Title
Topics in Dynamic Portfolio Choice Problems

Author
Wimonkittiwat, Poomyos

Publication Date
2013

Peer reviewed|Thesis/dissertation

Topics in Dynamic Portfolio Choice Problems

by

Poomyos Wimonkittiwat

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Associate Professor Andrew E.B. Lim, Chair
Professor Zuo-Jun Shen
Professor Steven N. Evans

Spring 2013

Topics in Dynamic Portfolio Choice Problems

Abstract

Topics in Dynamic Portfolio Choice Problems

by

Poomyos Wimonkittiwat

Doctor of Philosophy in Engineering - Industrial Engineering and Operations Research

University of California, Berkeley

Associate Professor Andrew E.B. Lim, Chair

We study two important generalizations of dynamic portfolio choice problems: a portfolio choice problem with market impact costs and a portfolio choice problem under the Hidden Markov Model.

In the first problem, we allow the presence of market impact and illiquidity. Illiquidity and market impact refer to the situation where it may be costly or difficult to trade a desired quantity of assets over a desire period of time. In this work, we formulate a simple model of dynamic portfolio choice that incorporates liquidity effects. The resulting problem is a stochastic linear quadratic control problem where liquidity costs are modeled as a quadratic penalty on the trading rate. Though easily computable via Riccati equations, we also derive a multiple time scale asymptotic expansion of the value function and optimal trading rate in the regime of vanishing market impact costs. This expansion reveals an interesting but intuitive relationship between the optimal trading rate for the illiquid problem and the classical Merton model for dynamic portfolio selection in perfectly liquid markets. It also gives rise to the notion of a "liquidity time scale". Furthermore, the solution to our illiquid portfolio problem shows promising performance and robustness properties.

In the second problem, we study dynamic portfolio choice problems under regime switching market. We assume the market follows the Hidden Markov Model with unknown transition probabilities and unknown observation statistics. The main difficulty of this dynamic programming problem is its high-dimensional state variables. The joint probability density function of the hidden regimes and the unknown quantities is part of the state variables, and this makes the problem suffer from the curse of dimensionality. Though the problem cannot be solved by any standard fashions, we propose approximate methods that tractably solve the problem. The key is to approximate the value function by that of a simpler problem where the regime is not hidden and the parameters are observable (the C-problem). This approximation allows the optimal portfolio to be computed in a semi-explicit way. The approximate solution shares the same structure with the solution of C-problem, but at the same time it provides clear insight into the unobservable extension. In addition, the perfor-

mance of the proposed methods is reasonably close to the upper-bound obtained from the information relaxation problem.

Contents

Acknowledgments

I would like to express how grateful I am to be one of Andrew Lim's PhD students. I truly appreciate his constant guidance and support. Without Andrew's persistent help, this dissertation would not be possible. Besides my advisor, I would like to thank my dissertation committee, Zuo-Jun Shen and Steven Evans, for their interest in my work and constructive comments. Also many thanks to UC Berkeley instructors those taught me valuable knowledge which has been found very useful in this dissertation.

Above all, I would like give my deepest appreciation to my mother who is always supportive and loves me unconditionally. I also thank my caring relatives, the Wimonkittiwat's, the Tiyapibulchaiya's, the Chaochonphan's and the Prayongpan's. Special thanks to Jackie for being with me in good and bad times. I also thank my IEOR friends for the good time at school and all the help during my studies, my Thai friends in Berkeley for the fun time outside school. Finally, to the rest I forgot to mention, thank you.

Chapter 1

Introduction

The dynamic portfolio choice problem has been mentioned countless of times by both researchers and practitioners. It is considerably one of the most important problems in investment science. This is a problem of dynamically allocating assets among sources of risks and returns; an investor may re-balance his/her portfolio based on the current situation and the knowledge of the market to maximize his/her utility and/or consumption. The problem was first studied by Robert C. Merton in the 1960s [49, 50], we refer to this problem as *the Merton problem*. The Merton problem can be explicitly solved in many cases due to the help of a few simplifying assumptions; these assumptions include short positions are allowed, there is no bid-ask spread, no transaction costs, the market is perfectly liquid, all variables and parameters are observable, and risky asset prices are driven by the Wiener processes. However, these assumptions are unrealistic, the solution to the Merton problem is optimal only in an ideal market. But the problem still provides good understanding to dynamic portfolio choice problem and is used as a first step to unfold more complicated problems. There are a huge number of efforts those try to generalize the Merton problem. These efforts do so by relaxing the simplifying assumptions, and this dissertation is one of them.

Some examples of the works on extending the Merton problem are as follows. First, investors may have constraints on asset holdings that limit the amount of short-sale or limit the maximum amount of assets being held. The works on dynamic portfolio choice problems those allow these position constraints are such as Cvitanic and Karatzas [22], Tepla [59] and Li et al. [42]. Second, in the real markets, investors buy and sell at different prices, buy at the ask price and sell at the bid price. This situation is also referred to as an existence of proportional transaction costs. The literature that studies relaxation of the bid-ask spread assumption is such as Constantinides [20], Davis and Norman [23], Shreve and Soner [58] and Liu [44].

Under perfectly liquid market assumption, investors may trade any amount of assets at any time and these decisions do not affect the prices. But this is not true in reality, markets are illiquid. Illiquidity and market impact refer to a situation where investor decisions can affect price processes. This is due to many reasons: demand and supply of securities are limited, there are feedback effects on prices as investors' decisions are dependent. The main difficulties in dealing with illiquidity and market impact in portfolio problems are twofold: first, market impact is very complicated as it is a result of complex market structure and all investors participating. The market impact model is controversial, there are various models of market impact in the literature (Dufour and Engle [24], Almgren and Chriss [3], Cetin et al. [18], Bouchaud [12], Cont et al. [21], Alfonsi et al. [2]). Second, incorporating market impact into the portfolio choice problem leads to much harder problem and sometimes impossible to solve. The works those study market impact in the context of portfolio optimization include Longstaff [46], Grinold [33], Ma et al. [48], Rogers and Singh [55, 54], Garleanu and Pedersen [29] and this dissertation. It is worth noting that illiquidity and market impact are also important in a closely related problem, an execution problem. The papers on this issue are e.g. Bertsimas and Lo [10], Almgren and Chriss [3] and Alfonsi et al. [1].

Another important development in dynamic portfolio selection problems is how to deal with hidden variables and parameters uncertainty. It is assumed in the classical Merton problem that all the model parameters, such as expected asset returns and volatilities, are fixed and known to the investors. In reality, however, it is unclear how investors may acquire information on these parameters. Besides, the studies by Hamilton [36], Ang and Bekaert [5], Tu [60], Carvalho et al. [17] agree that these quantities may change over time and ignoring this fact is costly. There are many approaches to model unknown parameters: Gennotte [31] and Brennan [14] assumes that the parameters are fixed but unknown. Brendle [13] assumes linear Gaussian processes for the unobservable parameters. Honda [38], Rieder and Bauerle [53], Sass and Haussmann [56] and this dissertation assume the unknown parameters are driven by unobservable Markov switching process.

Finally, the assumption that risky assets are driven by the Wiener processes may not be suffice as abrupt changes or shocks appear occasionally in real markets. These sudden changes are a result of market crash, beginning of financial crisis, company default, natural disaster etc. To capture these events, Liu et al. [45], Bauerle and Rieder [8], Capponi and Figueroa-Lopez [16] and Lim and Watewai [43] modify the asset price processes by allowing jumps in assets prices as well as sudden changes in price processes. The second part of this dissertation can be considered in this category too as we assume that the market is driven a Markovian regime switching process.

This dissertation studies two important generalizations in dynamic portfolio choice problem as mentioned above. We dedicate the first part of this dissertation to study dynamic portfolio choice problems with the presence of market impact costs, and dedicate the second part to study dynamic portfolio choice problems under unobservable regime switching market.

The first part focuses on solving portfolio optimization problem when there exists market impact. There are two main approaches in research community to model market impact. The first approach models market impact as impact of trades on security prices (Bertsimas and Lo [10], Almgren and Chriss [3], He and Mamaysky [37], Ly Vath et al. [47], Rogers and Singh [55], Alfonsi et al. [2]). Although this is a natural and intuitive way to model market impact, the problem suffers from intractability issue. Associating impact of investor decisions into price dynamics substantially increases complexity of the problem. An alternative approach which we prefer is to model market impact as costs on fast trading and incorporate these costs into investor objective. Grinold [33] and Garleanu and Pedersen [29] use this approach in the context of active asset management while this dissertation is the first to use this in the context of the Merton problem. The main advantage of incorporating costs into the objective is the problem becomes tractable, and therefore insightful findings are obtained.

We present a detailed motivation and more complete literature review as well as a brief review of the classical Merton problem [50] in Chapter 2. Then we formulate two equivalent problems to solve the portfolio selection problem with market impact costs and derive their relationship in Chapter 3. The first formulation is an extension to the classical Merton problem with costs on fast trading. The second one is very intuitive in the sense that, with presence of market impact, it is also optimal to track the Merton problem solution at minimal costs. We analyze the solution of these two formulation in Chapter 4. This analysis reveals an intuitive relationship between the optimal trading strategies with and without presence of market impact. We demonstrate the performance and robustness properties of the portfolio obtained using our approach on the temporary price impact model of Almgren et al. [4] in Chapter 5. The proofs which are not provided right after the assertions can be found in Chapter 6.

The second part of this dissertation focuses on solving the dynamic portfolio choice problem under regime switching market. We relax the Merton problem assumptions that variables and parameters are observable by adopting the Hidden Markov Model (HMM) to model the market; the market regime which is unobservable follows a Markov chain, and this regime affects the price processes through the asset expected returns and volatilities. Although there are a few previous works those study the dynamic portfolio choice problem and model the market in this particular way [38, 56, 53, 34, 35, 11, 26], we make the most practical assumption, namely the market regime is hidden, the regime transition probabilities and the regime-dependent parameters are also unknown to the investors. In this problem, investors only observe the output generated from the market which includes asset prices, analyst recommendation, news etc. We describe the problem setting in detail in Chapter 7. We then formulate the investor problem in Chapter 8. This dynamic portfolio choice problem suffers from the curse of dimensionality and is considered unsolvable because the investor's knowledge on the hidden process and unknown parameters is represented by a highly complicated joint probability density function of all unknown quantities, and solving the problem requires doing backward recursion on this density function. Instead of using the standard method like the backward recursion to solve the problem, we propose approximate methods which allow us to tractably compute the solution without directly dealing with the complex joint probability distribution. We present the methods as well as the main result from the approximation in Chapter 9. Also in Chapter 9, we compare the performance of the approximate methods with the existing approaches under both real data from US equity market and simulated data. We summarize the Markov Chain Monte Carlo method which we used in problem approximation in Chapter 10. Moreover, we provide a mean to evaluate the proposed methods by comparing its performance to the upper-bound of the exact problem. The upper-bound is obtained from the dual problem via a technique called the information relaxation technique. The material on performance evaluation is explained in Chapter 11. And the proofs in the second part of this dissertation can be found in Chapter 12. Finally, we conclude the work done in this dissertation in Chapter 13

Part I

Dynamic Portfolio Choice Problems with Market Impact Costs

Chapter 2

Motivation

Illiquidity and market impact refer to the situation where it may be costly or difficult to trade a desired quantity of assets over a desire period of time. Much of the recent effort in modeling illiquidity and incorporating it in hedging, portfolio choice, and execution problems is motivated by the recognition that ignoring it is both risky and costly (e.g. Bertsimas and Lo [10], Almgren and Chriss [3], Grinold [33] and Garleanu and Pedersen [28], He and Mamaysky [37], Ly Vath et al. [47], Moallemi and Saglam [51], Rogers and Singh [55], Schied et al. [57]). While much of this literature accounts for illiquidity by explicitly modeling the impact of trades on asset prices ([28, 33, 51] are exceptions, though these papers focus on active portfolio choice problems), this approach is challenging for a number of reasons. Firstly, the mechanism by which trades impact prices is complex and not completely understood, so model mis-specification is a serious concern. Secondly, even simple models for the price impact of trades lead to substantially harder optimization problems. For this reason, much of the literature has been restricted to simple single-asset models, and is difficult to extend to multiple assets and to models in which stochastic factors impact liquidity.

In this work, we adopt a different approach in which liquidity-related costs are incorporated through a penalty in the objective that makes it costly to trade large quantities at a fast rate rather than an explicit model of how trading affects prices. While ignoring the impact of trading on prices is perhaps controversial, though the same can be said about the decision to ignore the impact of trading rate on alphas in [28, 33, 51], one argument in support is that optimal trading rates, by virtue of our objective, will be slow so price impact (even if it were to be included) will be minimal, “justifying” to some degree our decision to ignore them. More generally, we adopt the view that an optimization model is simply a way of generating plausibly good solutions by making the appropriate tradeoffs - in this case, between investment performance and market impact costs - and that an explicit model of the price impact mechanism may be unnecessary so long as the appropriate tradeoffs are being captured in the simpler model. One advantage of our approach is that it is tractable. Specifically, our formulation naturally handles multiple assets which is not the case for many of the models that directly model the price impact mechanism. Though we focus on constant liquidity costs, our results extend to problems with stochastic liquidity regimes, jumps in prices, regime-switching etc., with no essential difficulties. That said, we agree that it is important to assess the performance of solutions generated by our model on more complex systems under different models of the impact mechanism. This is the focus of our numerical study which shows favorable performance and robustness properties when applied to the price impact model of Almgren et al. [4].

Several highlights of our work are as follows. Firstly, we show that our generalization of the Merton problem is equivalent to another "illiquid problem" where the objective is to track the optimal holding of the perfectly liquid Merton problem with minimal error and market impact cost. This gives an intuitively appealing relationship between two approaches for accounting for illiquidity. Secondly, we also derive a multiple time scale asymptotic expansion of the value function and optimal trading rate for our problem which provides interesting insights into the structure of the optimal strategy. The asymptotic expansion shows that when the market impact costs are small (i) the investor trades in a direction that decreases the gap between his/her current position and optimal holding for the Merton problem, and (ii) that the trading rate is increasing in the volatility of the risky asset, the risk-aversion of the investor, the market depth, and the remaining trading time. It also leads to the introduction of a "liquidity time scale", an intuitive notion we have not seen in other contexts.

We first give a brief review to the classical Merton problem. The Merton problem will be served as the base model in the following chapters.

2.1 Portfolio Problem in Liquid Market

We recall the classical Merton problem [50] for frictionless markets.

Asset Dynamics

We model uncertainty using an n-dimension standard Brownian motion living on a filtered probability space $(\Omega, \mathcal{F}, P, \{\mathcal{F}_t\})$ over a finite time horizon $[0, T]$. There are n risky assets with prices $s(t) = [s_1(t), \cdots, s_n(t)]$ follow geometric Brownian motion

$$ds(t) = \operatorname{diag}(\mu)s(t)dt + \sum_{j=1}^{n} \operatorname{diag}(\sigma_{\cdot j})s(t)dw_j(t), \tag{2.1}$$

with expected returns $\mu \in \mathbb{R}^n$, diffusion matrix $\sigma = [\sigma_{\cdot 1}\ \sigma_{\cdot 2}\ \ldots\ \sigma_{\cdot n}] \in \mathbb{R}^{n\times n}$, and $\operatorname{diag}(x)$ is a diagonal matrix with vector x as its main diagonal. The risk-free asset price $s_0(t)$ satisfies

$$ds_0(t) = rs_0(t)dt, \tag{2.2}$$

with risk-free rate of return r.

The Merton Problem

Let $n(t) = [n_1(t), \cdots, n_n(t)]$ be an adapted process denoting the number of shares in each of the risky assets at time t. The wealth $x(t)$ of a self-financing investor satisfies

$$dx(t) = \Big\{rx(t) + (\mu - r)' \operatorname{diag}(n(t))s(t)\Big\}dt + \sum_{j=1}^{n} \operatorname{diag}(\sigma_{\cdot j}) \operatorname{diag}(n(t))s(t)\; dw_j(t)$$

We rewrite the wealth dynamics in terms of the dollar value $\pi(t) \triangleq \operatorname{diag}(n(t))s(t)$ of the risky asset holdings

$$dx(t) = \Big\{rx(t) + (\mu - r)'\pi(t)\Big\}dt + \sum_{j=1}^{n} \sigma'_{\cdot j}\pi(t)dw_j(t). \tag{2.3}$$

The classical Merton problem maximizes expected utility of terminal wealth

$$\begin{cases} \sup\limits_{\pi(\cdot)} \; \mathbb{E}\{\Phi(x(T))\} \\ \text{subject to:} \\ dx(t) = \Big\{rx(t) + (\mu - r)'\pi(t)\Big\}dt + \sum\limits_{j=1}^{n} \sigma'_{\cdot j}\pi(t)dw_j(t) \\ x(0) = x_0. \end{cases} \tag{2.4}$$

The value function $V(t, x)$ for (2.4) is the solution of the dynamic programming equations

$$\begin{cases} V_t + \sup\limits_{\pi} \Big\{xrV_x + (\mu - r)'\pi V_x + \dfrac{1}{2}\pi'\sigma\sigma'\pi V_{xx}\Big\} = 0 \\ V(T, x) = \Phi(x(T)). \end{cases} \tag{2.5}$$

Explicit solutions for the Merton problem and the associated dynamic programming equation can be found when the utility function is of power, exponential, logarithmic and quadratic type. In the case of quadratic utility we have the following result.

Proposition 1. *The value function of the Merton problem (2.4) with quadratic utility* $\Phi(x(T)) = x(T) - \frac{\eta}{2}x(T)^2$ *is*

$$V_M(t, x) = -\frac{\eta}{2}a_M(t)x^2 + b_M(t)x + c_M(t) \tag{2.6}$$

where

$$a_M(t) = e^{\big(2r - \theta\big)(T-t)}$$
$$b_M(t) = e^{\big(r - \theta\big)(T-t)}$$
$$c_M(t) = \frac{1}{2\eta}\big(1 - e^{-\theta(T-t)}\big).$$

Note that $\theta = (\mu - r)'\Sigma^{-1}(\mu - r)$ *and* $\Sigma = \sigma\sigma'$*. And the optimal risky asset holdings are*

$$\begin{aligned}\pi_M^*(t, x) &= \Sigma^{-1}(\mu - r)(\frac{1}{\eta}\frac{b_M(t)}{a_M(t)} - x) \\ &= \Sigma^{-1}(\mu - r)(\frac{1}{\eta}e^{-r(T-t)} - x).\end{aligned} \tag{2.7}$$

(Note that the subscript M *is added for later reference.)*

Proof. It can be shown directly that (2.6) is the solution of the dynamic programming equations (2.5) and that $\pi_M^*(t, x)$ is the associated maximizer. □

Chapter 3

Illiquid Portfolio Choice Problem Formulation

We propose two formulations of the illiquid portfolio choice problem. First formulation, we extend the Merton problem (2.4) to account for illiquidity using the trading costs. Second formulation, we track the Merton optimal portfolio with minimum trading cost. Then we show that these two formulations are equivalent.

3.1 Illiquid Portfolio Problem: Model 1

In this section, we account for liquidity effects by formulating a modification of the Merton problem in which fast trading is expensive. In this model, trading rates are the control variables while the wealth and the risky asset positions are the state variables.

Trading rate

Let $n(t) = [n_1(t), \cdots, n_n(t)]$ denote the number of shares in each of the risky assets and $\rho(t) = [\rho_1(t), \cdots, \rho_n(t)]$ denote the rates at which shares in each of the risky asset are being purchased at time t, which is controlled by the investor. The risky asset holdings satisfy

$$dn(t) = \rho(t)dt. \tag{3.1}$$

As in the Merton problem (2.4) let $\pi(t) \triangleq \operatorname{diag}(n(t))s(t)$ denote the dollar value of the risky asset holdings. In contrast to the Merton problem, where $\pi(t)$ are the control variables, we now assume that our asset holdings can only be controlled through the trading rates $\rho(t)$ so $\pi(t)$ needs to be included as part of the state. Ito's formula together with (2.1) and (3.1) imply that

$$d\pi(t) = \{\operatorname{diag}(\mu)\pi(t) + \operatorname{diag}(\rho(t))s(t)\}dt + \sum_{j=1}^{n} \operatorname{diag}(\sigma_{\cdot j})\pi(t)dw_j(t).$$

Defining $\bar{\rho}(t) \triangleq \operatorname{diag}(\rho(t))s(t)$, it follows that

$$d\pi(t) = \{ \operatorname{diag}(\mu)\pi(t) + \bar{\rho}(t)\}dt + \sum_{j=1}^{n} \operatorname{diag}(\sigma_{\cdot j})\pi(t)dw_j(t). \tag{3.2}$$

In this equation, $\bar{\rho}(t)$ can be interpreted as a vector of trading rates, in dollars per unit time, of the risky assets. Equation (3.2) tells us that the change in the dollar value of the risky asset holdings is the sum of the change due to trade, which occurs at rate $\bar{\rho}(t) = \operatorname{diag}(\rho(t))s(t)$, and the change $\operatorname{diag}(\mu)\pi(t)dt + \sum_{j=1}^{n} \operatorname{diag}(\sigma_{\cdot j})\pi(t)dw_j(t)$ due to fluctuations in the value of assets already being held. The self-financing condition implies that the wealth process $x(t)$ remains unchanged from the Merton problem and is given by (2.3).

Liquidity costs & Model I

To incorporate the idea that it is costly to trade large quantities of assets in small periods of time, we consider the problem

$$\begin{cases} \sup\limits_{\bar{\rho}(\cdot)} \; \mathbb{E}\left\{ -\int_0^T \frac{\lambda}{2}\,\bar{\rho}(t)'\,R\,\bar{\rho}(t)\,dt + \Phi(x(T)) \right\} \\ \text{subject to:} \\ \begin{bmatrix} dx(t) \\ d\pi(t) \end{bmatrix} = \left\{ \tilde{A} \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + \tilde{B}\,\bar{\rho}(t) \right\} dt + \sum_{j=1}^{n} \tilde{C}_j \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} dw_j \\ \bar{\rho}(\cdot) \in \mathcal{A},\; x(0) = x_0,\; \pi(0) = \pi_0 \end{cases} \tag{3.3}$$

where

$$\tilde{A} = \begin{bmatrix} r & (\mu - r)' \\ \underline{0} & \operatorname{diag}(\mu) \end{bmatrix},\; \tilde{B} = \begin{bmatrix} \underline{0}' \\ I_n \end{bmatrix},\; \tilde{C}_j = \begin{bmatrix} 0 & \sigma_{\cdot j}' \\ \underline{0} & \operatorname{diag}(\sigma_{\cdot j}) \end{bmatrix},$$

$\underline{0}$ denotes an n-dimensional vector of zeros, I_n is an $n \times n$ identity matrix, and $\mathcal{A}$ is the class of admissible controls defined as

$$\mathcal{A} = \left\{ \bar{\rho} : [0,\, T] \times \Omega \to \mathbb{R}^n \;\middle|\; \text{such that } \bar{\rho}(\cdot) \text{ is } \{\mathcal{F}_t\}\text{-adapted and } \mathbb{E}\int_0^T \|\, \bar{\rho}(t) \,\|^2 \, dt < \infty \right\}.$$

Aside from the dynamics, which we have already discussed, a key modification relative to the Merton problem (2.4) is the introduction of the quadratic penalty on the trading rates

$$\mathbb{E}\int_0^T \frac{\lambda}{2}\,\bar{\rho}(t)'\,R\,\bar{\rho}(t)\,dt \tag{3.4}$$

in the objective. Optimizing (3.3) involves a trade-off between maximizing expected utility $\mathbb{E}[\Phi(x(T))]$ and minimizing the cost of trading. Assuming R is positive definite, the quadratic $\bar\rho'\, R\, \bar\rho$ in (3.4) means that marginal cost of trading increases in the trading rate. The *illiquidity coefficient* $\lambda > 0$ is large when illiquidity frictions are large. Similar models for liquidity costs have been proposed for dynamic models of active portfolio investment (Garleanu and Pedersen [29], Grinold [33]). To our knowledge, this work is the first in which such a model is used in the context of a Merton-type problem.

When the utility function is quadratic, the value function and the optimal trading rates can be characterized as follows.

Proposition 2. *Suppose that the illiquidity coefficient λ is positive, the matrix R is positive definite and the utility function is quadratic with risk-aversion parameter $\eta > 0$:*

$$\Phi(x(T)) = x(T) - \frac{\eta}{2} x(T)^2. \tag{3.5}$$

Then the value function is

$$V(t,x,\pi) = -\frac{1}{2} \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix}' a(t) \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + b(t)' \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + c(t), \tag{3.6}$$

and the optimal trading rates satisfy

$$\bar\rho^*(t,x,\pi) = \frac{1}{\lambda} R^{-1} \tilde{B}' \left(b(t) - a(t) \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} \right) \tag{3.7}$$

The functions $a : [0,T] \to \mathbb{R}^{(n+1)\times(n+1)}$, $b : [0,T] \to \mathbb{R}^{n+1}$ and $c : [0,T] \to \mathbb{R}$ are solutions of the differential equations

$$\begin{cases} 0 = & \dot a(t) + a(t)\tilde{A} + \tilde{A}' a(t) + \sum_{j=1}^{n} \tilde{C}_j' a(t) \tilde{C}_j - \frac{1}{\lambda} a(t)' \tilde{B} R^{-1} \tilde{B}' a(t) \\ a(T) = & \begin{bmatrix} \eta & \underline{0}' \\ \underline{0} & \underline{\underline{0}} \end{bmatrix} \end{cases} \tag{3.8}$$

$$\begin{cases} 0 = & \dot b(t) + \tilde{A}' b(t) - \frac{1}{\lambda} a(t)' \tilde{B} R^{-1} \tilde{B}' b(t) \\ b(T) = & \begin{bmatrix} 1 \\ \underline{0} \end{bmatrix} \end{cases} \tag{3.9}$$

$$\begin{cases} 0 = & \dot c(t) - \frac{1}{2\lambda} b(t)' \tilde{B} R^{-1} \tilde{B}' b(t) \\ c(T) = & 0 \end{cases} \tag{3.10}$$

where $\underline{\underline{0}}$ is an $n \times n$ matrix of zeros.

Proof. This result follows from standard analysis of linear-quadratic control problem. □

While Proposition 2 characterizes the value function and optimal trading rate for (3.3) which are numerically easy to compute when the utility function $\Phi(x)$ is quadratic, the expressions (3.6)-(3.10) are not particularly insightful. For instance, it is unclear how the presence of market impact penalty in (3.3) modifies the value function and optimal policy relative to the solution of the Merton problem. To clarify these issues we begin by introducing an alternative formulation of the market impact problem which provides us a simple and intuitive relationship between the market impact problem and the Merton problem.

3.2 Illiquid Portfolio Problem: Model 2

In this section, we introduce another formulation of the market impact problem. Though interesting in its own right, this model allows us to relate the market impact model (3.3) to the Merton problem.

Tracking Problem

Let $q(t)$ be an $n \times n$ deterministic positive definite matrix for all $t \in [0, T]$, $\xi(t)$ denote the trading rate, $x(t)$ the investor's wealth, and $\pi(t)$ his position in the risky assets. Consider the problem

$$\begin{cases} \inf_{\xi(\cdot)} \ \mathbb{E}\left\{ \int_0^T \frac{\lambda}{2}\, \xi(t)'\, R\, \xi(t) + \frac{1}{2}(\pi_M^*(t, x(t)) - \pi(t))' q(t) (\pi_M^*(t, x(t)) - \pi(t)) dt \right\} \\ \text{subject to:} \\ \begin{bmatrix} dx(t) \\ d\pi(t) \end{bmatrix} = \left\{ \tilde{A} \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + \tilde{B}\, \xi(t) \right\} dt + \sum_{j=1}^{n} \tilde{C}_j \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} dw_j \\ \xi(\cdot) \in \mathcal{A},\ x(0) = x_0,\ \pi(0) = \pi_0 \end{cases} \tag{3.11}$$

In this problem, the goal is to track the solution

$$\pi_M^*(t,\, x(t)) = \Sigma^{-1}(\mu - r) \left\{ \frac{1}{\eta} \frac{b_M(t)}{a_M(t)} - x \right\}$$

of the Merton problem (2.4) at minimal cost. Observe that the dynamics of this problem are identical to those of the illiquid problem (3.3).

Model II

First we define new state variables

$$\begin{cases} y(t) & \triangleq \dfrac{1}{\eta}\dfrac{b_M(t)}{a_M(t)} - x(t) \\ z(t) & \triangleq \pi_M^*(t, x(t)) - \pi(t). \end{cases} \tag{3.12}$$

Then rewrite the tracking problem (3.11) in terms of the new variables $y(t)$ and $z(t)$ with $q(t) = \eta\, a_M(t)\Sigma$.

$$\begin{cases} \inf_{\xi(\cdot)} \ \mathbb{E}\left\{\displaystyle\int_0^T \frac{\lambda}{2}\,\xi(t)'\,R\,\xi(t) + \frac{\eta a_M(t)}{2}\begin{bmatrix} y(t) \\ z(t) \end{bmatrix}' Q \begin{bmatrix} y(t) \\ z(t) \end{bmatrix} dt\right\} \\ \text{subject to:} \\ \begin{bmatrix} dy(t) \\ dz(t) \end{bmatrix} = \left\{A\begin{bmatrix} y(t) \\ z(t) \end{bmatrix} + B\xi(t)\right\}dt + \displaystyle\sum_{j=1}^n C_j \begin{bmatrix} y(t) \\ z(t) \end{bmatrix} dw_j \\ \xi(\cdot) \in \mathcal{A},\ y(0) = \dfrac{b_M(0)}{\eta a_M(0)} - x_0, \ \ z(0) = \pi_M^*(t, x_0) - \pi_0. \end{cases} \tag{3.13}$$

where

$$Q = \begin{bmatrix} 0 & \underline{0}' \\ \underline{0} & \Sigma \end{bmatrix}, \quad B = \begin{bmatrix} \underline{0}' \\ -I_n \end{bmatrix},$$

$$\begin{cases} A = & \begin{bmatrix} r-\theta & (\mu - r)' \\ A_{21} & A_{22} \end{bmatrix} \\ A_{21} = & \{(r-\theta)I_n - \operatorname{diag}(\mu)\}\Sigma^{-1}(\mu - r) \\ A_{22} = & \Sigma^{-1}(\mu - r)(\mu - r)' + \operatorname{diag}(\mu) \end{cases},$$

$$\begin{cases} C_j = & \begin{bmatrix} -\sigma_{\cdot j}'\Sigma^{-1}(\mu - r) & \sigma_{\cdot j}' \\ C_{21}^j & C_{22}^j \end{bmatrix} \\ C_{21}^j = & -\{\Sigma^{-1}(\mu - r)\sigma_{\cdot j}' + \operatorname{diag}(\sigma_{\cdot j})\}\Sigma^{-1}(\mu - r) \\ C_{22}^j = & \Sigma^{-1}(\mu - r)\sigma_{\cdot j}' + \operatorname{diag}(\sigma_{\cdot j}) \end{cases}$$

Note that the dynamics in (3.13) can be obtained directly by substituting the definition (3.12) into the dynamics in (3.11). The solution of (3.13) is stated in the following Proposition.

Proposition 3. *Suppose that the illiquidity coefficient λ and the investor risk-aversion parameter η are positive, and the matrix R and Σ are positive definite. The value function of*

(3.13) is

$$W(t,y,z) = \frac{1}{2}\begin{bmatrix} y \\ z \end{bmatrix}' \alpha(t) \begin{bmatrix} y \\ z \end{bmatrix}, \tag{3.14}$$

and the optimal trading rates are

$$\xi^*(t,y,z) = -\frac{1}{\lambda} R^{-1} B' \alpha(t) \begin{bmatrix} y \\ z \end{bmatrix}. \tag{3.15}$$

The coefficient $\alpha(t)$ is the solution to the Riccati equation:

$$\begin{cases} 0 = & \dot{\alpha}(t) + \alpha(t)A + A'\alpha(t) + \sum_{j=1}^{n} C_j'\alpha(t)C_j - \frac{1}{\lambda}\alpha(t)'BR^{-1}B'\alpha(t) + \eta a_M(t) Q \\ \alpha(T) = & \underline{\underline{0}} \end{cases} \tag{3.16}$$

Proof. This result follows from standard analysis of linear-quadratic control problem. □

3.3 Relationship between Model I, Model II and the Merton problem

We summarize the relationship between the two formulations (3.3) and (3.13) of portfolio selection problems with market impact costs and the Merton problem (2.4) in the following Theorem.

Theorem 4. *Let $V(t,\,x,\,\pi)$ and $W(t,\,y,\,z)$ denote the value functions for the market impact problems (3.3) and (3.13), where the relationship between $(x,\,\pi)$ and $(y,\,z)$ is stated in (3.12). $V_M(t,\,x)$ denotes the value function for the Merton problem (2.4). Then*

$$V(t,x,\pi) = V_M(t,x) - W\Big(t, \frac{1}{\eta}\frac{b_M(t)}{a_M(t)} - x,\, \pi_M^*(t,\,x) - \pi\Big) \tag{3.17}$$

Moreover, both problems (3.3) and (3.13) have identical optimal trading policies

$$\begin{aligned} \bar{\rho}(t,x,\pi) &= \xi\Big(t, \frac{1}{\eta}\frac{b_M(t)}{a_M(t)} - x,\, \pi_M^*(t,\,x) - \pi\Big) \\ &= -\frac{1}{\lambda}R^{-1}B'\alpha(t) \begin{bmatrix} \frac{1}{\eta}\frac{b_M(t)}{a_M(t)} - x \\ \pi_M^*(t,\,x) - \pi \end{bmatrix}. \end{aligned} \tag{3.18}$$

Proof. See Chapter 6. □

Chapter 4

Solution Approximation

4.1 Asymptotic Expansions

Although Theorem 4 establishes an interesting relationship between the Merton problem and two seemingly different formulations of portfolio selection problems that account for market impact costs, the relationship between the optimal trading rates (3.18) and the solution of the fully liquid Merton problem is not particularly clear. In this section, we utilize multiple time scale perturbation methods to derive an approximation of the Riccati equation (3.16) and the optimal trading rates (3.15) when the illiquidity coefficient λ is small. This analysis establishes an intuitive connection between the optimal trading rate and the value function of the illiquid problem and the frictionless Merton problem.

Intuitively, we expect that the value function $W(t, y, z)$ of the problem (3.13) will go to zero as the trading cost λ becomes small. This motivates us to derive asymptotic expansions of $W(t, y, z)$ in terms of λ. As a first step, we have the following result on the limiting behavior of $\alpha(t)$. We adopt the definition from Kokotovic et al. [40] that a real function $f(t, \epsilon)$ is $O(\epsilon)$ over an interval $[t_1, t_2]$ if there exist positive constants k and ϵ^* such that

$$|f(t,\epsilon)| \leq k\epsilon \quad \forall \epsilon \in [0, \epsilon^*],\ \forall t \in [t_1, t_2]$$

and that $f(t, \epsilon)$ is $o(\epsilon)$ as ϵ approaches ϵ_0 if $\lim_{\epsilon \to \epsilon_0} \frac{|f(t,\epsilon)|}{\epsilon} = 0$. For brevity, we denote that a matrix/vector is $O(\epsilon)$ or $o(\epsilon)$ if every entry in the matrix/vector is $O(\epsilon)$ or $o(\epsilon)$ respectively.

Lemma 5. *The matrix $\alpha(t)$ in (3.16) is $O(\sqrt{\lambda})$ over the interval $[0, T]$.*

Proof. See Chapter 6. □

When λ is small, this result allows us to expand $\alpha(t)$ of the form

$$\alpha(t) = \sqrt{\lambda}\alpha^0(t) + o(\sqrt{\lambda}), \tag{4.1}$$

where $\alpha^0(t)$ is $O(1)$ for all $t \in [0, T]$.

Proposition 6. *$\alpha^0(t)$ in (4.1) is*

$$\alpha^0(t) = \begin{bmatrix} 0 & \underline{0}' \\ \underline{0} & \sqrt{\eta a_M(t)}F(w(t)) \end{bmatrix} \tag{4.2}$$

where

$$w(t) \triangleq \frac{T-t}{\sqrt{\lambda}}\sqrt{\eta a_M(t)}, \tag{4.3}$$

and $F : [0, \infty) \to \mathbb{R}^{n\times n}$ is positive semi-definite and is the solution of the Riccati equation

$$\begin{cases} \dfrac{dF}{dw}(w) = -F(w)'R^{-1}F(w) + \Sigma \\ F(0) = \underline{\underline{0}} \end{cases} \tag{4.4}$$

Proof. First we assume that $\alpha(t)$ depends on the two different time scales which are defined by

$$u = T - t \quad \text{and} \quad v = \frac{T-t}{\sqrt{\lambda}}, \tag{4.5}$$

and write (4.1) as

$$\alpha(u, v) = \sqrt{\lambda}\alpha^0(u, v) + o(\sqrt{\lambda}).$$

Under this re-parameterization, the time derivative becomes

$$\frac{d}{dt}(\cdot) = -\frac{\partial}{\partial u}(\cdot) - \frac{1}{\sqrt{\lambda}}\frac{\partial}{\partial v}(\cdot).$$

and the differential equation (3.16) is transformed into a partial differential equation

$$\begin{cases} \sqrt{\lambda}\dfrac{\partial}{\partial u}\alpha^0(u, v) + \dfrac{\partial}{\partial v}\alpha^0(u, v) \\ = \sqrt{\lambda}\{\alpha^0(u, v)A + A'\alpha^0(u, v) + \displaystyle\sum_{j=1}^{n} C_j\alpha^0(u, v)C_j\} - \alpha^0(u, v)'BR^{-1}B'\alpha^0(u, v) \\ \quad + \eta a_M(t)Q + o(\sqrt{\lambda}) \end{cases}$$

Consider O(1) terms, the equation reduces to

$$\begin{cases} \dfrac{\partial}{\partial v}\alpha^0(u, v) & = -\alpha^0(u, v)'BR^{-1}B'\alpha^0(u, v) + \eta a_M(u)Q \\ \alpha^0(0, 0) & = 0 \end{cases}$$

Observe that $BR^{-1}B'$ and Q contain 0's in the first row and column,

$$BR^{-1}B' = \begin{bmatrix} 0 & \underline{0}' \\ \underline{0} & R^{-1} \end{bmatrix}, \quad Q = \begin{bmatrix} 0 & \underline{0}' \\ \underline{0} & \Sigma \end{bmatrix}$$

Hence starting at the initial condition $\alpha^0(0,0) = 0$, the first row and first column of $\alpha^0(u,v)$ remain zero for all positive u and v. $\alpha^0(u,v)$ can be written by

$$\alpha^0(u,v) = \begin{bmatrix} 0 & \underline{0}' \\ \underline{0} & \beta^0(u,v) \end{bmatrix}.$$

where $\beta^0(u,v)$ satisfies the following equation

$$\begin{cases} \frac{\partial}{\partial v}\beta^0(u,v) & = -\beta^0(u,v)'R^{-1}\beta^0(u,v) + \eta a_M(u)\Sigma \\ \beta^0(0,0) & = 0 \end{cases} \tag{4.6}$$

We can see that

$$\beta^0(u,\, v) = \sqrt{\eta a_M(u)} F\left(v\sqrt{\eta a_M(u)}\right)$$

is the solution of (4.6) where $F(v\sqrt{\eta a_M(u)})$ is the solution of the Riccati equation (4.4). Observe that $F(w)$ is positive semi-definite due to the fact that it is a Riccati equation and can be associated with the value function of a minimum cost optimal control problem. □

The following theorem gives asymptotic expansions of the value function $V(t,x,\pi)$ and $W(t,y,z)$ as well as the optimal trading policies for the problems (3.3) and (3.13).

Theorem 7. *With small λ, the value function of the market impact problem (3.3) satisfies*

$$\begin{aligned} V(t,x,\pi) =& V_M(t,x) - \frac{\sqrt{\lambda\eta a_M(t)}}{2}(\pi^*_M(t,x)-\pi)'F\Big(\frac{T-t}{\sqrt{\lambda}}\sqrt{\eta a_M(t)}\Big)(\pi^*_M(t,x)-\pi) \\ &+ o(\sqrt{\lambda}), \end{aligned}$$

and the optimal trading policy is

$$\bar{\rho}^*(t,x,\pi) = \frac{\sqrt{\eta a_M(t)}}{\sqrt{\lambda}} R^{-1} F\Big(\frac{T-t}{\sqrt{\lambda}}\sqrt{\eta a_M(t)}\Big)(\pi^*_M(t,x)-\pi) + o(1/\sqrt{\lambda}),$$

where $F(w)$ solves (4.4).

Proof. The expressions above are directly obtained by substituting the approximation (4.1)-(4.2) into (3.17)-(3.18). □

Theorem 7 establishes a close relationship between the portfolio selection problem in an illiquid market and the Merton problem. Several properties of the value function and the optimal trading rates are worth noting:

- $F(w)$ is a Riccati equation, and it can be shown that F is uniformly bounded and non-negative definite for every $w \in [0, \infty)$. If follows that the term
$$F\Big(\frac{T-t}{\sqrt{\lambda}}\sqrt{\eta a_M(t)}\Big)$$
is also uniformly bounded for every $\lambda > 0$ and $t < T$.
- As the illiquidity coefficient λ decreases the value function V increases, and when λ goes to zero, V converges to the value function of the Merton problem.
- As λ decreases, the optimal trading rate $\bar{\rho}^*$ increases and becomes arbitrarily large as λ goes to zero. Intuitively, the investor trades infinitely quickly when trading costs vanish, which allows him/her to track the optimal Merton portfolio with vanishing error.
- An investor trades to close the gap between his/her current positions π and the (ideal) positions π_M^* he/she would adopt if there were no liquidity costs.
- The value function V decreases in the distance $(\pi_M^*(t,x) - \pi)' F(w)(\pi_M^*(t,x) - \pi)$.

Single Risky Asset Case

In case of single risky asset, an explicit solution for $F(w)$ in (4.4) can be found. We summarize the findings as follows.

Theorem 8. *With small λ, the value function of the market impact problem (3.3) satisfies*

$$V(t,x,\pi) = V_M(t,x) - \sqrt{\lambda}\,\frac{\sigma\sqrt{\eta a_M(t)}}{2}\tanh\Big(\sigma\sqrt{\eta a_M(t)}\frac{T-t}{\sqrt{\lambda}}\Big)\Big\{\pi_M^*(t,x) - \pi\Big\}^2 + o(\sqrt{\lambda}). \tag{4.7}$$

The optimal trading policy is

$$\bar{\rho}^*(t,x,\pi) = \frac{\sigma\sqrt{\eta a_M(t)}}{\sqrt{\lambda}}\tanh\Big(\sigma\sqrt{\eta a_M(t)}\frac{T-t}{\sqrt{\lambda}}\Big)\Big\{\pi_M^*(t,x) - \pi\Big\} + o(1/\sqrt{\lambda}). \tag{4.8}$$

Let $x(t)$ and $\pi(t)$ be wealth and risky asset holding processes of an investor who adopts the trading rate (4.8). The difference between the holding $\pi(t)$ and that of the Merton investor is

$$E\left\{\int_0^T \frac{1}{2}\Big[\pi_M^*\Big(t,\, x(t)\Big) - \pi(t)\Big]^2 dt\right\} = \sqrt{\lambda}\frac{K}{2}(\pi_M^*(0,x(0)) - \pi(0))^2 + o(\sqrt{\lambda}) \tag{4.9}$$

where the coefficient

$$K = \frac{1}{2\sigma\sqrt{\eta a_M(0)}} \tanh(\sigma\sqrt{\eta a_M(0)}\frac{T}{\sqrt{\lambda}}).$$

Proof. See Chapter 6. □

In the single risky asset case, Theorem 8 provides additional insight into the relationship between the market impact problem and the Merton problem as follows.

- The liquidity time scale $(T-t)/\sqrt{\lambda}$ affects the optimal trading rate $\bar{\rho}^*$ and the value function only when the remaining time is comparable to $\sqrt{\lambda}$, since the hyperbolic tangent term will be significant if $T - t/\sqrt{\lambda}$ is small. When the liquidity time is small an investor will reduce his/her trading rate since market impact costs dominate potential improvements in utility from rebalalancing. In contrast to the solution of the classical Merton problem (Proposition 1) the investor needs to consider both the the liquidity as well as the original time scale when there are market impact costs.
- A larger volatility σ makes the optimal trade rate $\bar{\rho}^*$ increase. When the volatility is large, an investor trades more aggressively to close the gap between his/her position and the ideal position π_M^*.
- A larger risk-aversion parameter causes an investor to trade faster.

4.2 Accuracy

In this section we compare the numerical solutions of the coefficient $\alpha(t)$ in (3.16) and its approximation $\sqrt{\lambda}\alpha^0(t)$ in (4.1). To illustrate the solutions, we consider a market with one risk-free asset and three risky assets. The risk-free return r is 2% per year, the risky asset expected returns μ and diffusion matrix σ are

$$\mu = \begin{bmatrix} 0.06 \\ 0.08 \\ 0.06 \end{bmatrix}, \quad \sigma = \begin{bmatrix} 0.200 & 0 & 0 \\ 0.224 & 0.168 & 0 \\ 0.040 & 0.013 & 0.196 \end{bmatrix}. \tag{4.10}$$

The investor risk aversion factor η is 6.67×10^{-8}, chosen in according to \$10M initial wealth. The illiquidity coefficients λ are varied from is 1×10^{-12} to 1×10^{-8}, and the cost matrix R is I_3. The time horizon of the problem is one year (250 trading days).

Figure 4.1 compares the entries in $\alpha(t)$, solid lines, and $\beta(t) \triangleq \sqrt{\lambda}\alpha^0(t)$, dotted lines. We can see that the approximation error decreases as the illiquidity coefficient decreases.

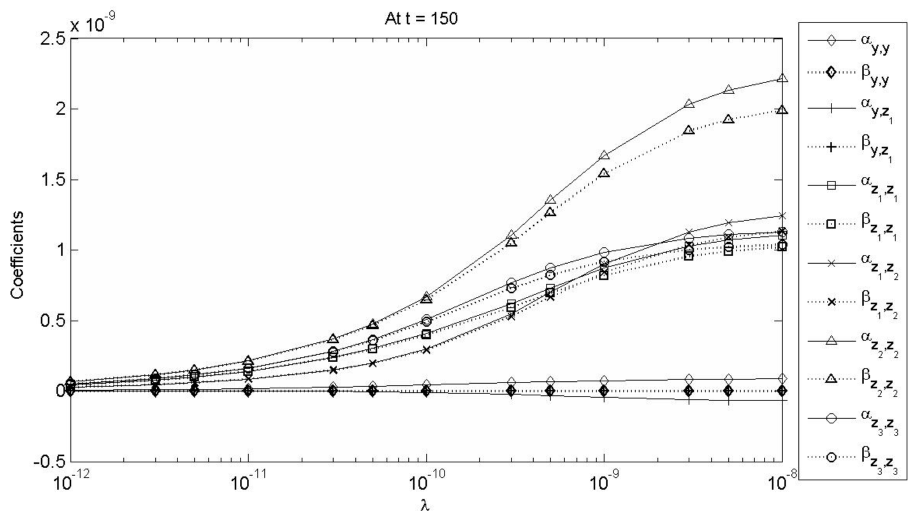

Figure 4.1: This figure compares the numerical solutions of the coefficient $\alpha(t)$ calculated from (3.16) and its approximation $\beta(t) \triangleq \sqrt{\lambda}\,\alpha^0(t)$ from Proposition 6. The coefficients are calculated at $t = 150$.

Chapter 5

Examples

In this chapter, we test the performance of the optimal solution of (3.3) on a temporary price impact model by Almgren et al. [4]. The price impact model by Almgren et al. is well-known among both researchers and practitioners as being mentioned in many places such as [2, 39, 30, 25]. Even though this price impact model is fairly straightforward, incorporating this impact into dynamic portfolio choice problem causes the problem impossible to solve. Instead of trying to solve the portfolio choice problem under the Almgren et al.'s price impact, we take another approach by adopting the solution to the proposed problem (3.3). Then we compare the performance of our approach to those of the cases where the market is perfectly liquid and where the price impact is ignore. We begin the chapter by describing the market impact model following by a description of the trading strategies which we will be testing.

Financial market with market impact costs

Throughout this chapter, we assume that investors re-balance daily over the course of one year (250 trading days). We denote re-balancing time and re-balancing interval by τ and Δ (in years) respectively. The risk-free asset follows the dynamics (2.2) and we assume a risk-free rate of $r = 2\%$ per year. For the risky assets, we assume a temporary impact model that builds on the model in Almgren et al. [4]. Specifically, the price of risky asset i is described using two components, an *observed price* $s_i(\tau)$ and an *execution price* $s_i^{exec}(\tau)$. Observed price follows geometric Brownian motion

$$ds_i(t) = \mu_i s_i(t)dt + \sum_{j=1}^{n} \sigma_{ij} s_i(t) dw_j(t),$$

while the execution price is given by

$$s_i^{exec}(\tau) = s_i(\tau)(1 + J_i(\rho_i(\tau))),$$

where the function

$$J_i(\rho_i(\tau)) = c_i \ \text{sgn}(\rho_i(\tau)) \left(\frac{\rho_i(\tau)\Delta}{V_i}\right)^{0.6} \tag{5.1}$$

is the *relative price impact*. Here, c_i and V_i are constants representing market depth and average daily volume of each asset, respectively. Observe that the sign of $J_i(\rho_i(\tau))$ depends on trading direction and its magnitude is proportional to the rate of trading to the power of 0.6 (which is the value fitted in [4] using data). Intuitively, $s_i^{exec}(\tau)$ is the average execution price of an order of size $\rho_i(\tau)\Delta$ submitted after observing a price of $s_i(\tau)$. Execution price $s_i^{exec}(\tau)$ is larger than $s_i(\tau)$ when buying an asset and smaller when selling and can be viewed as the cost of moving through the order book in order to execute a block trade.

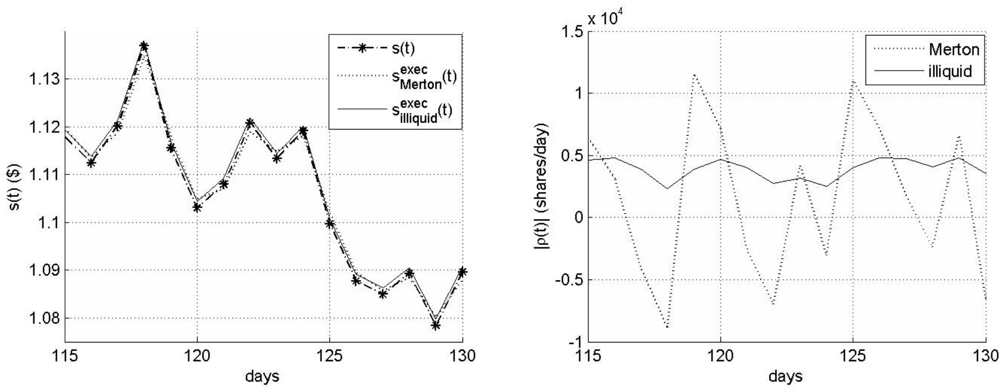

Figure 5.1: This figure illustrates the price impact. The left plot shows how execution prices deviate from the observed prices. The deviation is caused by the trades shown on the right plot.

Investors

We consider four investor types:

- the *Illiquid Investor*, who solves (3.3) for the optimal solution $\rho^*(t)$ (3.7). Conditional on the time τ prices $s(\tau)$ the illiquid investor trades quantities $\rho^*(\tau)\Delta$ at prices $s^{exec}(\tau)$ determined by the trading rates $\rho^*(\tau)$;

- the *Merton Investor* who at time τ, observes the prices $s(\tau)$ and chooses trading rates $\rho_M(\tau)$ such that his holdings after the trade coincide with the optimal holdings for (2.7) of the Merton problem: $\rho_M(\tau)\Delta = \pi_M^*(\tau) - \pi_M^*(\tau^-)$. The trade is executed at the prices $s^{exec}(\tau)$ determined by the trading rates $\rho_M(\tau)$;

- the *Riskless Investor*, who only holds the riskless asset, and

- the *Ideal Investor*, who like the Merton investor, re-balances to the Merton optimal holdings (2.7) each day, but unlike the Merton investor, is able to do so at the *observed prices* $s(\tau)$.

5.1 Examples: Single Risky Asset

First we consider the case where there is only one risky asset in the market. We assume the expected return and volatility of the risky asset to be $\mu = 6\%$ and $\sigma = 25\%$ per year, and the initial price to be one dollar. The average daily volume V is chosen such that the Merton investor trades on average 10% of daily volume.

We now compare the performance and trading behavior of these investors, as well as the sensitivity of the illiquid investor's performance to the choice of illiquidity coefficient λ.

Performance comparison

The parameters in the price impact model are consistent with the empirical analysis in Almgren et al. [4], namely the execution price for a trade as large as 10% of daily volume deviates by 0.2% from the observed price. The value of illiquidity coefficient λ used by illiquid investor is 10^{-10}.

Figure 5.2 show the portfolio return mean-standard deviation frontiers for the illiquid, Merton, and ideal investors, which are obtained by varying the risk aversion parameter η for all investors. The ideal investor's frontier can be regarded as a performance upper bound for the investors trading under the price impact.

When the risk aversion parameter is large (low portfolio return and standard deviation), the performance of all three investors' are roughly the same due to small exposure to risky asset, which leads to small amount of trade and hence small impact. On the other hand, for decreasing risk aversion (higher portfolio return and standard deviation), trading and the risky asset holding increase and the cost of price impact on P&L becomes substantial, and the frontier of both the illiquid and Merton investors are pushed to the right. We can clearly see, however, that the illiquid investor's frontier always dominates that of the Merton investor and that the deviation of the illiquid investor's frontier from that of the idealized investor is relatively small over all risk aversion levels.

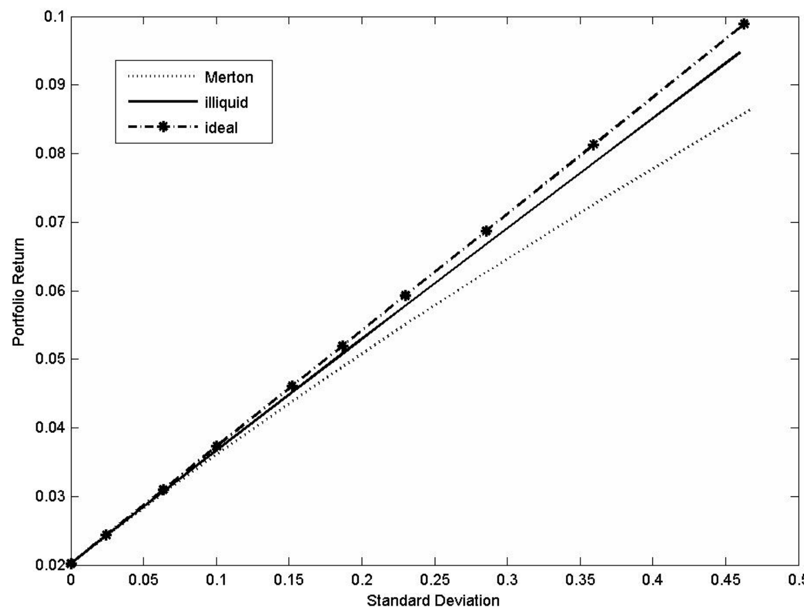

Figure 5.2: Mean-standard deviation frontiers of the Merton investor, illiquid investor and ideal investor. The illiquidity coefficient λ used by illiquid investor is 10^{-10}, and the relative price impact J is 0.2% if an investor trades 10% of the daily volume.

Figure 5.3 shows the histogram of the difference in final wealth of the Merton and illiquid investors, $x_{illiquid} - x_{Merton}$ when risk aversion η is 3.5×10^{-7}, which clearly shows that the illiquid investor outperforms the Merton investor on most occasions, and that the out performance is more substantial than the under performance.

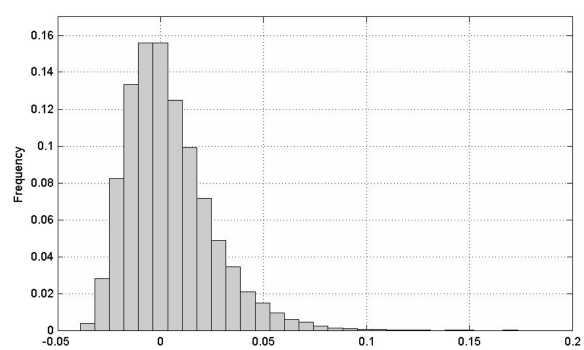

Figure 5.3: This figure shows the histogram of the different in final wealth of the Merton and illiquid investors. The illiquidity coefficient λ used by illiquid investor is 10^{-10}, and the relative price impact J is 0.2% if an investor trades 10% of the daily volume.

Trading Behavior

We now illustrate the major differences between the trading rate and risky asset holdings of the Merton and illiquid investors. To show these differences, we plot the sample paths of risky asset holdings, trading rates, risky asset price and the histogram of trading rate in Figure 5.4. The relative price impact, the illiquidity coefficient λ and the risk aversion η remain the same as in the previous example, namely J is 0.2% if an investor trades 10% of the daily volume, λ equals to 10^{-10} and η equals to 3.5×10^{-7}.

Consistent with our asymptotic result (Theorem 4), Figure 5.4 shows that the holdings of the illiquid investor tracks the holdings of the Merton investor. The third and fourth plots shows that the optimal trading rate of the illiquid investor is substantially smaller than that of the Merton investor.

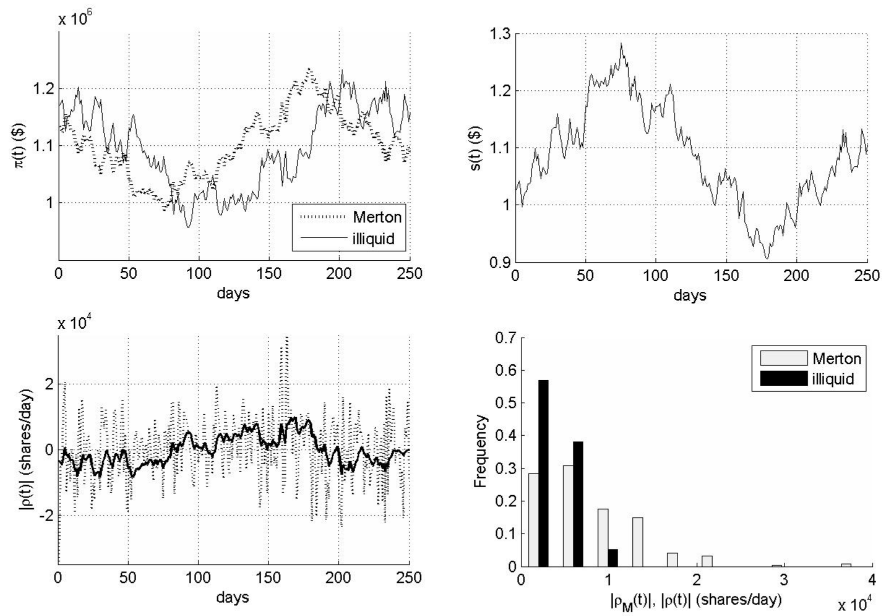

Figure 5.4: This figure compares the sample paths of the Merton and illiquid investors. The illiquidity coefficient λ used by illiquid investor is 10^{-10}, and the relative price impact J is 0.2% if an investor trades 10% of the daily volume.

Robustness

In this section, we examine the sensitivity of the performance of our model to the illiquidity coefficient λ and the price impact J. First, under a specific level of price impact, we show that implementing different values of λ's does not change the performance of the illiquid investor portfolio much. Second, under different levels of price impact, we show that a single value of λ selected from a particular range can perform considerably well compared to the Merton investor. Throughout this example, the risk aversion parameter η is set to 3.5×10^{-7}.

We vary the illiquidity coefficient λ from 5×10^{-12} to 5×10^{-8} while keeping the relative price impact J at 0.2%. Figure 5.5 plots the utilities of the illiquid, Merton, and riskless investors. While the optimal utility for the illiquid investor is achieved at $\lambda = 8 \times 10^{-10}$, a more important observation is that the utility curve is fairly flat around the optimal value. This suggests that illiquid investor performs well for a relatively large range of values for the illiquidity coefficient, and that one need not be concerned (at least for this example) that small changes in λ will lead to a large drop in performance. Indeed, this example suggests that being conservative and choosing a large value of λ will not hurt utility too much. From the second plot in Figure 5.5 it is interesting to note that while the utility of the illiquid investor does not change substantially over the range of λ being considered, his trading rate does. Intuitively, the savings from reduced market impact compensate for the loss in utility from slower trading.

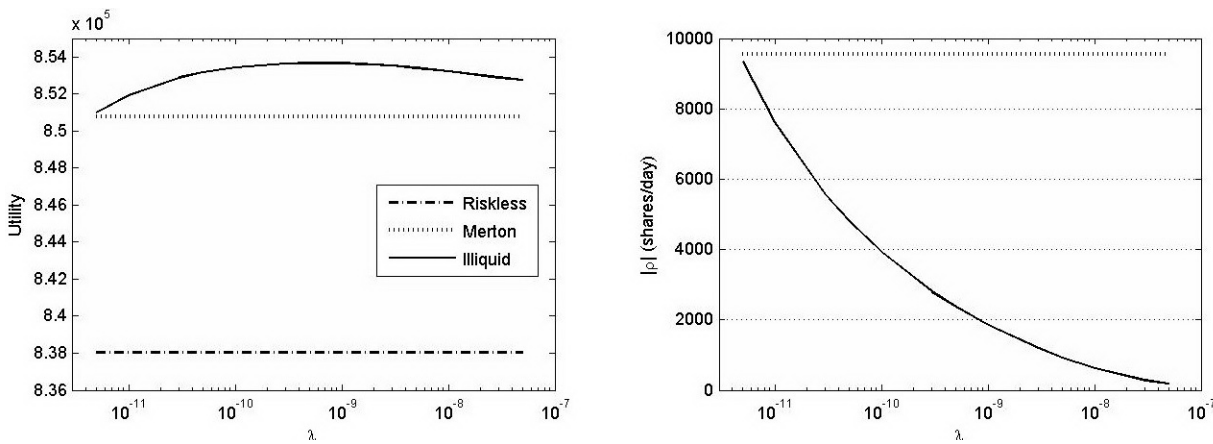

Figure 5.5: This figure shows the utility and the trading rates of the illiquid investor with different values of λ's. The relative price impact J is 0.2% if an investor trades 10% of the daily volume. The average utility of the Merton investor and riskless investor is added for comparison. It is interesting to note that while utility does not change substantially over the range of λ being considered, the trading rate changes significantly.

Next, we vary the relative price impact J between 0.02% to 1% at 10% of daily trading (note that the empirical work by [4] suggests J is around 0.2%), and plot the average utility of the illiquid, Merton and riskless investors. To demonstrate the effect of using different values of illiquid coefficients, we consider four illiquid investors with λ's equal to $5 \times 10^{-11}, 10^{-10}, 5 \times 10^{-10}$ and 1×10^{-9}. From the plot in Figure 5.6, we can see that performance of the illiquid investors (i.e. choices of λ) is comparable to that of the Merton investor when the impact is small, is considerably better when the impact is moderate to high, and is less sensitive to changes in price impact. We note that an impact of 0.12%[1] is approximately equal to paying the spread, so the Merton investor only outperforms the illiquid investors in a highly liquid market where buying and selling occur at almost the same price.

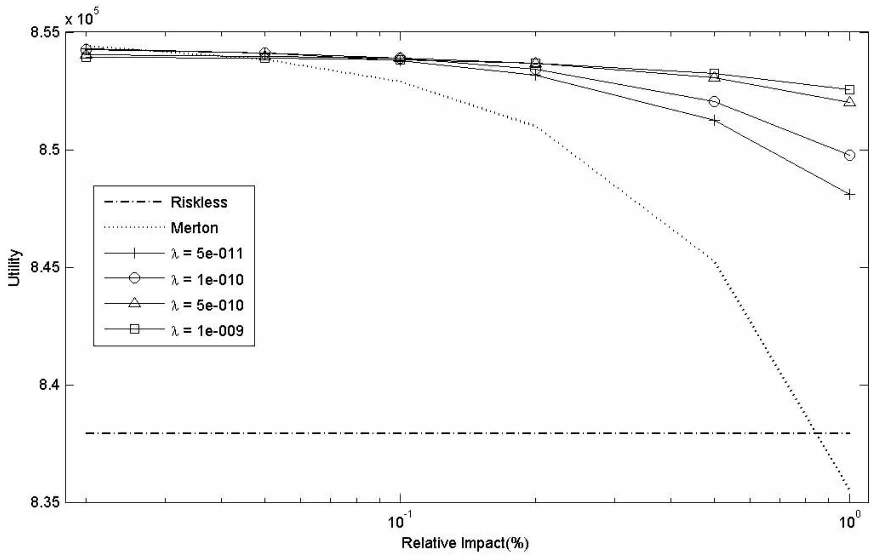

Figure 5.6: This figure compares the average utility of four illiquid investors with four different λ's under various illiquidity situations. The average utility of the Merton and riskless investor is added for comparison.

5.2 Examples: Multiple Risky Assets

We provide similar examples as above except that now there are multiple risky assets. We use the same expected returns and the diffusion matrix as shown in (4.10),

$$\mu = \begin{bmatrix} 0.06 \\ 0.08 \\ 0.06 \end{bmatrix}, \quad \sigma = \begin{bmatrix} 0.200 & 0 & 0 \\ 0.224 & 0.168 & 0 \\ 0.040 & 0.013 & 0.196 \end{bmatrix}. \tag{5.2}$$

We choose the initial prices to be one dollar, choose V_i's to be the same for all assets at 2M shares per day, choose the matrix R to be identity matrix, and let initial wealth x_0 be 10M for every investor. For asset i, the parameter c_i is chosen in consistent with the empirical analysis in Almgren et al. [4], namely, that the execution price for a trade as large as 10% of daily volume deviates by 0.2% from the observed price.

[1]The average spread of common stocks listed on the NYSE is 0.24%. For the stocks those are in the top five percent of dollar trading volume, the average spread comes down to 0.075%. The data is collected from 2001 to 2005 by Kyle and Obizhaeva [41].

Performance comparison

In this example, the illiquidity coefficient λ used by illiquid investor in his model (3.3) is 5×10^{-11}. Similar to Figure 5.2, Figure 5.7 shows the mean-standard deviation frontiers for the illiquid, Merton, and ideal investors. We can see that as the number of assets grows, the Merton investor becomes worse as he/she has more assets to re-balance and hence suffers more from price impacts. On the other hand, the illiquid investor continues to perform well. The mean-standard deviation frontier of the illiquid investor remains close to that of the ideal investor over all risk aversion levels.

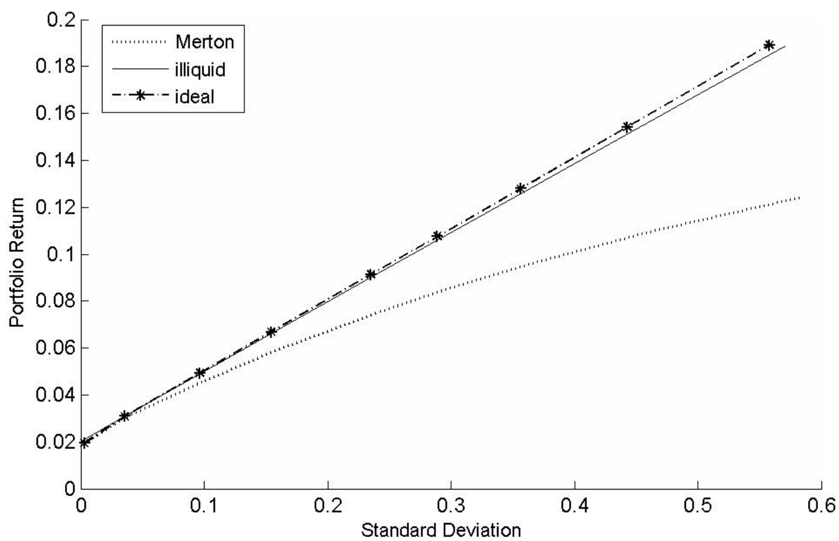

Figure 5.7: Mean-standard deviation frontiers of the Merton investor, illiquid investor and ideal investor.

Trading Behavior

Figure 5.8 shows the plot of the sample paths of risky asset holdings, trading rates and risky asset prices. The illiquidity coefficient λ and the risk aversion η equal to 5×10^{-11} and 6.67×10^{-8} respectively. We can still see that the trading rates of the illiquid investor remain substantially lower than those of the Merton investor. We can also see that the holdings of the illiquid investor track those of the Merton investor when there multiple risky assets.

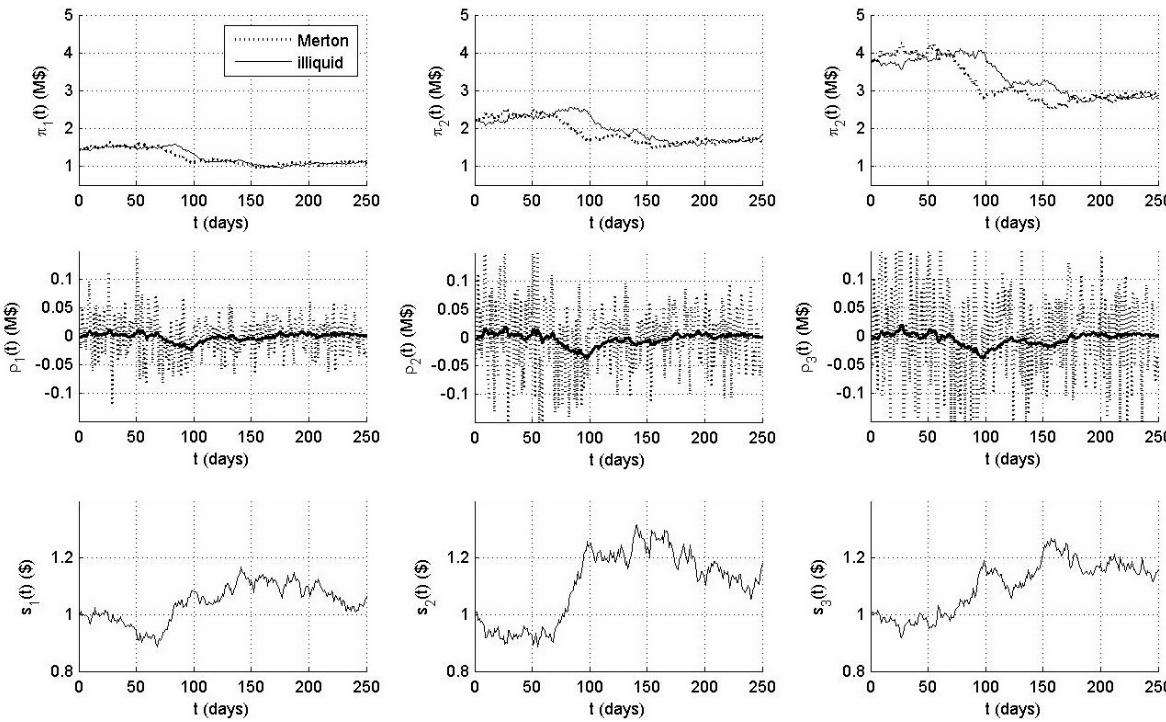

Figure 5.8: This figure compares the sample paths of risky asset holdings and trading rates of the Merton and illiquid investors. The illiquidity coefficient λ used by illiquid investor is 5×10^{-11}.

Robustness

The robustness of the proposed strategy (3.7) remains the same when there are multiple risky assets. In Figure 5.9, we vary the illiquidity coefficient λ from 10^{-12} to 10^{-8} while keeping the other parameters the same. The utility curve is fairly flat around the optimal value which suggests that illiquid investor can perform well for a relatively large range of values for the illiquidity coefficient.

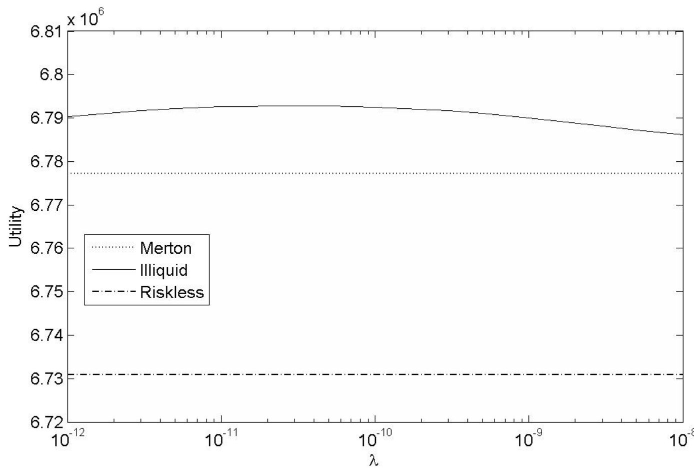

Figure 5.9: This figure shows the utility of illiquid investor using different values of λ's. The average utility of the Merton investor and riskless investor is added for comparison. It is interesting to note that illiquid investor's utility does not change substantially over the range of λ being considered.

Chapter 6

Proofs

6.1 Proof of Theorem 4

First we define the difference in value of the Merton problem (2.4) and the illiquid problem (3.3)

$$\tilde{W}(t,x,\pi) \triangleq V_M(t,x) - V(t,x,\pi).$$

From the quadratic structure of $V_M(t,x)$ in (2.6) and $V(t,x,\pi)$ in (3.6), $\tilde{W}(t,x,\pi)$ can be written down by

$$\tilde{W}(t,x,\pi) = -\frac{1}{2}\begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix}' \tilde{\alpha}(t) \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + \tilde{\beta}(t)' \begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + \tilde{\gamma}(t),$$

where $\tilde{\alpha}(t), \tilde{\beta}(t)$ and $\tilde{\gamma}(t)$ satisfy the following differential equations.

$$\begin{cases} 0 = & \dot{\tilde{\alpha}}(t) + \tilde{\alpha}(t)\tilde{A} + \tilde{A}'\tilde{\alpha}(t) + \sum_{j=1}^{n} \tilde{C}_j'\tilde{\alpha}(t)\tilde{C}_j - \frac{1}{\lambda}\tilde{\alpha}(t)'\tilde{B}R^{-1}\tilde{B}'\tilde{\alpha}(t) \\ & +\eta a_M(t)\begin{bmatrix} \theta & -(\mu-r)' \\ -(\mu-r) & -\Sigma \end{bmatrix} \\ \tilde{\alpha}(T) = & \underline{\underline{0}} \end{cases}$$

$$\begin{cases} 0 = & \dot{\tilde{\beta}}(t) + \tilde{A}'\tilde{\beta}(t) - \frac{1}{\lambda}\tilde{\alpha}(t)'\tilde{B}R^{-1}\tilde{B}'\tilde{\beta}(t) - b_M(t)\begin{bmatrix} \theta \\ \mu - r \end{bmatrix} \\ \tilde{\beta}(T) = & \underline{0} \end{cases}$$

$$\begin{cases} 0 = & \dot{\tilde{\gamma}}(t) - \frac{1}{2\lambda}\tilde{\beta}(t)'\tilde{B}R^{-1}\tilde{B}'\tilde{\beta}(t) + \frac{\theta}{2}\frac{b_M(t)^2}{\eta a_M(t)} \\ \tilde{\gamma}(T) = & 0 \end{cases}$$

It can be shown that $\tilde{W}(t,x,\pi)$ is the value function of the linear-quadratic control problem

$$\begin{cases} \inf_{\xi(\cdot)} \; \mathbb{E}\left\{\int_0^T \frac{\lambda}{2}\,\xi(t)'\,R\,\xi(t) + \frac{\eta a_M(t)}{2}\begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix}'\begin{bmatrix} \theta & -(\mu-r)' \\ -(\mu-r) & -\Sigma \end{bmatrix}\begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} \right. \\ \qquad\qquad \left. -b_M(t)\begin{bmatrix} \theta \\ \mu - r \end{bmatrix}'\begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + \frac{\theta}{2}\frac{b_M(t)^2}{\eta a_M(t)}\right\} \\ \text{subject to:} \\ \begin{bmatrix} dx(t) \\ d\pi(t) \end{bmatrix} = \left\{\tilde{A}\begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix} + \tilde{B}\,\xi(t)\right\}dt + \sum_{j=1}^n \tilde{C}_j\begin{bmatrix} x(t) \\ \pi(t) \end{bmatrix}dw_j \\ \xi(\cdot) \in \mathcal{A},\; x(0) = x_0,\; \pi(0) = \pi_0 \end{cases} \tag{6.1}$$

Now it is straightforward to see that the control problem (6.1) is exactly the same as the tracking problem (3.13) by the definition (3.12). Then the relationship (3.17) holds due to the fact that that the value functions of (6.1) and (3.13) are equal,

$$\tilde{W}(t,x,\pi) = W\Big(t, \frac{1}{\eta}\frac{b_M(t)}{a_M(t)} - x,\, \pi_M^*(t,\,x) - \pi\Big) = V_M(t,x) - V(t,x,\pi).$$

The relationship between the optimal trading rates (3.7) and (3.18) follows by similar analysis.

6.2 Proof of Lemma 5

To show that $\alpha(t)$ in (3.16) is $O(\sqrt{\lambda})$, we first construct an upper bound of the cost $W(t,y,z)$ of the tracking problem (3.13). The upper bound is the expected cost of a stochastic system which is obtained by modifying (3.13) as follows

- replace $a_M(t)$ in the objective function by $\bar{a}_M \geq a_M(t)$,
- adopt the suboptimal control policy

$$\bar{\xi}(t,z,y) = \frac{1}{\sqrt{\lambda}}R^{-1}\begin{bmatrix} 0 & \bar{\beta} \end{bmatrix}\begin{bmatrix} y(t) \\ z(t) \end{bmatrix} \tag{6.2}$$

where $\bar{\beta}$ is an $n \times n$ full rank $O(1)$ matrix.

The stochastic system is given by

$$\begin{cases} \mathbb{E}\left\{ \displaystyle\int_0^T \frac{\lambda}{2}\,\bar{\xi}(t)'\,R\,\bar{\xi}(t) + \frac{\eta \bar{a}_M}{2} \begin{bmatrix} y(t) \\ z(t) \end{bmatrix}' Q \begin{bmatrix} y(t) \\ z(t) \end{bmatrix} dt \right\} \\ \text{subject to:} \\ \begin{bmatrix} dy(t) \\ dz(t) \end{bmatrix} = \left\{ A \begin{bmatrix} y(t) \\ z(t) \end{bmatrix} + B\xi(t) \right\} dt + \displaystyle\sum_{j=1}^{n} C_j \begin{bmatrix} y(t) \\ z(t) \end{bmatrix} dw_j \\ y(0) = \dfrac{b_M(0)}{\eta a_M(0)} - x_0, \quad z(0) = \pi_M^*(t, x_0) - \pi_0. \end{cases} \tag{6.3}$$

Denoting the expected cost of (6.3)

$$\bar{W}(t,y,z) \triangleq \mathbb{E}\left\{ \int_t^T \frac{\lambda}{2}\,\bar{\xi}(s)'\,R\,\bar{\xi}(s) + \frac{\eta \bar{a}_M}{2} \begin{bmatrix} y(s) \\ z(s) \end{bmatrix}' Q \begin{bmatrix} y(s) \\ z(s) \end{bmatrix} ds \;\middle|\; y(t) = y, z(t) = z \right\},$$

it follows from the outlined construction that

$$\bar{W}(t,y,z) \geq W(t,y,z), \quad \forall t \in [0,T] \text{ and } (y,z) \in \mathbb{R}^{n+1}. \tag{6.4}$$

The linear-quadratic structure of (6.3) implies that $\bar{W}(t,y,z)$ is quadratic in (y,z)

$$\bar{W}(t,y,z) = \frac{1}{2} \begin{bmatrix} y \\ z \end{bmatrix}' M(t) \begin{bmatrix} y \\ z \end{bmatrix}$$

and the symmetric matrix $M(t) = \begin{bmatrix} M_{11}(t) & M'_{21}(t) \\ M_{21}(t) & M_{22}(t) \end{bmatrix}$ satisfies the linear differential equations

$$\begin{cases} -\dot{M}_{11}(t) = & 2M_{11}(t)A_{11} + 2M'_{21}(t)A_{21} \\ & +\sum_{j=1}^{n} \left\{ C^j_{11}M_{11}(t)C^j_{11} + 2\,C^j_{11}M'_{21}(t)C^j_{21} + C^{j\prime}_{21}M_{22}(t)C^j_{21} \right\} \\ M_{11}(T) = & 0 \end{cases} \tag{6.5}$$

$$\begin{cases} -\sqrt{\lambda}\dot{M}_{21}(t) = & \sqrt{\lambda}\big[A_{11}M_{21}(t) + M_{22}(t)A_{21} + A_{12}M_{11}(t) + A'_{22}M_{21}(t) \\ & +\sum_{j=1}^{n} \left\{ C^j_{12}M_{11}(t)C^j_{11} + C^{j\prime}_{22}M_{21}(t)C^j_{11} + C^j_{12}M'_{21}(t)C^j_{21} + C^{j\prime}_{22}M_{22}(t)C^j_{21} \right\} \big] \\ & -\bar{\beta}' R^{-1} M_{21}(t) \\ M_{21}(T) = & \underline{0} \end{cases} \tag{6.6}$$

$$\begin{cases} -\sqrt{\lambda}\dot{M}_{22}(t) = & \sqrt{\lambda}\big[M_{21}(t)A'_{12} + M_{22}(t)A_{22} + A_{12}M_{21}(t)' + A'_{22}M_{22}(t) \\ & +\bar{\beta}'R^{-1}\bar{\beta} + \eta\bar{a}_M\Sigma \\ & +\sum_{j=1}^{n} \left\{ C^j_{12}M_{11}(t)C^{j\prime}_{12} + C^{j\prime}_{22}M_{21}(t)C^{j\prime}_{12} + C^j_{12}M_{21}(t)C^j_{22} + C^{j\prime}_{22}M_{22}(t)C^j_{22} \right\} \big] \\ & -M_{22}(t)\,\bar{\beta}'\,R^{-1} - R^{-1}\,\bar{\beta}\,M_{22}(t) \\ M_{22}(T) = & \underline{\underline{0}}. \end{cases} \tag{6.7}$$

(We write an arbitrary $(n+1) \times (n+1)$ matrix using block notation $X = \begin{bmatrix} X_{11} & X'_{12} \\ X_{21} & X_{22} \end{bmatrix}$ where X_{11} is the $(1,1)$ entry of X, X_{12} is a vector containing the second to the $(n+1)^{th}$ entries of the first row, X_{21} is a vector containing the second to the $(n+1)^{th}$ entries of the first column, and X_{22} is an $n \times n$ matrix containing the rest.)

Observe that the system (6.5) - (6.7) is a standard singular perturbation model with small parameter $\sqrt{\lambda}$. We will show that $M_{11}(t), M_{21}(t)$ and $M_{22}(t)$ are $O(\sqrt{\lambda})$ for $t \in [0, T]$ by following standard analysis (e.g. [40]). First we introduce two time scales, the original (slow) time scale $u \triangleq T - t$ and the fast time scale $v \triangleq \dfrac{T-t}{\sqrt{\lambda}}$. The slow responses, $\bar{M}_{11}(u), \bar{M}_{21}(u)$

and $\bar{M}_{22}(u)$, solve the reduced-order model (the small parameter is set to be zero):

$$\begin{cases} \frac{\partial}{\partial u}\bar{M}_{11}(u) = & 2\bar{M}_{11}(u)A_{11} + 2\bar{M}_{21}'(u)A_{21} \\ & + \sum_{j=1}^{n} \left\{ C_{11}^{j}\bar{M}_{11}(u)C_{11}^{j} + 2\, C_{11}^{j}\bar{M}_{21}'(u)C_{21}^{j} + C_{21}^{j\prime}\bar{M}_{22}(u)C_{21}^{j} \right\} \\ \bar{M}_{11}(0) = & 0 \end{cases} \tag{6.8}$$

$$\underline{0} = -\bar{\beta}' R^{-1}\bar{M}_{21}(u) \tag{6.9}$$

$$\underline{\underline{0}} = -\bar{M}_{22}(t)\,\bar{\beta}'\,R^{-1} - R^{-1}\,\bar{\beta}\,\bar{M}_{22}(t) \tag{6.10}$$

The fast responses, $\hat{M}_{21}(v)$ and $\hat{M}_{22}(v)$, satisfy the boundary layer system:

$$\begin{cases} -\frac{\partial}{\partial v}\hat{M}_{21}(v) = & -\bar{\beta}' R^{-1}(\bar{M}_{21}(0) + \hat{M}_{21}(v)) \\ \hat{M}_{21}(0) = & -\bar{M}_{21}(0) \end{cases} \tag{6.11}$$

$$\begin{cases} -\frac{\partial}{\partial v}\hat{M}_{22}(v) = & -(\bar{M}_{22}(0) + \hat{M}_{22}(v))\,\bar{\beta}'\,R^{-1} - R^{-1}\,\bar{\beta}\,(\bar{M}_{22}(0) + \hat{M}_{22}(v)) \\ \hat{M}_{22}(0) = & -\bar{M}_{22}(0). \end{cases} \tag{6.12}$$

Under an assumption that $\bar{\beta}$ is full-rank, the solutions to the reduced-order and boundary layer systems (6.8) - (6.12) are that every entry of $\bar{M}_{11}(u)$, $\bar{M}_{21}(u)$, $\bar{M}_{22}(u)$, $\hat{M}_{21}(v)$ and $\hat{M}_{22}(v)$ are zero. It follows from Theorem 3.1 of [40] that the solutions of (6.5) - (6.7) can be written as

$$\begin{aligned} M_{11}(t) &= \bar{M}_{11}(u) + O(\sqrt{\lambda}) = O(\sqrt{\lambda}) \\ M_{11}(t) &= \bar{M}_{21}(u) + \hat{M}_{21}(v) + O(\sqrt{\lambda}) = O(\sqrt{\lambda}) \\ M_{11}(t) &= \bar{M}_{22}(u) + \hat{M}_{22}(v) + O(\sqrt{\lambda}) = O(\sqrt{\lambda}) \end{aligned}$$

We now conclude that $\alpha(t)$ is $O(\sqrt{\lambda})$ by the property (6.4).

6.3 Proof of Theorem 8

In case of single risky asset, $n = 1$, the solution of $F(w)$ in 4.4 is

$$F(w) = \tanh(\sigma\, w).$$

Then the result in (4.7)-(4.8) is obtained by direct substitution.

To obtain the result in (4.9), we consider the following expected cost problem

$$\begin{cases} E\left\{\int_0^T \frac{1}{2}\Big[\pi_M^*\Big(t,\, x(t)\Big) - \pi(t)\Big]^2 dt\right\} \\ \text{subject to:} \\ dx(t) = \{x(t)r + \pi(t)(\mu - r)\}dt + \pi(t)\sigma dw(t) \\ d\pi(t) = \{\bar{\rho}^*(t, x(t), \pi(t)) + \pi(t)\mu\}dt + \pi(t)\sigma dw(t) \\ x(0) = x_0, \quad \pi(0) = \pi_0 \end{cases}$$

Substituting for $\bar{\rho}^*(t,\, x(t),\, \pi(t))$ (see (4.8)), this can be rewritten as

$$\begin{cases} E\left\{\int_0^T \frac{1}{2} X(t)' Q\, X(t) dt\right\} \\ \text{subject to:} \\ dX(t) = \{\Phi + \frac{1}{\sqrt{\lambda}}\Psi(t)\}X(t)dt + \Omega\, X(t)dW(t) \\ X(0) = X_0 \end{cases} \tag{6.13}$$

where

$$X(t) = \begin{bmatrix} \frac{b_M(t)}{\eta a_M(t)} - x(t) \\ \pi(t) \end{bmatrix},\; Q = \begin{bmatrix} \left(\frac{\mu - r}{\sigma^2}\right)^2 & -\frac{\mu - r}{\sigma^2} \\ -\frac{\mu - r}{\sigma^2} & 1 \end{bmatrix},\; \Phi = \begin{bmatrix} r & -(\mu - r) \\ 0 & \mu \end{bmatrix},$$

$$\Psi(t) = \begin{bmatrix} 0 & 0 \\ \frac{\mu - r}{\sigma^2} f(t) & -f(t) \end{bmatrix},\; \Omega = \begin{bmatrix} 0 & -\sigma \\ 0 & \sigma \end{bmatrix},$$

$$f(t) = \sigma\sqrt{\eta a_M(t)}\tanh\Big(\sigma\sqrt{\eta a_M(t)}\frac{T - t}{\sqrt{\lambda}}\Big).$$

The linear quadratic structure of (6.13) implies that the expected cost is quadratic in $X(t)$

$$E\left\{\int_t^T \frac{1}{2} X(s)' Q\, X(s) ds\right\} = \frac{1}{2} X(t)' P(t)\, X(t) \tag{6.14}$$

where the coefficient matrix $P(t)$ satisfies the Lyapunov equation

$$\begin{cases} \dot{P}(t) + Q + P(t)\Phi + \Phi' P(t) + \frac{1}{\sqrt{\lambda}}(P(t)\Psi(t) + \Psi(t)' P(t)) + \Omega' P(t)\,\Omega &= 0 \\ P(T) &= 0 \end{cases} \tag{6.15}$$

An approximation of $P(t)$ can be found using a multiple time scale perturbation method. Similar to Proposition 6, two different time scales are defined by

$$u \triangleq T - t \;\text{ and }\; v \triangleq \frac{T - t}{\sqrt{\lambda}}. \tag{6.16}$$

With these new time scales, we approximate $P(t)$ by expanding it over the small parameter λ

$$P(u,v) = P^0(u,v) + \sqrt{\lambda}P^1(u,v) + o(\sqrt{\lambda}).$$

The Lyapunov equation (6.15) becomes

$$\begin{aligned}\frac{\partial}{\partial u}P^0(u,\, v)+&\frac{\partial}{\partial v}P^1(u,v) + \frac{1}{\sqrt{\lambda}}\frac{\partial}{\partial v}P^0(u,v)\\ &= Q + P^0(u,\, v)\Phi + \Phi' P^0(u,\, v) + P^1(u,v)\Psi(u,v) + \Psi(u,v)'P^1(u,v)\\ &\quad + \Omega\,' P^0(u,\, v)\,\Omega + \frac{1}{\sqrt{\lambda}}\big\{P^0(u,\, v)\Psi(t) + \Psi(t)'P^0(u,\, v)\big\} + o(\sqrt{\lambda}).\end{aligned}$$

with initial conditions $P^0(0,0) = 0$ and $P^1(0,0) = 0$. The equations obtained by collecting the terms scaled by $\lambda^{-1/2}$ and λ^0 can be solved by

$$\begin{aligned}P^0(u,v) &= 0\\ P^1(u,v) &= \Big\{\frac{v}{2}\text{sech}^2(\sigma\sqrt{\eta a_M(u)}v) + \frac{1}{2\sigma\sqrt{\eta a_M(u)}}\tanh(\sigma\sqrt{\eta a_M(u)}v)\Big\}Q.\end{aligned}$$

Substituting for u and v, the expected cost (6.14) becomes

$$\begin{aligned}&E\left\{\int_0^T \frac{1}{2}\Big[\pi_M^*(t,\, x(t)) - \pi(t)\Big]^2 dt\right\}\\ =\;& \frac{T}{2}\text{sech}^2(\sigma\sqrt{\eta a_M(0)}\frac{T}{\sqrt{\lambda}})(\pi_M^*(0,\, x(0)) - \pi(0))^2\\ &+\frac{\sqrt{\lambda}}{2\sigma\sqrt{\eta a_M(0)}}\tanh(\sigma\sqrt{\eta a_M(0)}\frac{T}{\sqrt{\lambda}})(\pi_M^*(0,\, x(0)) - \pi(0))^2 + o(\sqrt{\lambda}).\end{aligned}$$

The result in (4.9) follows from the observation that $\frac{T}{2}\text{sech}^2(\sigma\sqrt{\eta a_M(0)}\frac{T}{\sqrt{\lambda}})$ is $o(\sqrt{\lambda})$.

Part II

Dynamic Portfolio Choice Problems under the Hidden Markov Model

Chapter 7

Motivation

Regime switching is a phenomenon in financial market which is important and costly to ignore. Evidences of regime switching and benefits of modeling the markets in this particular way are presented in such as Hamilton [36], Ang and Bekaert [5, 6], Guidolin and Timmermann [34, 35], Tu [60] and Ang and Timmermann [7]. In general, market regime, or the market state, which is interpreted as an underlying factor that influences the asset returns is unobservable. This makes the well-known Hidden Markov Model (HMM) a natural candidate to model regime switching market. Under the HMM, investors only make an indirect observation of the market state through its noisy output such as asset prices, market analyst reports, news etc. Though the HMM is well studied and being used in various fields including financial economics (see such as Geweke et al. [32]), it is not yet fully incorporated into the portfolio choice problems or other dynamic decision problems; the existing literature either makes unrealistic assumptions regarding the knowledge on model parameters or ignores parameters uncertainties. In this work, we study dynamic portfolio choice problems when the market follows the HMM. The key difference from the previous literature is we do not make an impractical assumption that the transition probabilities of the hidden Markov chain and the observation statistics are known to the investors.

In dynamic portfolio optimization, the HMM assumption adds complications into the problem in many ways. First, an investor needs to learn about the hidden Markov chain and the stochastic processes those generate the observations. Under the Bayesian framework, this is done by updating the joint posterior distribution of the hidden states, the transition probabilities and state-dependent parameters such as asset expected returns, asset volatilities and other observation statistics. Besides, this learning process must be repeated when a new observation arrives. Second, for the HMM, the posterior distribution does not belong to any known distribution family and has no sensible ways to parameterize. Another difficulty which is largely ignored in the past literature is solving the dynamic portfolio choice problem knowing that there are uncertainties in the market states and the model parameters. This is an extremely complicated issue as the posterior distribution has no structure, which causes the dynamic optimization to be intractable and impossible to solve using any standard methods. In this work, we propose methods that handle all of these complications and solve the dynamic portfolio choice problems in a tractable fashion.

One approach to solve portfolio choice problem under the HMM is to separately learn the hidden state and model parameters then solve the portfolio choice problem. Guidolin and Timmermann [34, 35] first utilize the Expectation-Maximization (EM) algorithm to calculate the maximum likelihood estimate of the parameters, then solve the portfolio problem as if the estimates are the truth. The main shortage of this approach is when the unknown parameters are replaced by their point estimates the parameter uncertainties are ignored. Also this approach relies on numerical optimization which does not scale well with the problem size such as the number of hidden states or the number of assets.

A more elaborate approach in solving the dynamic portfolio selection problems based on the optimal filter theory. This approach gives us an explicit representation of optimal portfolio. The works done by this approach are Honda [38], Sass and Haussmann [56], Rieder and Bauerle [53], Bjork et al. [11] and Frey et al. [26]. However, the limitation of this approach is it assumes that the state transition probabilities are known to the investors. Since it is not clear how investors may acquire this set of information, the transition probabilities are replaced by point estimates and hence the uncertainties of these estimates are ignored. Moreover, due to limited tools in optimal filter framework, it cannot handle the case where the asset volatilities are state-dependent.

Another approach which we use does not assume the transition probabilities are observable and takes all parameter uncertainties into account. This approach applies the Bayesian framework to learn the hidden processes and utilizes the Markov Chain Monte Carlo (MCMC) technique to help solve optimization problem. Tu [60] is the first that uses this approach to the portfolio optimization problem under the HMM; however, the author ignores inter-temporal hedging and solves the single period mean-variance problems which greatly simplifies the optimization problem. The MCMC method enables Tu [60] to sample necessary model parameters for the mean-variance optimization from the posterior distribution.

We also adopt the Bayesian framework to learn the hidden processes as in Tu [60], but we are interested in a different case where an investor solves the dynamic utility-maximization problem. The investor re-balances his/her portfolio as new observation is made. The main roadblock to solve this problem is the joint posterior distribution which represents the investor knowledge on the hidden market regime, the unknown transition probabilities and the the regime-dependent parameters is part of the state variables, and performing backward recursion on this highly complicated posterior distribution is impossible. To avoid this, we solve this problem by first approximating the value function of the portfolio choice problem under the HMM (HMM-problem) by that of a complete information problem (C-problem). This approximation enables us to explicitly calculate the portfolio holdings in term of the weighted expected value of unknown parameters under the posterior distribution. We then utilize the MCMC technique to numerically computed this expectation.

Contributions of this work are: first, we propose tractable methods to solve the dynamic portfolio optimization problem when the market state is unobservable, and the transition probabilities and the state-dependent parameters are unknown. The methods allow us to calculate the optimal portfolio in a semi-explicit way.

Second, we provide a mean to evaluate the performance of this approximation. We use the information relaxation technique to construct a dual problem of the exact portfolio problem under the HMM. Then the performance of the dual problem is compared to that the proposed methods. The numerical experiments show that the performance of the approximate methods is reasonably close that of the dual problem which in theory is an upper-bound to the exact problem.

Finally, we are among the first to bring an advancement in MCMC to help solving the dynamic optimization problem with hidden processes. Although there are a lot of developments in sampling techniques from the posterior distribution, but there is not much utilization of these technique in solving optimization problems. In this work, we devise technique that utilizes the MCMC methods to calculate the optimal solution to the dynamic portfolio problem. And the proposed method is not limited to the portfolio problem under the HMM, it is ready to be applied to other contexts and other types of hidden processes.

7.1 Market Assumptions

We consider a financial market living on a filtered probability space $(\Omega, \overline{\mathcal{F}}, \{\overline{\mathcal{F}}_t\}, \mathbb{P})$ over a finite time horizon $[0, T]$. We model market uncertainties using two independent $\{\overline{\mathcal{F}}_t\}$-adapted processes, a J-dimensional standard Brownian motion w_t and a continuous-time Markov chain s_t. We assume that investors are price-takers. Elements of our market are as follows.

Asset price dynamics

The market can switch between S states/regimes. The market state at time t is denoted by $s_t \in \{1, 2, ..., S\}$, which evolves according to a continuous time Markov chain with transition rate matrix R. Specifically, state changes are characterized by

$$\mathbb{P}(s_{t+\delta} = k \,|\, s_t = j) = \begin{cases} \delta\, R_{jk} + o(\delta) & \text{if } k \neq j \\ 1 + \delta\, R_{jj} + o(\delta) & \text{if } k = j, \end{cases} \tag{7.1}$$

where $R_{jk} \geq 0$ for $j \neq k$ and $R_{jj} = -\sum_{k \neq j} R_{jk}$. There is one risk-free asset and I risky assets. The risk-free asset price p_t^0 grows continuously at rate r

$$dp_t^0 = p_t^0\, r\, dt. \tag{7.2}$$

The price of risky asset i is given by

$$dp_t^i = p_t^i \left(\mu_i(s_t) + r\right) dt + p_t^i \sum_{j=1}^{J} \sigma_{ij}(s_t)\, dw_t^j \tag{7.3}$$

where $\mu_i(s_t)$ and $\sigma_{i\cdot}(s_t) = [\sigma_{i1}(s_t), \cdots, \sigma_{iJ}(s_t)]$ are the state dependent return and volatility for asset i. We denote by

$$\mu(s) = [\mu_1(s), \mu_2(s), ..., \mu_I(s)]'$$

the state dependent $I \times 1$ vector of excess returns, and by

$$\sigma(s) = \begin{bmatrix} \sigma_{11}(s) & \cdots & \sigma_{1J}(s) \\ \vdots & \ddots & \vdots \\ \sigma_{I1}(s) & \cdots & \sigma_{IJ}(s) \end{bmatrix}$$

the state dependent $I \times J$ volatility matrix. We assume, as is standard, that $\sigma(s)\sigma(s)' \geq c\mathbf{I}$, where $\mathbf{I}$ is the $I \times I$ identity matrix, for some constant $c > 0$. We assume that the investor observes prices at discrete time points. For brevity, we also let $t \in \{0, 1, ..., T\}$ be a discrete time index where the length between period t and $t+1$ is some $\delta > 0$ (e.g. $\delta = 1/252$ for observations at the close of each day). At time t, the price of asset i is

$$p_t^i = p_{t-1}^i \exp\left\{ \big(\mu_i(s_{t-1}) + r - \frac{1}{2}\Sigma_{ii}(s_{t-1})\big)\delta + \sum_{j=1}^{J} \sigma_{ij}\,(w_t^j - w_{t-1}^j) \right\}. \tag{7.4}$$

We denote by $p^t = \{p_0, p_1, p_2, ..., p_t\}$ the history of prices up to time t.

Hidden Markov Model

In this work, we model the market regime as a continuous time Markov chain s_t. We are interested in the case when neither the market regime s_t nor the transition rate R are known by the investor, but need to be estimated using the history of asset prices and other noisy signals $y^t = \{y_0, y_1, \cdots, y_t\}$ of the hidden market state. The signals y^t could be subjective views about the regime that are periodically announced by market experts, though signals from other sources can be incorporated in a similar way.

There are various ways to model expert views. In this work, we assume for simplicity that the announcement at time t is multivariate normal with state-dependent mean and covariance

$$y_t \sim N(\mu_y(s_t), \Sigma_y(s_t)), \tag{7.5}$$

though other models are possible and can be handled using the methods we present in this work[1]. Observe that this model allows for multiple analysts with biased and dependent signals in that $\Sigma_y(s)$ need not be diagonal nor $\mu_y(s)$ be the same for all experts. We do not assume that the observation statistics $(\mu_y(\cdot), \Sigma_y(\cdot))$ are known to the investor; these need to be learned using the history of forecasts and returns.

[1]Specifically, we are able to handle the other models of analyst forecasts as long as we can sample the model parameters from the joint posterior distribution ρ_t.

Observations and learning

We have described dynamics of the market regime s_t, the observation process/analyst forecasts, and price process p^i_t, and their parameters R, $(\mu_y(s),\, \Sigma_y(s))$ and $(\mu_i(s),\, \sigma_{i1}(s), \cdots,\, \sigma_{iJ}(s))$, respectively. We assume throughout that the investor has the history of prices and forecasts

$$(p^t,\, y^t) = \Big\{(p_0,\, y_0),\, (p_1,\, y_1), \cdots,\, (p_t,\, y_t)\Big\},$$

and can use this data to estimate market regimes

$$s^t = \{s_0, s_1, s_2, ..., s_t\},$$

transition rate matrix R, and the regime dependent parameters

$$\xi \equiv \Big\{\mu(s), \sigma(s), \mu_y(s), \Sigma_y(s);\; s \in \{1, 2, ..., S\}\Big\}. \tag{7.6}$$

We denote by $\{\,\mathcal{F}_t\}$ the history of prices p_t and the analyst forecasts y_t available at time t to the investor

$$\mathcal{F}_t = \sigma\Big\{(p_0,\, y_0),\, (p_1,\, y_1), \cdots,\, (p_t,\, y_t)\Big\} \tag{7.7}$$

so $(s^t,\, R,\, \xi)$ need to be learned conditional on $\mathcal{F}_t$. One of the challenges of this work is to understand how these estimates should be used to make dynamic asset allocations.

Chapter 8

Portfolio Choice Problems under the Hidden Markov Model (*HMM-Problem*)

We summarize the Bayesian approach to learning the hidden market regime s^t, transition rates R, regime dependent price process parameters $(\mu(s), \Sigma(s))$ and observation statistics $(\mu_y(s), \Sigma_y(s))$, and formulate the associated dynamic portfolio choice problem. The dynamic programming equation for this problem is impossible to solve, due to the high dimensional state (which includes the posterior over (s^t, R, ξ)), which motivates subsequent discussion on approximations in the following chapter.

8.1 Bayesian Learning

With the Bayesian approach, the investor begins with prior views of the value of the hidden state s_0, the transition rate matrix R, and the state dependent return and observation parameters ξ (7.6), represented by a joint probability distribution

$$\rho_0(s_0, R, \xi).$$

The analyst observes asset returns/prices and expert signals over time and updates his prior using Bayes' rule. The resulting posterior distribution at time t is

$$\rho_t(s^t, R, \xi) \propto L(p^t, y^t; s^t, R, \xi)\, \rho_0(s_0, R, \xi), \tag{8.1}$$

where $L(\cdot)$ is the likelihood function. Observe that the history of states s^t, and not just the current state s_t, is also being learned. An estimate of the history of regimes is needed to estimate the current regime as well as the parameters (R, ξ).

In general, the posterior distribution ρ_t does not belong to any known distribution families and cannot be parameterized with a lower dimensional sufficient statistic. In such a situation, we can integrate under the posterior distribution by sampling from ρ_t using Markov Chain Monte Carlo (see Fruhwirth-Schnatter [27] and Chapter 10). Samples from this posterior are needed in our solution of the portfolio choice problem. In this work, we adopt a Dirichlet distribution for the prior on the transition probabilities, a normal-inverse-Wishart distribution for the prior on the observation parameters, and assume that they are independent. With this choice, the posterior can be sampled using Gibbs sampling.

8.2 Wealth dynamics and portfolio choice

The wealth process x_t of a self-financing investor with portfolio $\pi(t) = [\pi_1(t), \cdots, \pi_I(t)]'$ where $\pi_i(t)$ is the proportion of the investor's wealth in asset i at time t. The wealth process satisfies

$$dx_t = x_t\big\{r + \pi(t)'\mu(s_t)\big\}dt + x_t \sum_{j=1}^{J} \sigma'_{\cdot j}(s_t)\pi(t)dw_t^j, \tag{8.2}$$

where $\sigma_{\cdot j}(s)$ is the j^{th} column of $\sigma(s)$. The portfolio strategy $\pi(\cdot)$ is restricted to the admissible class

$$\mathcal{A} = \left\{\pi : [0,T]\times\Omega \to \mathbb{R}^I \;\middle|\; \pi \text{ is } \{\mathcal{F}_t\}\text{-adapted and } \mathbb{E}\left[\int_0^T \big\|x_s\pi(s)\big\|^2 ds < \infty\right]\right\}$$

where $\{\mathcal{F}_t\}$ is the history of prices and forecasts (7.7). The investor maximizes expected power utility of final wealth,

$$\max_{\pi(\cdot)\in\mathcal{A}} \; \mathbb{E}\left[\frac{x_T^\gamma}{\gamma}\right], \tag{8.3}$$

subject to wealth dynamics (8.2) and the dynamics of the posterior distribution (8.1). The state variables of the HMM-problem are $(x_t,\, \rho_t(\cdot))$ and the value function

$$\begin{aligned}
&V(s,\, x,\, \rho(\cdot)) = \max\nolimits_{\pi(\cdot)\in\mathcal{A}} \; \mathbb{E}\left[\tfrac{x_T^\gamma}{\gamma}\right],\\
&\text{Subject to:}\\
&\;x_t \text{ satisfies (8.2)},\\
&\;\rho_t(\cdot) \text{ satisfies (8.1)},\\
&\;x_s = x,\; \rho_s(\cdot) = \rho(\cdot),
\end{aligned}$$

which solves the dynamic programming recursion

$$V(t, x_t, \rho_t) = \max_{\pi_t} \mathbb{E}_{\rho_t}\Big[V(t+1, x_{t+1}, \rho_{t+1}) \,\Big|\, x_t\Big] \tag{8.4}$$

The subscript of the expectation represents the random variables associated with that expectation. The expectation in the Bellman equation is calculated over the hidden states s^t, the unknown parameters R and ξ, the wealth and the state of the next period x_{t+1} and s_{t+1}, and the observations of the next period p_{t+1} and y_{t+1}.

The state variable $(x_t,\ \rho_t(\cdot))$ of (8.3) is infinite dimensional since the posterior distribution $\rho_t(\cdot)$ is a continuous probability distribution. It follows that unless $\rho_t(\cdot)$ can be represented in terms of a finite dimensional sufficient statistic, which is not the case for the problem that we are interested in solving, the dynamic programming recursion (8.4) can not be solved exactly since it requires enumerating over all possible realizations of the posterior. (We encounter a similar problem even if we replace the continuous posterior with a discrete approximation since the dimension of the approximate state is still high). In the following chapter, we propose an approximate approach for solving (8.3) for which semi-analytic expressions of an approximately optimal dynamic policy can be written in terms of expectations under the posterior, and show that this policy can be computed using Gibbs sampling.

Chapter 9

Problem Approximation

In this chapter, we discuss an approximate method for solving (8.3). We begin by introducing the idea behind the approximation which involves the value function of a simpler variant of the HMM-problem in which the investor observes the market regime s_t and knows the model parameters. We derive the value function of this simpler problem and show how it can be used to construct an approximately optimal dynamic allocation policy that can be computed using Gibbs' sampling. We also illustrate this approximate method via examples at the end of this chapter.

9.1 Value function approximation

As noted in (8.4), the optimal policy for the HMM problem (8.3) is characterized by the solution of the optimization problem

$$\pi^*(t, x_t, \rho_t) = \arg\max_{\pi_t} \mathbb{E}_{\rho_t}\Big[\mathbb{E}_{x_{t+1}, s_{t+1}, p_{t+1}, y_{t+1}}\Big\{V(t+1, x_{t+1}, \rho_{t+1}) \,\Big|\, s^t, R, \xi\Big\}\Big]. \tag{9.1}$$

The challenge in solving this equation is not so much the optimization problem, but rather, that the objective function on the right hand side depends on the value function $V(t+1,\, x_{t+1},\, \rho_{t+1})$, which can only be computed by recursively solving (8.4) for every possible wealth-posterior pair $(x_t,\, \rho_t)$ for every time t. Indeed, if we had an oracle that could (magically) produce the value function for every desired choice of $(t,\, x_t,\, \rho_t)$, finding the optimal policy would be substantially easier. In particular, we would only need to solve (9.1) at wealth-posterior pairs $(x_t,\, \rho_t)$ that are encountered going forwards in time, and the computationally overwhelming task of implementing backwards recursion is not needed.

Of course, such an oracle does not exist, but the perspective just described suggests that backwards recursion can be avoided if an easily computable approximation to the value function, which we denote by $\widetilde{V}$, can be found. If this is possible, then then an easily computable approximately optimal policy can be obtained by replacing $V(t+1,\, \cdot)$ with

$\widetilde{V}(t+1,\,\cdot)$ in (9.1) and solving

$$\pi^*(t,\,x_t,\,\rho_t) = \arg\max_{\pi}\,\mathbb{E}_{\rho_t}[\widetilde{V}(t+1,\,\cdot)\,|\,x_t] \tag{9.2}$$

at values of $(x_t,\,\rho_t)$ that are encountered going forwards in time. (See Bertsekas [9] and Powell [52] for a description and applications of this approach in other applications).

While any approximation of V is (in principle) possible, the quality of the trading policy obtained using the approximation (9.2) depends on its quality. In this regard, a natural approximation $\widetilde{V}$ is the value function of a simpler but related problem to (8.3), which we now describe.

9.2 Complete Information Problem (*C-Problem*)

The complete information problem (C-problem) is similar to the HMM-problem except that an investor knows the transition rate R, can observe the market states s_t, and knows the return parameters $\{\mu(s),\sigma(s);\,s\in\{1,2,...,S\}\}$ for each regime. (Under these assumptions, the observations y_t are irrelevant). For this problem, the investor's goal is to maximize expected utility of terminal wealth

$$\max_{\pi(t)\in\overline{\mathcal{A}}}\,\mathbb{E}\left[\frac{x_T^\gamma}{\gamma}\right], \tag{9.3}$$

subject to the state dynamics of the Markov chain (7.1) and the wealth dynamics (8.2). His/her class of admissible trading policies is modified to reflect this additional information

$$\overline{\mathcal{A}} = \left\{\pi:[0,T]\times\Omega\to\mathbb{R}^I\,\middle|\,\pi\text{ is }\{\,\mathcal{F}_t^{W,\,s}\}\text{-adapted and }\mathbb{E}\left[\int_0^T\left\|x_s\pi(s)\right\|^2 ds<\infty\right]\right\}$$

where $\{\,\mathcal{F}_t^{W,\,s}\}$ is the filtration generated by the Brownian motion W_t and the Markov chain s_t. The solution to this problem is presented in the following proposition and forms the basis of the approximation to the HMM problem.

Proposition 9. *The value function for the C-problem is*

$$V_C(t,x_t,s_t) = \frac{x_t^\gamma}{\gamma}g_C(t,s_t) \tag{9.4}$$

where $g_C(t,s_t)$ solves the system of differential equations

$$\begin{cases} 0=\dot{g}_C(t,s)+\left\{\gamma r+\dfrac{1}{2}\dfrac{\gamma}{1-\gamma}\theta(s)\right\}g_C(t,s)+\displaystyle\sum_{k=1}^{S}R_{sk}\,g_C(t,k), & s\in\{1,2,...,S\},\\ g_C(T,s)=1, \end{cases} \tag{9.5}$$

$\theta(s)\equiv\mu'(s)\,\Sigma^{-1}(s)\,\mu(s)$ and $\Sigma(s)\equiv\sigma(s)\sigma'(s)$. The optimal portfolio is

$$\pi_C^*(t) = \frac{1}{1-\gamma}\,\Sigma^{-1}(s_t)\,\mu(s_t). \tag{9.6}$$

Proof. See Chapter 12. □

9.3 Approximate portfolio

With the value function approximation, $\widetilde{V} = V_C$, the approximate optimal policy (9.2) becomes

$$\begin{aligned}\pi^*(t,\, x_t,\, \rho_t) &= \arg\max_{\pi} \mathbb{E}_{\rho_t}\Big[V_C(t+1,\, x_{t+1},\, s_{t+1})\,\Big|\, x_t\Big] \\ &= \arg\max_{\pi} \mathbb{E}_{\rho_t}\Big[\mathbb{E}\Big\{V_C(t+1,\, x_{t+1},\, s_{t+1})\,\Big|\, x_t,\, s_t,\, R,\, \xi\Big\}\,\Big|\, x_t\Big]. \end{aligned} \tag{9.7}$$

Although it was introduced as an approximation to the recursive equation (9.7), there is another interpretation which we now describe. Specifically, suppose that at time t, $\{s^t,\, R,\, \xi\}$ are not known to the investor but he/she has a posterior distribution $\rho_t(s^t,\, R,\, \xi)$ on the set of possible values. The approximation (9.7) would be optimal for this investor if it was also the case that the hidden state and the true parameters are revealed to him the instant after his time t allocation has been decided (but before a state transition occurs), and the market state is observable from then on. By revealing the true states and parameters and making the state observable, the investor solves the C-problem instead of the HMM-problem from time $t+1$ onwards. In other words, the posterior dynamics (8.1) are no longer relevant from time $t+1$ onwards and the value function from $t+1$ onwards is V_C. Given information at time t, the probability that (s^t, R, ξ) is revealed to the investor after his time t decision is $d\rho_t(s^t, R, \xi)$, and the dynamic programming equations associated with this setup gives (9.7).

By conditioning on s^t, R and ξ, the inner expectation is easy to calculate. Specifically, conditioning on s^t, R, ξ, the next period wealth x_{t+1} and s_{t+1} are independent. The expectation over s_{t+1} can simply be calculated using the transition rate R (7.1), and the expectation over x_{t+1} is calculated by applying Ito's lemma. Specifically, we have

$$\begin{aligned}&\mathbb{E}\Big[\,V_C(t+1, x_{t+1}, s_{t+1})\,\Big|\, x_t,\, s^t,\, R, \xi\,\Big] \\ &= \mathbb{E}\,\Big[\,\mathbb{E}\Big\{\,V_C(t+1, x_{t+1}, s_{t+1})\,\Big|\, x_{t+1},\, s^t,\, R,\, \xi\,\Big\}\,\Big|\, x_t,\, s^t,\, R,\, \xi\,\Big] \\ &= \mathbb{E}\Big[V_C(t+1, x_{t+1}, s_t) + \delta\sum_{k=1}^{S} R_{s_t,k} V_C(t+1, x_{t+1}, k) + o(\delta)\,\Big|\, s^t, R, \xi\,\Big] \\ &= \mathbb{E}\Big[V_C(t, x_t, s_t) \\ &\quad + \delta\left\{\frac{\partial}{\partial t}V_C(t, x_t, s_t) + x_t\big(r + \pi'\mu(s_t)\big)\frac{\partial}{\partial x}V_C(t, x_t, s_t) + \frac{1}{2}x_t^2\,\pi'\,\Sigma(s_t)\,\pi\frac{\partial^2}{\partial x^2}V_C(t, x_t, s_t)\right\} \\ &\quad + \delta\sum_{k=1}^{S} R_{s_t,k} V_C(t, x_t, k)\Big] + o(\delta)\end{aligned}$$

Substituting into (9.7) gives

$$\pi^*(t,\, x_t,\, \rho_t) = \arg\max_{\pi}\, \mathbb{E}_{\rho_t}\left\{ x_t\pi'\mu(s_t)\frac{\partial}{\partial x}V_C(t,x_t,s_t) + \frac{1}{2}x_t^2\,\pi'\,\Sigma(s_t)\,\pi\frac{\partial^2}{\partial x^2}V_C(t,x_t,s_t) \,\Big|\, x_t\right\} + O(\delta)$$

where we have dropped terms that do not depend on π. Dropping higher order terms and noting that V_C depends on $(s_t,\, R,\, \xi)$ it follows that

$$\begin{aligned}\pi^*(t,\, x_t,\, \rho_t) &= \mathbb{E}_{\rho_t}\left[x_t\,\Sigma(s)\frac{\partial^2}{\partial x^2}V_C(t,x_t,s_t)\right]^{-1} \mathbb{E}_{\rho_t}\left[\mu(s_t)\frac{\partial}{\partial x}V_C(t,x_t,s_t)\right] \\ &= \frac{1}{1-\gamma}\,\mathbb{E}_{\rho_t}\Big[\Sigma(s)\, g_C(t,s;R,\xi)\Big]^{-1}\mathbb{E}_{\rho_t}\Big[\mu(s)\, g_C(t,s;R,\xi)\Big] \qquad (9.8)\end{aligned}$$

where the second line follows from (9.4) and $g_C(t,s;R,\xi)$ solves the system of differential equations (9.5). (We add R and ξ into the arguments of g_C as a reminder that g_C also depends on these unknown parameters). Observe that the structure of the optimal portfolio is very similar to that of the C-problem (9.6) except that the mean and covariance are replaced by their expectations under the posterior, weighted by $g_C(t,\, s;\, R,\, \xi)$.

Computing the approximately optimal policy

The expectation over $(s,\, R,\, \xi)$ in (9.8) is computed using Monte Carlo integration, with samples of $(s,\, R,\, \xi)$ under the posterior being generated by Gibbs' sampling. This can be summarized as follows.

1. Sample $\{s_m^t, R^m, \xi^m\}_{m=1,2,\ldots,M}$ from the posterior distribution ρ_t using the Gibbs sampler[1].

2. For each sample $\{s_m^t, R^m, \xi^m\}$, solve the differential equation (9.5) for $g_C(t, s_m^t; R^m, \xi^m)$.

3. Calculate the portfolio $\pi^*(t)$ by

$$\pi^*(t) \approx \frac{1}{1-\gamma}\left[\sum_{i=1}^{M}\Sigma^m(s)\, g_C(t,s_m^t;R^m,\xi^m)\right]^{-1}\left[\sum_{i=1}^{M}\mu^m(s_m^t)\, g_C(t,s_m^t;R^m,\xi^m)\right]. \quad (9.9)$$

9.4 Examples

In this section, we illustrate the use of the previously outlined approximation using the real data from US equity market and simulated data. The US equity data used in the first example is the Fama-French 2×3 Portfolios from 1980 to 2012. In the next two examples, we simulate the regime-switching markets with market analysts who provide the investor the noisy observations of the hidden market states.

[1] See details in Chapter 10

Example : Fama-French 6 Portfolios

In this example, we investigate the value-weighted Fama-French 2×3 portfolios from 1980-2012. The portfolios are denoted by SmLo, SmMed, SmHi, BigLo, BigMed and BigHi. We model the market using the HMM with two regimes, bull and bear regimes. The bear regime is characterized by the regime with lower expected returns and higher volatility. We denoted the bull and bear regimes by $s = 1$ and $s = 2$ respectively. The investor observes weekly returns of the six portfolios, and rebalances weekly. There are no experts/analysts in this example. We first describe the prior distribution as follows.

Prior parameters

The prior parameters of this example are obtained from the older data set from 1960-1979.[2] Since the observation is discrete, we assume that the hidden Markov chain is a discretization of the chain defined in (7.1). The relationship between the rate matrix R of the continuous-time MC and the transition matrix Q of the discrete-time MC is standard as shown in (10.1). We assume the prior distribution of each row of the transition matrix Q is Dirichlet,

$$Q_{j\cdot} \sim Dir(e_{j1}, e_{j2}).$$

The parameter e_{jk} can be interpreted as the prior number of switching from regime j to regime k. We set the e_{jk}'s to equal the average number of transitions within τ weeks of a Markov chain with rate matrix

$$R_{prior} = \begin{bmatrix} -2.14 & 2.14 \\ 5.84 & -5.84 \end{bmatrix}.$$

τ reflects the investor confidence in this prior distribution. With $\tau = 104$ weeks, the Dirichlet distribution parameters are

$$\begin{bmatrix} e_{11} & e_{12} \\ e_{21} & e_{22} \end{bmatrix} = \begin{bmatrix} 72.98 & 3.13 \\ 3.13 & 24.76 \end{bmatrix}.$$

The prior of the annualized expected returns obtained from 1960-1879 data set are:

$$\mu_{prior}(1) = \begin{bmatrix} 0.23 \\ 0.22 \\ 0.23 \\ 0.15 \\ 0.14 \\ 0.17 \end{bmatrix}, \quad \mu_{prior}(2) = \begin{bmatrix} -0.24 \\ -0.17 \\ -0.13 \\ -0.13 \\ -0.07 \\ -0.08 \end{bmatrix}.$$

[2]We setup the prior of parameters of this 1960-1979 data set by first manually identifying the bull and bear regimes and estimating the parameters in each regime. Then with very flat prior distribution, we update the model parameters using Bayesian approach. The posterior parameters of the 1960-1979 data set is calculated and used as the prior parameters for the 1980-2012 data set.

The prior covariance matrices $\Sigma_{prior}(1)$ and $\Sigma_{prior}(2)$ are calculated from the following standard deviation and correlation:

$$SD(1) = \begin{bmatrix} 0.13 \\ 0.10 \\ 0.10 \\ 0.11 \\ 0.09 \\ 0.11 \end{bmatrix}, \quad Corr.(1) = \begin{bmatrix} 1 & 0.94 & 0.90 & 0.81 & 0.78 & 0.77 \\ \cdot & 1 & 0.95 & 0.80 & 0.81 & 0.82 \\ \cdot & \cdot & 1 & 0.76 & 0.79 & 0.83 \\ \cdot & \cdot & \cdot & 1 & 0.87 & 0.82 \\ \cdot & \cdot & \cdot & \cdot & 1 & 0.87 \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix},$$

$$SD(2) = \begin{bmatrix} 0.28 \\ 0.23 \\ 0.23 \\ 0.24 \\ 0.20 \\ 0.21 \end{bmatrix}, \quad Corr.(2) = \begin{bmatrix} 1 & 0.96 & 0.93 & 0.86 & 0.86 & 0.87 \\ \cdot & 1 & 0.97 & 0.83 & 0.88 & 0.90 \\ \cdot & \cdot & 1 & 0.79 & 0.86 & 0.90 \\ \cdot & \cdot & \cdot & 1 & 0.88 & 0.87 \\ \cdot & \cdot & \cdot & \cdot & 1 & 0.93 \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix}.$$

From the price dynamics (7.4), the one-period asset returns

$$ret_t = [ret_t^1, ..., ret_t^I]' \equiv \left[\ln(\frac{p_{t+1}^1}{p_t^1}), ..., \ln(\frac{p_{t+1}^I}{p_t^I})\right]'$$

have a normal distribution,

$$ret_t \sim N(\mu_{ret}(s_t), \Sigma_{ret}(s_t)),$$

where $\mu_{ret}(s)$ and $\Sigma_{ret}(s)$ are defined in (10.2)-(10.3). We assume normal-inverse-Wishart (NIW) for the return parameters $\mu_{ret}(s)$ and $\Sigma_{ret}(s)$,

$$\Sigma_{ret}(s) \sim IW\Big(\Lambda_{ret}^0(s), \nu_{ret}^0(s)\Big) \quad \text{and} \quad \mu_{ret}(s) \,\Big|\, \Sigma_{ret}(s) \sim N\Big(\eta_{ret}^0(s), \frac{\Sigma_{ret}(s)}{\kappa_{ret}^0(s)}\Big), \tag{9.10}$$

where

$$\begin{aligned} &\kappa_{ret}^0(1) = \kappa_{ret}^0(2) = 260, \quad \nu_{ret}^0(1) = \nu_{ret}^0(2) = 20, \\ &\eta_{ret}^0(s) = \Big\{\mu_{prior}(s) - \frac{1}{2}\operatorname{diag}(\Sigma_{prior}(s))\Big\}\,\delta, \\ &\Lambda_{ret}^0(s) = \nu_{ret}^0(s)\,\Sigma_{prior}(s)\,\delta; \quad s \in \{1,2\},\ \delta = 1/52. \end{aligned}$$

$\kappa_{ret}^0(s)$ and $\nu_{ret}^0(s)$ reflect the confidence in the prior mean and the prior covariance respectively. $\kappa_{ret}^0(s)$ can be interpreted as the number of prior observations with sample mean $\eta_{ret}^0(s)$, while $\nu_{ret}^0(s)$ can be interpreted as the number of prior observations with covariance $\Sigma_{prior}(s)\,\delta$. Note that, in this example, the only data available to the investor is the returns.

Result

Figure 9.1 shows the Fama-French portfolios and the probability in bear regime. The bull regime is characterized by the regime with higher expected returns (15%-25% per year) and lower volatility (10%-15% per year) while the bear regime has lower expected returns (-30%-10%) and higher volatility (25%-30%). The summary statistics of the prior and the posterior distribution are show in Figure 9.2 and 9.3. From Figures 9.2 and 9.3, we can see that the prior distribution is fairly flat. The range of parameters sampled from the prior distribution is much wider compared to that of the posterior distribution. This flat prior is a result of choosing small values of τ, $\kappa^0_{ret}(s)$ and $\nu^0_{ret}(s)$.

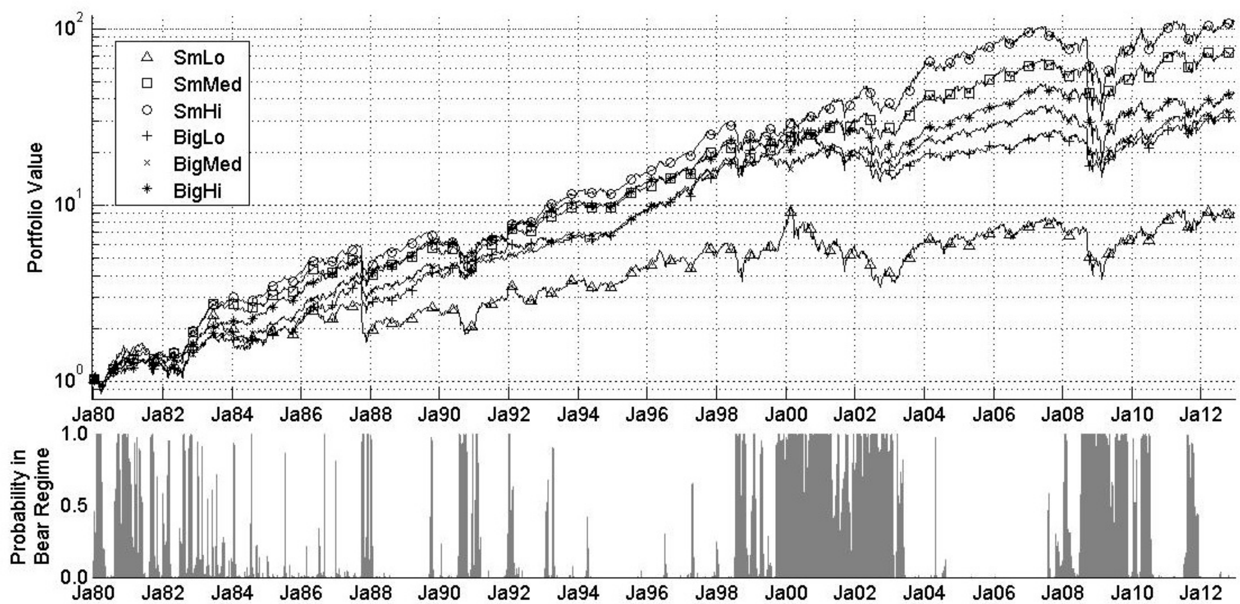

Figure 9.1: The top plot shows the Fama-French portfolios from 1980 to 2012. The bottom plot shows the probability of being in bear regime.

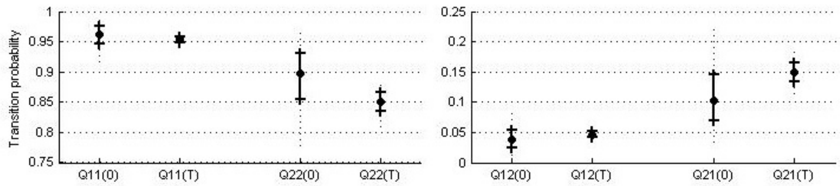

Figure 9.2: We plot the summary statistics of the prior distribution (denoted by $t = 0$) and the posterior distribution (denoted by $t = T$) of the transition matrix Q. The posterior distribution is updated from observing the whole data set. For each vertical line, the circular marker represents the 50th percentile, the solid line connects the 25th and the 75th percentiles, and the dotted lines are extended to the 5th and the 95th percentiles.

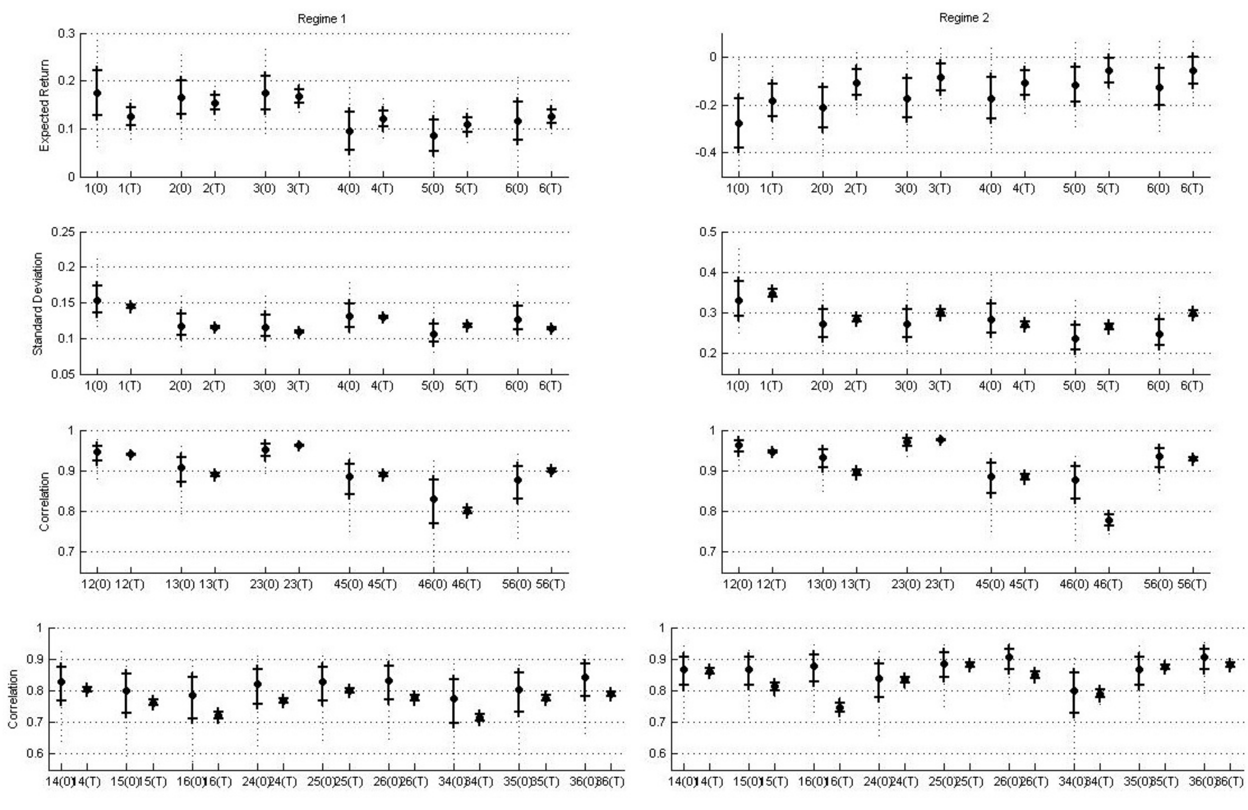

Figure 9.3: Similar to Figure 9.2, we plot the summary statistics of the annualized returns, standard deviation and correlation. The regime 1 (bull) parameters are in the left column the regime 2 (bear) parameters are in the right column. The number 1-6 represent the portfolio SmLo, SmMed, SmHi, BigLo, BigMed and BigHi respectively.

To illustrate the use of our approximate method, we let an investor solves the HMM-problem where he/she can invest either in the Fama-French portfolios or in risk-free asset[3]. The prior distribution of the model is described in the previous section. When the weekly data is revealed to the investor, the investor updates his/her knowledge then adjusts the asset holdings using the approximate method we proposed.

We show the investor portfolios and wealth in Figure 9.4-9.6. Figure 9.4 shows ten-year investment from 1980-1989. We can see that the investor is able to identify the bad periods and responds by reducing the asset holdings, as a result, the investor wealth does not affect much from the crash. Similarly, Figure 9.5 and 9.6 show the result from 1992-2001 and 2002-2012. We also show the hedging demand which is the difference in portfolio holdings obtained from the proposed method and from solving the single-period mean-variance problem. Specifically, the hedging demand $\Delta\pi$ is defined by

$$\Delta\pi(t) = \pi^*(t) - \frac{1}{1-\gamma}\Sigma_{pred}(t)^{-1}\,\mu_{pred}(t). \tag{9.11}$$

The term $\Sigma_{pred}(t)^{-1}\mu_{pred}(t)/(1-\gamma)$ is the optimal portfolio for the single period problem. $\mu_{pred}(t)$ and $\Sigma_{pred}(t)$ are the mean and covariance of the predictive density which can be calculated using the MCMC method. We can see that the hedging demand is larger at the beginning of the investment horizon as the intertemporal hedging is ignored in the single-period problem.

[3]For simplicity, we let an investor lend and borrow at the same rate which is averaged over investment horizon. And this risk-free rate is assumed to be known to the investor.

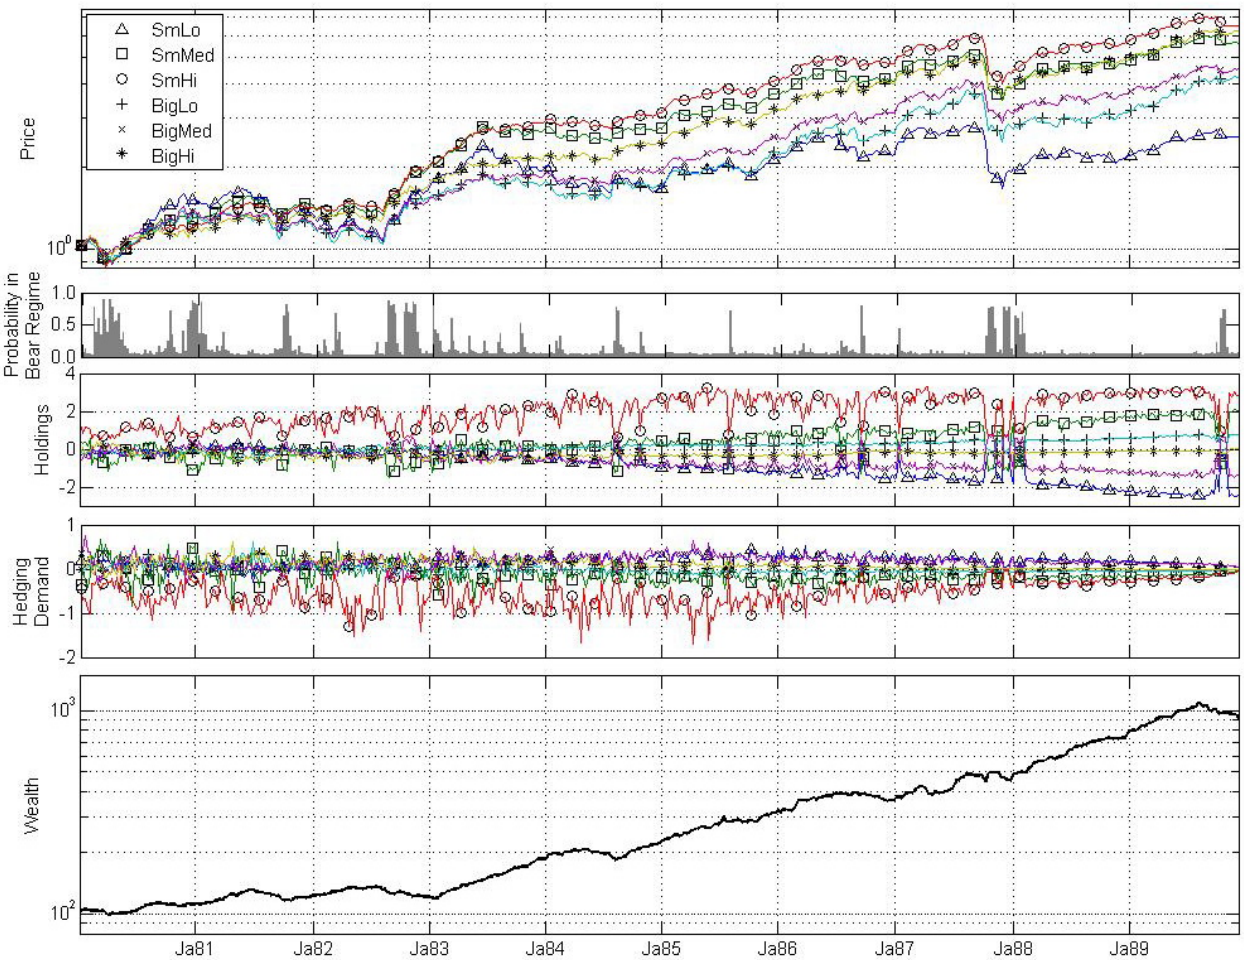

Figure 9.4: This figure shows, from top to bottom, the Fama-French portfolios from 1980 to 1989, the probability in bear regime, the risky asset holdings selected by the approximate method, the hedging demand, and the investor wealth. The investor initial wealth is \$100, and the risk aversion parameter γ is -8.

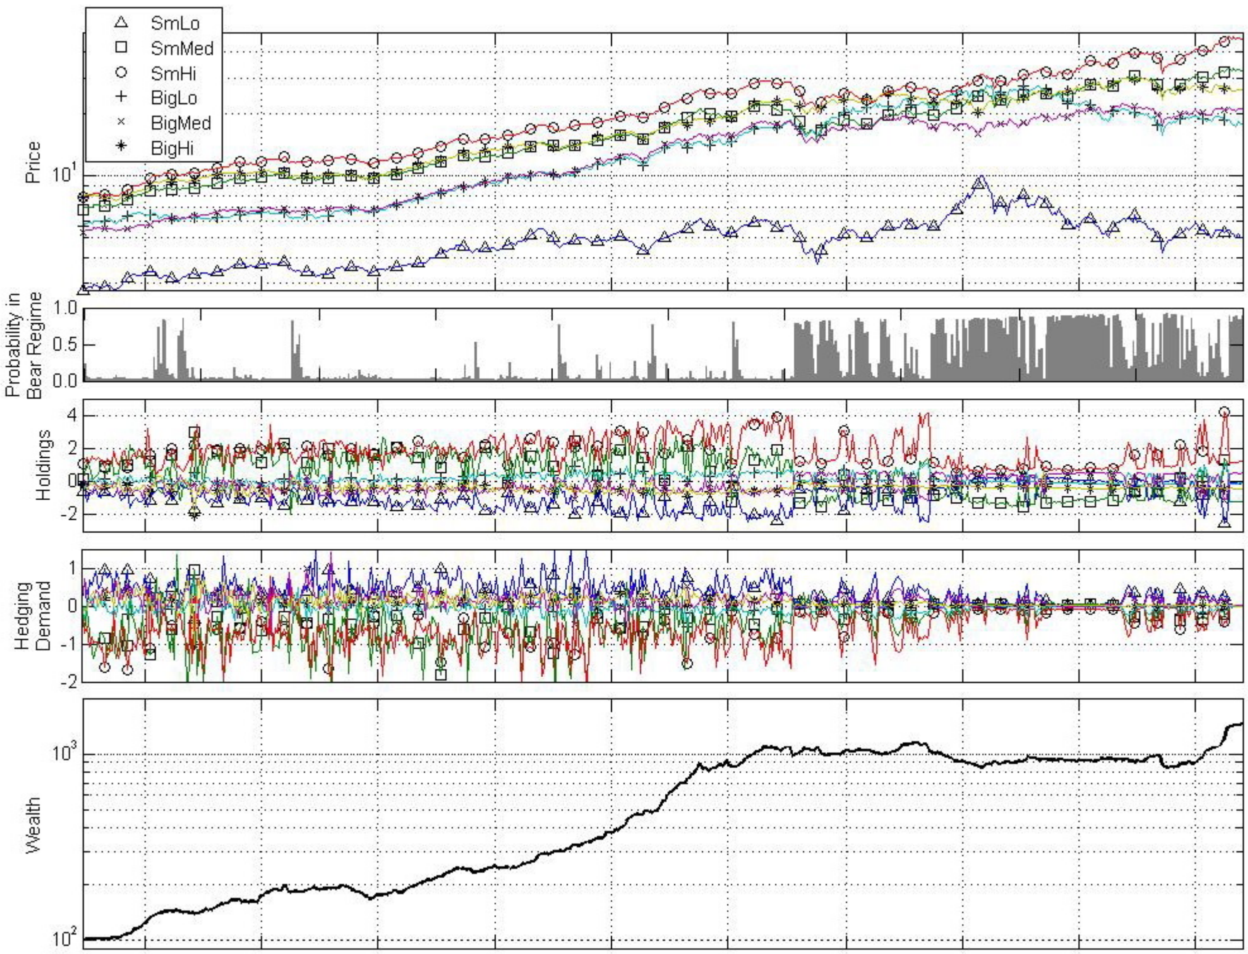

Figure 9.5: Similar to Figure 9.4, this figure shows the result from 1992 to 2001

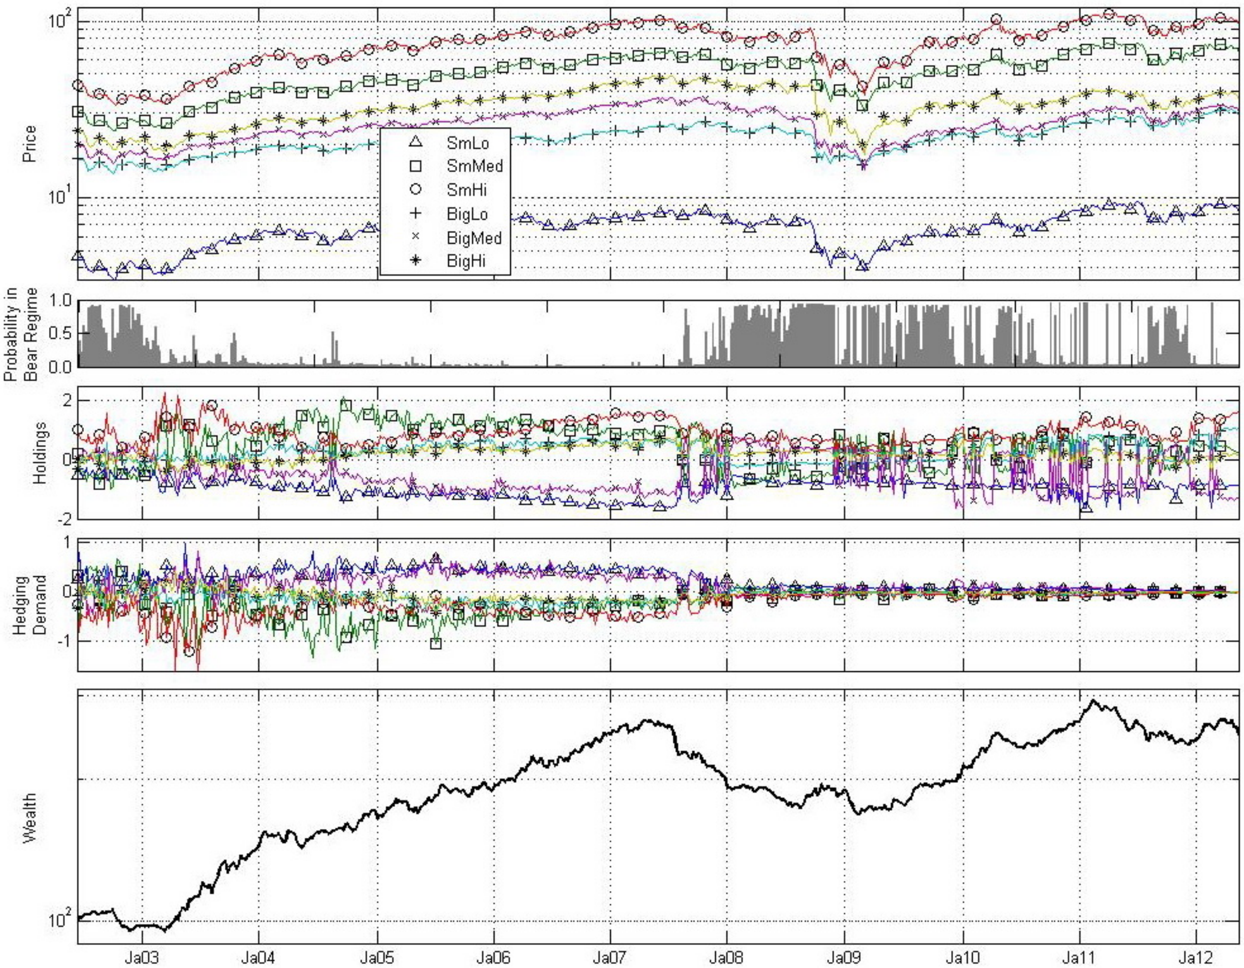

Figure 9.6: Similar to Figure 9.4, this figure shows the result from 2002 to 2012

Next, we compare the performance of the using six different approaches to solve the portfolio choice problem. The investor statistics are calculated by letting the investors solve a series one-year-horizon problems during 1980-2012. The mean-standard deviation frontiers and the certainty equivalence are plotted in Figure 9.7. Note that the investors are allowed to adjust their holdings every week. These six approaches are:

1. An investor solves the HMM-problem using the proposed dynamic approximation.

2. An investor repeatedly solves the single-period (one-week) mean-variance problem but still assumes that market follows the HMM with two regimes. In this case, the investor first computes the predictive mean and variance, then solves the mean-variance problem.

3. An investor repeatedly solves the single-period mean-variance problem without regime switching. The prior parameters for the asset return are obtained from 1960-1979,

$$\mu_{prior} = \begin{bmatrix} 0.11 \\ 0.13 \\ 0.16 \\ 0.077 \\ 0.095 \\ 0.13 \end{bmatrix}, \quad SD = \begin{bmatrix} 0.18 \\ 0.15 \\ 0.15 \\ 0.15 \\ 0.13 \\ 0.14 \end{bmatrix}, \quad Corr. = \begin{bmatrix} 1 & 0.96 & 0.93 & 0.85 & 0.83 & 0.83 \\ \cdot & 1 & 0.97 & 0.82 & 0.86 & 0.87 \\ \cdot & \cdot & 1 & 0.79 & 0.84 & 0.88 \\ \cdot & \cdot & \cdot & 1 & 0.87 & 0.85 \\ \cdot & \cdot & \cdot & \cdot & 1 & 0.90 \\ \cdot & \cdot & \cdot & \cdot & \cdot & 1 \end{bmatrix}.$$

   Similar to the HMM-problem, the prior distribution is assumed to be NIW with parameters

$$\kappa^0_{ret} = 260, \quad \nu^0_{ret} = 20, \quad \eta^0_{ret} = \left\{\mu_{prior} - \frac{1}{2}\operatorname{diag}(\Sigma_{prior})\right\}\delta, \quad \Lambda^0_{ret} = \nu^0_{ret}\,\Sigma_{prior}\,\delta.$$

4. An investor repeatedly solves the single-period mean-variance problem without regime switching and parameters uncertainty. The return parameters are the same as the previous approach but they are assumed to be the truth.

5. An investor assigns equal weight for all assets.

6. An investor invests all in the risk-free asset.

We can see in Figure 9.7 that dynamic hedging, regime switching and parameter uncertainty help improve the investor performance, and hence our approximate approach performs the best measured by both the portfolio mean-standard deviation plot and the investor utility. Though the improvement of the dynamic approximation is not substantial in the mean-standard deviation plot, it shows significant improvement in terms of the certainty equivalence as a result of its preferable higher-order risk.

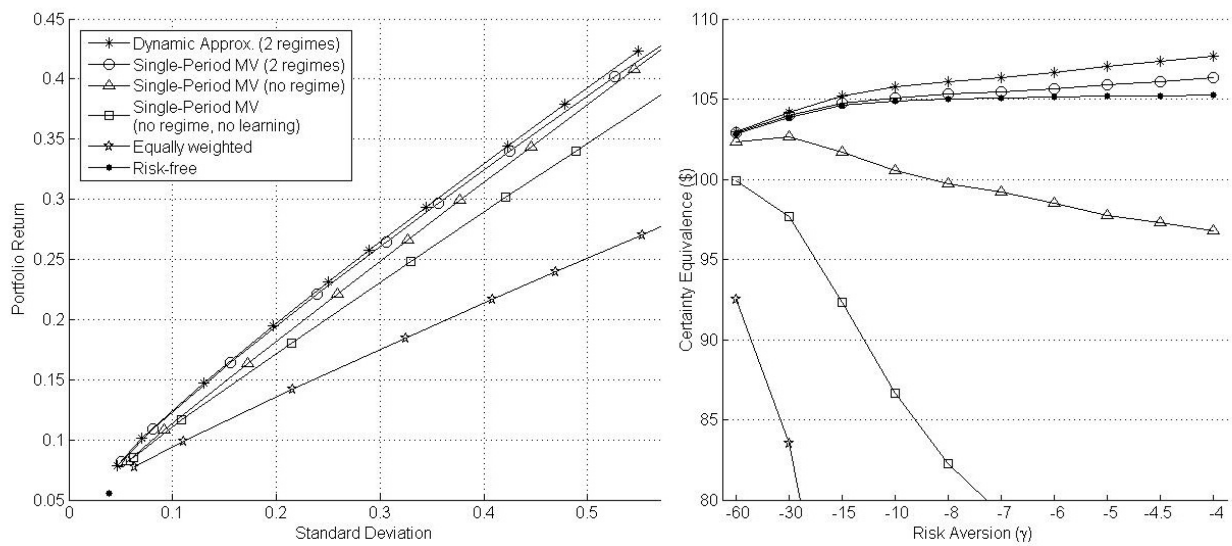

Figure 9.7: The left plot compares the mean-standard deviation frontiers obtained from six different approaches. The right plot compares the certainty equivalence at different risk-aversion levels. Note that the initial wealth is \$100.

Example : Setting 1

In this example, we simulate a market with three regimes: bull, calm and bear regimes. The regimes are named in according to the return characteristic in each regime. We denote the bull, calm and bear regimes by numbers 1, 2 and 3 respectively. The observations on the asset prices and the analyst forecasts are made daily. We first describe the true parameters used in the simulation.

Data generating process

The generated hidden Markov chain is a discretization of the chain defined in defined in (7.1) with rate matrix

$$R = \begin{bmatrix} -6.0 & 4.5 & 1.5 \\ 1.5 & -3.0 & 1.5 \\ 3.0 & 3.0 & -6.0 \end{bmatrix} \text{ per year.}$$

The risky asset price processes defined in (7.3) has the excess return vectors $\mu(s)$ and the covariance matrices $\Sigma(s)$

$$\mu(1) = \begin{bmatrix} 0.20 \\ 0.15 \end{bmatrix}, \qquad \mu(2) = \begin{bmatrix} 0.05 \\ 0.03 \end{bmatrix}, \qquad \mu(3) = \begin{bmatrix} -0.10 \\ -0.09 \end{bmatrix},$$

$$\Sigma(1) = \begin{bmatrix} 0.0225 & 0.0090 \\ 0.0090 & 0.0144 \end{bmatrix}, \quad \Sigma(2) = \begin{bmatrix} 0.0225 & 0.0045 \\ 0.0045 & 0.0144 \end{bmatrix}, \quad \Sigma(3) = \begin{bmatrix} 0.0225 & 0.0135 \\ 0.0135 & 0.0144 \end{bmatrix}.$$

Apart from the differences in excess returns across regimes, we also assume different correlation structures, specifically the correlations are 0.5, 0.25 and 0.75 respectively. Throughout the example, we assume that this risk-free rate r is 0.

The investor observes two analyst forecasts on the market state on a daily basis. The model of the forecasts is defined in (7.5) with parameters:

$$\mu_y(1) = \begin{bmatrix} 1.0 \\ 1.0 \end{bmatrix}, \quad \mu_y(2) = \begin{bmatrix} 2.0 \\ 2.0 \end{bmatrix}, \quad \mu_y(3) = \begin{bmatrix} 3.0 \\ 3.0 \end{bmatrix},$$
$$\Sigma_y(s) = v^2 \begin{bmatrix} 1.0 & 0.5 \\ 0.5 & 1.0 \end{bmatrix}; \quad s \in \{1, 2, 3\}.$$

The standard deviation of the forecasts denoted by v will be varied from 1 to 5 to test the performance of the investor. Next we describe the prior distribution of the parameters.

Prior parameters

We assume Dirichlet for the prior of the the transition matrix Q,

$$Q_{j\cdot} \sim Dir(e_{j1}, e_{j2}, e_{j3})$$

We set the value of e_{jk}'s by let them be the average number of transitions within τ days of a Markov chain with rate matrix

$$R\,(1 + \Delta R).$$

We use the parameter ΔR to characterize the precision of prior transition rate. In this example, we let τ to be 63 days and choose ΔR from $\{\,-0.2,\, 0.0,\, 0.2\,\}$. When $\Delta R = 0$, for example, the Dirichlet distribution parameters are

$$\begin{bmatrix} e_{11} & e_{12} & e_{13} \\ e_{21} & e_{22} & e_{23} \\ e_{31} & e_{32} & e_{33} \end{bmatrix} = \begin{bmatrix} 14.71 & 0.269 & 0.091 \\ 0.209 & 34.68 & 0.209 \\ 0.150 & 0.150 & 12.31 \end{bmatrix}.$$

Similar to the previous example, we assume NIW for the return parameters $\mu_{ret}(s)$ and $\Sigma_{ret}(s)$,

$$\Sigma_{ret}(s) \sim IW\Big(\Lambda^0_{ret}(s), \nu^0_{ret}(s)\Big) \quad \text{and} \quad \mu_{ret}(s)\,\big|\,\Sigma_{ret}(s) \sim N\Big(\eta^0_{ret}(s), \frac{\Sigma_{ret}(s)}{\kappa^0_{ret}(s)}\Big) \tag{9.12}$$

where

$$\nu^0_{ret}(1) = \kappa^0_{ret}(1) = 10000, \quad \nu^0_{ret}(2) = \kappa^0_{ret}(2) = 25000, \quad \nu^0_{ret}(3) = \kappa^0_{ret}(3) = 10000,$$
$$\eta^0_{ret}(s) = \big\{\mu(s)(1 + \Delta\mu) - \frac{1}{2}\,\mathrm{diag}(\Sigma(s))\big\}\delta, \quad \Lambda^0_{ret}(s) = \nu^0_{ret}(s)\,\Sigma(s)\,\delta; \quad s \in \{1, 2, 3\}.$$

The parameter $\Delta\mu$ allows us to characterize how precise the prior distribution is; in this example, we choose $\Delta\mu$ between -0.2 and 0.2.

We also assume NIW for the forecasts parameters $\mu_y(s)$ and $\Sigma_y(s)$ defined in (7.5),

$$\Sigma_y(s) \sim IW\Big(\Lambda_y^0(s), \nu_y^0(s)\Big) \quad \text{and} \quad \mu_y(s) \,\big|\, \Sigma_y(s) \sim N\Big(\eta_y^0(s), \frac{\Sigma_y(s)}{\kappa_y^0(s)}\Big) \tag{9.13}$$

where

$$\nu_y^0(1) = \kappa_y^0(1) = 36, \qquad \nu_y^0(2) = \kappa_y^0(2) = 36, \qquad \nu_y^0(3) = \kappa_y^0(3) = 36.$$

$$\eta_y(1) = \begin{bmatrix} 1.0 - 0.2v \\ 1.0 + 0.2v \end{bmatrix}, \qquad \eta_y(2) = \begin{bmatrix} 2.0 - 0.2v \\ 2.0 + 0.2v \end{bmatrix}, \qquad \eta_y(3) = \begin{bmatrix} 3.0 - 0.2v \\ 3.0 + 0.2v \end{bmatrix},$$

$$\Lambda_y^0(s) = 1.2^2\, \nu_y^0(s)\, \Sigma_y(s); \qquad s \in \{1, 2, 3\}.$$

We assume that the prior distribution on the expected forecasts is not quite accurate, an investor underestimates the mean the first analyst and overestimates that of the second analyst. And the prior variance is assumed to be 1.2^2 times the true variance. As a result of this imprecise prior parameters, the investor assigns relatively small value for the prior number of observations $\nu_y^0(s)$ and $\kappa_y^0(s)$.

Result

Figure 9.8 shows the result of Bayesian learning on model parameters. We draw the transition matrix Q, return parameters $\mu_{ret}(s)$ and forecasts parameters $\mu_y(s)$ from the prior distribution, denoted by $t = 0$, and from the posterior distribution after 500 days of observations, denoted by $t = T$. The summary statistics of these parameters are shown in the plot. Generally, we can see improvement at $t = 500$ days; the sampled parameters have smaller biased and variance. Note that this is a result of a single sample path, hence we might not be able to seen any improvement for some parameters.

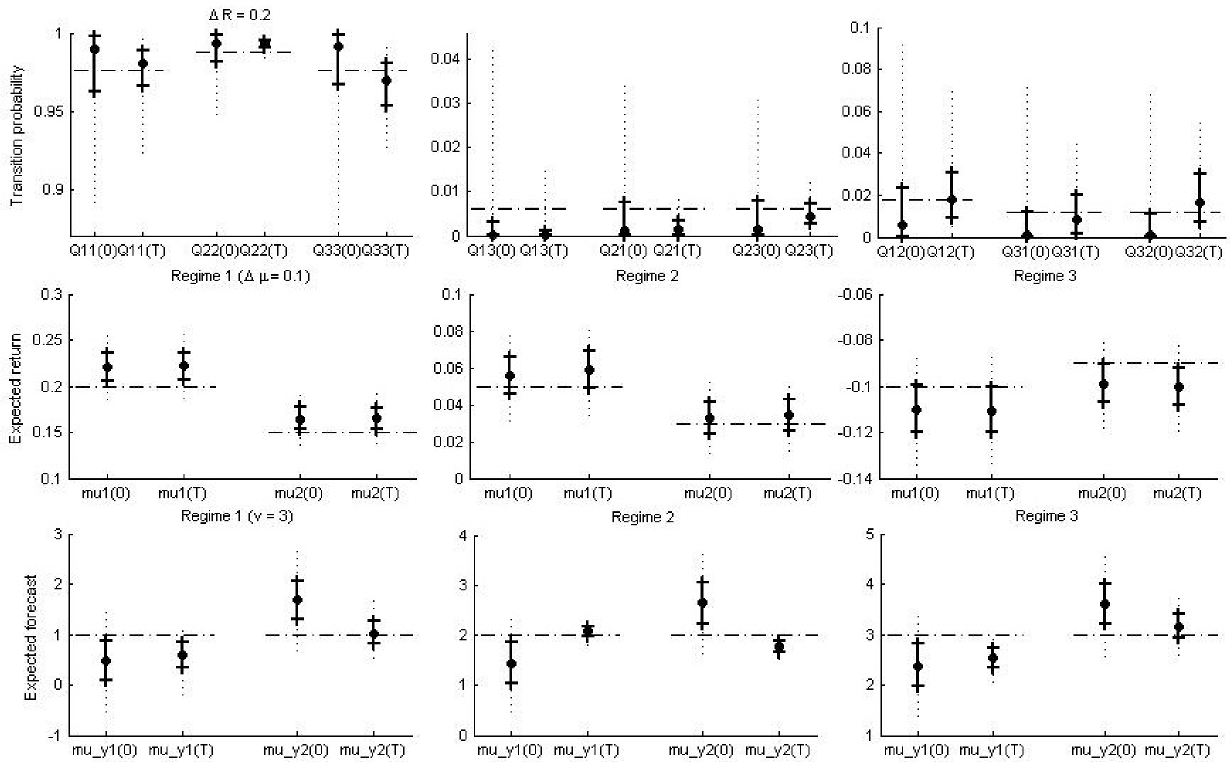

Figure 9.8: At forecast standard deviation $v = 3$, we plot the summary statistics of the transition matrix Q (first row), return parameters $\mu_{ret}(s)$ (second row) and forecasts parameters $\mu_y(s)$ (third row). For each vertical line, the circular marker represents the 50th percentile, the solid line connects the 25th and the 75th percentiles, and the dotted lines are extended to the 5th and the 95th percentiles. These parameters are drawn from the prior distribution at $t = 0$ and the posterior at $t = T = 500$ days. The horizontal dashed lines which are the true values of the parameters are added for comparison.

We plot the true market state, the inferred state probability as well as the portfolios obtained from solving the HMM-problem and the C-problem in Figure 9.9. This also compares the result from different market condition namely in the left column the analyst forecasts noise v is 1 while v is 3 for the right one. We can see that smaller forecast noise provides a more accurate state probability, and the difference between π_C and π_{HMM} depends on the state probability. When the forecast noise is low, an investor can easily infer the hidden states and the unknown parameters, and therefore π_{HMM} is closer to π_C.

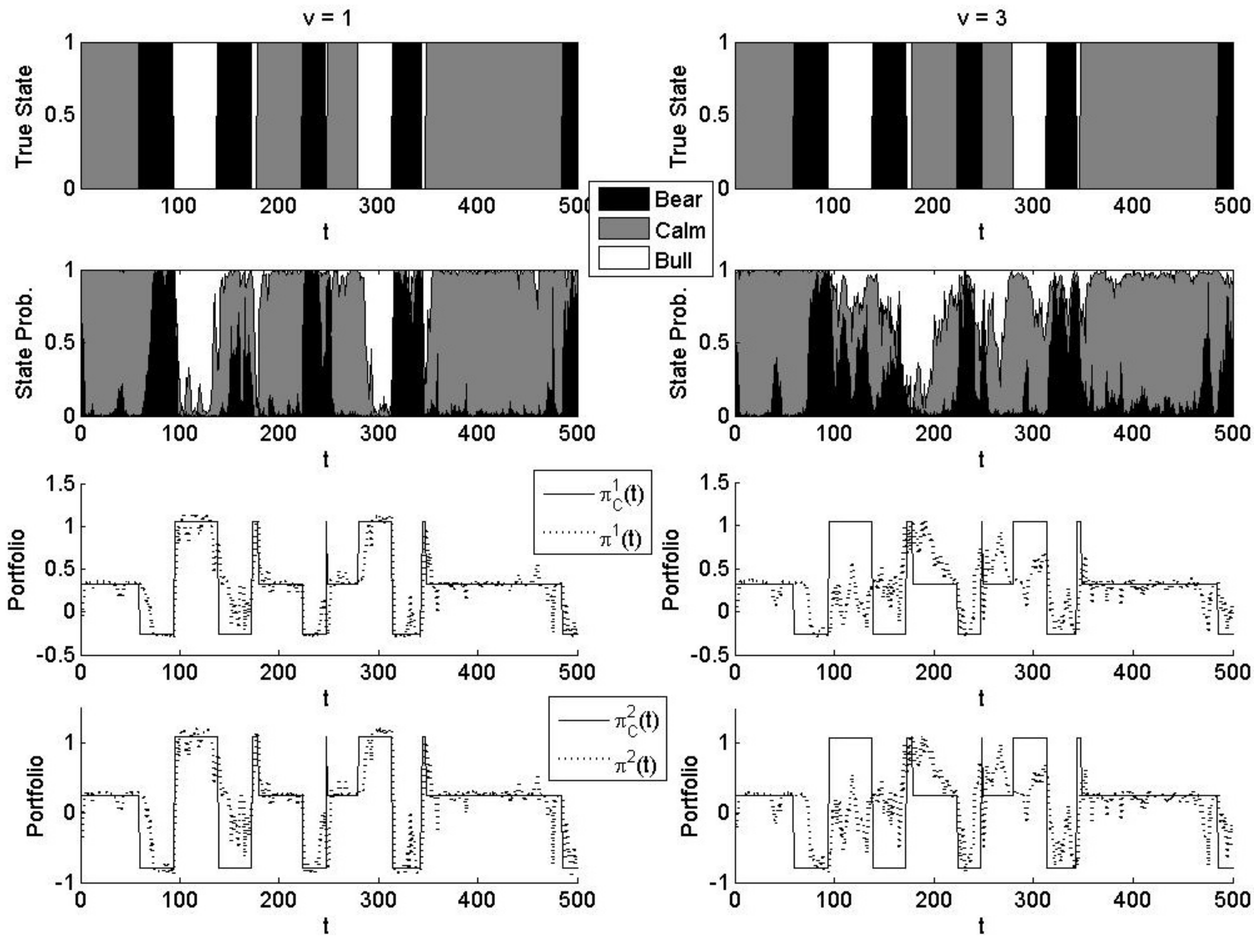

Figure 9.9: This figure compares the state probabilities and the portfolios from two market conditions when the analyst forecast noises v are 1 (left column) and 3 (right column). The plots in the first row show the true state. The plots in the second row show the probability of being in each state. The plots in the last two rows compares the risky asset holdings obtained from solving the HMM-problem (the dotted line) and from solving the C-problem (the solid line). The risk aversion parameter γ used is -5.

Next we compare the performance of the investors who solve the C-problem, the HMM-problem and the single-period mean-variance problem in Figure 9.10. For the single-period mean-variance problem, an investor also assumes the market follows the HMM with three states but instead of solving a dynamic problem, he/she repeatedly solves single-period problems. We can see superiority of the approximate methods over the single-period mean-variance approach in every set of parameters.

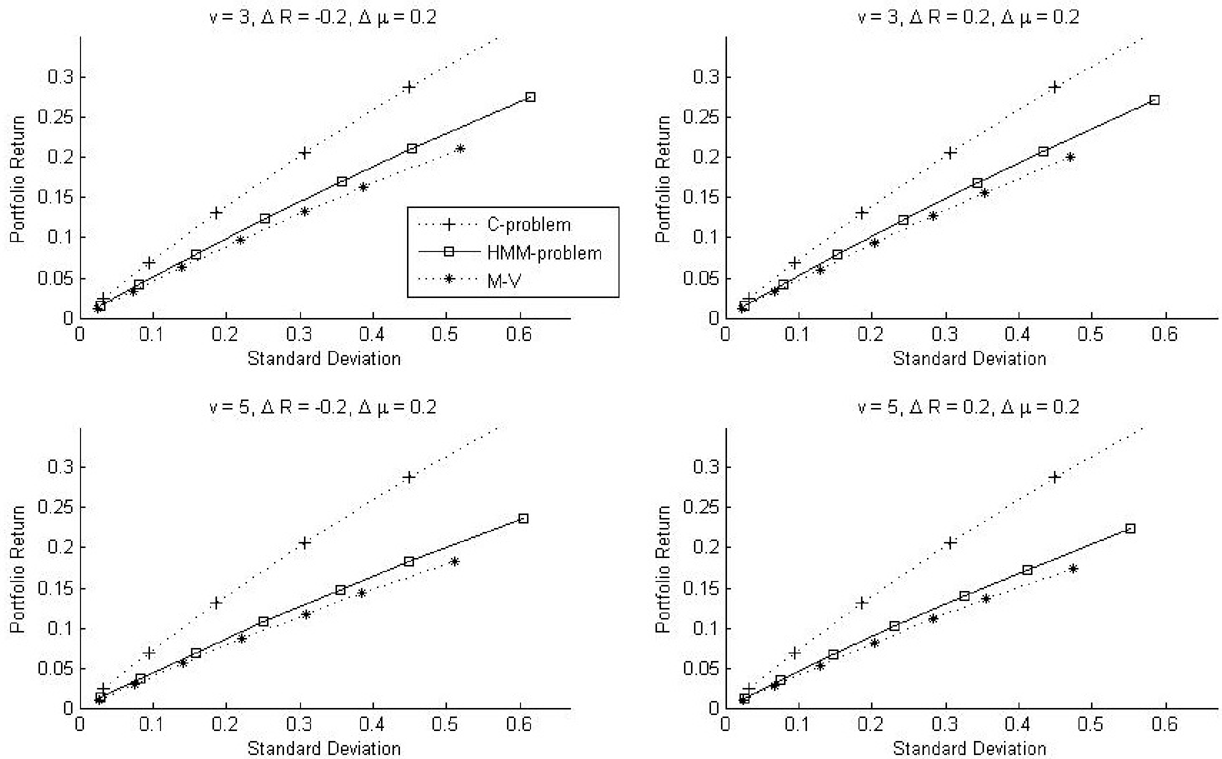

Figure 9.10: This shows the mean-standard deviation frontier of the portfolio constructed by solving the C-problem (dotted line with plus markers), the HMM-problem (solid line with square markers) and solving the single-period mean-variance problem (dotted line with star markers) under different sets of parameters. The investment horizon is 250 days.

Example : Setting 2

In this setting, we use the same parameters as the Setting 1 except the asset return parameters. The excess return vectors $\mu(s)$ and the covariance matrices $\Sigma(s)$ are

$$\mu(1) = \begin{bmatrix} 0.17 \\ 0.03 \end{bmatrix}, \qquad \mu(2) = \begin{bmatrix} 0.05 \\ 0.03 \end{bmatrix}, \qquad \mu(3) = \begin{bmatrix} -0.05 \\ 0.03 \end{bmatrix},$$

$$\Sigma(1) = \begin{bmatrix} 0.0225 & -0.0045 \\ -0.0045 & 0.0144 \end{bmatrix}, \quad \Sigma(2) = \begin{bmatrix} 0.0225 & 0.0045 \\ 0.0045 & 0.0144 \end{bmatrix}, \quad \Sigma(3) = \begin{bmatrix} 0.0225 & 0.0169 \\ 0.0169 & 0.0225 \end{bmatrix}.$$

The correlations in this setting are -0.25, 0.25 and 0.75 for regime 1, 2 and 3 respectively. We maintain the same structure of the prior parameters as in the Setting 1.

Similar to the previous section, we compare the performance of the investors who solve the C-problem, the HMM-problem and the single-period mean-variance problem in Figure 9.11. Again, the superiority of the approximate methods over the mean-variance approach can be seen in all settings.

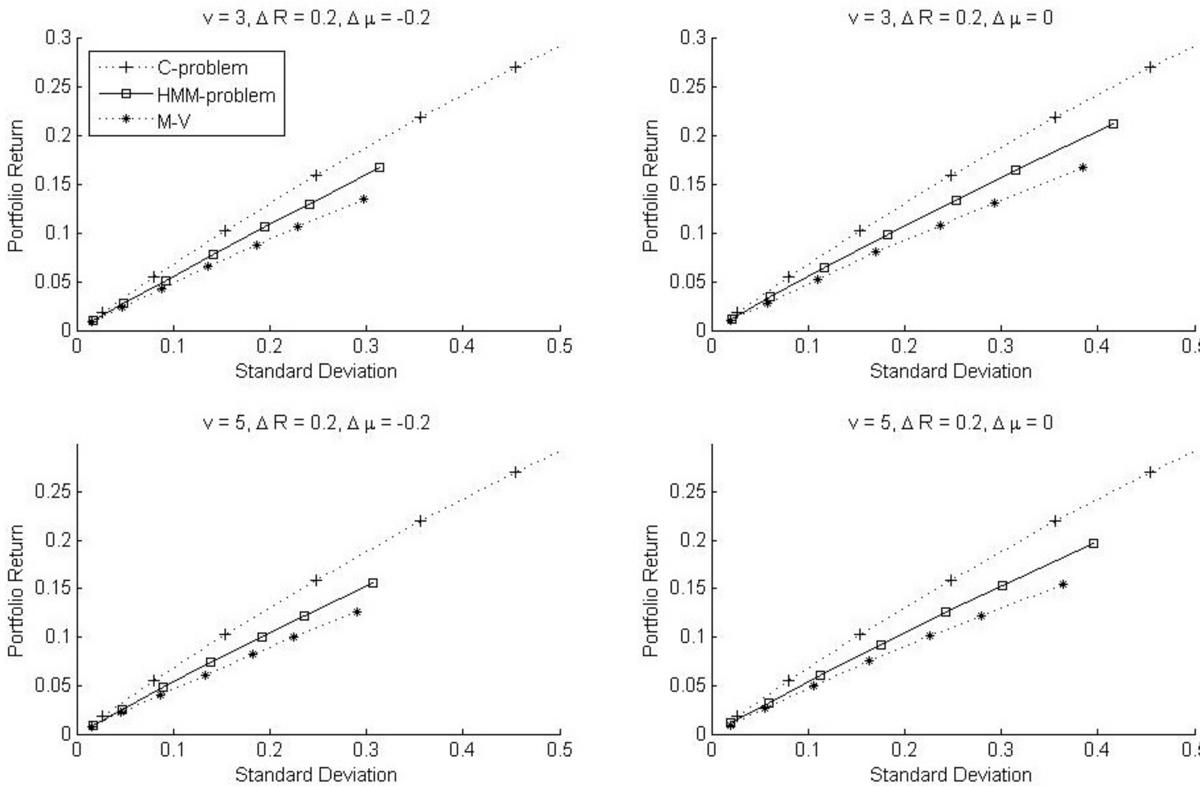

Figure 9.11: This shows the mean-standard deviation frontier of the portfolio constructed by solving the C-problem (dotted line with plus markers), the HMM-problem (solid line with square markers) and solving the single period mean-variance problem (dotted line with star markers) under different sets of parameters. The investment horizon is 250 days.

Chapter 10

Sampling from Posterior Distribution

In this chapter, we present the Gibbs sampler which is used to sample the state sample path s^t, the rate matrix R and the state-dependent parameters ξ from the posterior distribution ρ_t. We begin by stating the assumptions on the parameters distribution then we summarize the Gibbs sampler from Fruhwirth-Schnatter [27].

10.1 Prior Distribution

Transition matrix

We first discretize the Markov chain in (7.1) to make it compatible with discrete observations by defining the regime transition probability matrix Q as follows.

$$Q_{jk} = \begin{cases} R_{jk}\,\delta & \text{if} \quad j \neq k \\ 1 + R_{jk}\,\delta & \text{otherwise} \end{cases} \tag{10.1}$$

We assume Dirichlet as a prior distribution for each row of the transition matrix Q,

$$Q_{j\cdot} \sim Dir(e_{j1}, e_{j2}, ..., e_{jS}), \quad j = 1, 2, ..., S.$$

The parameter e_{jk} can be viewed as the prior number of switching from regime j to k. Condition on knowing the state sample path s^t, the posterior distribution of Q is

$$Q_{j\cdot} \sim Dir(e_{j1} + N_{j1}(s^t), e_{j2} + N_{j2}(s^t), ..., e_{jK} + N_{jK}(s^t)),$$

where $N_{jk}(s^t)$ is the number of transitions from state j to k along the sample path s^t.

Returns and forecasts

Condition on the hidden state and price parameters, asset returns which are defined by

$$ret_t^i \equiv \ln(\frac{p_{t+1}^i}{p_t^i}), \quad i = 1, 2, ..., I \tag{10.2}$$

are normally distributed,

$$ret_t = [ret_t^1, ..., ret_t^I]' \sim N(\mu_{ret}(s_t), \Sigma_{ret}(s_t)).$$

The relationship between $(\mu(s), \Sigma(s))$ and $(\mu_{ret}(s), \Sigma_{ret}(s))$ are

$$\mu_{ret}^i(s) = \left(\mu_i(s) + r - \frac{1}{2}\Sigma_{ii}(s) \right) \delta \quad \text{and} \quad \Sigma_{ret}(s) = \Sigma(s)\delta. \tag{10.3}$$

We also recall that the analyst forecast is normally distributed, $y_t \sim N(\mu_y(s_t), \Sigma_y(s_t))$.

We assume normal-inverse-Wishart (NIW) as a prior distribution for both $(\mu_{ret}, \Sigma_{ret})$ and (μ_y, Σ_y). Since the returns and the forecasts have exactly the same distribution, we omit writing down the distribution of the forecast parameters.

$$\Sigma_{ret}(s) \sim IW\Big(\Lambda_{ret}^0(s), \nu_{ret}^0(s)\Big) \quad \text{and} \quad \mu_{ret}(s) \,\Big|\, \Sigma_{ret}(s) \sim N\Big(\eta_{ret}^0(s), \frac{\Sigma_{ret}(s)}{\kappa_{ret}^0(s)}\Big)$$

for $s \in \{1, 2, ..., S\}$, and $\Lambda_{ret}^0, \nu_{ret}^0, \eta_{ret}^0$ and κ_{ret}^0 are prior distribution parameters. $\nu_{ret}^0(s)$ and $\kappa_{ret}^0(s)$ can be interpreted as the number of prior observations which have sample mean $\eta_{ret}^0(s)$ and sum of squared $\Lambda_{ret}^0(s)$.

Condition on knowing the state sample path s^t and the return realization $\{ret_1, ret_2, ..., ret_{n(s)}\}$, where $n(s)$ is the number of observations from state s, the posterior distribution of $(\mu_{ret}, \Sigma_{ret})$ is also NIW with parameters

$$\begin{aligned}
\Lambda_{ret}^n(s) &= \Lambda_{ret}^0(s) + \frac{\kappa_{ret}^0(s)\, n(s)}{\kappa_{ret}^0(s) + n(s)} \big(\overline{ret}(s) - \eta_{ret}^0(s)\big)\big(\overline{ret}(s) - \eta_{ret}^0(s)\big)' \\
&\quad + \sum_{j=1}^{n(s)} \big(ret_j - \overline{ret}(s)\big)\big(ret_j - \overline{ret}(s)\big)' && (10.4) \\
\nu_{ret}^n(s) &= \nu_{ret}^0(s) + n(s) && (10.5) \\
\eta_{ret}^n(s) &= \frac{\kappa_{ret}^0(s)\, \eta_{ret}^0(s) + n(s)\, \overline{ret}(s)}{\kappa_{ret}^0(s) + n(s)} && (10.6) \\
\kappa_{ret}^n(s) &= \kappa_{ret}^0(s) + n(s), && (10.7)
\end{aligned}$$

where $\overline{ret}(s)$ is a sample mean of the returns in state s.

10.2 Gibbs sampling

Gibbs sampling is a natural method to sample the hidden variables/parameters in the HMM as it allows us to sample from the joint posterior distribution given that we can sample from conditional distributions. The assumptions on the prior distributions in the previous section provide us easy ways to conditionally sample $Q, (\mu_{ret}, \Sigma_{ret})$ and (μ_y, Σ_y). Also the states

s^t are straightforward to sample given the model parameters are know. We summarize the Gibbs sampler for the HMM as follows.

Let m be an index of the sample. We start by generating a sample path s_0^t. Then for $m \in \{1, 2, ..., M\}$,

- Given s_{m-1}^t, sample the transition matrix Q from its posterior distribution

$$Q_{j\cdot}^m \sim Dir(e_{j1} + N_{j1}(s_{m-1}^t), e_{j2} + N_{j2}(s_{m-1}^t), ..., e_{jK} + N_{jK}(s_{m-1}^t)),$$

- Given s_{m-1}^t, calculate the parameters of NIW posterior distribution using (10.4) - (10.7). Then sample the return parameters from

$$\Sigma_{ret}^m(s) \sim IW\Big(\Lambda_{ret}^n(s), \nu_{ret}^n(s)\Big)$$
$$\mu_{ret}^m(s) \,\Big|\, \Sigma_{ret}^m(s) \sim N\Big(\eta_{ret}^n(s), \frac{\Sigma_{ret}(s)}{\kappa_{ret}^n(s)}\Big)$$

  for all state $s \in \{1, 2, ..., S\}$. And repeat this for the forecast parameters, $\mu_y^m(s)$ and $\Sigma_y^m(s)$.

- Given Q^m and $\big\{\mu_{ret}^m(s), \Sigma_{ret}^m(s), \mu_y^m(s), \Sigma_y^m(s); s = 1, 2, ..., S\big\}$, sample a sample path s_m^t using the multi-move sampling. First, we need to calculate filtered state probabilities $\mathbb{P}(s_\tau \,|\, ret^{\tau-1}, y^\tau)$ for all $\tau = 1, 2, .., t$. This can be done recursively using the two following equations:

$$\begin{cases} \mathbb{P}(s_\tau \,\big|\, ret^{\tau-1}, y^\tau) = \dfrac{\mathbb{P}(y_\tau \,\big|\, s_\tau)\, \mathbb{P}(s_\tau \,\big|\, ret^{\tau-1}, y^{\tau-1})}{\sum_{k=1}^S \mathbb{P}(y_\tau \,\big|\, s_\tau = k)\, \mathbb{P}(s_\tau = k \,\big|\, ret^{\tau-1}, y^{\tau-1})}, & \text{if } \tau = t \\ \mathbb{P}(s_\tau \,\big|\, ret^{\tau}, y^\tau) = \dfrac{\mathbb{P}(ret_\tau, y_\tau \,\big|\, s_\tau)\, \mathbb{P}(s_\tau \,\big|\, ret^{\tau-1}, y^{\tau-1})}{\sum_{k=1}^S \mathbb{P}(ret_\tau, y_\tau \,\big|\, s_\tau = k)\, \mathbb{P}(s_\tau = k \,\big|\, ret^{\tau-1}, y^{\tau-1})}, & \text{if } \tau \le t \end{cases}$$

  and

$$\begin{aligned} \mathbb{P}(s_\tau \,\big|\, ret^{\tau-1}, y^{\tau-1}) &= \sum_{k=1}^S \mathbb{P}(s_\tau, s_{\tau-1} = k \,\big|\, ret^{\tau-1}, y^{\tau-1}) \\ &= \sum_{k=1}^S Q_{ks_\tau}\, \mathbb{P}(s_{\tau-1} = k \,\big|\, ret^{\tau-1}, y^{\tau-1}). \end{aligned}$$

  Next the sample path s_m^t is sampled by first sample the state at time t, $s_{t,m}$, from the filtered state probability

$$\mathbb{P}(s_t \,\big|\, ret^{t-1}, y^t).$$

  Then for $\tau = t-1, t-2, ..., 1$, sample $s_{\tau,m}$ from the conditional distribution

$$\mathbb{P}(s_\tau \,\big|\, s_{\tau+1}, ret^\tau, y^\tau) = \frac{Q_{s_\tau, s_{\tau+1}}\, \mathbb{P}(s_\tau \,\big|\, ret^\tau, y^\tau)}{\sum_{k=1}^S Q_{k, s_{\tau+1}}\, \mathbb{P}(s_\tau = k \,\big|\, ret^\tau, y^\tau)}.$$

By following this procedure, the samples $\left\{s_m^t, Q^m, \mu_{ret}^m(s), \Sigma_{ret}^m(s), \mu_y^m(s), \Sigma_y^m(s)\right\}_{m=1,2,\ldots,M}$ have distribution ρ_t.

Chapter 11

Performance Evaluation

In this chapter, we discuss an approach to evaluate the performance of the approximate methods in Chapter 9. First, we formulate a dual problem using the information relaxation technique. (See Brown et al. [15] for details.) The dual problem value function will serve us as an upper-bound to the $\underline{\text{HMM}}$-problem. The key idea of the dual problem is two-fold. First, we relax (allow investors to have a better filtration) the HMM-problem such that the relaxation problem is easy to solve. Then we penalize the relaxation problem with penalties those are tractable and reflect the differences between the original and dual problem filtrations. As the HMM-problem suffers from the curse of dimensionality, we are interested in the dual problem that does not. To achieve this, the posterior distribution dynamics (8.1) must not be part of the dual problem dynamics. Then we show by numerical experiments that the performance of the proposed approximate methods is reasonably close to the dual upper-bound.

$\underline{\textbf{HMM}}$-problem

Solely for evaluation purpose, we introduce a slightly different version of the HMM-problem, we called it the $\underline{\text{HMM}}$-problem. The reason that we raise the $\underline{\text{HMM}}$-problem is we are not be able to formulate a dual problem of the original HMM-problem which does not suffer from the curse of dimensionality, i.e. the dual problem is as hard as the HMM-problem. However, if we consider a mildly inefficient version of the HMM-problem, we will be able to find an easy dual problem. Specifically, an investor that solves the $\underline{\text{HMM}}$-problem only utilizes the analyst forecasts to update the posterior distribution,

$$\underline{\rho}_t(s^t, R, \xi) \propto L(y^t; s^t, R, \xi)\, \rho_0(s_0, R, \xi). \tag{11.1}$$

This is in contrast to the HMM-problem where an investor use both prices and forecasts to update his/her belief. We shall now discuss the dual problem of the $\underline{\text{HMM}}$-problem as follows.

11.1 Information Relaxation

Let $\underline{\mathcal{F}}_t$ denote a σ-algebra that describes information available at time t to an investor who solves the $\underline{\text{HMM}}$-problem. By setting, $\{\underline{\mathcal{F}}_t\}$ is generated only by observing y_t. Next, we introduce two sets of information relaxation of the $\underline{\text{HMM}}$-problem:

- $\overline{\mathcal{G}}_t$ describes information available at time t to an investor who observes $\{s^T, p^t, y^T\}$ and the parameters $\{R, \xi\}$. This relaxation allows the investor to see the future sample paths of the state s^T and analyst forecast y^T. The investor also knows the true rate matrix R and the state-dependent parameters ξ.
- $\mathcal{G}_t$ describes information available at time t to an investor who observes only the forecasts y^T. This set of information allows the investor to see the future sample path analyst forecast y^T, and therefore he/she is able to construct the posterior distribution $\underline{\rho}_T(s^T, R, \xi)$. The reason we introduced $\mathcal{G}_t$ this way is to avoid the dynamics of posterior distribution, i.e.
$$\underline{\rho}_0 = \underline{\rho}_1 = ... = \underline{\rho}_T.$$

It is trivial to see that $\underline{\mathcal{F}}_t \subseteq \mathcal{G}_t \subseteq \overline{\mathcal{G}}_t$. Now the reader may see why we consider the relaxation of the $\underline{\text{HMM}}$-problem instead of the relaxation of the HMM-problem. This is to prevent an investor to incorporate the future asset prices into the decision making process.

Now, we consider a dynamic portfolio choice problem where an investor maximizes the expected power utility of the final wealth,

$$\max_{\pi(\cdot)\in\mathcal{A}^{\overline{\mathcal{G}}}} \mathbb{E}\left[\frac{x_T^\gamma}{\gamma}\right], \tag{11.2}$$

subject to the wealth dynamics (8.2), and $\mathcal{A}^{\overline{\mathcal{G}}}$ is a set of admissible portfolio strategies those are adapted to $\{\overline{\mathcal{G}}_t\}$. This problem is a trivial upper-bound of the $\underline{\text{HMM}}$-problem. The value function of this problem will be used in the dual problem formulation.

Proposition 10. *The value function of the relaxation problem* (11.2)*, denoted by V_R, is*

$$V_R(t, x_t; s_T, R, \xi) = \frac{x_t^\gamma}{\gamma} g_R(t; s^T, R, \xi) \tag{11.3}$$

where $g_R(t; s^T, R, \xi)$ solves the following DE,

$$\begin{cases} \quad 0 &= \dot{g}_R(t) + \left\{\gamma r + \frac{1}{2}\frac{\gamma}{1-\gamma}\theta(s_t)\right\} g_R(t) \\ g_R(T) &= 1. \end{cases} \tag{11.4}$$

Recall that $\theta(s) = \mu(s)'\Sigma^{-1}(s)\mu(s)$.

Proof. See Chapter 12. □

Penalties

To obtain a tight upper-bound to the HMM-problem, we introduce penalties on the extra-information: the hidden states s^T and the model parameters R and ξ. At period t, the penalty is constructed as follows.

$$Z(t, x_t, \pi) = \mathbb{E}\Big[V_R(t+1, x_{t+1}) - V_R(t, x_t) \,\Big|\, \overline{\mathcal{G}}_t\Big] - \mathbb{E}\Big[V_R(t+1, x_{t+1}) - V_R(t, x_t) \,\Big|\, \mathcal{G}_t\Big] \tag{11.5}$$

Intuitively, the penalty reflects the difference in value between knowing $\overline{\mathcal{G}}_t$ and $\mathcal{G}_t$. The first term on the RHS is easy to calculate by solving (11.4) since $\{s^T, R, \xi\}$ are known to the investor. The second term can be calculated using the Monte Carlo integration. We first sample $\{s^T_m, R^m, \xi^m\}_{m=1,2,\ldots,M}$ from the posterior distribution $\underline{\rho}_t$ then solve (11.4) for $g_R(t; s^T_m, R^m, \xi^m)$ then calculate the expectation.

Lemma 11. *The penalty $Z(t, x_t, \pi)$ in (11.5) is dual feasible.*

Proof. This follows directly from Proposition 2.2 in Brown et al. [15]. □

11.2 Dual Problem

In the dual problem, an investor maximizes the utility of the final wealth subtracted by the penalties on the extra-information,

$$\max_{\pi(\cdot)\in\mathcal{A}^{\overline{\mathcal{G}}}} \mathbb{E}\left[\frac{x_T^\gamma}{\gamma} - \sum_{t=0}^{T-1} Z(t, x_t, \pi(t))\right], \tag{11.6}$$

subject to the wealth dynamics (8.2). We can see that the posterior distribution $\underline{\rho}_t$ does not change over time in the dual problem, and therefore it is no longer a state variable. This allows us to solve the dual problem. The solution to the dual problem is summarized in the following Proposition.

Proposition 12. *The value function to the dual problem* (11.6)*, denoted by V_D, is*

$$V_D(t, x_t) = \frac{x_t^\gamma}{\gamma} h(t). \tag{11.7}$$

where $h(t)$ solves the following differential equation

$$\begin{cases} 0 &= \dot{h}(t) + \gamma r h(t) + \dfrac{1}{2}\dfrac{\gamma}{1-\gamma}\Big\{\theta(s_t) g_R(t) - \widetilde{\theta}(t)\Big\} \\ & \quad + \dfrac{1}{2}\dfrac{\gamma}{1-\gamma}\Big\{\mu(s_t)\big(h(t) - g_R(t)\big) + \widetilde{\mu}(t)\Big\}'\Big[\Sigma(s_t)(h(t) - g_R(t)) + \widetilde{\Sigma}(t)\Big]^{-1} \\ & \quad \times\Big\{\mu(s_t)\big(h(t) - g_R(t)\big) + \widetilde{\mu}(t)\Big\} \\ h(T) &= 0, \end{cases} \tag{11.8}$$

and

$$\widetilde{\theta}(t) = \mathbb{E}_{\underline{\rho}_T}[\theta(s)g_R(t)],\ \ \widetilde{\mu}(t) = \mathbb{E}_{\underline{\rho}_T}[\mu(s_t)g_R(t)]\ \ \textit{and}\ \ \widetilde{\Sigma}(t) = \mathbb{E}_{\underline{\rho}_T}[\Sigma(s)g_R(t)]. \tag{11.9}$$

Proof. See Chapter 12. □

The differential equation (11.8) can be solved in two steps:

1. Calculate $\{\widetilde{\theta}(t), \widetilde{\mu}(t), \widetilde{\Sigma}(t)\}_{t=0,1,\ldots,T}$ using the Monte Carlo integration. First, sample $\{s_m^T, R^m, \xi^m\}_{m=1,2,\ldots,M}$ from the posterior distribution $\underline{\rho}_T(s^T, R, \xi)$ and calculate $g_R(t; s_m^T, R^m, \xi^m)$ using (11.4) then compute the expectation .

2. Solve for $h(t)$ using numerical method.

Now we summarize the relationship between the HMM-problem and the dual problem (11.6) in the following Proposition.

Proposition 13. *The value function V_D is an upper-bound to that of the HMM-problem.*

Proof. This follows from the fact that $\underline{\mathcal{F}}_t \subseteq \mathcal{G}_t \subseteq \overline{\mathcal{G}}_t$ and Proposition 2.3 in Brown et al. [15]. □

11.3 Examples

In this section, we illustrate the use of the dual bound in evaluating performance of the approximate methods. We first simulate market with various conditions then adopt the approximate methods in Chapter 9 to solve the HMM-problem. To evaluate the investor performance, we compare the realization of the final wealth utility with three types of bounds.

1. *Complete information bound*
   This is a naive upper-bound obtained by assuming that the information on market state and the model parameters is available to the investor. Therefore the bound is the value function of the C-problem (9.4),

$$\text{C-bound} = \mathbb{E}[V_C(t=0, x_0, s_0)].$$

2. *Dual bound*
   This is a tighter upper-bound obtained by solving the dual problem (11.6),

$$\text{D-bound} = \mathbb{E}[V_D(t=0, x_0)].$$

3. *Risk-free bound*
   This a lower-bound obtained by investing all in the risk-free asset. The utility associated with this risk-free investment is $\frac{1}{\gamma}x_0^\gamma e^{\gamma r T}$.

Example : Setting 1

With the same true and prior parameters as in the Setting 1 in Section 9.4, Figure 11.1 compares the performance of the investor who solves the HMM-problem with the C-bound, D-bound and risk-free bound under various analyst forecast noises, different errors in prior return $\Delta\mu$ and difference errors in prior transition rate ΔR. We can see that the performance of the investor using the proposed approximate methods is fairly close to the dual upper-bound. The investor begins with \$100 in initial wealth and has 250 days of investment horizon.

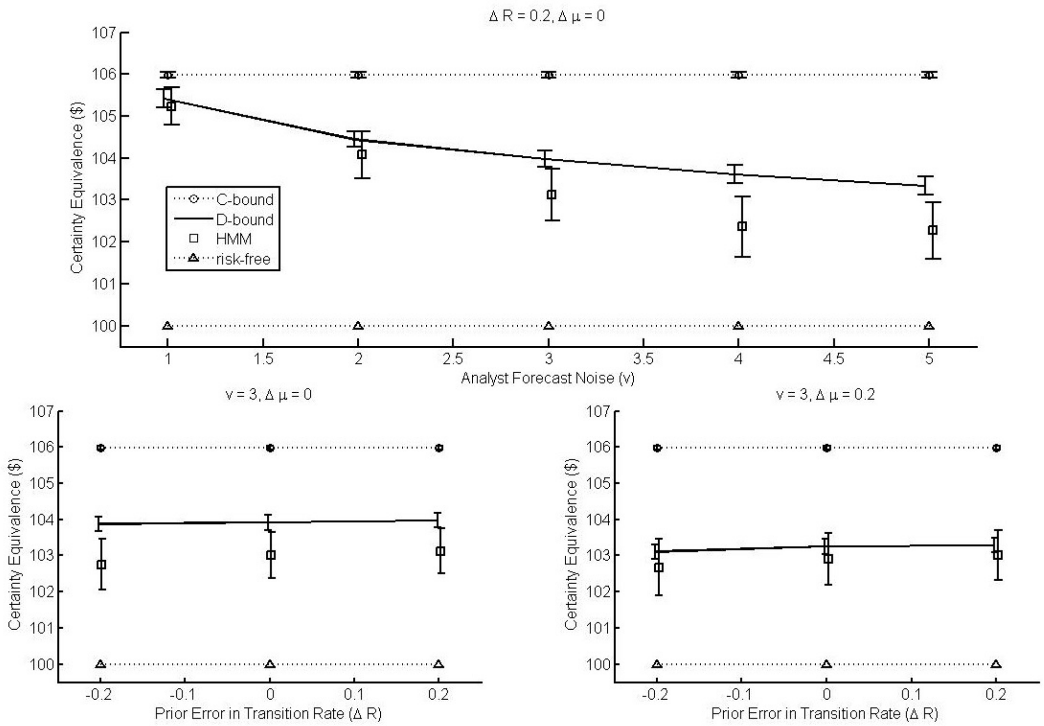

Figure 11.1: This compares the certainty equivalence of the investor who solves the HMM-problem (square markers with error bars) with the C-bound (dotted line with circular markers), the D-bound (solid line) and the risk-free bound (dotted line with triangular markers). The investment horizon is 250 days and the risk-aversion parameter is -5. The error bars represent ± two standard deviations. The certainty equivalence is computed from 1000 simulations.

Example : Setting 2

In this example, we use the same true and prior parameters as in the Setting 2 in Section 9.4. We compare the the performance of the investor who solves the HMM-problem with the C-bound, D-bound and risk-free bound under various settings in Figure 11.2.

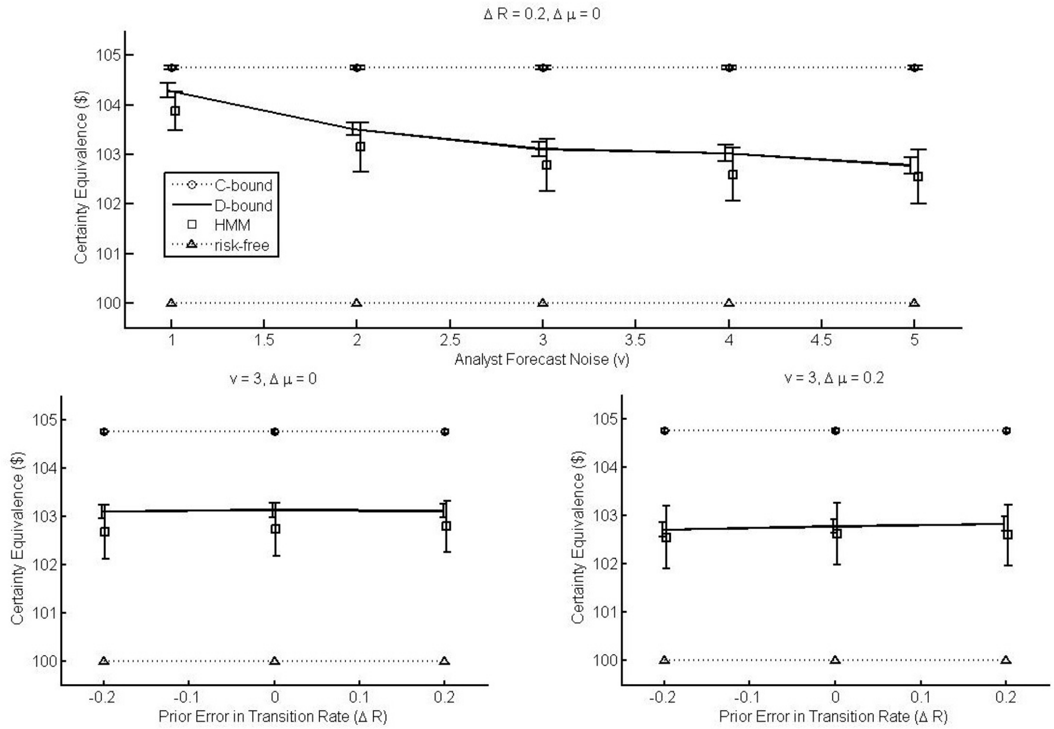

Figure 11.2: This compares the certainty equivalence of the investor who solves the HMM-problem (square markers with error bars) with the C-bound (dotted line with circular markers), the D-bound (solid line) and the risk-free bound (dotted line with triangular markers). The investment horizon is 250 days and the risk-aversion parameter is -5. The error bars represent ± two standard deviations. The certainty equivalence is computed from 1000 simulations.

Example : Setting 3

We consider a three-regime market similarly to that in Setting 1 and 2 but with different transition rate and return parameters. The transition rate of the Markov chain is

$$R = \begin{bmatrix} -24.0 & 16.8 & 7.2 \\ 3.0 & -6.0 & 3.0 \\ 19.2 & 4.8 & -24.0 \end{bmatrix} \text{ per year.}$$

The excess return vectors $\mu(s)$ and the covariance matrices $\Sigma(s)$ are

$$\mu(1) = \begin{bmatrix} 0.12 \\ 0.16 \end{bmatrix}, \quad \mu(2) = \begin{bmatrix} 0.04 \\ 0.05 \end{bmatrix}, \quad \mu(3) = \begin{bmatrix} -0.10 \\ -0.14 \end{bmatrix},$$

$$\Sigma(1) = \begin{bmatrix} 0.0400 & 0.0200 \\ 0.0200 & 0.0625 \end{bmatrix}, \quad \Sigma(2) = \begin{bmatrix} 0.0256 & 0.0173 \\ 0.0173 & 0.0324 \end{bmatrix}, \quad \Sigma(3) = \begin{bmatrix} 0.0400 & 0.0375 \\ 0.0375 & 0.0625 \end{bmatrix}.$$

The correlations across assets are 0.4, 0.6 and 0.75 respectively.

The prior parameters for the transition matrix Q are

$$\begin{bmatrix} e_{11} & e_{12} & e_{13} \\ e_{21} & e_{22} & e_{23} \\ e_{31} & e_{32} & e_{33} \end{bmatrix} = \begin{bmatrix} 19.80 & 2.02 & 0.68 \\ 1.25 & 75.87 & 1.25 \\ 1.45 & 0.48 & 22.19 \end{bmatrix}.$$

This matrix is calculated from the average number of state transitions within 125 days of a Markov chain with rate

$$\begin{bmatrix} -30.0 & 22.5 & 7.5 \\ 4.0 & -8.0 & 4.0 \\ 15.0 & 5.0 & -20.0 \end{bmatrix} \text{ per year.} \tag{11.10}$$

Observe that the rate matrix used to constructed the prior distribution is quite different from the true one. In this example, we let the precision in prior asset return $\Delta\mu$ to be 0 while we vary the forecast standard deviation v from 0.5 to 5.

Figure 11.3 compares the performance of the investor who solves the HMM-problem with the C-bound, D-bound and risk-free bound. The investor starts with \$100 and has 125 days to invests in the market. We can see that the performance of the investor using the proposed approximate methods is very close to the theoretical dual bound.

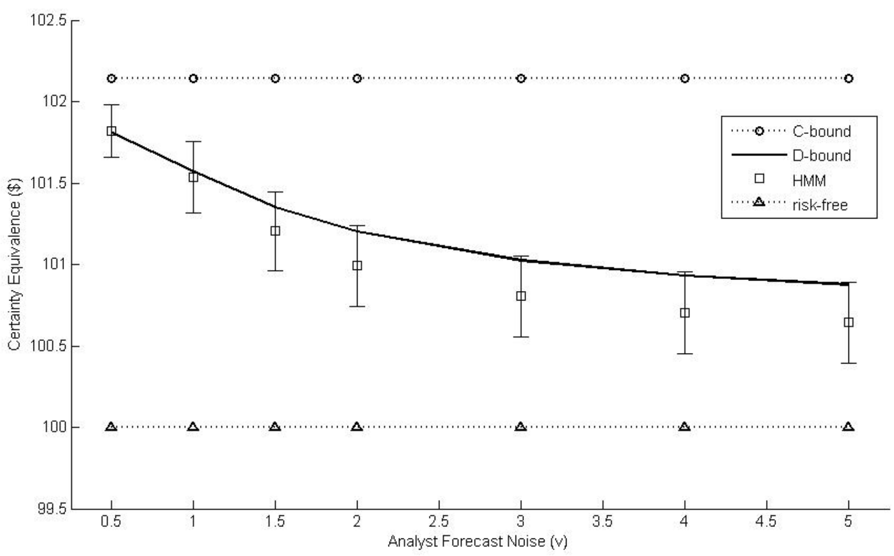

Figure 11.3: This compares the certainty equivalence of the investor who solves the HMM-problem (square markers with error bars) with the C-bound (dotted line with circular markers), the D-bound (solid line) and the risk-free bound (dotted line with triangular markers). The investment horizon is 125 days and the risk-aversion parameter is -5. The error bars represent $\pm$ two standard deviations. The certainty equivalence of the investor is computed from 3000 simulations.

Chapter 12

Proofs

12.1 Proof of Proposition 9

Let $V_C(t, x, s)$ be the value function of the C-problem. The Bellman equation writes

$$\begin{aligned}
&V_C(t, x_t, s_t) \\
&= \max_{\pi} \mathbb{E}_{x_{t+\delta}, s_{t+\delta}} \Big[V_C(t+\delta, x_{t+\delta}, s_{t+\delta}) \Big] \\
&= \max_{\pi} \mathbb{E}_{x_{t+\delta}} \Big[\mathbb{E}_{s_{t+\delta}} \Big[V_C(t+\delta, x_{t+\delta}, s_{t+\delta}) \Big] \Big] \\
&= \max_{\pi} \mathbb{E}_{x_{t+\delta}} \left[V_C(t+\delta, x_{t+\delta}, s_t) + \delta \sum_{k=1}^{S} R_{s_t k} V_C(t+\delta, x_{t+\delta}, k) + o(\delta) \right] \\
&= \max_{\pi} \left[\begin{array}{l} V_C(t, x_t, s_t) + \delta \dfrac{\partial}{\partial t} V_C(t, x_t, s_t) + \delta\, x_t \big(r + \pi' \mu(s_t) \big) \dfrac{\partial}{\partial x} V_C(t, x_t, s_t) \\ + \delta\, \dfrac{1}{2} x_t^2\, \pi' \, \Sigma(s_t)\, \pi \dfrac{\partial^2}{\partial x^2} V_C(t, x_t, s_t) + \delta \displaystyle\sum_{k=1}^{S} R_{s_t k} V_C(t, x_t, k) + o(\delta) \end{array} \right]
\end{aligned} \tag{12.1}$$

Neglect the $o(\delta)$ term, the optimal portfolio is

$$\pi_C^*(t) = - \left[x_t\, \Sigma(s_t)\, \frac{\partial^2}{\partial x^2} V_C(t, x_t, s_t) \right]^{-1} \mu(s_t)\, \frac{\partial}{\partial x} V_C(t, x_t, s_t). \tag{12.2}$$

Substitute $\pi_C^*(t)$ back into the Bellman equation (12.1). Then we can see that (12.1) is solved by the value function V_C defined in (9.4) - (9.5).

12.2 Proof of Proposition 10

Let V_R be the value function of this problem. The Bellman equation of this problem is

$$\begin{aligned} V_R(t, x_t) &= \max_{\pi} \mathbb{E}_{x_{t+1}}\Big[V_R(t+1, x_{t+1}) \,\Big|\, \mathcal{G}_t^C\Big] \\ &= V_R(t, x_t) + \delta \frac{\partial}{\partial t} V_R(t, x_t) \\ &\quad + \delta \max_{\pi} \left\{ x_t\big(r + \pi' \mu(s_t)\big) \frac{\partial}{\partial x} V_R(t, x_t) + \frac{1}{2} x_t^2 \, \pi' \, \Sigma(s_t) \, \pi \frac{\partial^2}{\partial x^2} V_R(t, x_t) \right\} + o(\delta). \end{aligned} \tag{12.3}$$

The optimal π^* is

$$\pi^* = -\left[x_t \Sigma(s_t) \frac{\partial^2}{\partial x^2} V_R(t, x_t) \right]^{-1} \mu(s_t) \frac{\partial}{\partial x} V_R(t, x_t). \tag{12.4}$$

Substitute π^* back into the Bellman equation (12.3). Then we can see that (12.3) is solved by the value function V_R defined in (11.3) - (11.4).

12.3 Proof of Proposition 12

Let V_D be the value function of the dual problem. The Bellman equation of the dual problem is

$$V_D(t, x_t) = \max_{\pi} \left\{ -Z(t, x_t, \pi) + \mathbb{E}\Big[V_D(t+1, x_{t+1}) \,\Big|\, \mathcal{G}_t^C \Big] \right\} \tag{12.5}$$

Apart from the argument t and x_t, the value function V_D also depends on s^T, R, ξ and $\underline{\rho}_T$; however, there parameters do not change over time, hence we omit writing them.

We first consider the terms from $Z(t, x_t, \pi(t))$,

$$\begin{aligned} &\mathbb{E}_{x_{t+1}}\Big[V_R(t+1, x_{t+1}) - V_R(t, x_t) \,\Big|\, \mathcal{G}_t^C \Big] \\ &\quad = \delta \, g_R(t) \, \frac{x_t^\gamma}{\gamma} \left\{ -\frac{1}{2} \frac{\gamma}{1-\gamma} \, \theta(s_t) + \gamma \pi' \, \mu(s_t) + \frac{1}{2} \gamma(\gamma - 1) \pi' \, \Sigma(s_t) \, \pi \right\} + o(\delta) \end{aligned} \tag{12.6}$$

$$\begin{aligned} &\mathbb{E}_{s^T, R, \xi, x_{t+1}}\Big[V_R(t+1, x_{t+1}) - V_R(t, x_t) \,\Big|\, \mathcal{G}_t \Big] \\ &\quad = \mathbb{E}_{\underline{\rho}_T}\Big[\mathbb{E}_{x_{t+1}}\Big[V_R(t+1, x_{t+1}) - V_R(t, x_t) \,\Big|\, s^T, R, \xi \Big]\Big] \\ &\quad = \delta \, \frac{x_t^\gamma}{\gamma} \left\{ -\frac{1}{2} \frac{\gamma}{1-\gamma} \, \widetilde{\theta}(t) + \gamma \pi' \, \widetilde{\mu}(t) + \frac{1}{2} \gamma(\gamma - 1) \pi' \, \widetilde{\Sigma}(t) \, \pi \right\} + o(\delta), \end{aligned} \tag{12.7}$$

and the value function at period $t+1$,

$$\mathbb{E}_{x_{t+1}}\Big[V_D(t+1, x_{t+1}) \,\Big|\, \mathcal{G}_t^C\Big]$$
$$= V_D(t, x_t) + \delta \left\{ \frac{\partial}{\partial t} V_D(t, x_t) + x_t\big(r + \pi' \mu(s_t)\big) \frac{\partial}{\partial x} V_D(t, x_t) + \frac{1}{2} x_t^2 \, \pi' \, \Sigma(s_t) \, \pi \frac{\partial^2}{\partial x^2} V_D(t, x_t) \right\}$$
$$+ o(\delta). \tag{12.8}$$

Substitute (12.6)-(12.8) into the Bellman equation (12.5), and solve for the optimal portfolio,

$$\pi^* = \frac{1}{1-\gamma} \Big[\Sigma(s_t)\big(h(t) - g_R(t)\big) + \widetilde{\Sigma}(t)\Big]^{-1} \Big[\mu(s_t)\big(h(t) - g_R(t)\big) + \widetilde{\mu}(t)\Big]. \tag{12.9}$$

Then by straightforward substitution, the Bellman equation is solved by V_D defined in (11.7)-(11.8).

Chapter 13

Conclusion

This dissertation studies two important topics of dynamic portfolio choice problems: a portfolio choice problem with market impact costs and a portfolio choice problem under the Hidden Markov Model.

In the first part, we study a portfolio selection problem with market impact by extending the well-known Merton portfolio selection problem. The key idea of our model is we capture market impact using the penalty term which is quadratic in the trading rate. In the case of quadratic utility function, our problem becomes a linear-quadratic control problem which is straightforward to solve. We formulate two illiquid portfolio problems; the first one, we incorporate the penalty into the investor objective. The second one which is very intuitive is to track the Merton solution with minimum trading costs. We show that these two formulations deliver the same solution; when there is market impact, it is optimal to either solve the Merton problem with market impact penalty or track the Merton solution with minimum trading costs. Then we derive the approximation of the optimal trading policy in term of the optimal solution of perfectly liquid market case. The result clearly shows the intuitive relationship between our solution and the solution to the Merton problem, and how the states variables and model parameters affect the optimal trading policy. Our analysis also gives rise to the notion of a *liquidity time scale*. Numerical experiments show promising performance and robustness properties. Though this work focuses on quadratic utility, we also believe that non-quadratic utilities can be handled by applying similar ideas to the associated dynamic programming equations. Other interesting extensions are also possible, including models with stochastic market impact costs that jump whenever asset prices jump.

In the second part, we study a portfolio selection problem under the Hidden Markov Model. The main difference from the previous literature is we assume that the market regime, the transition probabilities and the regime-dependent parameters are unknown to the investor. The difficulty in solving this problem is the investor knowledge on the unknown state variables and parameters is represented by a high-dimensional joint probability distribution with no way to parameterize. As a result, the problem cannot be solved due to the curse of dimensionality. In this work, we propose approximate methods that allow us to compute the optimal portfolio without doing backward recursion on the joint probability distribution. The approximate methods rely on an approximation of value function by a value function of the complete information problem, the C-problem. The optimal portfolio share similar structure to that of the C-problem except that the unknown parameters are replaced by their weighted expectations under the joint probability distribution. Under the HMM, calculating this expectation can be done by the Gibbs sampling. The proposed methods show promising result as we compare it to the exact-problem upper-bound obtained from the dual problem. Though the problem setting is the market under the HMM, the idea of these approximate methods can be generalized to solve other dynamic optimization problems when there are hidden processes.

Bibliography

[1] A. Alfonsi, A. Fruth, and A. Schied. Constrained portfolio liquidation in a limit order book model. Advances in Mathematics of Finance, 83:9–25, 2008.

[2] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. Quantitative Finance, 10(2):143–157, 2010.

[3] R. Almgren and N. Chriss. Optimal execution of portfolio transaction. Journal of Risk, 3:5–39, 2000.

[4] R. Almgren, C. Thum, E. Hauptmann, and H. Li. Equity market impact. Risk, 18(4): 57–62, 2005.

[5] A. Ang and G. Bekaert. International asset allocation. The Review of Financial Studies, 15(4):1137–1187, 2002.

[6] A. Ang and G. Bekaert. How regimes affect asset allocation. Financial Analysts Journal, 60(2):86–99, 2004.

[7] A. Ang and A. Timmermann. Regime changes and financial markets. Forthcoming in Annual Review of Financial Economics, 2011.

[8] N. Bauerle and U. Rieder. Portfolio optimization with jumps and unobservable intesity process. Mathematical Finance, 17(2):205–224, 2007.

[9] D. Bertsekas. Dynamic Programming and Optimal Control, volume 1. Athena Scientific, 3 edition, 2005.

[10] D. Bertsimas and A. Lo. Optimal control of execution costs. Journal of Financial Markets, 1:1–50, 1998.

[11] T. Bjork, M. Davis, and C. Landen. Optimal investment with partial information. Mathematical Methods of Operations Research, 71:371–399, 2010.

[12] J. Bouchaud. Price impact. Encyclopedia of Quantitative Finance, 2010.

[13] S. Brendle. Portfolio selection under incomplete information. Stochastic Processes and their Applications, 116:701–723, 2006.

[14] M. Brennan. The role of learning in dynamic portflio decisions. European Finance Review, 1:295–306, 1998.

[15] D. Brown, J. Smith, and P. Sun. Information relaxations and duality in stochastic dynamic programs. Operations Research, 58(4):785–801, 2010.

[16] A. Capponi and J. Figueroa-Lopez. Dynamic portfolio optimization with a defaultable security and regime-switching. Forthcoming in Mathematical Finance, 2013.

[17] C. Carvalho, H. Lopes, and O. Aguilar. Dynamic stock selection strategies: a structured factor model framework. In Bayesian Statistics 9, 2011.

[18] U. Cetin, R. Jarrow, and P. Protter. Liquidity risk and arbitrage pricing theory. Finance and Stochastics, 8(3):311–341, 2004.

[19] G. Chen and S. Hsu G. Chen. Linear Stochastic Control Systems. CRC Press, 1 edition, 1995.

[20] G. Constantinides. Capital market equilibrium with transaction costs. Journal of Political Economy, 94(4):842–862, 1986.

[21] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. Operations Research, 58:549–563, 2010.

[22] J. Cvitanic and I. Karatzas. Convex duality in constrained portfolio optimization. Annals of Applied Probability, 2(4):767–818, 1992.

[23] M. Davis and A. Norman. Portfolio selection with transaction costs. Mathematics of Operations Research, 15(4):676–713, 1990.

[24] A. Dufour and R. Engle. Time and the price impact of a trade. Journal of Finance, 55 (6):2467–2498, 2000.

[25] P. Forsyth, J. Kennedy, S. Tse, and H. Windcliff. Optimal trade execution: A mean quadratic variation approach. Journal of Economic Dynamics and Control, 36:1971–1991, 2012.

[26] R. Frey, A. Gabih, and R. Wunderlich. Portfolio optimization under partial information with expert opinions. working paper, 2011.

[27] S. Fruhwirth-Schnatter. Finite Mixture and Markov Switching Models. Springer, 2006.

[28] N. Garleanu and L. Pedersen. Liquidity and risk management. The American Economic Review Papers and Proceedings, 97:193–197, 2007.

[29] N. Garleanu and L. Pedersen. Dynamic trading with predictable returns and transaction costs. Forthcoming in Journal of Finance, 2013.

[30] J. Gatheral. No-dynamic-arbitrage and market impact. Quantitative Finance, 10:749–759, 2010.

[31] G. Gennotte. Optimal portfolio choice under incomplete information. Journal of Finance, 41(3):733–746, 1986.

[32] J. Geweke, G. Koop, and H. van Dijk. The Oxford Handbook of Bayesian Econometrics. Oxford University Press, 2011.

[33] R. Grinold. A dynamic model of portfolio management. Journal of Investment Management, 4:5–22, 2006.

[34] M. Guidolin and A. Timmermann. Asset allocation under multivariate regime switching. Journal of Economic Dynamics and Control, 31:3503–3544, 2007.

[35] M. Guidolin and A. Timmermann. International asset allocation under regime switching, skew, and kurtosis preferences. Unknown Journal, 21(2):889–935, 2008.

[36] J. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. Econnometrica, 57(2):357–384, 1989.

[37] H. He and H. Mamaysky. Dynamic trading with price impact. Journal of Economic Dynamics and Control, 29:891–930, 2005.

[38] T. Honda. Optimal portfolio choice for unobservable and regime-switching mean returns. Journal of Economic Dynamics and Control, 28:45–78, Mar.-Apr. 2003.

[39] I. Kharroubi and H. Pham. Optimal portfolio liquidation with execution cost and risk. SIAM Jounal of Financial Mathematics, 1:897–931, 2010.

[40] P. Kokotovic, H. Khalil, and J. O'Reilly. Singular Perturbation Methods in Control: Analysis and Design. Academic Press, 1986.

[41] A. Kyle and A. Obizhaeva. Market microstructure invariants: Empirical evidence from portfolio transitions. Available at SSRN: http://ssrn.com/abstract=1978943, 2011.

[42] X. Li, X. Y. Zhou, and A. Lim. Dynamic mean-variance portfolio selection with no-shorting constraints. SIAM Journal on Control and Optimization, 40(5):1540–1555, 2002.

[43] A. Lim and T. Watewai. Optimal investment and consumption when regime transitions cause price shocks. working paper, 2013.

[44] H. Liu. Optimal consumption and investment with transaction costs and multiple risky assets. Journal of Finance, 59(1):289–338, 2004.

[45] J. Liu, F. Longstaff, and J. Pan. Dynamic asset allocation with event risk. Journal of Finance, 58:231–259, 2003.

[46] F. Longstaff. Optimal portfolio and the valuation of illiquid securities. Review of Financial Studies, 14(2):407–431, 2001.

[47] V. Ly Vath, M. Mnif, and H. Pham. A model of optimal portfolio selection under liquidity risk and. Finance and Stochastics, 11:51–90, 2007.

[48] J. Ma, Q. Songy, J. Xuz, and J. Zhang. Impulse control and optimal portfolio selection with general transaction cost. Preprint, 2009.

[49] R. Merton. Lifetime portfolio selection under uncertainty: The continuous-time case. Review of Economics and Statistics, 51:247257, 1969.

[50] R. Merton. Optimum consumption and portfolio rules in a continuous-time model. Journal of Economic Theory, 3:373–413, 1971.

[51] C. Moallemi and M. Saglam. Dynamic portfolio choice with linear rebalancing rules. Preprint, 2012.

[52] W. Powell. Approximate Dynamic Programming: Solving the Curses of Dimensionality. John Wiley & Sons, 2007.

[53] U. Rieder and N. Bauerle. Portfolio optimization with unobservable markov-modulated drift process. Journal of Applied Probability, 42:362–378, 2005.

[54] L. Rogers and S. Singh. Modelling liquidity and its effects on price. Technical report, Cambridge University, 2004.

[55] L. Rogers and S. Singh. The cost of illiquidity and its effects on hedging. Mathematical Finance, 20(4):597–615, 2010.

[56] J. Sass and U. Haussmann. Optimizing the terminal wealth under partial information: The drift process as a continuous time markov chain. Finance and Stochastics, 8:553–577, 2004.

[57] A. Schied, T. Schoneborn, and M. Tehranchi. Optimal basket liquidation for cara investors is deterministic. Applied Mathematical Finance, 17:471–489, 2010.

[58] S. Shreve and H.M. Soner. Optimal investment and consumption with transaction costs. The Annals of Applied Probability, 4(3):609–692, 1994.

[59] L. Tepla. Optimal portfolio policies with borrowing and shortsale constraints. Journal of Economic Dynamics and Control, 24:1623–1639, 2000.

[60] J. Tu. Is regime switching in stock returns important in portfolio decisions? Management Science, 56(7):1198–1215, July 2010.

[61] J. Yong and X. Y. Zhou. Stochastic Controls: Hamiltonian Systems and Hjb Equations. Springer, 1 edition, 1999.