldsymbol{\phi}^T \boldsymbol{\Sigma} \boldsymbol{\phi} \\ & \text{subject to} && \boldsymbol{\phi}^T \boldsymbol{\mu} \geq \mu_0, \\ & && \mathbf{1}^T \boldsymbol{\phi} = 1. \end{aligned} \tag{1}$$

The formulation suggests that, given a lower bound on the expected return the investor is willing to accept, the optimal portfolio can be obtained by minimizing the variance. If parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are known, the solution to (1) can be found analytically[29]. It can also be solved numerically under additional constraints, such as no short-selling and position limits in any single asset.

For any μ_0 such that (1) is feasible, the optimal solution $\boldsymbol{\phi}^*$ gives the portfolio allocation among n risky assets which minimizes variance. In the return-variance space, the relationship between constraint μ_0 and the optimal value of objective function, $\boldsymbol{\phi}^{*T} \boldsymbol{\Sigma} \boldsymbol{\phi}^*$, can be represented by a curve known as the efficient frontier. Intuitively, an investor should always allocate the capital to stay on the efficient frontier. Based on Markowitz's work, Treynor[36, 37], Sharpe[34], and Lintner[25] developed the Capital Asset Pricing Model (CAPM) for pricing an individual security or portfolio. CAPM has a profound impact in the academia and practitioners. It has been generalized into a family of factor-based investing models such as the Fama-French models[11, 12], which are widely used in today's financial industry.

The first part of this dissertation addresses the issue of uncertainty in Markowitz model. Since in real-world applications, parameters such as $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are not known exactly and are usually estimated using historical data of financial securities and statistical regression technique, which introduce noise and error to estimated values. Because the quadratic optimization problem in Markowitz model is sensitive to small perturbation of $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$, the result could be problematic[4] and many practitioners turn away from it. To address this issue, several approaches have been introduced. Goldfarb and Iyengar[16] developed a robust optimization framework to take possible ranges of parameters into consideration. Black and Litterman[7, 6] introduced a Bayesian framework to modify statistical estimations with investors' subjective view. Meucci[31] latter unified two approaches into a robust Bayesian model.

The second part of the dissertation studies portfolio optimization problems from a dynamic perspective. As the financial market is ever evolving, investors need to adjust

portfolio allocations according new data and opinion. Thus, portfolio selection can be viewed as dynamics optimization problems in discrete time or continuous time. Discrete portfolio optimization requires less mathematical foundation and can usually be solved using discrete dynamic programming. Continuous portfolio optimization, which was first studied by Merton[28, 30] in the 1970s, sits on stochastic calculus and partial differential equation. We will review both types of problems and focus on a problem which is in the intersection of stochastic control and Bayesian theory.

The rest of the dissertation is organized as follow. In chapter 2, we review robust portfolio optimization models developed to address drawbacks of Markowitz's model. In chapter 3, we look at optimization models that borrow ideas from on Bayesian statistics. In chapter 4, we go back an review the classic continuous-time portfolio optimization problem proposed by Merton. In chapter 5 and 6, we introduce and study dynamic portfolio optimization models in discrete-time and in continuous-time.

Chapter 2

Robust Portfolio Optimization

Although Markowitz's mean-variance model provides a convenient and useful theoretical framework for portfolio optimization, it has several drawbacks. Broadie[9] has shown that the optimal portfolio is sensitive to perturbations of market variables $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. Since $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are subject to statistical estimation error, the result of its following optimization problem is not reliable. In addition, the model can and often lead to poorly diversified portfolio by allocating capital to only a few assets. To reduce the sensitivity of the Markowitz's optimal portfolios to input error, Goldfarb and Iyengar[16] proposed portfolio selection models that are robust to parameter uncertainty and estimation errors, and can be efficiently solved as convex optimization problems.

2.1 Robust Portfolio Selection Model

Assume that the market opens for trading at discrete time and has n traded assets. The vector of asset returns over a single market period is denoted by a column vector $\mathbf{r} \in \mathbb{R}^n$. The returns on the assets in different market periods are assumed to be independent. The single period asset return $\mathbf{r} \in \mathbb{R}^n$ is assumed to be a random variable driven by the factor model

$$\mathbf{r} = \boldsymbol{\mu} + \mathbf{V}^T \mathbf{f} + \boldsymbol{\epsilon}, \quad (2)$$

where $\boldsymbol{\mu} \in \mathbb{R}^n$ is the mean return, $\mathbf{f} \sim \mathcal{N}(\mathbf{0}, \mathbf{F}) \in \mathbb{R}^m$ denotes the normally distributed return of factors that drives the market with mean $\mathbf{0}$ and covariance \mathbf{F} , $\mathbf{V} \in \mathbb{R}^{m \times n}$ is the

deterministic factor loading matrix, and $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{D}) \in \mathbb{R}^n$ is a normally distributed residual vector with mean $\mathbf{0}$ and covariance \mathbf{D} . In addition, we assume $\boldsymbol{\epsilon}$ is independent of \mathbf{f} , \mathbf{F} is positive definite, and \mathbf{D} is a positive semi-definite diagonal matrix with diagonal elements $\mathbf{d} \in \mathbb{R}^n$ (i.e. $\mathbf{D} = \text{diag}(\mathbf{d})$). According to the factor model, the asset returns is also normally distributed with $\mathbf{r} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{V}^T \mathbf{F} \mathbf{V} + \mathbf{D})$. Denote the position of investment as $\boldsymbol{\phi} \in \mathbb{R}^n$, where the i th component ϕ_i represents the fraction of total wealth invested in the i th asset. The return of the portfolio $\boldsymbol{\phi}$ is then normally distributed with distribution

$$\mathbf{r}_\phi \sim \mathcal{N}(\boldsymbol{\phi}^T \boldsymbol{\mu}, \boldsymbol{\phi}^T (\mathbf{V}^T \mathbf{F} \mathbf{V} + \mathbf{D}) \boldsymbol{\phi}). \quad (3)$$

In contrast to Markowitz's model where mean return and the covariance matrix are constant, Goladfarb and Iyengar (GI) introduced uncertainty sets for \mathbf{D} , \mathbf{V} , and $\boldsymbol{\mu}$ in order to deal with statistical errors. In GI model, each diagonal element of the covariance matrix \mathbf{D} can take any value in an closed interval and the uncertainty set S_d for \mathbf{D} is

$$S_d = \{\mathbf{D} : \mathbf{D} = \text{diag}(\mathbf{d}), d_i \in [\underline{d}_i, \bar{d}_i] \text{ for } i = 1, \dots, n\}. \quad (4)$$

Each element of mean return vector $\boldsymbol{\mu}$ can take values from an interval centered at the i th element of a constant vector $\boldsymbol{\mu}_0 \in \mathbb{R}^n$

$$S_m = \{\boldsymbol{\mu} : \boldsymbol{\mu} = \boldsymbol{\mu}_0 + \boldsymbol{\xi}, |\xi_i| < \gamma_i \text{ for } i = 1, \dots, n\}, \quad (5)$$

where γ_i 's are constant real numbers. The uncertainty set of loading matrix \mathbf{V} is given by restricting each column to an elliptic ball w.r.t. a symmetric positive definite matrix $\mathbf{G} \in \mathbb{R}^{m \times m}$, i.e.

$$S_v = \{\mathbf{V} : \mathbf{V} = \mathbf{V}_0 + \mathbf{W}, \|\mathbf{W}_i\|_g \leq \rho_i \text{ for } i = 1, \dots, n\}, \quad (6)$$

where \mathbf{W}_i denotes the i th column of \mathbf{W} , and $\|\mathbf{w}\|_g = \sqrt{\mathbf{w}^T \mathbf{G} \mathbf{w}}$ for any $\mathbf{w} \in \mathbb{R}^m$. The choices of uncertainty sets is motivated by the fact that parameters where uncertainty sets are established around are estimated by linear regression between historical time series of \mathbf{r} and \mathbf{f} . Other parameters such as \underline{d}_i , \bar{d}_i , $\boldsymbol{\mu}_0$, and \mathbf{V}_0 can be obtained by multivariate linear regression and by specifying confidence levels.

Combining (3) and uncertainty sets, GI proposed a robust version of Markowitz's portfolio optimization problem

$$\begin{aligned}
& \underset{\boldsymbol{\phi} \in \mathbb{R}^n}{\text{minimize}} && \max_{\{\mathbf{V} \in S_v, \mathbf{D} \in S_d\}} \boldsymbol{\phi}^T (\mathbf{V}^T \mathbf{F} \mathbf{V} + \mathbf{D}) \boldsymbol{\phi} \\
& \text{s.t.} && \min_{\{\boldsymbol{\mu} \in S_m\}} \boldsymbol{\phi}^T \boldsymbol{\mu} \geq \alpha \\
& && \mathbf{1}^T \boldsymbol{\phi} = 1,
\end{aligned} \tag{7}$$

where $\alpha \in \mathbb{R}$ stands for the worst-case return level. The objective of the robust minimum variance portfolio selection problem (7) is to minimize the worst-case variance of the portfolio subject to the constraint that the worst-case expected return on the portfolio is at least α . Since it takes into account all possible values of \mathbf{V} , \mathbf{D} , and $\boldsymbol{\mu}$ in uncertainty sets, it is less sensitive to parameter perturbation than the original formulation of Markowitz.

With uncertainty sets S_d , S_m , and S_v described in (4), (5), and (6), the robust optimization problem can be converted into a Second Order Cone Programming (SOCP) problem and solved by numerical convex optimization software such as CVX[17].

2.2 Linear Regression and Uncertainty Sets

GI provided a way of getting statistical estimations of uncertainty sets S_d , S_m , and S_v via multivariate linear regression. Suppose the market data consists of asset returns $\{\mathbf{r}^t : t = 1, \dots, p\}$ for $p \in \mathbb{N}$ periods, and the corresponding factor returns $\{\mathbf{f}^t : t = 1, \dots, p\}$. Then the factor model $\mathbf{r} = \boldsymbol{\mu} + \mathbf{V}^T \mathbf{f} + \boldsymbol{\epsilon}$ implies that

$$r_i^t = \mu_i + \sum_{j=1}^m V_{ji} f_j + \epsilon_i^t \tag{8}$$

for $1 \leq i \leq n$ and $1 \leq t \leq p$. In linear regression analysis, it is standard to assume the vector of residual returns $\boldsymbol{\epsilon}^t$'s are independent and identically distributed normal random variables, i.e. $\epsilon_i^t \sim \mathcal{N}(0, \sigma_i^2)$ for $t = 1, 2, \dots, p$. Let $\mathbf{S} = [\mathbf{r}^1, \dots, \mathbf{r}^p] \in \mathbb{R}^{n \times p}$ be the matrix of observed asset returns and $\mathbf{B} = [\mathbf{f}^1, \dots, \mathbf{f}^p] \in \mathbb{R}^{m \times p}$ be the matrix of observed factor returns. The factor model implies the following linear model for each asset

$$\mathbf{y}_i = \mathbf{A} \mathbf{x}_i + \boldsymbol{\epsilon}_i, \tag{9}$$

where $\mathbf{y}_i = [r_i^1, r_i^2, \dots, r_i^p]^T$, $\mathbf{A} = [\mathbf{1}, \mathbf{B}^T]$, $\mathbf{x}_i = [\mu_i, V_{1i}, V_{2i}, \dots, V_{mi}]^T$, and $\epsilon_i = [e_i^1, e_i^2, \dots, e_i^p]$ is the residual returns corresponding to asset i . If we assume $\mathbf{A}^T \mathbf{A}$ is invertible, which is usually the case when p is large enough, the least squares estimation $\bar{\mathbf{x}}_i$ of \mathbf{x}_i is

$$\bar{\mathbf{x}}_i = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{y}_i. \quad (10)$$

Substituting (9) into (10) yields

$$\bar{\mathbf{x}}_i - \mathbf{x}_i = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \epsilon_i \sim \mathcal{N}(\mathbf{0}, \sigma_i (\mathbf{A}^T \mathbf{A})^{-1}). \quad (11)$$

As a result $\mathcal{X}_i = \frac{1}{\sigma_i^2} (\bar{\mathbf{x}}_i - \mathbf{x}_i)^T (\mathbf{A}^T \mathbf{A}) (\bar{\mathbf{x}}_i - \mathbf{x}_i)$ is a χ^2 random variable with $m+1$ degrees of freedom. Since the true value of σ_i^2 is unknown, we can replace σ_i^2 by its unbiased estimation $s_i^2 = \frac{\|\mathbf{y}_i - \mathbf{A} \bar{\mathbf{x}}_i\|_2^2}{p-m-1}$ and get the resulting random variable

$$\mathcal{Y}_i = \frac{1}{(m+1)s_i^2} (\bar{\mathbf{x}}_i - \mathbf{x}_i)^T (\mathbf{A}^T \mathbf{A}) (\bar{\mathbf{x}}_i - \mathbf{x}_i), \quad (12)$$

which is distributed according to the F -distribution with $m+1$ degrees of freedom[1] in the numerator and $p-m-1$ degrees of freedom in the denominator. Now for any number $\omega \in (0, 1)$, there is a ω -critical value $c(\omega)$ such that the probability of the event $\mathcal{Y}_i \leq c(\omega)$ is ω . Equivalently, the set $S_i(\omega)$ defined below is a ω -confidence set for \mathbf{x}_i

$$S_i(\omega) = \{\mathbf{x}_i : (\bar{\mathbf{x}}_i - \mathbf{x}_i)^T (\mathbf{A}^T \mathbf{A}) (\bar{\mathbf{x}}_i - \mathbf{x}_i) \leq (m+1)c(\omega)s_i^2\}. \quad (13)$$

By independent assumption, $S(\omega) = S_1(\omega) \times S_2(\omega) \times \dots \times S_n(\omega)$ is a ω^n -confidence set for $(\boldsymbol{\mu}, \mathbf{V})$. Based on $S(\omega)$, we can then get a natural description of an ω^n -confidence uncertainty set for $\boldsymbol{\mu}$ by projecting $S(\omega)$ along the direction of $\boldsymbol{\mu}$:

$$S_m(\omega) = \{\boldsymbol{\mu} : \boldsymbol{\mu} = \boldsymbol{\mu}_0 + \mathbf{v}, |v_i| \leq \gamma_i, i = 1, 2, \dots, n\}, \quad (14)$$

where $\boldsymbol{\mu}_0$ is sample average of historical returns, and $\gamma_i = \sqrt{(\mathbf{A}^T \mathbf{A})_{11}^{-1} c(\omega) s_i^2}$. Similarly, the uncertainty set for the loading matrix is:

$$S_v(\omega) = \{\mathbf{V} : \mathbf{V} = \mathbf{V}_0 + \mathbf{W}, \|\mathbf{W}_i\|_g \leq \rho_i, i = 1, 2, \dots, n\}, \quad (15)$$

where \mathbf{V}_0 comes from the least-square estimation in (10), $\|\mathbf{W}_i\|_g$ is the elliptic norm of the i th column of \mathbf{W} with respect to the matrix $\mathbf{G} = \mathbf{B}\mathbf{B}^T - \frac{1}{p}(\mathbf{B}\mathbf{1})(\mathbf{B}\mathbf{1})^T$, and $\rho_i = \sqrt{mc(\omega)s_i^2}$. GI didn't give a specific description of uncertainty set S_d . It is understood to be more flexible than S_m and S_v and can be treated as a known parameter and replaced by empirical estimations.

2.3 Joint Ellipsoidal Uncertainty Set

There are several criticisms of GI on the uncertainty sets. One of them is about the size of uncertain sets \mathbf{S}_v and \mathbf{S}_v . Because they are constructed by projecting $S(\omega)$ along the directions of $\boldsymbol{\mu}$ and \mathbf{V} . The inner maximization problem in (7) optimizes over the Cartesian product $S_m \times S_v$, which can be substantially larger than the original uncertainty set $S(\omega)$. In addition, the actual probability of $S_m \times S_v$ is unknown, and the optimal portfolio tends to be too conservative and poorly diversified.

Lu[26] introduced a joint ellipsoidal uncertainty set to address this issue within GI's framework. He also gave a natural description of the joint ellipsoidal uncertainty set based on multivariate linear regression and an equivalent formulation of problem (7). Instead of looking at each individual random variable \mathcal{Y}_i defined by equation (12), Lu considered their sum $\mathcal{Y} = \sum_{i=1}^n \mathcal{Y}_i$ and defined the joint ellipsoidal uncertainty set $S_{m,v}(\omega)$ as

$$S_{m,v}(\omega) = \{(\boldsymbol{\mu}, \mathbf{V}) : \mathcal{Y} = \sum_{i=1}^n \frac{(\bar{\mathbf{x}}_i - \mathbf{x}_i)^T (\mathbf{A}^T \mathbf{A})(\bar{\mathbf{x}}_i - \mathbf{x}_i)}{(m+1)s_i^2} \leq \tilde{c}(\omega)\}, \quad (16)$$

where $\tilde{c}(\omega)$ is a number depends on ω and the distribution of \mathcal{Y} such that the probability of $\mathcal{Y} \leq \tilde{c}(\omega)$ is ω . Although this new formulation avoids projecting sets and taking Cartesian product, one problem that comes up is the the transformed random variable \mathcal{Y} does not have an analytic cumulative distribution function (CDF). As a result, finding $\tilde{c}(\omega)$ given ω is not as easy as in GI model, where one can simply look up $c(\omega)$ from the CDF of an appropriate F -distribution. To handle this problem, Lu suggested that one can either find $\tilde{c}(\omega)$ by running numerical simulations if the number of assets n is not large, and approximate the CDF of \mathcal{Y} via the Central Limit Theorem if n is large. The corresponding robust optimization formulation using ellipsoidal uncertainty set is

$$\begin{aligned} & \underset{\boldsymbol{\phi} \in \mathbb{R}^n}{\text{minimize}} && \max_{\{(\boldsymbol{\mu}, \mathbf{V}) \in S_{m,v}\}} \boldsymbol{\phi}^T (\mathbf{V}^T \mathbf{F} \mathbf{V} + \mathbf{D}) \boldsymbol{\phi} \\ \text{s.t.} &&& \min_{\{(\boldsymbol{\mu}, \mathbf{V}) \in S_{m,v}\}} \boldsymbol{\phi}^T \boldsymbol{\mu} \geq \alpha, \\ &&& \mathbf{1}^T \boldsymbol{\phi} = 1. \end{aligned} \quad (17)$$

Using duality argument, Lu showed the above optimization problem has the same solution to the following problem for some positive constant θ :

$$\begin{aligned} & \underset{\boldsymbol{\phi} \in \mathbb{R}^n}{\text{maximize}} && \max_{\{(\boldsymbol{\mu}, \mathbf{V}) \in \mathcal{S}_{m,v}\}} \boldsymbol{\phi}^T \boldsymbol{\mu} - \theta \boldsymbol{\phi}^T (\mathbf{V}^T \mathbf{F} \mathbf{V} + \mathbf{D}) \boldsymbol{\phi} \\ & \text{s.t.} && \mathbf{1}^T \boldsymbol{\phi} = 1. \end{aligned} \tag{18}$$

Lu further proved that problem (18) can be efficiently solved after being transformed into a SOCP problem and claimed the resulting optimal portfolios are more diversified than the solutions based on GI model.

2.4 Numerical Study

We present some numerical study result to see how GI and Lu models perform with real financial data. In the study, we use historical price data of 13 assets to estimate their mean return vector, covariance matrix, loading matrix, and other model parameters discussed in previous sections. These assets range from fixed income products to stock indices, and we summarize their return, risk, and correlations matrix in Figure 1. Then we solve robust portfolio optimization problems using GI model and Lu model, respectively. Since the only difference between Lu model and GI is in the uncertainty set, we can make a meaningful comparison. Notice there is a hyper-parameter ω in both models that represents the level of confidence, and it must be specified by user. Consequently, we could have different optimal portfolio allocations based on different values of ω .

Asset	1	2	3	4	5	6	7	8	9	10	11	12	13
Mean Return (%)	0.98	1.89	2.64	3.73	3.3	2.51	3.14	3.33	3.78	6.09	4.79	6.13	7.00
Standard Deviation (%)	0.45	4.03	10.37	5.36	4.57	4.54	7.45	6.07	10.14	7.51	8.44	10.75	15.28

(a) Annualized return and risk of assets

Asset	1	2	3	4	5	6	7	8	9	10	11	12	13
1	1												
2	-0.2	1											
3	-0.16	0.35	1										
4	-0.17	0.75	0.5	1									
5	0.08	0.08	0.05	0.26	1								
6	-0.13	0.5	0.31	0.51	0.36	1							
7	-0.2	0.53	0.35	0.52	0.21	0.95	1						
8	-0.16	0.54	0.49	0.88	0.5	0.56	0.52	1					
9	-0.15	0.4	0.48	0.77	0.58	0.53	0.5	0.95	1				
10	-0.2	0.74	0.67	0.89	0.14	0.55	0.59	0.77	0.69	1			
11	-0.22	0.63	0.64	0.83	-0.01	0.4	0.46	0.67	0.56	0.88	1		
12	-0.22	0.58	0.72	0.77	-0.09	0.33	0.41	0.62	0.51	0.84	0.97	1	
13	-0.17	0.22	0.48	0.43	-0.17	0.07	0.16	0.34	0.27	0.43	0.66	0.73	1

(b) Correlation matrix of assets

Figure 1: Summary statistics of asset returns

Based on GI model, Portfolio allocations corresponding to different ω can be found in Figure 2(a). When $\omega = 0$, each uncertainty set shrinks to a single point, and the robust optimization problem produces the same result as Markowitz's model. The optimal portfolio says one should invest in 5 out of 13 assets. As ω increases, GI model suggest that one allocate more in the first and sixth asset, and allocate less in the second and third asset. Eventually, we see a 78% concentration in the 1st asset. In terms of diversification, GI model doesn't produce superior result in our study. One possible explanation is uncertainty sets in GI model is too large due the projection methods used in the model. As a result, the model tend to prefer few conservative assets with better risk and return trade-off.

Portfolio allocations based on Lu model is shown in Figure 2(b). Notice the portfolios displayed correspond to different value of $c(\omega)$, which is the critical value for ω based on the distribution of joint ellipsoidal uncertainty sets. Since the distribution is difficult to sample and optimized portfolios are not sensitive to changing ω , we use $c(\omega)$ as a proxy of ω (see Section 2.5 for more on this issue). As $c(\omega)$ increases, we see similar trend as in GI model, more allocation in the first and sixth asset and less in the second and the third. However, Lu model picks up the seventh and the eleventh asset before completely

abandoning the third asset and put over 70% weight in the first asset. Based on our observation, Lu model does give slightly better diversification than GI model. The observation is consistent with the claim that Lu model gives more accurate description of the uncertainty sets.

$\omega \backslash \text{Asset}$	1	2	3	4	5	6	7	8	9	10	11	12	13
0	40%	39%	15%			2%							5%
0.1	49%	32%	12%			3%							4%
0.3	55%	27%	10%			4%							4%
0.5	61%	22%	8%			6%							4%
0.7	67%	16%	5%			8%							4%
0.9	78%	7%				11%					1%		3%

(a) Portfolio allocations for different ω using GI model

$c(\omega) \backslash \text{Asset}$	1	2	3	4	5	6	7	8	9	10	11	12	13
0	40%	39%	15%			2%							5%
0.4	53%	24%	11%			7%							4%
0.8	59%	18%	8%			9%					1%		4%
1.2	63%	15%	7%			9%					2%		4%
1.6	66%	13%	6%			9%	1%				2%		4%
2	68%	11%	5%			8%	1%				2%		4%

(b) Portfolio allocations for different $c(\omega)$ using Lu model

Figure 2: Portfolio allocations using GI and Lu models

2.5 Discussion

To deal with statistical error and parameter perturbation, GI and Lu optimized the objective function over uncertainty sets constructed around the least-square estimation of $\boldsymbol{\mu}$ and \mathbf{V} . Because the objective is to optimize outcomes in worst scenarios, the optimal solutions tend to be conservative. We observed quite conservative portfolio allocations in numerical studies using both real market data and randomly generated data. In addition to the conservativeness, numerical issues came up in our experimental studies of these methods. For example, the inverse CDF of random variable $\mathcal{Y} = \sum_{i=1}^n \mathcal{Y}_i$ in Lu model, which is the sum of n i.i.d. F -distributed random variable, is extremely steep near the origin. As a result, the uncertainty set $S_{m,v}(\omega)$ and the optimal solution is not sensitive to the confidence threshold ω , which is counterproductive to the primary objective of robust optimization. To put these robust optimization models in practice,

we need a more accurate description of underlying uncertainty sets and better numerical technique to simulate complicated distributions.

Chapter 3

Bayesian Methods

The classic approach to portfolio optimization relies on estimations of market parameters such as mean return and covariance matrix. Yet many assumptions about the distribution of asset returns are not realistic. Moreover, statistical methods used to estimate these parameters can be problematic. For example, estimating covariance matrix when the sample size is small compared to the number of assets. As the optimization problem is sensitive to input parameters, the resulting portfolio allocation can be sub-optimal[24].

Bayesian theory provides a way to reduce the sensitivity of optimal portfolios to input parameters. In Bayesian framework, we are able to consider not only statistical estimations based on historical data, but also to incorporate prior information such as investors' subjective opinion or expert knowledge. By adjusting statistical estimations with investors' views, more reasonable and well-diversified portfolio allocations can be achieved. In following sections, we focus on two models. The Black-Litterman model was prominent for applying Bayesian method to estimation mean return and covariance matrix. The robust Bayesian allocation approach extends the Black-Litterman model by adding robust optimization elements to it.

3.1 The Black-Litterman Model

Black and Litterman (BL)[7, 6] introduced a portfolio optimization framework which incorporates more inputs than Markowitz's model. Besides statistical estimation based on market data, an investor is allowed to use her views on asset returns as input to the underlying optimization problem. Consequently, portfolio allocations produced by BL model naturally reflect investors' opinion[32] in a consistent manner and do not need post-processing afterwards.

Consider a market of $N \in \mathbb{N}$ securities or asset classes, whose returns $\mathbf{X} \in \mathbb{R}^N$ are normally distributed with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, i.e. $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Similar to Markowitz's model, $\boldsymbol{\Sigma}$ is estimated from exponential smoothing of historical time series. However, BL treat $\boldsymbol{\mu}$ as a random variable since its exact value is unknown. In addition, BL assume distribution of $\boldsymbol{\mu}$ is also normal with $\boldsymbol{\mu} \sim \mathcal{N}(\boldsymbol{\pi}, \tau\boldsymbol{\Sigma})$ where $\tau \in \mathbb{R}$ is a small constant. Intuitively, $\boldsymbol{\pi}$ stands for the best estimation of $\boldsymbol{\mu}$ and $\tau\boldsymbol{\Sigma}$ represents the uncertainty on this guess. Note the covariance matrix of $\boldsymbol{\mu}$ is exactly the covariance matrix of \mathbf{X} scaled by the constant τ , reflecting the fact that $\boldsymbol{\mu}$ is less volatile than \mathbf{X} .

If there is no estimation error, then $\boldsymbol{\mu} = \boldsymbol{\pi}$ and $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\pi}, \boldsymbol{\Sigma})$. In the mean-variance framework, an investor tries to find the optimal portfolio allocation from solving the optimization problem

$$\begin{aligned} \boldsymbol{\omega}_\lambda = & \underset{\boldsymbol{\omega} \in \mathbb{R}^N}{\operatorname{argmax}} \boldsymbol{\omega}^T \boldsymbol{\pi} - \lambda \boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega} \\ \text{s.t.} \quad & \mathbf{1}^T \boldsymbol{\omega} = 1, \end{aligned} \tag{19}$$

where the constant $\lambda \geq 0$ represents the risk-aversion level which depends on each individual investor. Assume there is an average risk-aversion level $\bar{\lambda}$ that applies to the whole market, the quadratic optimization problem can be solved analytically. According to the first order condition, the resulting market optimal portfolio $\tilde{\boldsymbol{\omega}}$ satisfies

$$\boldsymbol{\pi} = 2\bar{\lambda}\boldsymbol{\Sigma}\tilde{\boldsymbol{\omega}} \tag{20}$$

and $\boldsymbol{\pi}$ is called the market equilibrium estimation. Thus we can set $\boldsymbol{\pi}$ in terms of $\tilde{\boldsymbol{\omega}}$ once the exogenous parameter $\bar{\lambda}$ is chosen (BL set it to be around 1.2). It is worth pointing out that historical return information does not play a direct role in determining

π . Assume historical returns $\{\mathbf{X}_t\}_{t=1}^{t=T}$ in different time periods are independent and identically distributed (with the common distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the standard sample estimator of $\boldsymbol{\mu}$ is

$$\hat{\boldsymbol{\mu}} = \frac{1}{T} \sum_{t=1}^{t=T} \mathbf{X}_t \sim \mathcal{N}\left(\boldsymbol{\mu}, \frac{1}{T} \boldsymbol{\Sigma}\right), \quad (21)$$

Above equation suggests that it is reasonable to set $\tau = \frac{1}{T}$. As discussed in previous chapters, the sample estimator $\hat{\boldsymbol{\mu}}$ tends to be inaccurate and its perturbation could lead to dramatically different optimal portfolio in Markowitz's model. One way to deal with this issue is to apply Stein's shrinkage estimator and set $\boldsymbol{\mu}$ as a convex combination of $\hat{\boldsymbol{\mu}}$ and market equilibrium estimation $\boldsymbol{\pi}$:

$$\boldsymbol{\mu}^s = (1 - s)\hat{\boldsymbol{\mu}} + s\boldsymbol{\pi}, \quad (22)$$

where $0 \leq s \leq 1$.

BL also allows investors to express their views on the expected return $\boldsymbol{\mu}$. An investor can specify a "pick" matrix $\mathbf{P} \in \mathbb{R}^{K \times N}$, whose k th row describes a linear relative performance relationship among assets. In addition, an investor can express level of uncertainty in any of her views through a normal model

$$\mathbf{P}\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{v}, \boldsymbol{\Omega}). \quad (23)$$

In short, the mean vector $\mathbf{v} \in \mathbb{R}^K$ describes relative performance among assets, and the covariance matrix $\boldsymbol{\Omega}$ quantifies levels of uncertainty of views. If an investor has a specific view on asset returns, components of \mathbf{v} can be chosen accordingly. However, if an investor has only qualitative view, BL suggest \mathbf{v} can be set in terms of market volatility:

$$v_k = (\mathbf{P}\boldsymbol{\pi})_k + \eta_k \sqrt{(\mathbf{P}\boldsymbol{\Sigma}\mathbf{P}^T)_{k,k}}, \quad k = 1, \dots, K, \quad (24)$$

where $\eta_k \in \{-\beta, -\alpha, \alpha, \beta\}$. Here α and β are positive constants and different values of η_k represent "bullish" or "bearish" views to some extent. Meucci[31] suggests that for convenience one can set

$$\boldsymbol{\Omega} = \frac{1}{c} \mathbf{P}\boldsymbol{\Sigma}\mathbf{P}^T \quad (25)$$

among other choices, where c is a positive constant representing investor's confidence.

Using Bayes formula, given \mathbf{P} , \mathbf{v} , and $\mathbf{\Omega}$, the posterior distribution of $\boldsymbol{\mu}$ can be obtained:

$$\boldsymbol{\mu}|\mathbf{v}, \mathbf{\Omega} \sim \mathcal{N}(\boldsymbol{\mu}_{BL}, \boldsymbol{\Sigma}_{BL}), \quad (26)$$

where

$$\begin{aligned} \boldsymbol{\mu}_{BL} &= \boldsymbol{\pi} + \tau \boldsymbol{\Sigma} \mathbf{P}^T (\tau \mathbf{P} \boldsymbol{\Sigma} \mathbf{P}^T + \mathbf{\Omega})^{-1} (\mathbf{v} - \mathbf{P} \boldsymbol{\pi}) \\ \boldsymbol{\Sigma}_{BL} &= (1 + \tau) \boldsymbol{\Sigma} - \tau^2 \boldsymbol{\Sigma} \mathbf{P}^T (\tau \mathbf{P} \boldsymbol{\Sigma} \mathbf{P}^T + \mathbf{\Omega})^{-1} \mathbf{P} \boldsymbol{\Sigma} \end{aligned} \quad (27)$$

The distribution of \mathbf{X} can be computed using the fact that $\mathbf{X} \stackrel{d}{=} \boldsymbol{\mu} + \mathbf{Z}$ where $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$. The posterior distribution of \mathbf{X} is:

$$\mathbf{X}|\mathbf{v}, \mathbf{\Omega} \sim \mathcal{N}(\boldsymbol{\mu}_{BL}, \boldsymbol{\Sigma}_{BL} + \boldsymbol{\Sigma}). \quad (28)$$

Thus, an investor can use the posterior estimations $\boldsymbol{\mu}_{BL}$ and $(\boldsymbol{\Sigma}_{BL} + \boldsymbol{\Sigma})$ as the expected return and covariance matrix in an mean-variance optimization model. BL compared the resultant portfolios based on posterior estimations displayed meaningful differences from the market portfolio, while preserving overall allocation characteristics.

3.2 Extension of Black-Litterman Model

Meucci considers how does an investor's confidence in her views affect posterior parameters $\boldsymbol{\mu}_{BL}$ and $\boldsymbol{\Sigma}_{BL}$. He argues that when the confidence is full, the posterior distribution $\mathcal{N}(\boldsymbol{\mu}_{BL}, \boldsymbol{\Sigma}_{BL})$ should converge the posterior distribution of $\mathcal{N}(\boldsymbol{\pi}, \boldsymbol{\Sigma})$ conditioned on the specific views \mathbf{v} , that is

$$\begin{aligned} \lim_{\mathbf{\Omega} \rightarrow \mathbf{0}} \boldsymbol{\mu}_{BL} &= \boldsymbol{\mu}|\mathbf{v} \\ \lim_{\mathbf{\Omega} \rightarrow \mathbf{0}} \boldsymbol{\Sigma}_{BL} &= \boldsymbol{\Sigma}|\mathbf{v}, \end{aligned} \quad (29)$$

where

$$\begin{aligned} \boldsymbol{\mu}|\mathbf{v} &= \boldsymbol{\pi} + \boldsymbol{\Sigma} \mathbf{P}^T (\mathbf{P} \boldsymbol{\Sigma} \mathbf{P}^T)^{-1} (\mathbf{v} - \mathbf{P} \boldsymbol{\pi}) \\ \boldsymbol{\Sigma}|\mathbf{v} &= \boldsymbol{\Sigma} - \boldsymbol{\Sigma} \mathbf{P}^T (\mathbf{P} \boldsymbol{\Sigma} \mathbf{P}^T)^{-1} \mathbf{P} \boldsymbol{\Sigma}. \end{aligned} \quad (30)$$

However, above relations are not completely satisfied in the original BL model. To make the model more consistent with investors' intuition, Meucci modifies BL model by directly considering views on the asset return vector $X \in \mathbb{R}^N$. Instead of considering

$\boldsymbol{\mu}$ as a random variable, Meucci set it to be the market equilibrium estimation, i.e. $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\pi}, \boldsymbol{\Sigma})$. In addition, the view is modelled as a normally distributed random variable \mathbf{V} with mean $\mathbf{P}\mathbf{X}$ and covariance $\boldsymbol{\Omega}$. In particular, given any realization \mathbf{x} of \mathbf{X} , the conditional distribution of \mathbf{V} is

$$\mathbf{V}|\mathbf{x} \sim \mathcal{N}(\mathbf{P}\mathbf{x}, \boldsymbol{\Omega}). \quad (31)$$

By Bayesian formula, the posterior distribution of \mathbf{X} given a realization \mathbf{v} of \mathbf{V} and $\boldsymbol{\Omega}$ is

$$\mathbf{X}|\mathbf{v}, \boldsymbol{\Omega} \sim \mathcal{N}(\boldsymbol{\mu}_{BL}^m, \boldsymbol{\Sigma}_{BL}^m), \quad (32)$$

where

$$\begin{aligned} \boldsymbol{\mu}_{BL}^m &= \boldsymbol{\pi} + \boldsymbol{\Sigma}\mathbf{P}^T(\mathbf{P}\boldsymbol{\Sigma}\mathbf{P}^T + \boldsymbol{\Omega})^{-1}(\mathbf{v} - \mathbf{P}\boldsymbol{\pi}) \\ \boldsymbol{\Sigma}_{BL}^m &= \boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{P}^T(\mathbf{P}\boldsymbol{\Sigma}\mathbf{P}^T + \boldsymbol{\Omega})^{-1}\mathbf{P}\boldsymbol{\Sigma}. \end{aligned} \quad (33)$$

Notice parameter τ in original BL model doesn't appear in Meucci's modification and it's no longer necessary to add the original covariance matrix $\boldsymbol{\Sigma}$ to the posterior. The modified model can now perfectly match both the market and investors' view. On the one hand if $\boldsymbol{\Omega} \rightarrow \infty$ (no confidence in the view), the posterior distribution $\mathcal{N}(\boldsymbol{\mu}_{BL}^m, \boldsymbol{\Sigma}_{BL}^m)$ converges to the prior distribution $\mathcal{N}(\boldsymbol{\pi}, \boldsymbol{\Sigma})$. On the other hand if $\boldsymbol{\Omega} \rightarrow \mathbf{0}$ (full confidence in the view), the posterior distribution converges to $\mathcal{N}(\boldsymbol{\mu}|\mathbf{v}, \boldsymbol{\Sigma}|\mathbf{v})$ that represents investors' view.

3.3 Robust Bayesian Allocation

So far, we have seen two approaches that aim to deal with statistical errors and portfolios' sensitivity. Robust optimization works by optimizing the worst-case outcome over a set of possible values centered around estimations. Bayesian methods work by adjusting statistical estimations toward investors' views on asset returns. Meucci[31] introduced a Robust Bayesian Allocation (RBA) model that combines robust optimization and Bayesian methods. In RBA model, Bayesian methods are used not only to combine the view and market data, but also to identify suitable uncertainty ranges for the parameters such as mean return and covariance matrix. Then the robust optimization problem is solved via convex optimization.

Consider a market with N assets where trading happens at a set of discrete times. An equivalent version of the classical mean-variance portfolio optimization problem is

$$\begin{aligned} \max_{\phi \in \mathbb{R}^N} \quad & \phi^T \boldsymbol{\mu} \\ \text{s.t.} \quad & \phi^T \boldsymbol{\Sigma} \phi \leq \beta, \\ & \mathbf{1}^T \phi = 1, \end{aligned} \tag{34}$$

where $\phi \in \mathbb{R}^N$, $\boldsymbol{\mu} \in \mathbb{R}^N$, $\boldsymbol{\Sigma} \in \mathbb{R}^{N \times N}$, and β are portfolio weights, mean return vector, covariance matrix, and maximal acceptable portfolio variance. Note this problem is the dual of Markowitz's formulation defined in chapter 1. A robust version of the problem can be set up similarly by considering sets of possible values for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$:

$$\begin{aligned} \max_{\phi \in \mathbb{R}^N} \quad & \min_{\boldsymbol{\mu} \in \Theta_\mu} \phi^T \boldsymbol{\mu} \\ \text{s.t.} \quad & \max_{\boldsymbol{\Sigma} \in \Theta_\Sigma} \phi^T \boldsymbol{\Sigma} \phi \leq \beta, \\ & \mathbf{1}^T \phi = 1, \end{aligned} \tag{35}$$

where Θ_μ and Θ_Σ denote corresponding uncertainty sets for $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$.

Assume at time $t \in \mathbb{N}$ an asset purchased $\tau \in \mathbb{N}$ period ago is being traded at price $P_t \in \mathbb{R}$. We can define the return of this asset between time $t - \tau$ and t as $R_{t,\tau} = \frac{P_t}{P_{t-\tau}} - 1$. In addition, denote the return data of all assets between time $t - \tau$ and t as $\mathbf{R}_{t,\tau} \in \mathbb{R}^n$. In many linear factor models, $\mathbf{R}_{t,\tau}$ is assumed to be normally distributed with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, i.e. $\mathbf{R}_{t,\tau} | \boldsymbol{\mu}, \boldsymbol{\Sigma} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Meucci models that the investor's prior knowledge about $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ with a normal-inverse-Wishart distribution:

$$\begin{aligned} \boldsymbol{\mu} | \boldsymbol{\Sigma} &\sim \mathcal{N}\left(\boldsymbol{\mu}_0, \frac{\boldsymbol{\Sigma}}{T_0}\right) \\ \boldsymbol{\Sigma}^{-1} &\sim \mathcal{W}\left(\nu_0, \frac{\boldsymbol{\Sigma}_0^{-1}}{\nu_0}\right), \end{aligned} \tag{36}$$

where $\boldsymbol{\mu}_0$ and $\boldsymbol{\Sigma}_0$ represent the investor's view about $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. T_0 and ν_0 represent the confidence in the view (in number of periods). The larger T_0 (or ν_0) is the more confidence an investor has in $\boldsymbol{\mu}_0$ (or $\boldsymbol{\Sigma}_0$). The posterior knowledge on $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ can be updated with sample mean $\hat{\boldsymbol{\mu}}$ and sample covariance $\hat{\boldsymbol{\Sigma}}$, which are estimated based on

$T \in \mathbb{N}$ periods of historical market data. In Bayesian framework, there are analytic formulas for the posterior estimation $(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$:

$$\begin{aligned}
T_1 &= T_0 + T \\
\nu_1 &= \nu_0 + T \\
\boldsymbol{\mu}_1 &= \frac{T_0 \boldsymbol{\mu}_0 + T \hat{\boldsymbol{\mu}}}{T_1} \\
\boldsymbol{\Sigma}_1 &= \frac{1}{\nu_1} [\nu_0 \boldsymbol{\Sigma}_0 + T \hat{\boldsymbol{\Sigma}} + \frac{(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})(\boldsymbol{\mu}_0 - \hat{\boldsymbol{\mu}})^T}{\frac{1}{T} + \frac{1}{T_0}}].
\end{aligned} \tag{37}$$

In above equation, T_1 and ν_1 can be understood as the total confidence in terms of the number of periods, and the posterior estimation $(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ is the weighted average of prior $(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$ and evidence $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})$.

Meucci's formulation not only gives the posterior estimation of $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ but also the uncertainty sets Θ_μ and Θ_Σ :

$$\begin{aligned}
\Theta_\mu &= \{\boldsymbol{\mu} : (\boldsymbol{\mu} - \boldsymbol{\mu}_{ce})^T \mathbf{S}_\mu^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}_{ce}) \leq q_\mu^2\} \\
\Theta_\Sigma &= \{\boldsymbol{\Sigma} : \text{vech}[\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{ce}]^T \mathbf{S}_\Sigma^{-1} \text{vech}[\boldsymbol{\Sigma} - \boldsymbol{\Sigma}_{ce}] \leq q_\Sigma^2\},
\end{aligned} \tag{38}$$

where

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_{ce} &= \boldsymbol{\mu}_1 \\
\hat{\boldsymbol{\Sigma}}_{ce} &= \frac{\nu_1}{\nu_1 + N + 1} \boldsymbol{\Sigma}_1 \\
\mathbf{S}_\mu &= \frac{1}{T_1} \frac{\nu_1}{\nu_1 - 2} \boldsymbol{\Sigma}_1 \\
\mathbf{S}_\Sigma &= \frac{2\nu_1^2}{(\nu_1 + N + 1)^3} (\mathbf{D}_N^T (\boldsymbol{\Sigma}_1^{-1} \otimes \boldsymbol{\Sigma}_1^{-1}) \mathbf{D}_N)^{-1},
\end{aligned} \tag{39}$$

where ‘‘vech’’ means stack the columns of a matrix skipping the redundant entries above the diagonal, \mathbf{D}_N is the duplication matrix that reinstates the redundant entries above the diagonal of a symmetric matrix, \otimes is the Kronecker product, and (q_μ, q_Σ) is a pair of hyper parameters which determine the size of uncertainty sets. In short, the uncertainty sets are ellipsoids centered at posterior estimations $\boldsymbol{\mu}_1$ and $\boldsymbol{\Sigma}_1$, which is quite similar to their counterparts in the GI model.

Meucci further proves that the robust optimization problem in equation (35) with uncertainty sets define in equations (37) and (38) is equivalent to

$$\begin{aligned} \max_{\phi \in \mathbb{R}^N} \quad & \phi^T \boldsymbol{\mu}_1 - \gamma_\mu \sqrt{\phi^T \boldsymbol{\Sigma}_1 \phi} \\ \text{s.t.} \quad & \phi^T \boldsymbol{\Sigma}_1 \phi \leq \gamma_\Sigma \\ & \mathbf{1}^T \phi = 1, \end{aligned} \tag{40}$$

where

$$\begin{aligned} \gamma_\mu &= \sqrt{\frac{q_\mu}{T_1} \frac{\nu_1}{\nu_1 - 2}} \\ \gamma_\Sigma &= \frac{\beta}{\frac{\nu}{\nu + N + 1} + \sqrt{\frac{2\nu_1^2 q_\Sigma^2}{(\nu + N + 1)^3}}}. \end{aligned} \tag{41}$$

Finally, the optimization problem in equation (40) can be treated as a SOCP and solved numerically using existing convex optimization softwares such as CVX[17].

3.4 Numerical Study

The RBA model combines features of Bayesian methods and robust optimization. Consequently, it requires input of both types of models we have previously discussed. For example, prior mean return μ_0 and prior covariance matrix $\boldsymbol{\Sigma}_0$, together with their confidence T_0 and ν_0 expressed in number of periods. In addition, we must also choose q_μ and q_Σ that govern the size of uncertainty set (similar to the parameter ω in GI model). Last but not least, another pair of mean and covariance- $(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})$ -needs to be estimated from historical price data so that the prior estimation can be updated. We present some summary statistics in Figure 3(a) including prior belief, estimations from data, and posterior adjusted values of monthly return and risk of assets, which are the same as in the study of GI or Lu models. In our numerical study, we assign equal confidence (weight) to parameters that represent investor's view and those estimated from historical data. By using different levels of q_μ and q_Σ that control the size of uncertainty sets, we study how portfolio allocations change accordingly.

Asset	1	2	3	4	5	6	7	8	9	10	11	12	13
Prior Return (%)	0.08	0.16	0.22	0.31	0.28	0.21	0.26	0.28	0.32	0.51	0.4	0.51	0.58
Data Return (%)	0.02	0.19	0.34	0.52	0.04	0.39	0.52	0.44	0.64	0.7	0.69	0.73	1.16
Posterior Return (%)	0.05	0.17	0.28	0.42	0.16	0.3	0.39	0.36	0.48	0.6	0.55	0.62	0.87
Prior Risk (%)	0.14	1.16	2.99	1.55	1.32	1.31	2.15	1.75	2.93	2.17	2.44	3.10	4.41
Data Risk (%)	0.00	0.46	0.65	1.35	1.30	1.18	1.89	1.22	2.36	1.74	1.52	1.76	3.51
Posterior Risk (%)	0.10	0.88	2.17	1.46	1.32	1.25	2.03	1.51	2.66	1.97	2.03	2.52	4.00

(a) Return and risk representing prior, data, and posterior estimations

Robustness\Asset	1	2	3	4	5	6	7	8	9	10	11	12	13
0							11%		54%	31%			3%
0.1					4%		13%		53%	29%			2%
0.15					13%		11%		50%	26%			1%
0.2					18%		9%		49%	23%			1%
0.25					22%		9%		48%	21%			1%
0.3					25%		8%		47%	19%			1%
0.35					28%		8%		46%	18%			1%
0.4					30%		7%		46%	17%			1%

(b) Portfolio allocations using RBA model

Figure 3: Summary statistics and portfolio allocation using RBA model

Portfolio allocations corresponding to different robustness (equal to q_μ and q_Σ) can be found in Figure 3(b). Overall, optimal portfolios are concentrated in 5 out of 13 assets. This might be caused by the risk-return feature of assets, and the fact that there is no allocation constraint. In particular, when the robustness is zero, there is more 50% allocation in the ninth asset. When the robustness increases, the optimal portfolio gradually picks up allocation in the fifth asset, while giving less weight on the other 4 assets. Eventually, the fifth asset get a 30% weight and the weight on the ninth asset is reduced to be less than 50%. Within this limited case, RBA model does gives slightly better diversification comparing to the classic model. However, we must emphasize how much more complicated RBA is than Markowitz's, or even the Black-Litterman's model. With many parameters that must be chosen by the investor, RBA definitely requires more than average investment expertise. Together with sophisticated elements in robust optimization, the model complexity is the greatest among all models we have discussed so far. In summary, one should be very careful and have sufficient confidence to use RBA model in practice.

Chapter 4

Dynamic Portfolio Optimization

So far, we have considered portfolio optimization in the static setting, where an investor makes investment decision at a specific moment in time without further update. In practice, new data and information are coming in every day, hour, and second. Statistical estimations and investors' views can change frequently. Thus, we can treat portfolio optimization as dynamic decision making problems. In this chapter we discuss two dynamic programming problems. The first is a general robust dynamic programming problem in discrete time. The second is a portfolio optimization problem which was first studied by Merton.

In the sequential setting, investment decisions are made in different stages in time. Each decision, besides generating a possible immediate reward, changes the ground of all future decisions and rewards associated. Due to the uncertain nature of financial markets, an investor only has limited information about the immediate reward of each decision and future decisions are dependent on previous ones. To achieve a good performance through all stages, the investor has to consider both the payoff of immediate decision and how is this decision going to influence future payoffs. Dynamic programming (DP) is a mathematical framework that allows the decision maker to compute a good overall strategy by succinctly encoding the evolving information state. It was originally developed by Richard Bellman[3] in the 1950s and has been widely applied in areas such as aerospace engineering, computer science, and economics. The fundamental idea behind DP is to break a complicated problem into nested sub-problems and solve

them recursively. In formal DP theory, uncertainty in the environment is modeled by a Markov process whose transition probability depends both on the information state and the action taken by the decision maker. It assumes the transition probability of each state-action pair is known exactly, and the goal is to choose a policy to maximize an objective function. The idea of dynamic programming can be extended to continuous time with moderate mathematical sophistication. As we will see in Merton's portfolio problem, using stochastic calculus and differential equation, a continuous-time dynamic programming can be elegantly formulated as a partial differential equation.

4.1 Robust Dynamic Programming

Iyengar[19] proposed a Robust Dynamic Programming (RDP) framework, where the transition probability of each state-action pair is no longer assumed to be known exactly (compare to classic dynamic programming problems).

Assume decisions are made at discrete point in time $t \in T = \{0, 1, \dots, N - 1\}$ where N is a fixed positive integer. At each $t \in T$, the decision maker observes a current state $s_t \in S_t$, where S_t is a discrete set of possible states at time t . Based the state s_t at time t , the decision maker needs to take an action $a \in A(s_t)$ - also a discrete set of all possible actions at s_t . For any discrete set B , let $M(B)$ be the set of all probability measures on B . Then a random action based on state s_t at time t naturally corresponds to a probability measure $q_{s_t} \in M(A(s_t))$. For $t < N$, the decision maker receives a reward $r_t(s_t, a_t, s_{t+1})$ if the action a_t is taken in state s_t at time t , and the state s_{t+1} is reached at time $t + 1$. The reward at the last stage r_N is assumed to be a function of s_N only.

To describe the concept that the transient probability is not known exactly, Iyengar introduced the notion of ambiguity. At each time $t \in T$, a set of probability measures $P_t(s, a) \subset M(S_{t+1})$ is assigned to a state-action pair (s, a) . This means at time t , if the decision maker in state s takes action a , then the transient probability to all possible states at $t + 1$ is an element of $p_{sa} \in P_t(s, a)$. However, there is no way to determine which element of $P_t(s, a)$ is actually governing the transition process.

Under the ambiguity of transition probability, a decision maker needs a rule to choose suitable actions based on historical states and actions, denoted as $h_t = (s_0, a_0, \dots, s_{t-1}, a_{t-1}, s_t)$ for $t < N$. Let H_t be the set of all possible historical states and actions up to time t . A randomized decision rule is a map $d_t : H_t \mapsto M(A(s_t))$, from H_t to the set of probability measures on possible action based on state s_t . A policy π of a robust dynamic programming problem is a collection of decision rule, i.e. $\pi = \{d_t | t \in T\}$. Given a decision rule d_t , the set of probability measures consistent with d_t is defined as

$$F^{d_t} = \{P : H_t \mapsto M(A(s_t) \times S_{t+1}) | \forall h \in H_t, P_h(a, s) = q_{d_t(h)}(a)p_{s_t a}(s), \\ p_{s_t a} \in P(s_t, a), a \in A(s_t), s \in S_{t+1}\}. \quad (42)$$

It is also assumed (for regularity) that the set of probability measures consistent with a policy π is just the product of sets of probability measures consistent with decision rule d_t at each time $t < N$:

$$F^\pi = F^{d_0} \times F^{d_1} \times \dots \times F^{d_{N-1}}. \quad (43)$$

In a RDP problem, the reward $V_0^\pi(s)$ generated by a policy π starting from the initial state $s_0 = s$ is defined as

$$V_0^\pi(s) = \inf_{P \in F^\pi} \mathbf{E}^P[r_N(s_N) + \sum_{0 \leq t < N} r_t(s_t, d_t(h_t), s_{t+1})], \quad (44)$$

where \mathbf{E}^P denotes the expectation w.r.t. the probability measure $P \in F^\pi$. The interpretation is that the reward of a policy π is the minimal expected reward over all probability measures consistent with π . The goal of a decision maker is to find an optimal policy π^* that maximize the robust value function

$$V_0^*(s) = \sup_{\pi \in \Pi} V_0^\pi(s), \quad (45)$$

where Π is the set of all admissible history-dependent policies. Similarly, the reward of policy π starting from history h_n with $n < N$ is defined as

$$V_n^\pi(h_n) = \inf_{P \in F_n^\pi} \mathbf{E}^P[r_N(s_N) + \sum_{n \leq t < N} r_t(s_t, d_t(h_t), s_{t+1})], \quad (46)$$

where $F_n^\pi = F^{d_n} \times F^{d_{n+1}} \times \dots \times F^{d_{N-1}}$. The optimal reward starting from h_n is defined as

$$V_n^*(h_n) = \sup_{\pi \in \Pi_n} V_n^\pi(h_n), \quad (47)$$

where Π_n denotes the set of all admissible history-dependent policies concerning $t \geq n$.

One remarkable result on RDP is the robust Bellman equation. It states that optimal reward functions $\{V_n^* | 0 \leq n \leq N\}$ satisfy the recursive relation

$$\begin{aligned} V_N^*(h_N) &= r_N(s_N), \\ V_n^*(h_n) &= \sup_{a \in A(s_n)} \left\{ \inf_{p \in P(s_n, a)} \mathbf{E}^p[r_n(s_n, a, s) + V_{n+1}^*(h_n, a, s)] \right\}, n < N, \end{aligned} \quad (48)$$

where $V_{n+1}^*(h_n, a, s)$ means the optimal reward function at time $n+1$ with history date obtained by concatenating a and s to h_n . The relation suggests that the optimal reward given history h_n can be achieved by considering all possible action in $A(s_n)$ and move one step forward in time. Although the relation gives a guideline of the connection between optimal reward functions at different times, it is not easy to apply due to fact that every state and action in all previous times come in to play. Fortunately, there is a significantly simplified equivalent version which states the following:

$$\begin{aligned} V_N^*(s_N) &= r_N(s_N), \\ V_n^*(s_n) &= \sup_{a \in A(s_n)} \left\{ \inf_{p \in P(s_n, a)} \mathbf{E}^p[r_n(s_n, a, s) + V_{n+1}^*(s)] \right\}, n < N, \end{aligned} \quad (49)$$

where $s \in S_{n+1}$. Compare to the history-dependent version, above relation says robust value functions $V_n^*(h_n)$ are functions of the most recent states only, i.e. in the form of $V_n^*(s_n)$. Instead of carrying a sequence of historical events, we only need to consider the current state and the immediate future. Thus, the robust dynamic programming is a Markov decision-making problem with ambiguity in transition probability.

When dealing with discrete states and actions, RDP can be a powerful tool. However, it is not easy to find a good, natural set of ambiguity sets. In addition, it's not completely clear how to adapt RDP to continuous decision-making problems. Although approximation by discretizing time, action, and underlying probability measures

is a possibility, the computational complexity and numerical issues that come with discretization can be difficult to deal with.

4.2 Merton's portfolio problem

Robert Merton[28, 30] examined the problem of optimal portfolio selection and consumption rules in continuous-time setting. Unlike discrete-time portfolio allocation problems, assets continuously generate returns in Merton's model and they are the only source of an investor's income. Likewise, consumption and portfolio allocations can also be continuously adjusted by the investor.

Assume there are $m \in \mathbb{N}$ assets in the market and the price of i th asset, $P_i(t)$ is generated by a Geometric Brownian Motion (GBM) with constant drift $a_i \in \mathbb{R}$ and constant volatility $\sigma_i \in \mathbb{R}$, i.e.

$$dP_i(t) = a_i P_i(t) dt + \sigma_i P_i(t) dB_i(t), \quad (50)$$

where $\{B_i(t)\}_{1 \leq i \leq m}$ are independent Brownian Motions (BM). At time $t \in \mathbb{R}$, denote the amount of total wealth as $W(t)$, consumption per unit time as $C(t)$, and the proportion of total wealth in the i th asset as $\omega_i(t)$ for $1 \leq i \leq m$. It is understood that the investor distribute all wealth among m assets, i.e. $\sum_{i=1}^m \omega_i(t) = 1$ for any t . It is easy to derive a Stochastic Differential Equation (SDE) that describes the evolution of $W(t)$ over time:

$$dW(t) = \left[\sum_{i=1}^m \omega_i(t) a_i W(t) - C(t) \right] dt + \sum_{i=1}^m \omega_i(t) \sigma_i W(t) dB_i(t). \quad (51)$$

Merton further studied a special case of above problem where there are only two assets in the market. One is a risky asset with drift $a \in \mathbb{R}$ and volatility $\sigma > 0$. The other is risk-free with drift $r \in \mathbb{R}$ and zero volatility. In the two-asset scenario, if we denote the proportion of total wealth in the risky asset as $\omega(t)$, then the proportion in the risk-free asset is $1 - \omega(t)$. The governing equation for the wealth process becomes

$$dW(t) = [(r + (a - r)\omega(t))W(t) - C(t)]dt + \omega(t)\sigma W(t)dB(t), \quad (52)$$

where $B(t)$ is a standard Brownian Motion (BM). Given the equation, an investor's goal is to find the optimal consumption rule $C(t)$ and portfolio allocation $\omega(t)$ which maximize the sum of expected utility and bequest at some future time:

$$\begin{aligned} \max_{C(t), \omega(t)} \quad & \mathbf{E}[Q(W(T), T) + \int_0^T e^{-\rho\tau} U(C(\tau)) d\tau] \\ \text{s.t.} \quad & C(t) \geq 0 \\ & \omega(t) \geq 0 \\ & W(0) = W_0 \geq 0. \end{aligned} \tag{53}$$

In above formulation, \mathbf{E} is short for the conditional expectation operator $E(0)$ at initial time 0, T stands for the investment horizon, ρ is a non-negative discount factor, $Q(\cdot, \cdot)$ is the bequest valuation function, and $U(\cdot)$ is the utility function. The first part inside the expectation represent the bequest at time T and the second part stands for the accumulated utility incurred by consumption between time 0 and T . In addition, standard utility theory assumes $U(\cdot)$ is concave and smooth enough with $U' > 0$ and $U'' < 0$.

To derive the optimal conditions for the weight $\omega(t)$ and consumption $C(t)$, Merton utilized the Bellman principle of optimality for Dynamic Programming. For $t \in [0, T]$, define

$$\begin{aligned} I(W(t), t) = \max_{C(s), \omega(s)} \quad & \mathbf{E}(t) \left[\int_t^T e^{-\rho\tau} U(C(\tau)) d\tau + Q(W(T), T) \right] \\ \text{s.t.} \quad & C(s) \geq 0 \\ & \omega(s) \geq 0, \end{aligned} \tag{54}$$

where $\mathbf{E}(t)$ is the conditional expectation operator at time t . $I(W(t), t)$ is understood as the maximal expected utility that can be achieved under any consumption $C(s)$ and weight $\omega(s)$ starting at time t . The principle of DP implies that for any $t_0 < t$, $I(W(t_0), t_0)$ should be determined by the following optimization problem

$$\begin{aligned} I(W(t_0), t_0) = \max_{C(s), \omega(s)} \quad & \mathbf{E}(t_0) \left[\int_{t_0}^t e^{-\rho\tau} U(C(\tau)) d\tau + I(W(t), t) \right] \\ \text{s.t.} \quad & C(s) \geq 0 \\ & \omega(s) \geq 0. \end{aligned} \tag{55}$$

Assume $I(\cdot, \cdot)$ has bounded third partial derivatives, we can apply Taylor's theorem to $I(W(t), t)$ with $t = t_0 + h$ and let $h \rightarrow 0$, the governing equation for optimal consumption $C(t)$ and $\omega(t)$ at any $t \in [0, T]$ can be derived:

$$0 = \max_{C(t), \omega(t)} \left\{ e^{-\rho t} U(C(t)) + \frac{\partial I}{\partial t} + \frac{\partial I}{\partial W} (\omega(t) [(a-r) + r] W(t) - C(t)) \right. \\ \left. + \frac{1}{2} \frac{\partial^2 I}{\partial W^2} \omega^2(t) \sigma^2 W^2(t) \right\}. \quad (56)$$

Denote the objective function to be maximized in the above equation as $\Phi(\omega, C, W, t)$. At a fixed time t , the first-order conditions for a regular interior maximum are

$$\begin{cases} \frac{\partial \Phi}{\partial C} = 0 = e^{-\rho t} U'(C) - \frac{\partial I}{\partial W} \\ \frac{\partial \Phi}{\partial \omega} = 0 = \frac{\partial I}{\partial W} (a-r) + \frac{\partial^2 I}{\partial W^2} \omega \sigma^2 W. \end{cases} \quad (57)$$

For the second order conditions, note $\frac{\partial^2 \Phi}{\partial C \partial \omega} = \frac{\partial^2 \Phi}{\partial \omega \partial C} = 0$, and $\frac{\partial^2 \Phi}{\partial C^2} = e^{-\rho t} U''(C) < 0$ since U is assumed to be strictly concave. Merton further assumed $I(W, t)$ is strictly concave in W . Consequently, the second order conditions are automatically satisfied. Thus, the original problem of determining optimal consumption and portfolio allocation is reduced to solving the following system of equations

$$\begin{cases} \Phi(\omega^*, C^*, W, t) = 0 \\ \Phi_\omega(\omega^*, C^*, W, t) = 0 \\ \Phi_C(\omega^*, C^*, W, t) = 0 \end{cases} \quad (58)$$

under the boundary condition $I(W(T), T) = Q(W(T), T)$ at the terminal time.

In general, the system of equations in (58) are difficult to solve due to the coupling between partial differential equation and algebraic equations. However, it can be solved explicitly if the utility has constant relative risk-aversion, i.e.

$$U(C) = \frac{C^\gamma}{\gamma}, \quad (59)$$

where $\gamma \in (0, 1)$. If we further assume that the bequest at the terminal time is given by

$Q(W(T), T) = \epsilon^{1-\gamma} e^{-\rho T} U(W(T))$ ($0 < \epsilon \ll 1$), the optimality conditions become

$$\begin{cases} C^*(t) = [e^{\rho t} \frac{\partial I_t}{\partial W}]^{\frac{1}{1-\gamma}} \\ \omega^*(t) = \frac{-(a-r) \frac{\partial I_t}{\partial W}}{\sigma^2 W^2 \frac{\partial^2 I_t}{\partial W^2}} \\ 0 = e^{\frac{-\rho t}{1-\gamma}} \left[\frac{1-\gamma}{\gamma} \frac{\partial I_t}{\partial W} \right]^{\frac{\gamma}{\gamma-1}} + \frac{\partial I_t}{\partial t} + \frac{\partial I_t}{\partial W} r W - \frac{(a-r)^2}{2\sigma^2} \frac{[\frac{\partial I_t}{\partial W}]^2}{\frac{\partial^2 I_t}{\partial W^2}}, \end{cases} \quad (60)$$

subject to $I(W(T), T) = \epsilon^{1-\gamma} e^{-\rho T} \frac{[W(T)]^\gamma}{\gamma}$. To solve the last equation in (52), substitute a trial solution $\bar{I}_t(W(t), t) = \frac{b(t)}{\gamma} e^{-\rho t} [W(t)]^\gamma$ and it turns out $b(t)$ must satisfy the ordinary differential equation

$$\begin{aligned} b'(t) &= \mu b(t) - (1-\gamma)[b(t)]^{\frac{-\gamma}{1-\gamma}} \\ \text{s.t. } b(T) &= \epsilon^{1-\gamma}, \end{aligned} \quad (61)$$

where the constant $\mu = \rho - \gamma \left[\frac{(a-r)^2}{2\sigma^2(1-\gamma)} + r \right]$. The governing equation of $b(t)$ has solution $b(t) = \left[\frac{1+(v\epsilon-1)e^{vt-T}}{v} \right]^{1-\gamma}$, where $v = \frac{\mu}{1-\gamma}$. Based on $b(t)$, the optimal solutions for $\omega(t)$ and $C(t)$ are

$$\begin{aligned} \omega^*(t) &= \frac{a-r}{\sigma^2(1-\gamma)} \\ C^*(t) &= \begin{cases} \frac{v}{1+(v\epsilon-1)e^{vt-T}} W(t) & \text{if } v \neq 0 \\ \frac{1}{T-t+\epsilon} W(t) & \text{if } v = 0. \end{cases} \end{aligned} \quad (62)$$

There is a very important implication from the expression of $\omega^*(t)$. Notice $\omega^*(t)$ is a constant independent of time. The solution to Merton's problem suggest that with any constant relative risk-aversion utility function, an investor should put a constant portion of wealth into the risky asset to maximize the current expectation of utility at terminal time. Notably, Merton also studied a wide range of related problems which involved other types of utility function, multiple asset classes, and finite/infinite time horizon. His pioneering work had tremendous impacts on portfolio allocation, asset pricing, and the study of financial derivatives. In 1997, Robert Merton and Myron Scholes were awarded the Nobel Memorial Prize in Economic Sciences for their work in derivative pricing in collaboration with the late Fischer Black (1938-1995).

Nowadays, Merton problem has become a classic example in stochastic control[33] and we do not need to go through all the technical steps went through by Merton. In

general, there are mainly three steps to solve a stochastic control problem. The first step is to transform the problem into a Partial Differential Equation (PDE) using Dynamic Programming Principle. The second step is to solve the PDE, usually by further converting it to a system of Ordinary Differential Equations. Finally, if there is a nice solution to the PDE, we come back and try to verify it is actually a solution to the original stochastic control problem.

Chapter 5

Dynamic Portfolio Optimization with Bayesian Update - Discrete Model

We consider a discrete portfolio optimization problem that is in the intersection of dynamic programming, robust optimization, and Bayesian theory. To be specific, an investor can adjust portfolio allocation at the beginning of each period, which makes the problem dynamic. However, a crucial parameter that governs the transitional probability measure of price movement is unknown to us. Consequently, we cannot solve the problem directly using Dynamic Programming principle. To deal with the parameter uncertainty, we model the unknown parameter via a probability distribution. Together with Bayesian theory, we are able to update the distribution of unknown parameter with price movement we have observed. Eventually, we will see how to solve it with dynamic programming and Bayesian update.

5.1 A Discrete Model with an Uncertain Parameter

We consider a discrete-time market where there are $T \in \mathbb{N}$ investment periods (indexed by $0, 1, \dots, T-1$). In this market, there is a risky asset and a risk-free asset. The risk-free

asset offers a constant rate of return $r \in \mathbb{R}$ in each period. In contrast, the price of risky asset can either go up by a factor e^a or go down by a factor e^{-a} , where $a \in \mathbb{R}^+$ is a constant. However, the probability of price increase is not known exactly. We treat the probability of price increase as a continuous random variable and use the beta distribution to model its probability distribution. The beta distribution[21] is a family of continuous probability distributions defined on the interval $[0, 1]$ and parameterized by two positive shape parameters $\alpha \in \mathbb{R}$ and $\beta \in \mathbb{R}$. The probability density function given α and β is

$$B(x, \alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}, \quad (63)$$

where $\Gamma(\cdot)$ is the gamma function[13]. Thus, the probability that the price of risky asset will increase in a certain period could be any number in $[0, 1]$ with respect to the beta distribution $B(x, \alpha, \beta)$. We assign a Bayesian-type method for the investor to make inference about the probability distribution. Initially, we set $\alpha = \beta = 1$. For $0 \leq t < T$, denote i_t as the number of periods where the price of risky asset moves up in the previous t periods, and j_t as the number of periods where the price moves down. A simple counting argument implies us $i_t + j_t = t$. Based on i_t and j_t , the investor then update her belief about the distribution as $B(x, 1 + i_t, 1 + j_t)$. Under this mechanism, the investor believes the distribution is likely to tilts toward 1 if $j_t > i_t$ and vice versa.

With the price dynamics and belief update mechanism laid out, we consider the investment problem. At the beginning of each period $0 \leq t < T$, the investor decides to put $\pi_t \in [0, 1]$ portion of her wealth in the risky asset, and rest $1 - \pi_t$ portion goes to the risk-free asset. The investor then hold the portfolio until the beginning of next period when a reallocation is allowed. Depending on the previous decision and actual price movement, the value of portfolio could go up or down. Similar to the Merton problem, the objective of the investor is to find the optimal portfolio allocation π_t^* so that the expected utility of at some future time is maximized. In our problem, we do not consider consumption of portfolio. The utility at the end of the T th period is a strictly concave function of portfolio value. For any time period $t < T$, the utility is determined by maximizing the expected utility of immediate future over all possible investment decisions. Formally, define the utility function at the end of the T th period

as

$$U_T(y) = \frac{y^{1-\gamma}}{1-\gamma}, \quad (64)$$

where y denotes the total portfolio value at the end of T th period, and $\gamma \neq 1$ is a positive constant. For $0 \leq t < T$, denote the portfolio value at the beginning of t th period as $v_t \in \mathbb{R}$. If the investor decide to put $\pi_t \in [0, 1]$ portion of v_t in the risky asset, her portfolio value at the beginning of next period (or the end of this period, equivalently) could be $v_t(1 - \pi_t)e^r + v_t\pi_t e^a$ if the price of risky asset increase during t th period, or $v_t(1 - \pi_t)e^r + v_t\pi_t e^{-a}$ if the price decrease. Since the actual probability of price change is only known as a distribution, we define the value function associated with portfolio value v_t and allocation π_t as

$$V_t(v_t, \pi_t) = \int_0^1 \left[U_{t+1}(v_t(1 - \pi_t)e^r + v_t\pi_t e^a)x + U_{t+1}(v_t(1 - \pi_t)e^r + v_t\pi_t e^{-a})(1 - x) \right] B(x, 1 + i_t, 1 + j_t) dx. \quad (65)$$

In other words, $V_t(v_t, \pi_t)$ is the expected utility of immediate future w.r.t. the distribution of price move that represents investor's current belief. Based on $V_t(v_t, \pi_t)$, we can define the utility of v_t as

$$U_t(v_t) = \max_{\pi_t \in [0, 1]} V_t(v_t, \pi_t), \quad (66)$$

which is the maximal value function among all possible allocations. The optimal portfolio allocation is naturally given by

$$\pi_t^* = \operatorname{argmax}_{\pi_t \in [0, 1]} V_t(v_t, \pi_t). \quad (67)$$

In summary, the investor chooses π_t^* that maximize the value function of v_t for each $0 \leq t < T$ based on her current belief of the probability distribution. In particular, at the very beginning of the investment process she wants to find

$$\pi_0^* = \operatorname{argmax}_{\pi_0 \in [0, 1]} V_0(v_0, \pi_0). \quad (68)$$

5.2 Transformation to a Dynamic Programming

It turns out the portfolio allocation problem described in section 5.1 can be reformulated into an ordinary dynamic programming problem. We describe details in the following theorems.

Theorem 1. For $0 \leq t \leq T$, utility function $U_t(v_t)$ takes the form

$$U_t(v_t) = c_t(i_t, j_t)v_t^{1-\gamma}, \quad (69)$$

where the coefficient c_t is a function of i_t and j_t which govern the beta distribution $B(x, 1 + i_t, 1 + j_t)$.

Proof. We proceed by backward induction. First note the utility function after the last investment period has the form $U_T(y) = \frac{y^{1-\gamma}}{1-\gamma}$, which satisfy the claim. Now assume $U_{n+1}(v_{n+1}) = c_{n+1}(i_{n+1}, j_{n+1})v_{n+1}^{1-\gamma}$ satisfy the claim for some $n \leq T - 1$. Consider the value function for any pair (v_n, π_n)

$$\begin{aligned} V_n(v_n, \pi_n) = & \int_0^1 \left[U_{n+1}(v_n(1 - \pi_n)e^r + v_n\pi_n e^a)x \right. \\ & \left. + U_{n+1}(v_n(1 - \pi_n)e^r + v_n\pi_n e^{-a})(1 - x) \right] B(x, 1 + i_n, 1 + j_n) dx. \end{aligned} \quad (70)$$

By the inductive assumption about U_{n+1} , we can rewrite the above integral as

$$\begin{aligned} V_n(v_n, \pi_n) = & \int_0^1 \left[c_{n+1}(i_n + 1, j_n)v_n^{1-\gamma}(e^r - \pi_n e^r + \pi_n e^a)^{1-\gamma}x \right. \\ & \left. + c_{n+1}(i_n, j_n + 1)v_n^{1-\gamma}(e^r - \pi_n e^r + \pi_n e^{-a})^{1-\gamma}(1 - x) \right] B(x, 1 + i_n, 1 + j_n) dx, \end{aligned} \quad (71)$$

where the difference in c_{n+1} reflects different price movement in two different scenarios. After pulling $v_n^{1-\gamma}$ outside, the integral is then completely determined by i_n , j_n , and π_n . Note when evaluating the integral, i_n and j_n have been determined from past price movement. Thus we can write

$$V_n(v_n, \pi_n) = f_n(i_n, j_n, \pi_n)v_n^{1-\gamma}, \quad (72)$$

where

$$\begin{aligned} f_n(i_n, j_n, \pi_n) = & \int_0^1 \left[c_{n+1}(i_n + 1, j_n)(e^r - \pi_n e^r + \pi_n e^a)^{1-\gamma}x \right. \\ & \left. + c_{n+1}(i_n, j_n + 1)(e^r - \pi_n e^r + \pi_n e^{-a})^{1-\gamma}(1 - x) \right] B(x, 1 + i_n, 1 + j_n) dx. \end{aligned} \quad (73)$$

Now the utility function can be written as

$$U_n(v_n) = v_n^{1-\gamma} \max_{\pi_n \in [0,1]} f_n(i_n, j_n, \pi_n), \quad (74)$$

where the last term depend only on i_n and j_n by maximizing over π_n . Thus the inductive step is complete and claim is valid for all $0 \leq t \leq T$. \square

Consequently, the theorem suggest that the amount of portfolio value v_t does not affect the optimal allocation, which is consistent with our intuition since we talk about portion π_t rather than amount. Moreover, it is the historical price movement i_t and j_t that determines the optima allocation through the beta distribution distribution. After separating v_t , at each state (i_t, j_t) , we only need to solve a one dimensional optimization problem over $\pi_t \in [0, 1]$. However, the expression of $f_n(i_n, j_n, \pi_n)$ makes it very difficult to find a closed-form solution to optimal, mainly due to the fractional exponents and the fact we restrict π_t within $[0, 1]$. Fortunately, in practice we don't need to consider every possible allocation between 0 and 1, and a discrete approximation should serve the purpose, e.g. all integer multiples of 1% or 0.1%.

When π_t is restricted to be chosen from a finite action set A , the maximization problem can be solved numerically by enumerating all choices of π_t . At time t , the portfolio optimization problem with the pair (i_t, j_t) becomes

$$\operatorname{argmax}_{\pi_t \in A} f_t(i_t, j_t, \pi_t), \quad (75)$$

where $f_t(i_t, j_t, \pi_t)$ is defined in the proof of the previous theorem. Combining the process presented in the proof, we summarize the complete problem-solving process in the algorithm below.

Algorithm 1: Discrete Dynamic Programming with Bayesian Update

Input: number of periods T , action set A ,

1 For $t = T - 1, T - 2, \dots, 1, 0$

(i) For every possible combination of i_t and j_t , solve the discrete optimization problem

$$\operatorname{argmax}_{\pi_t \in A} f_t(i_t, j_t, \pi_t)$$

to find the the optimal π_t^* and coefficient $c_t(i_t, j_t)$.

(ii) Stop if $t = 0$. Otherwise continue to next value of t .

Note in above algorithm, we must solve the optimization problem for all possibilities of i_t and j_t for a given period, which requires $c_{t+1}(i_{t+1}, j_{t+1})$ at the next time period. That is why we must solve the dynamic programming problem backward in time. Due to the constraint $i_t + j_t = t$, there are $t + 1$ different pairs of (i_t, j_t) . The connections between states (i_t, j_t) at different times can be conveniently represented by a binomial tree, which can be efficiently modeled in many numerical computing software. We will present some result of numerical experiment in the next section.

5.3 Numerical Study

We consider two instances of numerical studies: one does not allow short sell and the other does. In both cases we set the total number of periods $T = 8$, the return of risk-free asset $r = 0.1$, price jump constant $a = 0.1$, and the constant in utility function $\gamma = 2$. The dynamic programming algorithm described earlier is implemented in Python.

In the case where no short sell is allowed, the admissible control set is

$$A = \left\{ \frac{n}{1000} \mid n \in \mathbb{N}, 0 \leq n \leq 1000 \right\}. \quad (76)$$

The result of numerical study is presented in Figure 4. In Figure 4, each branch represents a possible price movement and each node represents a point when the investor must make decision about π_t . The number attached to each branch is the expected probability (according to the distribution $B(x, i_t, j_t)$) of price moving up or down, depending on whether the branch leads to a higher or lower node. If the branch is green (or red), it leads to a higher (or lower) price in the next period. The probability of price movement can also be visualized in the width of each branch. For example, if the price keeps decreasing, the probability of having an even lower price in the next period will increase, which is reflected by the increasing width of red branches as well as the increasing values associated with them. The number attached to each node directs how much (in percentage) the investor should put in the risky asset. The general pattern is that if the price drops more than it moves up, the investor would avoid risking the

portfolio at all. For example in period 7, the optimal allocation is to completely avoid the risky asset if one were at 4 nodes from the bottom. On the other hand, if the price moves up more often than it drops, the investor starts to invest in the risky asset, e.g. 4 nodes from the top in period 7. In general, this model leads to very conservative investment decision and the reason is that there is no “drift” term that drives the return of risky assets, i.e. the expected return of risky asset has no obvious advantage over the risk-free asset.




Figure 4: Optimal action tree (no short sell allowed)

Next we consider the situation where short sell is allowed. While everything else is unchanged, the admissible control set becomes

$$A = \left\{ \frac{n}{1000} \mid n \in \mathbb{Z}, -1000 \leq n \leq 1000 \right\}, \tag{77}$$

which means an investor can choose to short the risky asset up to 100%. The result of numerical study is presented in Figure 5. The general pattern is similar to what we can

observe from Figure 5, i.e. optimal allocations depends how does price move in previous periods. However, since short sell is allowed, it becomes beneficial to short when price of the risky asset becomes really low, e.g. bottom nodes at each time period.




Figure 5: Optimal action tree (short sell allowed)

Chapter 6

Dynamic Portfolio Optimization with Bayesian Update - Continuous Model

Models we have discussed in previous chapters can be seen as different extensions to Markowitz's model. Merton's portfolio problem studies portfolio optimization in continuous-time where an investor can continuously adjust portfolio allocations. The expected return and covariance in Markowitz's model correspond to the drift and volatility of the underlying price process in Merton's model. Bayesian methods allows statistical estimations to be modified according to investor's opinions or additional observations. Consequently, model parameters can also be updated in a dynamic fashion, as we have seen in discrete Bayesian models. Thus, it is natural to consider the combination of Bayesian methods and continuous-time models. In this chapter, we shall develop a continuous Bayesian model and study a portfolio optimization problem based on the model.

6.1 Conjugate Distributions

In Bayesian statistics, we often want to infer the distribution of a parameter θ based on observed data x . According to Bayes' formula[8]

$$p(\theta|x) = \frac{p(x|\theta)p(\theta)}{p(x)}, \quad (78)$$

the posterior distribution $p(\theta|x)$ is determined by the product of prior distribution $p(\theta)$ and likelihood function $p(x|\theta)$. If the posterior distributions $p(\theta|x)$ are in the same probability distribution family as the prior probability distribution $p(\theta)$, the prior and posterior are called conjugate distributions, and the prior is called a conjugate prior for the likelihood function.

Let X be a normal random variable with distribution $X \sim \mathcal{N}(\mu, \sigma^2)$, where the variance σ^2 is a known constant but the mean μ is not known for sure. Assume μ is also a normal random variable with (prior) distribution $\mu \sim \mathcal{N}(\mu_0, \sigma_0^2)$. In the Bayesian framework, we can update the (posterior) distribution of μ from independent observations of X . If there are $n \in \mathbb{N}$ independent observations $D = (x_1, \dots, x_n)$ of X , then the likelihood function is

$$p(D|\mu) = \prod_{i=1}^n p(x_i|\mu) = (2\pi\sigma^2)^{-\frac{n}{2}} e^{-\frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2}. \quad (79)$$

According to Bayes's formula, the posterior probability density of μ is a constant multiple of

$$e^{-\frac{1}{2\sigma_0^2}(\mu - \mu_0)^2 - \frac{1}{2\sigma^2} \sum_{i=1}^n (x_i - \mu)^2}. \quad (80)$$

Notice the posterior distribution of μ is also normal with $\mu|D \sim \mathcal{N}(\mu_p, \sigma_p^2)$, where $\frac{1}{\sigma_p^2} = \frac{1}{\sigma_0^2} + \frac{n}{\sigma^2}$, and $\mu_p = \sigma_p^2 \left(\frac{\mu_0}{\sigma_0^2} + \frac{\sum_{i=1}^n x_i}{\sigma^2} \right)$. As we can see, the prior and posterior distribution of μ are both Gaussian, hence conjugated distributions. With the posterior distribution of μ computed, we can further update the posterior predictive distribution of X by

$$\begin{aligned} p(x|D) &= \int p(x|\mu)p(\mu|D)d\mu \\ &= \int (2\pi\sigma^2)^{-\frac{1}{2}} e^{-\frac{1}{2\sigma^2}(x-\mu)^2} (2\pi\sigma_p^2)^{-\frac{1}{2}} e^{-\frac{1}{2\sigma_p^2}(x-\mu_p)^2} d\mu. \end{aligned} \quad (81)$$

A simple computation shows the posterior predictive distribution of X is again normal, i.e. $X|D \sim \mathcal{N}(\mu_p, \sigma_p^2 + \sigma^2)$.

6.2 Stochastic Market Model

In this section, we will extend the idea of conjugate distributions to a continuous stochastic process with unknown drift. Consider a stochastic process X_t evolves according to the stochastic differential equation (SDE)

$$dX_t = \mu dt + \sigma dW_t, \quad (82)$$

where W_t is a standard one dimensional Brownian motion, σ is a positive constant, and μ is a constant but its value is not known for sure. Based on information available at time $t = 0$, our belief about μ is a normal distribution with mean μ_0 and standard deviation σ_0 , i.e. $\mu \sim \mathcal{N}(\mu_0, \sigma_0^2)$. This distribution represents our belief about μ without observing X_t . Later when we observe values of X_t , the distribution of μ can be updated using the conjugate technique described earlier.

For example, if we have observed X_t on the time interval $[0, s]$, we can apply the following approach to update the distribution of μ . First for a positive integer N , partition $[0, s]$ into N sub-interval with $\{s_i = i\Delta s | 0 \leq i \leq N, \Delta s = \frac{s}{N}\}$. Then define random variables $\Delta X_i = X_{s_i} - X_{s_{i-1}}$ for $1 \leq i \leq N$. Due to the stationary independent increment property of Brownian motion, we know $\{\Delta X_i\}_{1 \leq i \leq N}$ are independent and identically distributed with a common normal distribution $\mathcal{N}(\mu\Delta s, \sigma^2\Delta s)$. By our prior knowledge of μ , the random variable $\mu\Delta s$ also has a (prior) normal distribution $\mathcal{N}(\mu_0\Delta s, \sigma_0^2\Delta s^2)$. Now we can repeat the conjugate distributions procedure to update the posterior distribution of $\mu\Delta s$ (hence μ via scaling) given $\{\Delta X_i\}_{1 \leq i \leq N}$, that is

$$\mu\Delta s | \{\Delta X_i\}_{1 \leq i \leq N} \sim \mathcal{N}\left(\frac{\Delta s \mu_0 \sigma^2 + \Delta s \sigma_0^2 X_s}{(\sigma^2 + s\sigma_0^2)}, \frac{\sigma^2 \sigma_0^2 (\Delta s)^2}{\sigma^2 + s\sigma_0^2}\right), \quad (83)$$

where we have used the identity $\sum_{i=1}^N \Delta X_i = X_s$. Similarly, the posterior predictive distribution of the increment $\Delta X = X_{s+\Delta s} - X_s$, becomes

$$\Delta X | \{\Delta X_i\}_{1 \leq i \leq N} \sim \mathcal{N}\left(\frac{\Delta s \mu_0 \sigma^2 + \Delta s \sigma_0^2 X_s}{(\sigma^2 + s\sigma_0^2)}, \frac{\sigma^2 \sigma_0^2 (\Delta s)^2}{\sigma^2 + s\sigma_0^2} + \sigma^2 \Delta s\right). \quad (84)$$

Finally let $N \rightarrow \infty$ and normalize the posterior predictive distribution over Δs , we obtain a normal distribution with following mean and variance

$$\lim_{\Delta s \rightarrow 0} \frac{\frac{\Delta s \mu_0 \sigma^2 + \Delta s \sigma_0^2 X_t}{(\sigma^2 + s\sigma_0^2)}}{\Delta s} = \frac{\mu_0 \sigma^2 + \sigma_0^2 X_s}{(\sigma^2 + s\sigma_0^2)}, \quad (85)$$

and

$$\lim_{\Delta s \rightarrow 0} \frac{\frac{\sigma^2 \sigma_0^2 (\Delta s)^2}{\sigma^2 + s \sigma_0^2} + \sigma^2 \Delta s}{\Delta s} = \sigma^2. \quad (86)$$

Above equations imply that if we update the distribution of μ using observations up to time s , the posterior predictive distribution of the normalized increment of X_t starting at time s , i.e. $\frac{X_{s+\Delta s} - X_s}{\Delta s}$, is $\mathcal{N}\left(\frac{\mu_0 \sigma^2 + \sigma_0^2 X_s}{(\sigma^2 + s \sigma_0^2)}, \sigma^2\right)$. If we further draw an analogy between the discrete updating process and the incremental distribution of the stochastic process described by equation (82), we obtain the following SDE corresponding to the posterior predictive distribution

$$dX_t = \frac{\mu_0 \sigma^2 + \sigma_0^2 X_t}{(\sigma^2 + t \sigma_0^2)} dt + \sigma dW_t. \quad (87)$$

Notice equation (82) cannot be used directly to set a stochastic control problem since the value of μ is unknown. However, we were able to derive equation (87) in the Bayesian framework which avoids direct use of μ by combining the prior distribution of μ and observations of the stochastic process. Thus, it becomes possible to study stochastic control problems based on equation (87). It is easy to verify

$$X_t = \mu_0 t + (\sigma^2 + t \sigma_0^2) \int_0^t \frac{\sigma}{\sigma^2 + s \sigma_0^2} dW_s \quad (88)$$

assuming $X_0 = 0$, from which we see X_t is a Gaussian process with normal distribution $\mathcal{N}\left(\mu_0 t, \frac{\sigma^2}{\sigma_0^2} (\sigma^2 + t \sigma_0^2)^2 \left(\frac{1}{\sigma^2} - \frac{1}{\sigma^2 + t \sigma_0^2}\right)\right)$ at any time t .

Remark. From equation (83), we can derive the posterior distribution of μ based on observations up to time t , denoting as μ_t ,

$$\mu_t \sim \mathcal{N}\left(\frac{\mu_0 \sigma^2 + \sigma_0^2 X_t}{(\sigma^2 + t \sigma_0^2)}, \frac{\sigma^2 \sigma_0^2}{\sigma^2 + t \sigma_0^2}\right). \quad (89)$$

In stochastic filtering theory, the conditional expectation of μ over the σ -algebra generated by observations of $\{X_s\}_{0 \leq x \leq t}$ has the same distribution as in equation (89). For a more rigorous derivation within the measure-theoretic framework, one can refer to Bain and Crisan[2].

6.3 Stochastic Control Problem

Consider a market where there is one risky asset, and one risk-free asset. The price process of the risky asset follows the stochastic process $Y_t = e^{X_t}$, where X_t is defined in

equation (87). By Itô's lemma, we can derive a stochastic differential equation for Y_t

$$dY_t = \left(f(t, Y_t) + \frac{1}{2}\sigma^2\right)Y_t dt + \sigma Y_t dW_t, \quad (90)$$

where

$$f(t, y) = \frac{\mu_0\sigma^2 + \sigma_0^2 \log y}{\sigma^2 + t\sigma_0^2}. \quad (91)$$

Assume the risk-free asset provides a constant (continuously compounding) rate of return r , which leads to the ordinary differential equation for its price process

$$dB_t = rB_t dt. \quad (92)$$

Based on market dynamics, an investor can continuously decide the portion of her money to put in the risky asset (the rest goes to the risk-free asset). Her objective is to maximize the expected utility of pay-off at some future time T . Denote the total value of her investment at time t as P_t , then P_t should satisfy the following stochastic differential equation if α_t portion of P_t were invested in the risky asset

$$dP_t = \alpha_t P_t \frac{dY_t}{Y_t} + (1 - \alpha_t) P_t \frac{dB_t}{B_t}, \quad (93)$$

which is equivalent to

$$dP_t = \left[\alpha_t \left(f(t, Y_t) + \frac{1}{2}\sigma^2 - r\right) + r\right] P_t dt + \alpha_t \sigma P_t dW_t. \quad (94)$$

We further set the investor's pay-off at a fixed future time T to be a constant relative risk aversion (CRRA) utility function of the portfolio value $P_T \in \mathbb{R}^+$

$$U(P_T) = \frac{P_T^{1-\gamma}}{1-\gamma}, \quad (95)$$

where $\gamma > 0$ is a constant that is not equal to 1 (in this thesis we will mainly focus on the case where $\gamma > 1$). For any $t \leq T$ the investor has the freedom to choose $\alpha_t \in \mathcal{A}$ so that the expected pay-off can be maximized, where \mathcal{A} is a pre-defined action set which describes possible actions the investor can take. For example, $\mathcal{A} = \mathbb{R}$ means the investor can take arbitrary long/short positions in both assets, $\mathcal{A} = \{0.5\}$ means the investment must split the portfolio evenly into two assets. Formally, define the value function $V(t, y, p)$ as the conditional expectation given the price of risky asset and portfolio value at time t

$$V(t, y, p) = \sup_{\alpha_t \in \mathcal{A}} \mathbb{E}[U(P_T) | Y_t = y, P_t = p], \quad (96)$$

where $0 \leq t \leq T$, $y, p \in \mathbb{R}^+$.

The standard approach to solve a stochastic control problem can be decomposed into several steps. The first step is to set up the Hamilton-Jacobi-Bellman (HJB) equation that the value function should satisfy. The second step is to obtain a solution to the HJB equation, e.g. by solving a system of ODE. The final step is to verify the solution to the HJB equation is indeed the value function of the stochastic control problem. Under some regularity conditions, the value function is a smooth solution to the HJB equation, for example Friedman[14], Gilbarg[15], and Krylov[23]. However, those conditions are not met in our problem. Yet we will see a smooth solution can be found when T satisfies some conditions.

The HJB equation of this stochastic control problem is

$$0 = V_t + \sup_{\alpha_t \in \mathcal{A}} \left\{ V_y \left(f + \frac{1}{2} \sigma^2 \right) y + V_p \left[\alpha_t \left(f + \frac{1}{2} \sigma^2 - r \right) + r \right] p + \frac{1}{2} V_{yy} \sigma^2 y^2 + \frac{1}{2} V_{pp} \alpha_t^2 \sigma^2 p^2 + V_{yp} \alpha_t \sigma^2 yp \right\}. \quad (97)$$

The process of deriving HJB equation in a stochastic control problem typically involves Itô's Lemma, and we include details in the Appendix. In the coming sections, we will look at three different cases where the action set \mathcal{A} is the set of all real numbers, contains a single real number, and is a closed interval. To make it notation-wise easier, we denote the expression inside the curly bracket in equation (97) as

$$G(\alpha_t) = V_y \left(f + \frac{1}{2} \sigma^2 \right) y + V_p \left[\alpha_t \left(f + \frac{1}{2} \sigma^2 - r \right) + r \right] p + \frac{1}{2} V_{yy} \sigma^2 y^2 + \frac{1}{2} V_{pp} \alpha_t^2 \sigma^2 p^2 + V_{yp} \alpha_t \sigma^2 yp, \quad (98)$$

emphasizing that we need to solve a maximization problem within equation (97). Notice that $G(\alpha_t)$ is a quadratic function of α_t and we can obtain a closed-form expression in some cases.

6.4 Stochastic Control Problem - Unconstrained Action

We first look at the case where the action set \mathcal{A} is the entire real line, i.e. $\mathcal{A} = \mathbb{R}$. Clearly the function G attains its maximum at

$$\alpha_t^* = \frac{V_p(r - f - \frac{1}{2}\sigma^2) - V_{yp}\sigma^2 y}{V_{pp}\sigma^2 p}. \quad (99)$$

Substitute α_t^* back to the HJB equation (97) yields a second order non-linear partial differential equation

$$0 = V_t + V_y(f + \frac{1}{2}\sigma^2)y + prV_p + \frac{1}{2}V_{yy}\sigma^2 y^2 - \frac{1}{2V_{pp}\sigma^2}[V_p(f + \frac{1}{2}\sigma^2 - r) + V_{yp}y\sigma^2]^2. \quad (100)$$

Notice the stochastic differential equation (94) is linear in the process P_t , meaning if P_t solve the equation so does any constant multiple of P_t . Together with the utility function (17) at time T , we can assume the value function V has the following form

$$V(t, y, p) = A(t, y)p^{1-\gamma}. \quad (101)$$

Substituting equation (101) into equation (100) gives us a non-linear partial differential equation which does not involve p

$$0 = A_t + (1 - \gamma)rA + A_y(f + \frac{1}{2}\sigma^2)y + \frac{1}{2}A_{yy}\sigma^2 y^2 + \frac{A(1 - \gamma)(f + \frac{\sigma^2}{2} - r)^2}{2\gamma\sigma^2} + \frac{y^2 A_y^2 (1 - \gamma)\sigma^2}{2A\gamma} + \frac{yA_y(f + \frac{\sigma^2}{2} - r)(1 - \gamma)}{\gamma}. \quad (102)$$

At this point, it is not obvious how to solve equation (102) due to its non-linearity. Instead, we will solve it numerically to gain some insights into it.

In our initial experiment, we would like to solve for $A(t, y)$ numerically in the domain $(t, y) \in [0, T] \times (0, C)$ where C is a positive constant. A natural terminal condition to be applied is

$$A(T, y) = \frac{1}{1 - \gamma} \quad (103)$$

according to the utility function in equation (95). To better analyze the behavior as the price of risky asset approaches to 0, we apply the change of variable formula $x = \log(y)$

and rewrite equation (102) accordingly as

$$0 = A_t + (1 - \gamma)rA + A_x(f + \frac{1}{2}\sigma^2) + \frac{1}{2}(A_{xx} - A_x)\sigma^2 + \frac{A(1 - \gamma)(f + \frac{\sigma^2}{2} - r)^2}{2\gamma\sigma^2} + \frac{A_x^2(1 - \gamma)\sigma^2}{2A\gamma} + \frac{A_x(f + \frac{\sigma^2}{2} - r)(1 - \gamma)}{\gamma}, \quad (104)$$

where A is understood as $A(t, e^x)$. We then numerically solve equation (104) using an explicit finite difference scheme. Since our focus here is not the numerical algorithm, we postpone the detail description of it to later sections. We set $T = 0.5$ where A is constant according to equation (103), and iteratively compute the value of A for $t < T$ and. Plots of A at some earlier times are given in Figure 6 (labeled with different colors). We observe that for a fixed x , A increases quickly as we stepped back in time. This can also be visualized by the fact the a curve corresponding to a smaller time is higher than a curve corresponding to a larger time. At a fixed point in time, represented by each single curve, A would also increases as x deviates from 0. It is worth mentioning that when we try to carry out the experiment for larger value of T , numerical blow-up occurs.




Figure 6: Numerical solution of A at some specific times

Our initial observations motivates us to reduce the scale of function A so that we can deal with moderate numerical values. To that purpose, we define a new function B

as

$$B(t, x) = \log A(t, e^x), \quad (105)$$

where $x = \log y$ as defined previously. This transformation leads to the a partial differential equation equivalent to equation (104) in terms of t , x , and $B(t, x)$

$$0 = B_t + (1 - \gamma)r + B_x \left(f + \frac{1}{2}\sigma^2 \right) + \frac{1}{2}(B_{xx} + B_x^2 - B_x)\sigma^2 + \frac{(1 - \gamma)(f + \frac{\sigma^2}{2} - r)^2}{2\gamma\sigma^2} + \frac{B_x^2(1 - \gamma)\sigma^2}{2\gamma} + \frac{B_x(f + \frac{\sigma^2}{2} - r)(1 - \gamma)}{\gamma}, \quad (106)$$

where $f = \frac{\mu_0\sigma^2 + \sigma_0^2 x}{\sigma^2 + t\sigma_0^2}$ after transforming from y to x , together with the terminal condition $B(T, x) = -\log(1 - \gamma)$. Equation (106) has two advantages over equation (104). The first is the range of $B(t, x)$ is significantly smaller compare to $A(t, y)$ in equation (104). The second is the only non-linear term in equation (106) is B_x^2 while $A(t, y)$ also appears in a denominator in equation (104).

A numerical study of equation (106) is conducted using the same algorithm which was used in the study of equation (104) (the details are postponed to later sections). Figure 7 displays plots of numerical solution at several point in time. As we can see, the range of $B(t, x)$ is significantly reduced and we can plot for a much wider range of x .




Figure 7: Numerical solution of $B(t, x)$ at some specific times

In addition to function $B(t, x)$, we also examined first order numerical derivative w.r.t. x , which are computed using upwind difference scheme. As shown in Figure 8 (and verified by second order numerical derivatives), the numerically value of $B_x(t, x)$ changes linearly w.r.t to x . This finding strongly suggests that if we fix the value of t , then $B(t, x)$ becomes a quadratic function of x . Thus, we are motivated to use a natural ansatz to solve equation (106) which is illustrated in the next theorem.




Figure 8: Numerical derivative of $B(t, x)$ w.r.t x at some specific times

Theorem 2. Let a , b , and c be functions that map \mathbb{R} to \mathbb{R} . Assume they solve the following system of ordinary differential equations on $[0, T]$ with boundary conditions $a(T) = 0$, $b(T) = 0$, and $c(T) = -\log(1 - \gamma)$. Then $B(t, x) = a(t)x^2 + b(t)x + c(t)$ solves

the partial differential equation (106) with terminal condition $B(T, x) = -\log(1 - \gamma)$.

$$\begin{cases} \dot{a} + \frac{2\sigma_0^2}{\gamma(\sigma^2+t\sigma_0^2)}a + \frac{2\sigma^2}{\gamma}a^2 + \frac{(1-\gamma)\sigma_0^4}{2\gamma\sigma^2(\sigma^2+t\sigma_0^2)^2} = 0, \\ \dot{b} + \left[\frac{\sigma_0^2}{\gamma(\sigma^2+t\sigma_0^2)} + \frac{2\sigma^2}{\gamma}a\right]b + 2a\left(\frac{\mu_0\sigma^2}{\gamma(\sigma^2+t\sigma_0^2)} + \frac{\sigma^2-2r}{2\gamma} + r - \frac{\sigma^2}{2}\right) + \\ \frac{(1-\gamma)\sigma_0^2}{\gamma\sigma^2(\sigma^2+t\sigma_0^2)}\left(\frac{\mu_0\sigma^2}{(\sigma^2+t\sigma_0^2)} - r + \frac{\sigma^2}{2}\right) = 0, \\ \dot{c} + (1-\gamma)r + \left(\frac{\mu_0\sigma^2}{\gamma(\sigma^2+t\sigma_0^2)} + \frac{0.5\sigma^2-r}{\gamma} + r - \frac{\sigma^2}{2}\right)b + \\ \sigma^2a + \frac{\sigma^2}{2\gamma}b^2 + \frac{1-\gamma}{2\gamma\sigma^2}\left(\frac{\mu_0\sigma^2}{(\sigma^2+t\sigma_0^2)} - r + \frac{\sigma^2}{2}\right)^2 = 0, \\ a(T) = 0, b(T) = 0, c(T) = -\log(1 - \gamma). \end{cases} \quad (107)$$

Proof. Substitute $B(t, x) = a(t)x^2 + b(t)x + c(t)$ into equation (106) to obtain a quadratic polynomial of x where the coefficients are functions of t . Set the coefficients to be zeros and match the degree-zero term to the terminal condition of the partial differential equation: $B(T, x) = -\log(1 - \gamma)$. We then obtain precisely the system of ordinary differential equations. \square

The system of equations (107) looks more complicated than it is. Notice only the first equation is a non-linear equation due to the term a^2 . Once the first equation is solved, the second and third equations can be solved sequentially as first-order linear ordinary differential equations. It turns out the first equation in (107) is a Riccati equation which can be solved via standard methods. Under a mild assumption, system (107) admits a closed-form solution.

Theorem 3. *The system of differential equations (107) admits bounded, continuously differentiable solutions a , b , and c on $[0, T]$ if $\sigma_0^2(1 - \gamma)T < \sigma^2\gamma$.*

Proof. Once the first equation in (107) is solved, the second and third equations follow naturally. For the first equation, we can solve it via a standard technique[18] that transform it into a linear ordinary differential equation. The solution is

$$a(t) = \frac{-\sigma_0^2}{2\sigma^2} \frac{1}{\sigma^2 + t\sigma_0^2} + \frac{1}{\frac{2\sigma^2}{\gamma}t + C_T}, \quad (108)$$

where C_T is a constant determined by the boundary condition at $t = T$. In this case, $a(T) = 0$ yields $C_T = \frac{1-\gamma}{4\gamma\sigma^4} \frac{\sigma_0^4}{(\sigma^2+T\sigma_0^2)^2}$. Thus, solution to the first equation is

$$a(t) = \frac{-\sigma_0^2}{2\sigma^2} \frac{1}{\sigma^2 + t\sigma_0^2} + \frac{1}{\frac{2\sigma^2}{\gamma}t + \frac{2\sigma^4}{\sigma_0^2} + 2\sigma^2T(1 - \frac{1}{\gamma})}. \quad (109)$$

While $\sigma^2 + t\sigma_0^2$ is positive for $t \geq 0$, the denominator of the second term on the right-hand-side of equation (109) can become 0 and leads to a blow-up. Denote $h(t) = \frac{2\sigma^2}{\gamma}t + \frac{2\sigma^4}{\sigma_0^2} + 2\sigma^2T(1 - \frac{1}{\gamma})$. Since $h(t)$ is monotone increasing in t , it assumes 0 on $[0, T]$ if and only if $h(0) \leq 0$, which happens exactly when $\sigma_0^2(1 - \gamma)T \geq \sigma^2\gamma$. Thus, as long as $\sigma_0^2(1 - \gamma)T < \sigma^2\gamma$, there exist a bounded continuously differentiable solution to a . Otherwise, a blows up to infinity in finite time. In particular, the inequality is always satisfied when $\gamma > 1$, and we only need to worry about this criteria if $\gamma < 1$. Solutions to the second and third follow naturally due to standard linear ODE theory. The boundedness is a simple consequence of the fact that a , b , and c are continuous functions on the closed interval $[0, T]$. \square

According to theorem 1 and theorem 2, we can derive a solution to the HJB equation, together with an expression for the optimal control α_t^* . We summarize the result in the following theorem.

Theorem 4. *There is a closed-form solution $V(t, y, p)$ to the HJB equation(97) when $\sigma_0^2(1 - \gamma)T < \sigma^2\gamma$,*

$$V(t, y, p) = p^{1-\gamma}e^{a(t)(\log y)^2 + b(t)\log y + c(t)}, \quad (110)$$

where $a(t)$, $b(t)$, and $c(t)$ are described in Theorem 2. Moreover, the optimal choice of α_t in the HJB equation is given by

$$\alpha_t^* = \left(\frac{2a(t)}{\gamma} + \frac{\sigma_0^2}{\gamma\sigma^2\sigma_0^2t + \gamma\sigma^4} \right) \log y - \left(\frac{r}{\gamma\sigma^2} - \frac{1}{2\gamma} - \frac{b(t)}{\gamma} - \frac{\mu_0}{\gamma\sigma^2 + \gamma\sigma_0^2t} \right). \quad (111)$$

In particular, α_t^* is a linear function of $\log y$ for any fixed t .

Proof. Notice $B(t, x)$ and $A(t, y)$ are the same function expressed in different scales. According to Theorem 2, $B(t, x) = a(t)x^2 + b(t)x + c(t)$ solves equation (106). Using the coordinate transforms $x = \log y$, $B(t, x) = \log A(t, e^x)$, and equation (101), we obtain the expression for $V(t, y, p)$. Then the expression for α_t^* follows from equation (99). \square

Now that we have figured out the solution to the HJB equation, what's left is to prove the solution is indeed the value function of the original stochastic control problem. In addition, we also want to verify α_t^* is the optimal control which produced the value

function. For the proof, we will use the HJB equation and adapt a classic Change-of-Numeraire method[35].

Theorem 5. *Among all Gaussian control processes with bounded mean and variance on the interval $[0, T]$ that are also jointly Gaussian with the log-price process X_t , and T satisfies the criteria in Theorem 3. Then $V(t, y, p)$ and α_t^* in Theorem 4 are the value function and the optimal control of the stochastic control problem, respectively.*

Proof. Consider at time $t \in [0, T]$, the price of risky asset is $y > 0$ and the portfolio value is $p > 0$. For $s \in [t, T]$, denote the process $Y_s^{t,y}$ the price process whose value is y at time t . Also, denote the process $P_s^{t,y,p,\alpha}$ the portfolio process where its value is p at time t , the price of risky asset is y at time t , and the control process α is applied on the interval $[t, T]$. Apply Itô's lemma to the process $V(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})$ and collect terms with ds and dW_s , we get

$$dV(s, Y_s^{t,y}, P_s^{t,y,p,\alpha}) = L_s ds + H_s dW_s, \quad (112)$$

where

$$\begin{aligned} L_s = & \frac{\partial V}{\partial t}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha}) + \frac{\partial V}{\partial y}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})(f + \frac{1}{2}\sigma^2)Y_s^{t,y} \\ & + \frac{\partial V}{\partial p}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\left[\alpha_s(f + \frac{1}{2}\sigma^2 - r) + r\right]P_s^{t,y,p,\alpha} + \frac{1}{2}\frac{\partial^2 V}{\partial y^2}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\sigma^2(Y_s^{t,y})^2 \\ & + \frac{1}{2}\frac{\partial^2 V}{\partial p^2}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\sigma^2\alpha_t^2(P_s^{t,y,p,\alpha})^2 + \frac{\partial^2 V}{\partial p \partial y}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\alpha_t\sigma^2 Y_s^{t,y} P_s^{t,y,p,\alpha}, \end{aligned}$$

and

$$H_s = \frac{\partial V}{\partial y}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\sigma Y_s^{t,y} + \frac{\partial V}{\partial p}(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\alpha_s\sigma P_s^{t,y,p,\alpha}.$$

Notice the terms in L_s are those in the original HJB equation (97) only with stochastic arguments and we always have $L_s \leq 0$ and the equality is achieved if $\alpha_s = \alpha_s^*$.

Next we perform a Change-of-Numeraire using H_s and then apply Girsanov theorem[20]. Using the expression of $V(t, y, p)$ given in Theorem 4, we can deduce

$$H_s = V(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})\left(2a(s)\log Y_s^{t,y} + b(s) + (1 - \gamma)\alpha_s\right)\sigma. \quad (113)$$

Since the logarithm of $Y_s^{t,y}$ and α_s are jointly Gaussian with bounded mean and variance on $[0, T]$, so is their sum. Set

$$J_s = \left(2a(s) \log Y_s^{t,y} + b(s) + (1 - \gamma)\alpha_s \right) \sigma. \quad (114)$$

By Fernique's theorem[10], there is $\epsilon > 0$ such that

$$\mathbb{E} \left[\exp \int_u^{(u+\epsilon) \wedge T} \frac{1}{2} J_\tau^2 d\tau \right] < +\infty \quad (115)$$

for any $u \in [0, T]$. By dividing $[0, T]$ into finite many sub-intervals with length $\frac{\epsilon}{2}$, the Novikov condition is satisfied according to Corollary 3.5.14 in Karatzas and Shreve[20]. Consequently, the Doleans exponential of J_s is a martingale for $s \in [t, T]$

$$\eta_s^t = \exp \int_t^s J_\tau dW_\tau - \frac{1}{2} \int_t^s J_\tau^2 d\tau. \quad (116)$$

By applying Itô's lemma to $\frac{V(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})}{\eta_s^t}$ (use η_s^t as the Numeraire) we get

$$d \left(\frac{V(s, Y_s^{t,y}, P_s^{t,y,p,\alpha})}{\eta_s^t} \right) = \frac{L_s}{\eta_s^t} ds, \quad (117)$$

which is non-positive and is (always) zero when α_s^* is applied. By comparing $s = t$ and $s = T$ we see

$$V(t, y, p) \eta_T^t \geq V(T, Y_T^{t,y}, P_T^{t,y,p,\alpha}). \quad (118)$$

Now take expectation on both side, using the fact that $V(t, y, p)$ is non-random, η_T^t is a martingale, and $V(T, y, p) = \frac{p^{1-\gamma}}{1-\gamma}$ we see

$$V(t, y, p) \geq \mathbb{E} \left[\frac{(P_T^{t,y,p,\alpha})^{1-\gamma}}{1-\gamma} \right], \quad (119)$$

where the right hand side is the expected utility function at T .

Notice when α^* is applied on $[t, T]$, L_s will always be zero and so is the right-hand-side of equation (117). In this case, equality is achieved in equation (118). Finally note α_s^* is a Gaussian process with bounded mean and variance on $[0, T]$ thanks to the fact $X_s^{t,y} = \log(Y_s^{t,y})$ and equation (88). As the deterministic functions in equation (114) are bounded and continuous, the proof is concluded. \square

From Theorem 4 and Theorem 5 we see the value function V exist on $[0, T]$ for a suitable value T . If T is too large, V will reach infinity in finite time due to the blow-up of $a(t)$. For a fix time t , V is unbounded and grows quickly as the log-price deviates from zero. An intuitive explanation consists of two parts. The first being the drift of price process is determined by current price through equation (87) in the Bayesian framework. That is, if the price increases then we get a larger drift, and the price is likely to get even larger due to the increased drift. The second is an investor is allowed to take arbitrary long/short positions in both assets. We will see that the value function will behave differently if we restrict the investment action one can take.

6.5 Stochastic Control Problem - Single Action

In this section, we consider a special stochastic control problem with the same market dynamics described in equations (90) through (94), and the value function defined in equations (95) and (96). The only difference is that we restrict the the investment action to be a single point, i.e. $\mathcal{A} = \{\alpha\}$ for a fixed $\alpha \in \mathbb{R}$. Obviously, the HJB equation becomes simpler as we do not need to maximize a quadratic function within the equation. Yet this problem provides useful insight regarding the boundary behavior when the price is very large or close to zero. In later section we will see how it helps us specify the correct boundary conditions in numerical simulations.

Since there is only one choice of $\alpha_t \equiv \alpha$, the HJB equation (97) is reduced to

$$0 = V_t + V_y(f + \frac{1}{2}\sigma^2)y + V_p[\alpha(f + \frac{1}{2}\sigma^2 - r) + r]p + \frac{1}{2}V_{yy}\sigma^2y^2 + \frac{1}{2}V_{pp}\alpha^2\sigma^2p^2 + V_{yp}\alpha\sigma^2yp. \quad (120)$$

We assume the value function has the same form as defined in equation (101) (for the same reason we used to write out equation (101))

$$V(t, y, p) = A(t, y)p^{1-\gamma}. \quad (121)$$

After applying transformations $x = \log y$ and $B(t, x) = \log A(t, e^x)$, we obtain an equivalent formulation of equation (120) in terms of x and $B(t, x)$

$$0 = B_t + B_x(f + \frac{1}{2}\sigma^2) + \frac{1}{2}\sigma^2[B_{xx} + (B_x)^2 - B_x] + (1 - \gamma)r + \frac{1}{2}(1 - \gamma)(-\gamma)\sigma^2\alpha^2 + B_x(1 - \gamma)\sigma^2\alpha + (1 - \gamma)(f + \frac{1}{2}\sigma^2 - r)\alpha. \quad (122)$$

As in the unconstrained control problem, we conduct a numerical experiment to get some idea about how does the solution behave. The experiment is based on equation (122) and explicit finite difference scheme. We fix $\alpha = 0.5$ in this experiment. The underlying region is $(t, x) \in [0, 0.5] \times [-100, 100]$. The terminal condition applied at $t = 0.5$ is $B(0.5, x) = -\log(1 - \gamma)$. Plots of numerically calculated $B(t, x)$ can be found in Figure 9. It suggests that for a fixed t the function $B(t, x)$ is a linear function of x .




Figure 9: Numerical solution of $B(t, x)$ at some specific times ($\alpha = 0.5$)

It turns out the numerical study again provides us with useful insight about the value function in the singleton control problem. We summarize the finding in the following theorem.

Theorem 6. *Let b , and c be functions that map \mathbb{R} to \mathbb{R} . Assume they solve the following system of ordinary differential equations on $[0, T]$ with boundary conditions $b(T) = 0$, and $c(T) = -\log(1 - \gamma)$. Then $B(t, x) = b(t)x + c(t)$ solves the partial differential*

equation (122) with terminal condition $B(T, x) = -\log(1 - \gamma)$.

$$\begin{cases} \dot{b} + \frac{\sigma_0^2}{\sigma^2 + \sigma_0^2 t} b + (1 - \gamma) \alpha \frac{\sigma_0^2}{\sigma^2 + \sigma_0^2 t} = 0, \\ \dot{c} + b \left[\frac{\mu_0 \sigma^2}{\sigma^2 + \sigma_0^2 t} + (1 - \gamma) \sigma^2 \alpha \right] + \frac{1}{2} \sigma^2 b^2 + (1 - \gamma) r + \\ \frac{1}{2} (\gamma - 1) \gamma \sigma^2 \alpha^2 + (1 - \gamma) \left(\frac{\mu_0 \sigma^2}{\sigma^2 + \sigma_0^2 t} + \frac{1}{2} \sigma^2 - r \right) \alpha = 0, \\ b(T) = 0, c(T) = -\log(1 - \gamma). \end{cases} \quad (123)$$

Proof. Substitute $B(t, x) = b(t)x + c(t)$ into equation (122) to obtain a linear function of x where the coefficients are functions of t . Set the coefficients to be zeros and match the degree-zero term to the terminal condition of the partial differential equation: $B(T, x) = -\log(1 - \gamma)$. We then obtain the system of ordinary differential equations. \square

Notice the first equation in equation (123) is a first order linear ordinary differential equation of b , which can be easily solved. Once we obtain b , c can be solved by integration according the second equation. Thus, the existence of solution to the system is established.

Theorem 7. *The system of differential equations (123) admits bounded, continuously differentiable solutions b , and c on $[0, T]$.*

Proof. The existence and uniqueness of solution is guaranteed by the Picard–Lindelöf Theorem[22]. Being continuous functions on an closed interval, the solutions must be bounded as well. \square

As a result of Theorem 6 and Theorem 7, we confirm that the function $B(t, x)$ is indeed a linear function in x for a fixed time t . In addition, we are able to obtain a closed-form solution to the HJB equation and the solution is the value function in the singleton control problem.

Theorem 8. *For any $T \geq 0$, there is a closed-form solution $V(t, y, p)$ to the HJB equation(36)*

$$V(t, y, p) = p^{1-\gamma} e^{b(t) \log y + c(t)}, \quad (124)$$

where $b(t)$ and $c(t)$ are as described in theorem 5. In addition, $V(t, y, p)$ is the value function in the singleton control problem with $\mathcal{A} = \{\alpha\}$.

Proof. We can follow the same verification argument of Theorem 5, or a direct application of Feynman-Kac formula[33] will lead to the conclusion. \square

We conclude this section by discussing how the value function behave and compare it to the value function in the unconstrained control problem. To make the discussion easier, we factor out $p^{1-\gamma}$ and work in the $(t, x, B(t, x))$ system.

The first thing we can do is to go back and check the accuracy of the numerical solution illustrated in Figure 9. It turns out the numerical solution provides good approximations to the close-form solution. In particular, for a fixed time t , $B(t, x)$ is linearly increasing with the log-price of the risky asset. This can be partly explained by the rule that an investor must always puts a fixed portion of wealth in the risky asset. Since the drift of risky asset is positively related to its price, the expected value drops/increases when the price drops/increases. In comparison, for the unconstrained control problem, investors can expect large value function no matter high or low the observed prices are, since they are allowed to take arbitrary positions in both assets.

The second concerns the limit behavior of $B(t, x)$ when x approaches $+\infty$ or $-\infty$. From Theorem 2, we see function $B(t, x)$ grows quadratically in x for any fixed time t in the unconstrained control problem (Figure 7). In addition, the optimal control α_t^* is a linear function of x for a fixed t , describing how an investor adjusts the position in risky asset according to the underlying price. In the singleton control problem, an invest always put a fixed portion of wealth in the risky asset and the function $B(t, x)$ is a linear function of x . Of interest is a the case in-between: a control problem where an investor has some but limited flexibility. For example, we can set the action set \mathcal{A} to be a closed interval $[l, h]$. Intuitively, an investor can adjust the optimal decision α_t^* within the interval based on X_t , the log-price of risky asset. For a large positive X_t , α_t^* will be close to or even equal to h since the risky asset is expected to give a good return. For a large negative X_t , α_t^* is likely to be close or equal to l for a similar reason. Once h (or l) has been reached, even one observes larger (or smaller) values of X_t , no further actions are allowed to be taken due to the constraint $\mathcal{A} = [l, h]$. Thus, for extreme values of x , the investor is essentially taking a constant position (l or h) in the risky asset. Based on our study of the singleton control problem. We can expect the value functions to be

asymptotically linear in x for a fixed time t .

Remark. *In the case where \mathcal{A} is a set of finite cardinality or closed interval, we do not expect to be able to derive a closed-form solution. Instead, we will rely on numerical study to obtain a numerically smooth solution. From this perspective, previous discussion helps us determine the proper boundary conditions.*

6.6 Stochastic Control Problem - Finite Action Set

Based on the discussion on the single action problem, we can make a generalization by allowing the action set \mathcal{A} to contain finite number of values. The main purpose is to build a connection to the constrained action problems (section 6.7) where the action set is a closed interval. Intuitively, when we approximate an closed interval with a (fine) finite set, the resulting constrained stochastic control problem should have similar solution to its finite action counterpart. We will see that is indeed the case based on result of numerical study, but in this section we focus on the finite action problem.

When the action set is finite, an investor can choose how much wealth to allocate in the risky asset but the choices are limited and discrete (compare to the unconstrained problem). For example, if we set $\mathcal{A} = \{0, 1\}$, then an investor can choose to invest only in bond or only in stock. Of course \mathcal{A} can contain values greater than 1 or smaller than 0, representing potential short positions in two assets. As we shall see, the solutions are not fundamentally different from the single action case.

Assume the action set contains k distinct real numbers, i.e. $\mathcal{A} = \{a_i | 1 \leq i \leq k, a_i \in \mathbb{R}\}$. Since the actions set is finite, the HJB equation (97) is equivalent to

$$0 = V_t + \max_{\alpha_t \in \mathcal{A}} \left\{ V_y \left(f + \frac{1}{2} \sigma^2 \right) y + V_p \left[\alpha_t \left(f + \frac{1}{2} \sigma^2 - r \right) + r \right] p + \frac{1}{2} V_{yy} \sigma^2 y^2 + \frac{1}{2} V_{pp} \alpha_t^2 \sigma^2 p^2 + V_{yp} \alpha_t \sigma^2 yp \right\}. \quad (125)$$

On the right hand side of equation (125), there is a discrete optimization problem over a finite set. We do not expect to obtain a closed-form solution to the the optimal choice

α_t . However, in numerical study it can be solved by simply comparing among all possible values. To reduce the dimension and scale of the problem, apply transformations $V(t, y, p) = A(t, y)p^{1-\gamma}$, $x = \log y$, and $B(t, x) = \log A(t, e^x)$, we reach the following equation

$$0 = B_t + B_x(f + \frac{1}{2}\sigma^2) + \frac{1}{2}\sigma^2[B_{xx} + (B_x)^2 - B_x] + (1 - \gamma)r + (1 - \gamma)\max_{\alpha_t \in \mathcal{A}} \left\{ -\frac{1}{2}\gamma\sigma^2\alpha_t^2 + (B_x\sigma^2 + f + \frac{1}{2}\sigma^2 - r)\alpha_t \right\}. \quad (126)$$

The optimal choice α_t^* is determined by comparison among all possible values:

$$\alpha_t^* = \operatorname{argmax}_{\alpha \in \mathcal{A}} \left\{ -\frac{1}{2}\gamma\sigma^2\alpha^2 + (B_x\sigma^2 + f + \frac{1}{2}\sigma^2 - r)\alpha \right\}. \quad (127)$$

Equation (127) is useful when it come to solve equation (126) numerically since it reduces an optimization problem into finding the maximal element in an array, which can be integrated into classic finite difference schemes.

Although the nature of the discrete optimization problem makes it difficult to find a closed-form solution to equation (126), we can still pursue numerical solutions thanks to the finiteness of the action set \mathcal{A} . The idea is to apply simple explicit finite difference methods while solving optimization problems via equation (127) at each grid point. Before proceeding to numerical study, we discuss the terminal and boundary conditions. It is obvious that at time T , $B(T, x) = -\log(1 - \gamma)$. The boundary conditions as $x \rightarrow \pm\infty$ deserve some discussion. As seen in the single action problem, $B(t, x)$ is a linear function of x if t is fixed. In addition, it is intuitive to allocation more (less) in the risky asset if x is large (small) due to the market dynamics. Thus, in finite action problems, we expected the optimal control to be constant for extreme values of x , and the function $B(t, x)$ is asymptotically linear in x in the limit (which agrees with the result obtained without boundary conditions).

We outline the numerical algorithm before presenting result of numerical studies. Start by evenly discretizing the time domain $[0, T]$ and price domain $[-X, X]$ into $M + 1$ and $2N + 1$ grid points, respectively. Then we assign terminal condition at time T by setting $B(T, x) = -\log(1 - \gamma)$. After that, an iteration is carried out to compute

$B(t, x)$, using explicit finite difference scheme based on equation (126). In addition, the optimal control α_t^* at each grid point is computed via equation (127). For grid points corresponding to $x = X$ or $x = -X$, a linear boundary condition is applied to determine $B(t, x)$. The complete process is summarized in the algorithm below.

Algorithm 2: Finite Difference Scheme for Finite Action Problems

Input: $T, X, M, N, \mathcal{A} = \{0, 1\}, \gamma = 0.5, \mu_0 = 0.03, r = 0.03, \sigma = 0.3$, and

$$\sigma_0 = 0.03.$$

1 Discretize $[0, T]$ and $[-X, X]$ into M and $2N$ sub-intervals.

(i) Denote $t_i = iT/M$ for $0 \leq i \leq M$.

(ii) Denote $x_j = (j - N)X/N$ for $0 \leq j \leq 2N$.

2 Assign terminal condition:

$$B(t_M, x_j) = -\log(1 - \gamma) \text{ for } 0 \leq j \leq 2N.$$

3 For each j between $M - 1$ and 0:

(i) For each i between 1 and $2N - 1$:

(i-a) Compute $B_x(t_j, x_i)$ and $B_{xx}(t_j, x_i)$ using $B(t_{j+1}, x_{i-1}), B(t_{j+1}, x_i)$, and

$$B(t_{j+1}, x_{i+1}) \text{ via finite difference.}$$

(i-b) Compute α_t^* via equation (127).

(i-c) Compute $B(t_j, x_i)$ via equation (126).



(ii) Apply linear boundary conditions:

$$B(t_j, x_0) = 2B(t_j, x_1) - B(t_j, x_2), \text{ and}$$

$$B(t_j, x_{2N}) = 2B(t_j, x_{2N-1}) - B(t_j, x_{2N-2}).$$

We present some case studies of the finite action control problem. For the purpose of comparison, we restrict numerical studies in $(t, x) \in [0, 2] \times [-1, 1]$. We start with the simplest case where $\mathcal{A} = \{0, 1\}$, meaning an investor can invest only in the risky asset or only in the risk-free asset at a given time. Plots of numerical solutions to $B(t, x)$ at different times are shown in Figure 10, together with the optimal actions. In Figure 10(a) we see plots of $B(t, x)$ tend to be flat for small (negative) values of x , and the corresponding optimal controls are $\alpha_t^* = 0$. This phenomenon agrees with earlier discussion which suggests that when the price of risky asset is small it tends to be even smaller in the near future. Consequently, it is safer to invest all the wealth in bond rather than stock. In addition, $B(t, x)$ behaves like its linear counterpart in the single action problem where an investor invest in the bond all the time. On the other hand, as

x increases, $B(t, x)$ becomes upward-sloping and the optimal controls switch to 1, which agrees with our discussion for exactly the same reason. We also observe that $B(t, x)$ tend to be linear when x approaches to 1.

(a) Numerical solution of $B(t, x)$ 

(b) Optimal control

Figure 10: Numerical solution of $B(t, x)$ and optimal control

We can examine the asymptotic behavior of $B(t, x)$ in more detail by comparing plots of numerical solution with the plots of analytic solutions where $\mathcal{A} = \{0\}$ and $\mathcal{A} = \{1\}$. In Figure 11, color plots represent the numerical solution we obtained from Algorithm 2, where $\mathcal{A} = \{0, 1\}$. The smaller t is, the higher the curve corresponding to t lies in the figure. The dashed black lines represent analytical solutions obtained following Theorem 6. More specifically, flat dashed black lines represent analytic solutions where $\mathcal{A} = \{0\}$, and only left halves are shown to make the comparison more clear. Dashed black lines with positive slopes represent analytic solutions where $\mathcal{A} = \{1\}$, and only right halves are displayed. We can clearly see in Figure 11 that the numerical solutions almost coincide with analytic solutions when x becomes very large or small. In addition, as x increase from -1 to 1 , numerical solutions change gradually from horizontal lines on the left to upward-sloping lines on the right.


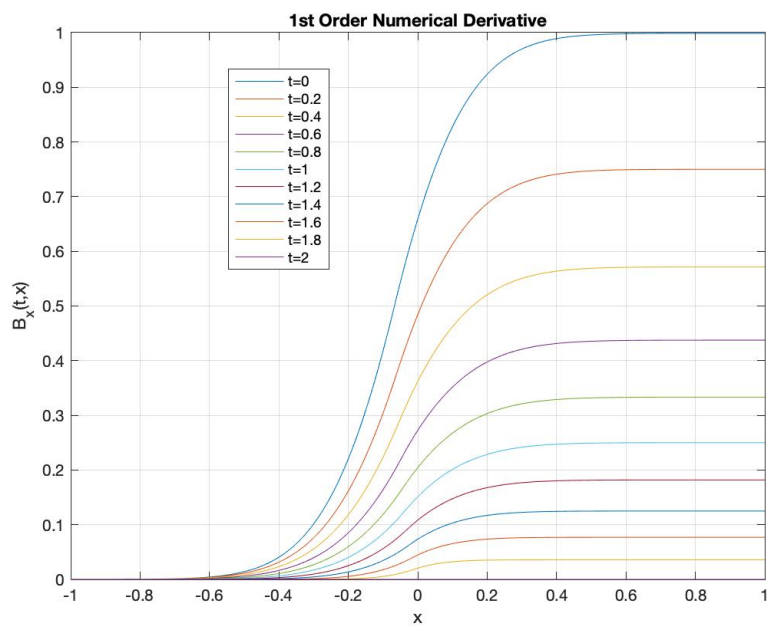



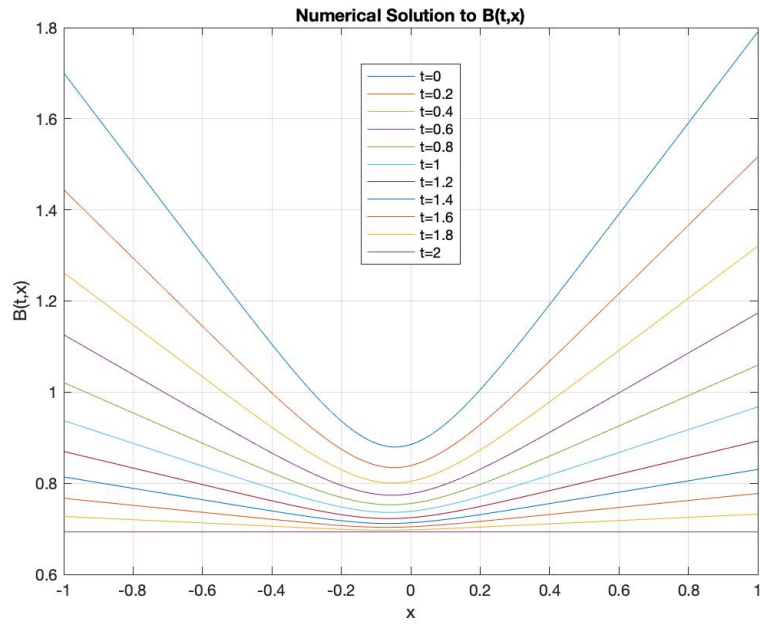
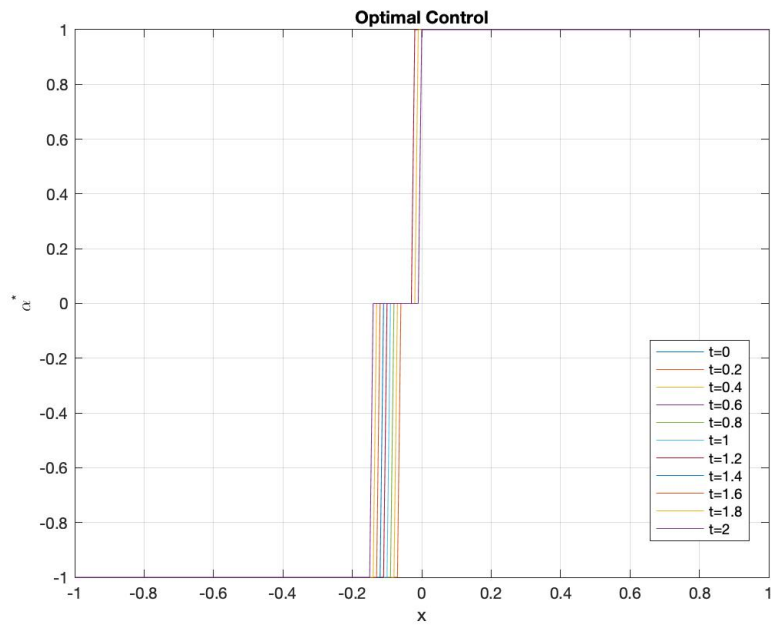
Figure 11: Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{0\}$ and $\mathcal{A} = \{1\}$)

In addition to asymptotic properties, we are also interested in the smoothness of the numerical solutions. Based on numerical solution to $B(t, x)$, we can compute its first-order and second-order derivatives w.r.t x . In our studies, numerical derivatives are computed using central finite difference schemes. As for boundary points, the derivatives are implied by linear boundary condition. Figure 12(a) shows plots of first-order numerical derivatives at different times. It suggests that $B_x(t, x)$ gradually increases from 0 which represents all-bond investment, to different positive levels that represent all-stock investment. The further a moment is from terminal time, the higher the level on the right-hand-side is. This agrees with earlier result in the section of single action problems. Figure 12(b) shows plots of second-order numerical derivatives at corresponding times. We see that the slope of $B(t, x)$ is always increasing but the rate it changes is not a constant. Starting from $x = -1$, the first-order derivative increases at a growing rate until a point near zero. After that the rate it increases keeps declining. In summary, our analysis shows how the numerical solutions to $B(t, x)$ are connected to the solutions to the single action problems, and suggests they are numerically twice continuously differentiable w.r.t x , which is a property implicitly assumed by the HJB equation (97).

(a) 1st order derivative $B_x(t, x)$ (b) 2nd order derivative $B_{xx}(t, x)$ Figure 12: Numerical derivatives of $B(t, x)$

Next we add some flexibility to the previous finite action problem by setting $\mathcal{A} = \{-1, 0, 1\}$. In this case, an investor can choose to take a short position in the risky asset with $\alpha_t = -1$. Consequently, the expected return can be substantial even for smaller (negative) values of x by simply shorting the risky asset and invest proceedings in the risk-free asset. Notice with this modification, equation (126), equation (127), and the numerical procedure outlined in Algorithm 2 remain valid. In fact, the only difference in numerical study is that the computer needs to make extra comparisons due to the increased size of \mathcal{A} .

We present the result of numerical studies with the same setup except for the admissible action set \mathcal{A} . First we examine the numerical solutions to $B(t, x)$ and corresponding optimal actions. Numerical solutions to $B(t, x)$ at different times are plotted in Figure 13(a). The right half of Figure 13(a) is similar to that of Figure 8(a), but the left part is quite different. Since an investor can choose to short the stock when x is small, the left halves of plots are no longer flat. We can get a better idea about this from Figure 13(b) which shows optimal actions one would take at different times w.r.t. different x . For small or large x , e.g. $x < -0.2$ or $x > 0.1$, an investor would take a -100% short position or a 100% long position in the stock. Note for a small interval in between -0.2 and 0.1 , one would put everything in the bond, i.e. choosing $\alpha_t^* = 0$. The transition is by no means smooth since the admissible action set itself is not continuous.

(a) Numerical solution of $B(t, x)$ 

(b) Optimal control

Figure 13: Numerical solution of $B(t, x)$ and optimal control

Next we compare the numerical solution to analytic solutions of single action problems where the control is set to be $\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$, respectively. As the optimal control takes either the smallest or largest possible value in \mathcal{A} when x is small enough or large enough, we expect the numerical solution to $B(t, x)$ to be asymptotically similar to solutions of those single action control problems. We can visualize that effect in Figure 14. Similar to Figure 11, numerically calculated solutions are plotted in solid color lines and analytic solutions are plotted in dashed black lines. For large or small x , curves with color are not distinguishable with dashed line segments which represent analytical solutions where $\mathcal{A} = \{-1\}$ or $\mathcal{A} = \{1\}$, respectively.



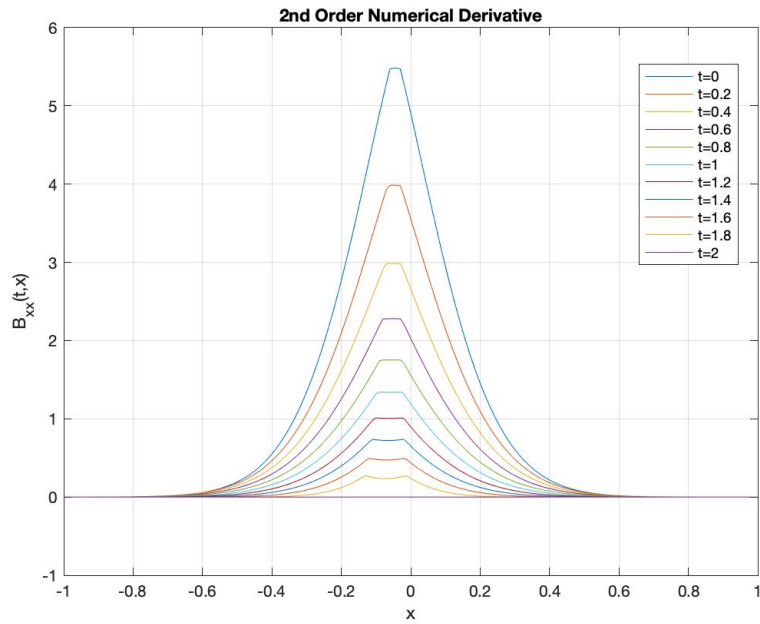



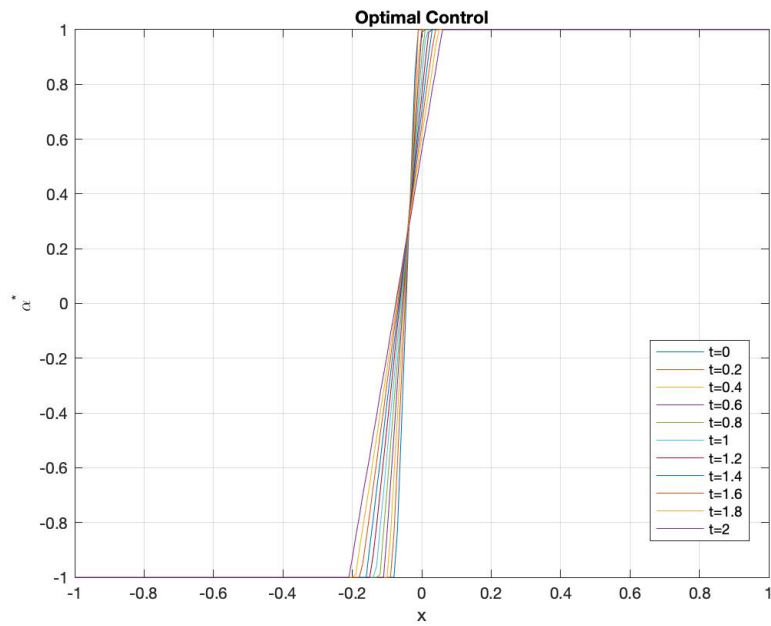
Figure 14: Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$)

Numerical derivatives w.r.t x are displayed in Figure 15. Figure 15(a) shows how do the slopes of $B(t, x)$ gradually change from lower levels corresponding to single action problem with $\mathcal{A} = \{-1\}$ to higher levels corresponding to single action problem with

$\mathcal{A} = \{1\}$. Figure 15(b) shows the corresponding second order derivatives. Notice there are relatively flat regions on top of curves(comparing to those in Figure 12(b)). Overall, the transitional pattern w.r.t x is similar to what we have seen in Figure 7. The result again suggests the numerical solution is well-behaved w.r.t. x .

(a) 1st order derivative $B_x(t, x)$ (b) 2nd order derivative $B_{xx}(t, x)$ Figure 15: Numerical derivatives of $B(t, x)$

We conclude this section with a case where the admissible control set is further refined. Define $\mathcal{A} = \{0.01i - 1 | 0 \leq i \leq 200, i \in \mathbb{N}\}$ and it's a superset of $\{-1, 0, 1\}$ we used in the previous case study. We want to check what happens when a much "larger" control set is put into use. In addition, the refined set can be considered as an discrete approximation to the closed interval $[-1, 1]$. We hope to connect it with the constrained stochastic control problems in the next section. Without reiterating the numerical process, we present results in Figure 16, 17, and 18, respectively. From Figure 16(a) and Figure 17, we observe that the numerical solution to $B(t, x)$ is very similar to what displayed in Figure 13(a) and Figure 14. However, optimal control plotted in Figure 16(b) is distinct from its counterpart in Figure 13(b). As an investor now has many more choices with the refined control set, the optimal control changes gradually from -1 to 1 in a linear fashion. The derivative plots in Figure 18 are also similar to those in Figure 15. The only visible difference is perhaps in the second order numerical derivatives. Notice in Figure 18(b), the plots tend to be smoother "bell-shaped" curves, and there is no flat regions on top as in Figure 15.

(a) Numerical solution of $B(t, x)$ 

(b) Optimal control

Figure 16: Numerical solution of $B(t, x)$ and optimal control



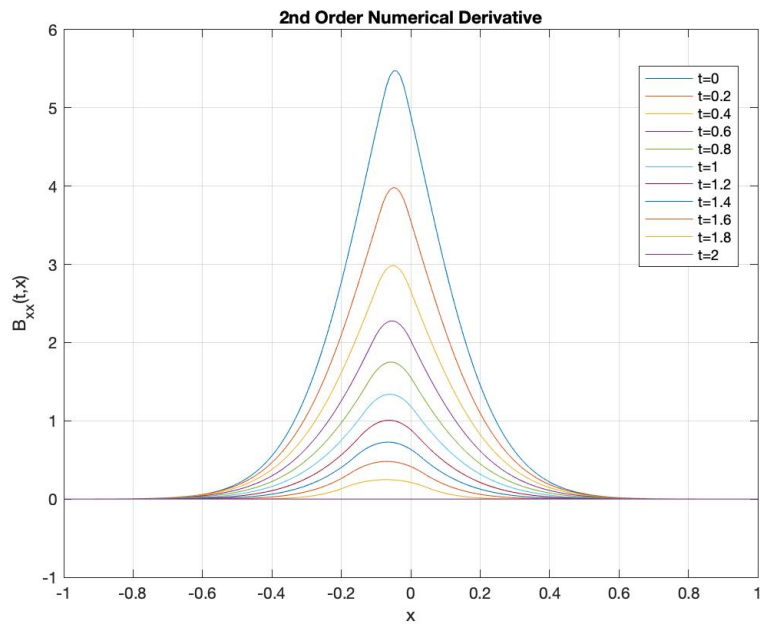


Figure 17: Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$)

(a) 1st order derivative $B_x(t, x)$ (b) 2nd order derivative $B_{xx}(t, x)$ Figure 18: Numerical derivatives of $B(t, x)$

6.7 Stochastic Control Problem - Constrained Action

In this section, we consider a stochastic control problem with the same market dynamics and utility function as defined in previous sections. The only difference is the control set becomes a closed interval, i.e. $\mathcal{A} = [m, n] \subset \mathbb{R}$ where m and n are constant real numbers. In particular, when $m < 0$ or $n > 1$, an investor can take a limited short position in the risky or risk-free asset.

Substitute $\mathcal{A} = [m, n]$ into equation (97) yields

$$0 = V_t + \sup_{\alpha_t \in [m, n]} \left\{ V_y \left(f + \frac{1}{2} \sigma^2 \right) y + V_p \left[\alpha_t \left(f + \frac{1}{2} \sigma^2 - r \right) + r \right] p \right. \\ \left. + \frac{1}{2} V_{yy} \sigma^2 y^2 + \frac{1}{2} V_{pp} \alpha_t^2 \sigma^2 p^2 + V_{yp} \alpha_t \sigma^2 yp \right\}. \quad (128)$$

Notice the second term on the right-hand-side involves a constrained quadratic optimization problem, and we do not expect to find a closed-form solution for $V(t, y, p)$. To reduce the number of variables and make it easier to carry out numerical experiments, we again apply transformations $V(t, y, p) = A(t, y)p^{1-\gamma}$, $x = \log y$, and $B(t, x) = \log A(t, e^x)$. Then an equivalent formulation in terms of x , t , and $B(t, x)$ is

$$0 = B_t + B_x \left(f + \frac{1}{2} \sigma^2 \right) + \frac{1}{2} \sigma^2 [B_{xx} + (B_x)^2 - B_x] + (1 - \gamma)r + \\ (1 - \gamma) \sup_{\alpha_t \in [m, n]} \left\{ -\frac{1}{2} \gamma \sigma^2 \alpha_t^2 + (B_x \sigma^2 + f + \frac{1}{2} \sigma^2 - r) \alpha_t \right\}. \quad (129)$$

The quadratic function of α_t in the curly bracket attains its maximum when $\alpha_t = \frac{B_x \sigma^2 + f + 0.5 \sigma^2 - r}{\gamma \sigma^2}$. But due to the constraint $\alpha_t \in [m, n]$, the optimal α_t^* should be determined as

$$\alpha_t^* = \min \left(n, \max \left(m, \frac{B_x \sigma^2 + f + 0.5 \sigma^2 - r}{\gamma \sigma^2} \right) \right). \quad (130)$$

Thus, we can set up numerical experiments similar to the finite-action cases based on equation (130). More specifically, at any point $(t_0, x_0) \in [0, T] \times \mathbb{R}$, α_t^* can be determined by the first order numerical derivative w.r.t x , plus comparisons with m and n . In terms of boundary conditions, since the admissible control set is bounded, asymptotically the optimal control should be either m or n . Thus, $B(t, x)$ should still be asymptotically linear in x for any fixed time. The rest of the numerical algorithm is the same as the

finite-action case and we outline the algorithm below.

Algorithm 3: Finite Difference Scheme for Constrained Action Problems

Input: T, X, M, N . $\mathcal{A} = \{0, 1\}$, $\gamma = 0.5$, $\mu_0 = 0.03$, $r = 0.03$, $\sigma = 0.3$, and $\sigma_0 = 0.03$.

- 1 Discretize $[0, T]$ and $[-X, X]$ into M and $2N$ sub-intervals.
 - (i) Denote $t_i = iT/M$ for $0 \leq i \leq M$.
 - (ii) Denote $x_j = (j - N)X/N$ for $0 \leq j \leq 2N$.
 - 2 Assign terminal condition:

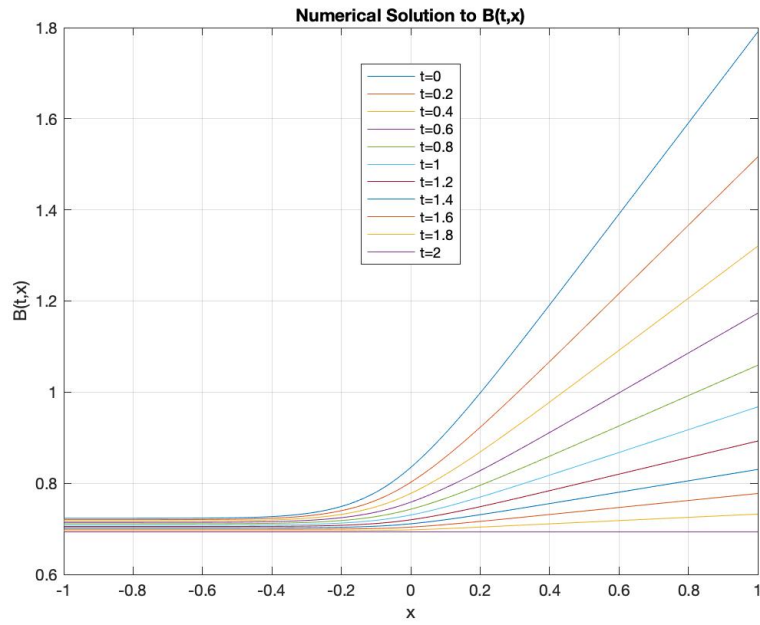

$$B(t_M, x_j) = -\log(1 - \gamma) \text{ for } 0 \leq j \leq 2N.$$
 - 3 For each j between $M - 1$ and 0:
 - (i) For each i between 1 and $2N - 1$:
 - (i-a) Compute $B_x(t_j, x_i)$ and $B_{xx}(t_j, x_i)$ using $B(t_{j+1}, x_{i-1})$, $B(t_{j+1}, x_i)$, and $B(t_{j+1}, x_{i+1})$.
 - (i-b) Compute α_i^* via equation (130).
 - (i-c) Compute $B(t_j, x_i)$ via equation (129).
 - (ii) Apply linear boundary conditions:

$$B(t_j, x_0) = 2B(t_j, x_1) - B(t_j, x_2), \text{ and}$$

$$B(t_j, x_{2N}) = 2B(t_j, x_{2N-1}) - B(t_j, x_{2N-2}).$$
-

We present some result of numerical studies based on Algorithm 3. For the purpose of comparison, we restrict numerical studies to $(t, x) \in [0, 2] \times [-1, 1]$, which is the same as the finite-action problems.

First we consider the case where $\mathcal{A} = [0, 1]$. Plots of numerical solutions to $B(t, x)$ and optimal actions are shown in Figure 19. In Figure 19(a), we see the plots gradually transit from flat to positive-slope. The corresponding optimal controls are mostly 0 on the left part and 1 on the right part. However, in the middle part of Figure 19(b) we observe optimal controls transit from 0 to 1. Although the interval where the changes take place is small, it makes a fundamental difference from what we observe in Figure 10(b) where optimal controls jump from 0 to 1.

(a) Numerical solution of $B(t, x)$ 

(b) Optimal control

Figure 19: Numerical solution of $B(t, x)$ and optimal control

In addition, we plot numerical solutions against analytic solutions to single-action problems where the actions are held constant, i.e. $\mathcal{A} = \{0\}$ or $\mathcal{A} = \{1\}$. In Figure 20, color plots represent our numerical solutions and dashed black plots represent analytic solutions where $\mathcal{A} = \{0\}$ and $\mathcal{A} = \{1\}$, respectively. Clearly they match very well for large or small x .




Figure 20: Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{0\}$ and $\mathcal{A} = \{1\}$)

We can further examine numerical derivatives of numerical solution w.r.t. x . As shown in Figure 21(a), the first order derivatives gradually transit from 0 to a constant values, which are slopes of solutions to the single-action problems. From Figure 21(b) we see the second order derivatives are bell-shaped curves which are similar to Figure 18(b) in the finite-action section.







Figure 21: Numerical derivatives of $B(t, x)$

Next we look at the numerical experiment where the admissible control set $\mathcal{A} = [-1, 1]$. In this case, an investor is allowed to take not only any long-only positions, but also any short position in the risky asset up to -1. The results are plotted in Figure 22, 23, and 24. Since short positions are allowed, $B(t, x)$ is no longer flat when x approaches -1, as shown in Figure 22(a). Together with Figure 23, we see the numerical solutions asymptotically match analytical solutions to constant control problem where the action is always -1 or 1. In terms of optimal controls, Figure 22(b) shows they change gradually from -1 to 1 in an interval between -0.2 and 0.2, and assume either -1 or 1 otherwise. In Figure 24, we see first order numerical derivatives transit smoothly from a set of negative values to a set of positive value, corresponding to two constant control problems.

(a) Numerical solution of $B(t, x)$ 

(b) Optimal control

Figure 22: Numerical solution of $B(t, x)$ and optimal control

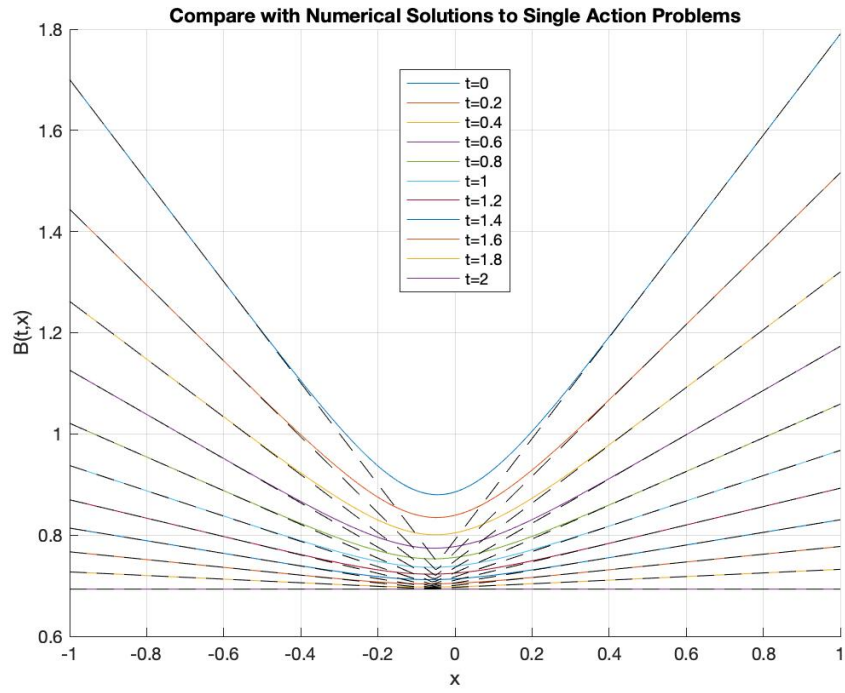
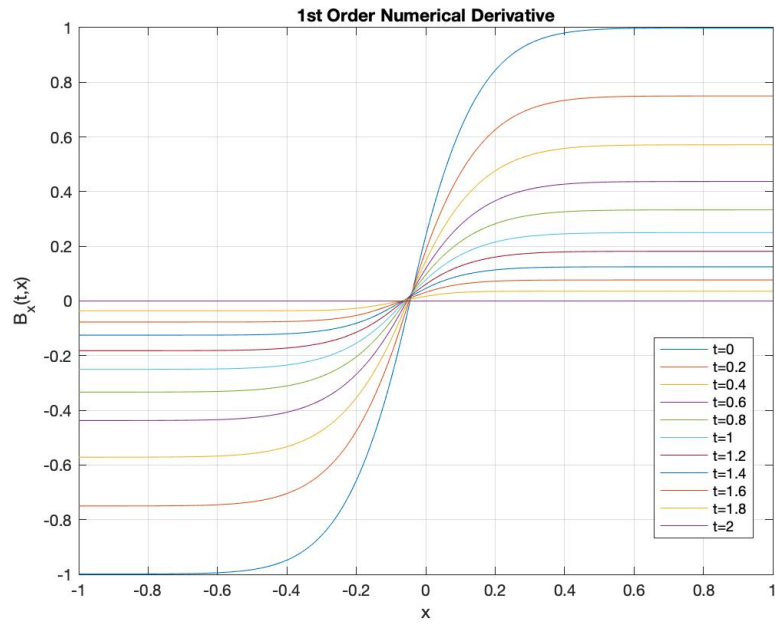



Figure 23: Comparison of the numerical solution with analytical solutions to single action problems ($\mathcal{A} = \{-1\}$ and $\mathcal{A} = \{1\}$)

(a) 1st order derivative $B_x(t, x)$ (b) 2nd order derivative $B_{xx}(t, x)$ Figure 24: Numerical derivatives of $B(t, x)$

We end this section by discussing how the the optimal control changes over time. In Figure 22(b) (or similarly Figure 19(b)), the curves representing the change from -1 to 1 (or from 0 to 1) are getting steeper as we step back in time, and the intervals where the change actually happen is getting smaller. An investor would choose either -1 or 1 outside those shrinking intervals. This effect can be better visualized by looking at the color map in Figure 25 which corresponding to the case where $\mathcal{A} = [-1, 1]$. In Figure 25, the horizontal axis is the log-price of risky asset, the vertical axis is time, and the color of each pixel represents the optimal control corresponding to a specific price and a specific time. The pure blue region on the left means an investor will choose $\alpha^* = -1$, the pure yellow region on the right means an investor will choose $\alpha^* = 1$. Between the pure blue and pure yellow region, the optimal control changes from -1 to 1 . As we can see, the transitional region shrinks when we move from $t = 2$ to $t = 0$. The observation agrees with equation (130) where the fractional term is a linear function of B_x and x with positive coefficients for a fixed t , and the coefficients become larger if t is smaller (see equation (91) for detail). Thus, if we fix x and move back in time, it is easier for the optimal control to hit upper or lower bound since due to the expanding value of B_x (e.g. Figure 18(a)).

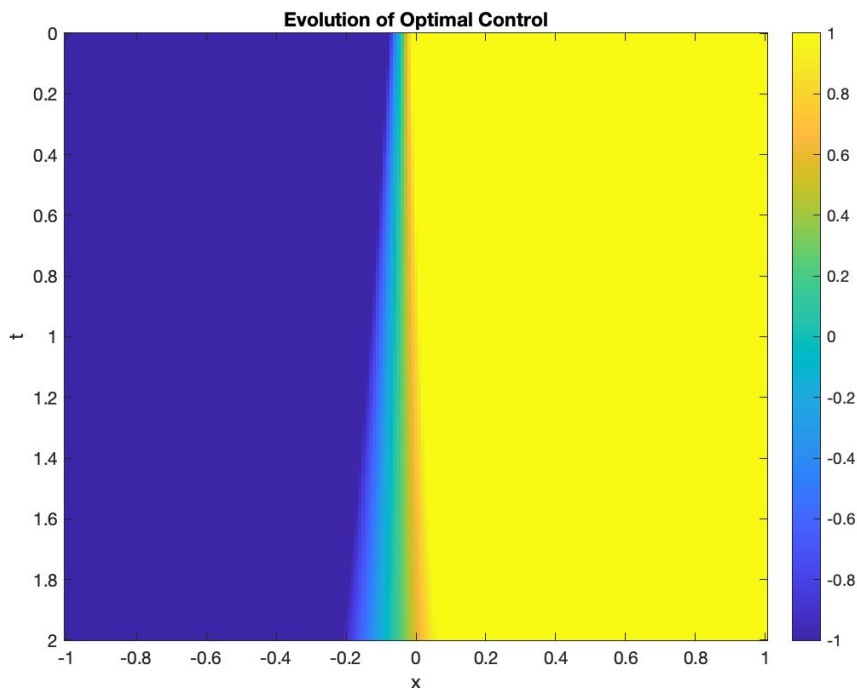


Figure 25: Evolution of optimal control in time ($\mathcal{A} = [-1, 1]$)

Will the transitional interval shrink to a single point? We don't have a complete answer to that. However, transitional interval can become very small in some cases. For example, when we conduct a study on $(t, x) \in [0, 2] \times [-5, 5]$ with $\mathcal{A} = [-30, 30]$ and other parameters the same as previous studies. As we can see in Figure 26, the transitional region is approximately $[-4, 4]$ near $t = 2$, but it almost shrinks to a single point when t is close to 0. As a matter of fact, we discretize the price range $[-5, 5]$ into 501 evenly separated grid points, and there are only four consecutive points (out of 500) where the optimal control is neither -30 nor 30 when $t < 0.1$.

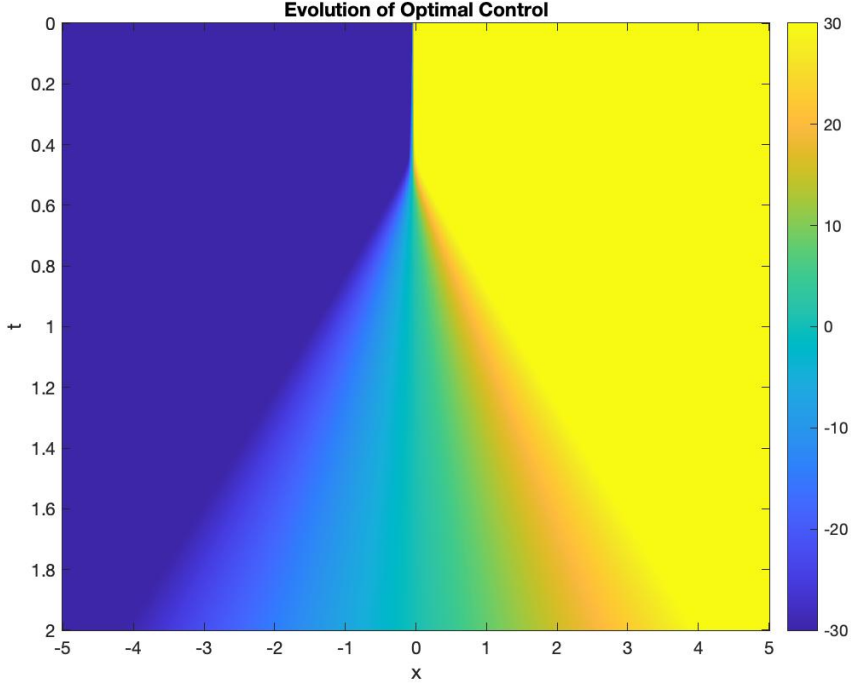


Figure 26: Evolution of optimal control in time ($\mathcal{A} = [-30, 30]$)

6.8 Discussion

In this chapter, we have studied 4 different cases of the portfolio allocation problem based on one common stochastic differential equation with Bayesian update. The main difference among these cases is the action set \mathcal{A} . For the unconstrained-action and single-action problems, we were able to establish theoretical results and find closed-form solutions. We notice that a problem similar to the unconstrained-action problem has been studied and one can refer to Bismuth[5] for a more general framework in the multi-dimensional case. The single-action problem provides us with important information about the asymptotic behavior of the value function when the price is extreme, which gives us the correct boundary conditions in the numerical studies of the finite-action and constrained-action problems. We then took a deep dive into the finite-action and constrained-action problems with several interesting cases of action set, and obtain

some interesting numerical results.

In section 6.6, we have conducted experiment for a finite-action problem where the admissible control set is $\mathcal{A} = \{0.01i - 1 | 0 \leq i \leq 200, i \in \mathbb{N}\}$. The results are presented in Figure 16, 17, and 18. When compare those results to what we obtained from the constrained-action problem using $\mathcal{A} = [-1, 1]$, we do not see much difference. In other words, numerical solutions in two cases are very close and share the same asymptotic behavior. In fact, the maximal difference in $B(t, x)$ is negligible based on our computation. In addition, optimal controls and numerical derivatives also behave very similarly in those two cases. We did notice that the actions in the finite-action problem must be integer multiples of 0.01 and there is no such constraint in the constrained-action problem. Despite the difference, the finite-action problem is qualified to be a very close approximation to the constrained-action problem according to our result.

To conclude this chapter, we compare the unconstrained-action problem with the classic Merton's portfolio problem. According to Merton, if the drift of risky asset μ is known, an investor should always hold a constant portion of risky asset (equation (62)). In our model, we do not know the exact value of μ . Under Bayesian belief system, we are able to reach equation equation (87) which suggests that the drift of the risky asset changes dynamically with its price. If we substitute the drift term of equation (87) into equation (62), we get the "optimal control"

$$\omega^* = \frac{\mu_0\sigma^2 + \sigma_0^2x}{(\sigma^2 + t\sigma_0^2)\sigma^2\gamma} - \frac{r}{\sigma^2\gamma}, \quad (131)$$

where x stands for the realization of X_t at time t . Using above equation and equation (111), we can directly compute

$$\alpha^* - \omega^* = \frac{2a(t)}{\gamma}x + \frac{b(t)}{\gamma} + \frac{1}{2\gamma}, \quad (132)$$

where the relation $x = \log y$ is used. Thus, we see that applying Merton's classic result with dynamic drift only gives sub-optimal investment choice. In particular, α^* and ω^* differ by a linear function of x for a fixed t .

Chapter 7

Conclusion

We have studied a variety of portfolio optimization models in this thesis. Started with the classic Markowitz model, we went through several modifications related to robust optimization and Bayesian methods. Then we took a dynamic perspective and examined portfolio optimization models in multiple periods and continuous-time. By combining a continuous-time model with Bayesian update, we derived a stochastic differential equation whose drift changes dynamically according observed price. Based on the equation, we further studied different cases of optimal portfolio allocation problems, and obtain theoretical as well as numerical results.

There are two directions worth further exploration. The first is to study the existence of classic solutions to HJB equations in finite-action and constrained-action cases. Although we have obtained numerically smooth solutions in both cases, the study is not completely rigorous due to the lack of theoretical support. The second direction is to apply effective numerical methods for HJB equations. We have relied on special coordinate transformations that make HJB equations easier to be studied, with linear or quadratic boundary conditions. However, such transformation may be hard to find for more general stochastic control problems and utility functions. As existing numerical methods often requires strong regularity conditions and limited boundary conditions , a general numerical framework about HJB equation will be much appreciated.

References

- [1] ANDERSON, T. W. *An introduction to multivariate statistical analysis*. John Wiley & Sons, New York, 1984.
- [2] BAIN, A., AND CRISAN, D. *Fundamentals of stochastic filtering*, vol. 60. Springer Science & Business Media, 2008.
- [3] BELLMAN, R. The theory of dynamic programming. *Bulletin of the American Mathematical Society* 60, 6 (1954), 503–516.
- [4] BEST, M. J., AND GRAUER, R. R. On the sensitivity of mean-variance-efficient portfolios to changes in asset means: some analytical and computational results. *The review of financial studies* 4, 2 (1991), 315–342.
- [5] BISMUTH, A., GUÉANT, O., AND PU, J. Portfolio choice, portfolio liquidation, and portfolio transition under drift uncertainty. *Mathematics and Financial Economics* 13, 4 (2019), 661–719.
- [6] BLACK, F., AND LITTERMAN, R. Global portfolio optimization. *Financial Analysts Journal* 48 (1992), 28–43.
- [7] BLACK, F., AND LITTERMAN, R. B. Asset allocation: Combining investor views with market equilibrium. *The Journal of Fixed Income* 1, 2 (1991), 7–18.
- [8] BOLSTAD, W. M., AND CURRAN, J. M. *Introduction to Bayesian statistics*. John Wiley & Sons, 2016.
- [9] BROADIE, M. Computing efficient frontiers using estimated parameters. *Annals of operations research* 45, 1 (1993), 21–58.

- [10] DA PRATO, G., AND ZABCZYK, J. *Stochastic equations in infinite dimensions*. Cambridge university press, 2014.
- [11] FAMA, E. F., AND FRENCH, K. R. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33 (1993), 3–56.
- [12] FAMA, E. F., AND FRENCH, K. R. A five-factor asset pricing model. *Journal of financial economics* 116, 1 (2015), 1–22.
- [13] FOLLAND, G. B. *Real analysis: modern techniques and their applications*, vol. 40. John Wiley & Sons, 1999.
- [14] FRIEDMAN, A. Stochastic differential equations and applications. In *Stochastic differential equations*. Springer, 2010, pp. 75–148.
- [15] GILBARG, D., AND TRUDINGER, N. S. *Elliptic partial differential equations of second order*, vol. 224. springer, 2015.
- [16] GOLDFARB, D., AND IYENGAR, G. Robust portfolio selection problems. *Mathematics of operations research* 28, 1 (2003), 1–38.
- [17] GRANT, M., AND BOYD, S. Cvx: Matlab software for disciplined convex programming, version 2.1.
- [18] HILLE, E. *Ordinary differential equations in the complex domain*. Courier Corporation, 1997.
- [19] IYENGAR, G. N. Robust dynamic programming. *Mathematics of Operations Research* 30, 2 (2005), 257–280.
- [20] KARATZAS, I., AND SHREVE, S. *Brownian Motion and Stochastic Calculus*, vol. 113. Springer Science & Business Media, 2012.
- [21] KEEPING, E. S. *Introduction to statistical inference*. Courier Corporation, 1995.
- [22] KNAPP, A. W. *Basic real analysis*. Springer Science & Business Media, 2005.
- [23] KRYLOV, N. V. *Nonlinear elliptic and parabolic equations of the second order*, vol. 7. Springer, 1987.

- [24] LEE, W. *Theory and methodology of tactical asset allocation*, vol. 65. John Wiley & Sons, 2000.
- [25] LINTNER, J. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics* (1965), 13–37.
- [26] LU, Z. Robust portfolio selection based on a joint ellipsoidal uncertainty set. *Optimization Methods & Software* 26, 1 (2011), 89–104.
- [27] MARKOWITZ, H. Portfolio selection. *The Journal of Finance*, 07 (1952), 77–91.
- [28] MERTON, R. C. Lifetime portfolio selection under uncertainty: The continuous-time case. *The review of Economics and Statistics* (1969), 247–257.
- [29] MERTON, R. C. An analytic derivation of the efficient portfolio frontier. *Journal of financial and quantitative analysis* (1972), 1851–1872.
- [30] MERTON, R. C., ET AL. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3, 4 (1971), 373–413.
- [31] MEUCCI, A. *Risk and Asset Allocation*. Springer Science & Business Media, 2007.
- [32] MEUCCI, A. The black-litterman approach: Original model and extensions. *Shorter version in, THE ENCYCLOPEDIA OF QUANTITATIVE FINANCE*, Wiley (2010).
- [33] PHAM, H. *Continuous-time stochastic control and optimization with financial applications*, vol. 61. Springer Science & Business Media, 2009.
- [34] SHARPE, W. F. Capital asset prices: A theory of market equilibrium under conditions of risk. *The journal of finance* 19, 3 (1964), 425–442.
- [35] SHREVE, S. E. *Stochastic calculus for finance II: Continuous-time models*, vol. 11. Springer Science & Business Media, 2004.
- [36] TREYNOR, J. L. Market value, time, and risk. *Time, and Risk (August 8, 1961)* (1961).

- [37] TREYNOR, J. L. *Toward a theory of market value of risky assets*. 1961.

Appendix A

A.1 Derivation of HJB Equation

We follow the notations used in Chapter 3 of Pham[33]. Consider a n -dimensional stochastic process $X_s \in \mathbb{R}^n$ defined by

$$dX_s = b(s, X_s, \alpha_s)ds + \sigma(s, X_s, \alpha_s)dW_s, \quad (133)$$

where s stands for time, α_s is the control process, and W_s is a d -dimensional Brownian motion. At some future time $T \in \mathbb{R}$, a reward will be given according to a function $g(\cdot)$ with X_T as the input to $g(\cdot)$. Assume at time $t < T$ the observed value of X_s is $x \in \mathbb{R}^n$, we can define the value function of (t, x) as

$$v(t, x) = \sup_{\alpha \in \mathcal{A}} \mathbb{E}[g(X_T^{t,x})]. \quad (134)$$

In equation (134), $X_T^{t,x}$ stands for the stochastic process whose value is x at time t , with the control α applied between t and T , and \mathcal{A} is the set of all admissible control processes. The objective of a stochastic control problem is to find the optimal control process α^* so that the expectation on the right-hand-side of equation (134) is maximized.

Now consider applying a constant control $\alpha_s = a$ between t and $t + h$ where a is a proper choice w.r.t. the pair (t, x) . By definition of $v(t, x)$ and the dynamic programming principle (section 3.3 of Pham[33]), we have the following relation

$$v(t, x) \geq \mathbb{E}[v(t+h, X_{t+h}^{t,x})]. \quad (135)$$

Assume v is smooth enough, we may apply Ito's formula between t and $t + h$:

$$v(t + h, X_{t+h}^{t,x}) = v(t, x) + \int_t^{t+h} \left(\frac{\partial v}{\partial t} + \mathcal{L}^a v \right)(s, X_s^{t,x}) ds + M_s, \quad (136)$$

where

$$\mathcal{L}^a v = b(t, x, a) \cdot D_x v + \frac{1}{2} \text{tr}(\sigma(t, x, a) \sigma(t, x, a)^T D_{xx} v) \quad (137)$$

and M_s is a (local) martingale. In equation (137), D_x and D_{xx} stand for the first order and second order derivative operator, respectively. In general, we do not need to worry about M_s since eventually we will let $h \rightarrow 0$ and the expectation of M_s can be handled by localization. Thus, substituting equation (136) into equation (135) we have

$$0 \geq \mathbb{E} \left[\int_t^{t+h} \left(\frac{\partial v}{\partial t} + \mathcal{L}^a v \right)(s, X_s^{t,x}) ds \right]. \quad (138)$$

Divide both sides of equation (138) by h and let $h \rightarrow 0$, we obtain

$$0 \geq \frac{\partial v}{\partial t}(t, x) + \mathcal{L}^a v(t, x). \quad (139)$$

Now since a can be any admissible choice, we have

$$0 \geq \frac{\partial v}{\partial t}(t, x) + \sup_{a \in A} [\mathcal{L}^a v(t, x)], \quad (140)$$

where A represents the set of admissible choices of a w.r.t. (t, x) . However, if a is chosen according to the optimal control that achieves equation (134), using the same argument we obtain the equality

$$0 = \frac{\partial v}{\partial t}(t, x) + \sup_{a \in A} [\mathcal{L}^a v(t, x)] \quad (141)$$

assuming the supreme is finite. Finally, since (t, x) is arbitrary, we obtain the HJB equation

$$0 = \frac{\partial v}{\partial t}(t, x) + \sup_{\alpha \in A} [\mathcal{L}^a v(t, x)]. \quad (142)$$