Universidad Carlos III de Madrid

Doctoral Dissertation

Dynamic Portfolio Selection with Transaction Costs

Author: **Xiaoling Mei**

Advisors:
Francisco J. Nogales
Universidad Carlos III de Madrid
Victor DeMiguel
London Business School

Department of Statistics
Madrid, November 2015

Universidad
Carlos III de Madrid
www.uc3m.es

(a entregar en la Oficina de Posgrado, una vez nombrado el Tribunal evaluador , para preparar el documento para la defensa de la tesis)

TESIS DOCTORAL

Dynamic Portfolio Selection with Transaction Costs

Autor: *Xiaoling Mei*

Director/es: **Francisco Javier Nogales**
Victor DeMiguel

Firma del Tribunal Calificador:

		Firma
Presidente:	Alberto Suárez	
Vocal:	Katerina Papadaki	
Secretario:	Carlos Ruíz	

Calificación:

Leganés, 26 de Noviembre de 2015

“读万卷书，行万里路”

To my family
致我的家人

Acknowledgements

This thesis becomes a reality with the kind help and assistance of many individuals. I would like to take this opportunity to express my appreciation and thanks to all of them.

First and most of all, I would like to express my deepest gratitude towards my advisors Francisco Javier Nogales and Victor DeMiguel for their consistent guidance, patience and continuous support throughout all the phases of my PhD study and research life. I will always be grateful to you for giving me the chances to work at London Business School, for encouraging me every time when I feel depressed. Both of you have been great mentors for me and I will cherish forever the memories of this period.

I also would like to express my gratitude to my colleagues and friends from Universidad Carlos III de Madrid. Especially to Liu Ling for being great officemate and friend over all the years. Vahe Avagyan, Nicola, Ziyuan Tang for the laughter you have brought me during the coffee. And Ignacio Cascos, María Durban and Elisa Molanes for your help in teaching.

Thanks to Alba, Boris and Erix, my Spanish friends I met in London. I will never forget the moments we shared and nice conversations we had in London and in Spain. I also want to thank my best friends Zhou Juan, Xiuping, Zuo Heng, Qian Xi, Zhang Liang, You Sheng, Tang Yong and Hui Hui for the constant joy and fun you give me when I mostly need it. Thank you guys for letting me feel life is so beautiful.

Last but not the least, I would like to thank my parents, my brother, Yan Tianjiao and Mei Shiyi for their unconditional love and support. Thanks to my parents for letting me become the person I want to be. Thanks to my brother, Tianjiao and Shiyi for the encouragement throughout all these years. To end with, I would like to express my special thanks to my boyfriend Weixuan Zhu for his unconditional support. Thank you for being the most precious things in my life.

Abstract

The last few decades have witnessed a surge in research activity in the area of multiperiod portfolio optimization due to its theoretical and practical importance. One of the main issues that researchers have confronted in the implementation of successful portfolio strategies is the consideration of transaction costs. Characterizing the optimal portfolio policy in the presence of transaction costs is an open area of research in portfolio optimization. In this dissertation, we aim to investigate the portfolio policy for a multiperiod portfolio problem with multiple risky assets and different types of transaction costs under the utility maximization framework. In particular, we study and characterize the optimal portfolio rebalancing rules for a risk-averse investor who incurs transaction costs in a discrete-time setting.

We begin with an analytically tractable framework where the investor has mean-variance preference and constant investment opportunity set in the presence of different types of transaction costs. For small trades, which result in proportional transaction costs, we provide a closed-form expression for the no-trade region that confines the investor's optimal portfolio policy for each period, and we further analyze how the size of the no-trade region changes with relevant model parameters. For large trades, which can distort prices and thus result in strictly convex costs, we show analytically that there exists a state-dependent rebalancing region for each period, such that the optimal policy at each period is to trade to its boundary. We further study the utility losses associated with ignoring transaction costs and behaving myopically with an empirical dataset.

Then, we move to an analytically intractable framework where the investor has power utility preference and predictable risky asset returns in the presence of proportional transaction costs. We first consider an approximation to the model with power utility using a mean-variance problem, and we propose several approximate solutions that induce low utility losses for the mean-variance problem. Furthermore, we adapt these feasible trading strategies to the framework with power utility of intermediate consumption and study numerically the associated losses in certainty equivalent. We show that the multiperiod portfolio selection problem can be tackled through the utilization of a duality approach.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

How to invest optimally is one of the main concerns of rational investors. The first breakthrough work on this topic is the mean-variance portfolio model proposed by Markowitz [1952]. In this model, the investor accomplishes her optimal portfolio based on the trade-off between risk and expected return. Markowitz measures the risk of a portfolio through the expected portfolio variance. Specifically, let $R \in \mathbb{R}^N$ be the return vector of available securities. Given the expectation $\mu = \mathbb{E}[R]$ of the return vector and the covariance matrix $\Sigma = \mathbb{E}[(R-\mu)(R-\mu)']$, the investor allocates her capital among the N assets, forming a portfolio $w \in \mathbb{R}^N$ where each component w_i represents the fraction of total capital held in asset i. The expected return of a portfolio and its variance are given respectively by

$$\mu_p = w'\mu, \tag{1.1}$$

$$\sigma_p^2 = w'\Sigma w. \tag{1.2}$$

Markowitz's work suggests that, for a any given level of expected return μ^*, the investor seeks the vector of weights which minimizes the portfolio variance. It takes the form

$$\min_{w} \quad w'\Sigma w \tag{1.3}$$

$$\text{s.t.} \quad w'\mu = \mu^*, \tag{1.4}$$

$$w'\mathbb{1} = 1, \tag{1.5}$$

where $\mathbb{1}$ is a vector of ones with length N. In the absence of a risk-free asset (i.e., an asset with zero return variance and zero covariance with all other assets), the optimal solution to the above optimization for every given level of expected return μ^* defines the upper edge line (known as the *efficient frontier*) of a risk-return region which consists of all combinations of return and risk for every possible portfolio w in risk-return space. For any given point on the efficient frontier, it provides a rational investor with the lowest level of risk exposed to a given target expected return μ^* or equivalently, the highest expected return that can be reached for a fixed level of risk. Without imposing the expected return constraint (1.4), the optimal solution gives the global minimum-variance (MV) portfolio, which is depicted on Figure 1.1. When one

Figure 1.1: Efficient Frontier

This figure shows the efficient frontier obtained from three individual assets represented by green dots. The red dots represent the Minimum Variance (MV) portfolio and tangency portfolio respectively. The risk-free rate is 2%. The blue line which joins the risk-free asset and tangency portfolio is the Capital Market Line (CML.)

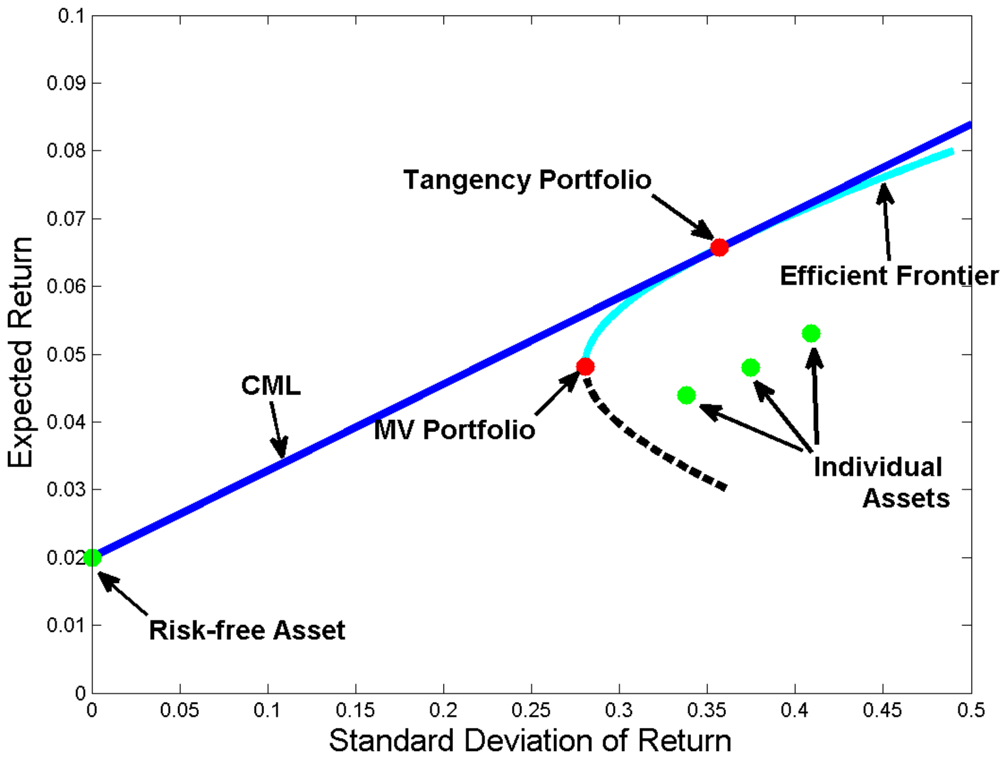

of the assets is risk-free, the diagram for the mean return and risk changes to the straight line, which is also known as the Capital Market Line (CML). The portfolios on the CML are formed by linear combinations of the risk-free asset and the *tangency portfolio* located on the efficient frontier. Problem (1.3)-(1.5) does not specify the the level of target return μ^* which should be decided by the investor depending on the risk she is willing to bear. Let γ be a parameter that represents the investor's *risk aversion*. Markowitz [1959] introduces a formulation where the investor maximizes her expected *mean-variance* utility, with the corresponding expected utility maximization problem as:

$$\max_{w} \quad w'\mu - \frac{\gamma}{2} w'\Sigma w \tag{1.6}$$

$$s.t. \quad w'\mathbb{1} = 1. \tag{1.7}$$

This problem is equivalent to problem (1.3)-(1.5), and it admits an explicit solution which can be attained through first-order conditions:

$$w = \frac{1}{\gamma}\Sigma^{-1}\mu - \frac{\iota}{\gamma}\Sigma^{-1}\mathbb{1}, \tag{1.8}$$

where ι is the Lagrange multiplier associated with constraint (1.7) and is defined as $\frac{\mu'\Sigma^{-1}\mathbb{1}-\gamma}{\mathbb{1}'\Sigma^{-1}\mathbb{1}}$.

1.1 Utility Function

Problem (1.6) defines the quadratic utility based on the mean and variance of a portfolio return. However, the investor in this model only considers an investment in one period and does not account for a multiperiod investment. Considerable work on multiperiod portfolio management has been conducted in subsequent years, in which a different framework of expected utility is used: instead of the mean and variance of a portfolio, the expected terminal wealth $\mathbb{E}(U(W))$ for a utility function is maximized. In particular, Mossin [1968] shows that the constant relative risk aversion (CRRA) function is the only type of utility function that leads to the investment of a fixed proportion in each of the risky assets. Specifically, for a utility function of the terminal wealth $U(W)$, the Arrow-Pratt coefficient of relative risk aversion is defined as

$$RRA(W) = \frac{-WU''(W)}{U'(W)}. \tag{1.9}$$

A CRRA utility function (which is also called power utility or isoelastic utility function) refers to one for which $RRA(W)$ is a constant. Mathematically speaking, a CRRA utility function is defined as:

$$U(W) = \frac{W^{1-\gamma}-1}{1-\gamma}, \tag{1.10}$$

with $\gamma \neq 1$, a constant referring to the relative risk aversion parameter. A special case of interest is the logarithm utility: when $\gamma = 1$, it has been shown that a multiperiod investor with logarithm utility on terminal wealth will perform myopically.

In addition to its utility for terminal wealth, the CRRA utility function defined over a stream of intermediate consumption is also a major concern of the in multiperiod portfolio optimization literature. In this thesis, we focus on both mean-variance utility and CRRA utility of intermediate consumption in Chapter 2 and Chapter 3, respectively.

1.2 Investment Opportunity Set

Let $R_{t+1} \in \mathbb{R}^N$ be a vector of random asset returns between t and $t+1$. A certain type of distribution of future returns is typically required for solving the multiperiod portfolio optimization models. This section reviews the models that are most commonly used in the literature, starting with the simple case where the investment opportunity set is constant.

1.2.1 Constant Investment Opportunity Set

In this case, the risky asset returns for each period R_{t+1} follow a log normal distribution. Specifically, let R_f be a constant denoting risk-free asset return. Define $r_f = \log(R_f)$ and $r_{t+1} = \log(R_{t+1})$. The risky asset return process follows

$$r_{t+1} - r_f = \mu_0 + \epsilon_{t+1}, \tag{1.11}$$

where $\mu_0 \in \mathbb{R}^N$ is a vector of constant, denoting the mean of log returns in excess of risk-free asset. ϵ_{t+1} is a multivariate normally distributed with mean zero and covariance matrix Σ_ϵ. With constant investment opportunity set, Mossin [1968] shows that, for a multiperiod investor, the *myopic* portfolio choice is optimal.

1.2.2 Factor Model

Factor models are used to explain the dynamics of asset returns by introducing pervasive factors that drive returns in common. Generally, the return of asset i over the period t is decomposed into the returns of K factor returns:

$$R_{it} = \alpha_{iR} + \sum_{j=1}^{K} \beta_{ij} f_{jt} + u_{it}, \tag{1.12}$$

where α_{iR} is a regression constant, f_{jt} is the j-th common factor at time t, and u_{it} is a zero-mean residual term uncorrelated with factor returns. The coefficients β_{ij} are the factor loadings of asset i on each of the K factors.

The main advantage of factor models is that they can considerably reduce the number of parameters needed to be estimated. Consider a portfolio of $N = 100$ assets as a simple example. There are $N(N+1)/2 = 5,000$ parameters to be estimated for the covariance matrix in Markowitz's mean-variance analysis. With $K = 3$ factors,[1] the covariance matrix of the asset returns is given by

$$\tilde{\Sigma} = \beta_R \Sigma_f \beta_R' + \Sigma_u, \tag{1.13}$$

where β_R is the $N \times K$ matrix of factor loadings, Σ_f is the $K \times K$ covariance matrix of the factors, and Σ_u is a diagonal matrix of N residual return variances. Then the number of parameters to be estimated reduces to $K + N(K+1) = 403$ terms.

1.2.3 VAR Model

In the context of return predictability, a factor model can provide some direction in forecasting expected asset returns by positing a dynamic model for making forecasts of the factors. Among

[1] Assuming that the residuals are mutually uncorrelated and are uncorrelated with the factors.

the dynamic models, the vector autoregression (VAR) model is the most common one for establishing the dynamics. Specifically, let F_t be the $K \times 1$ vector of factor returns over period t, and α_R be the vector of regression constant containing all the α_{iR}. The related VAR model is defined as follows:

$$R_t = \alpha_R + \beta_R F_t + u_t, \tag{1.14}$$

$$F_{t+1} = \alpha_K + \beta_K F_t + \nu_t, \tag{1.15}$$

where α_K is a vector of constant and β_K is a matrix of first-order autoregression factors. The ν_t is normally distributed white noise with mean zero. Regarding the predictive factors, Fama and French [1989] find that dividend yield and term spread predict return components for a cross-section of asset classes. As a consequence, the model which describes predictability of stock returns by a single dividend yield term has been widely employed in the literature. In Chapter 3, we will focus on this type of VAR model for log of asset returns.

1.3 Multiperiod Problems

Markowitz's original model considers an investment in one period: the investor determines the vector of portfolio weights at the beginning of the investment period and waits patiently until the end of the period. In subsequent years, there has been considerable work on multiperiod portfolio problems, pioneered by Samuelson [1969]. In the multiperiod setting, an investor makes a sequence of decisions. At each period, she rebalances her portfolios by taking into consideration a future changing opportunity set and the remaining investment horizon.

Specifically, consider the problem where the investor determines a vector of portfolio weights w_t at each time $t = 0, 1, 2, \cdots, T-1$. The objective is to find the allocation decisions that maximize an expected utility function. For instance, we can consider the utility of wealth at the final time T, given the initial wealth. That is,

$$\max_{w_0, w_1, \cdots, w_{T-1}} \quad \mathbb{E}_0 \left[U(W_T) \right], \tag{1.16}$$

subject to the *budget constraint* stating that, for each period, all the wealth is invested:

$$\begin{aligned} W_{t+1} &= W_t(1 + R_{p,t+1}) &(1.17)\\ &= W_t(1 + w_t' R_{t+1} + (1 - w_t' \mathbb{1}) R_f), &(1.18) \end{aligned}$$

where W_t is the investor's wealth at time t, and $R_{p,t+1}$ is the portfolio return at period $t+1$. Note that problem (1.16) is also equivalent to the problem where at each time step t, the investor tries to intertemporally maximize the expected utility of wealth at T, given the current wealth W_t,

$$\max_{w_0, w_1, \cdots, w_{T-1}} \quad \mathbb{E}_t \left[U(W_T) \right], \tag{1.19}$$

for $t = 0, 1, \cdots, T-1$. Moreover, constraint (1.17) describes the dynamics of wealth by specifying that the portfolio return for period $t+1$ arises from the allocation w_t to the risky assets and the rest $(1 - w_t'\mathbb{1})$ from the risk-free asset. The investment in the risk-free asset can be negative if we allow the investor borrow to from it. We also require the wealth at each period to be positive, $W_t \geq 0$.

Based on a formulation by dynamic programming, it is convenient to express (1.19) in terms of a value function which varies according to time period, current wealth and other state variables z_t,

$$
\begin{aligned}
V(t, W_t, z_t) &= \max_{\{w_s\}_{s=t}^{s=T-1}} \mathbb{E}_t[U(W_T)] && (1.20)\\
&= \max_{w_t} \mathbb{E}_t\Big[\max_{\{w_s\}_{s=t+1}^{s=T-1}} \mathbb{E}_{t+1}[U(W_T)]\Big] && (1.21)\\
&= \max_{w_t} \mathbb{E}_t\Big[V\left(t+1, W_{t+1} = W_t(1 + w_t'R_{t+1} + (1 - w_t'\mathbb{1})R_f), z_{t+1}\right)\Big] && (1.22)
\end{aligned}
$$

subject to the constraint for terminal wealth: $V(T, W_T, z_T) = U(W_T)$. Here, the expectations at time t are conditional on the information available at time t, which depends directly on the joint distribution of risky asset returns.

This problem does not admit, in general, an explicit solution. It can only be solved numerically using the system of nonlinear equations obtained from the first-order conditions. However, when an investor has the preference of power utility $U(W) = \frac{W^{1-\gamma}}{1-\gamma}$ $(\gamma \neq 1)$, some analytic progress can be achieved. In this case, the investor's decision problem can be established as follows:

$$
\begin{aligned}
V(t, W_t, z_t) &= \max_{w_t} \mathbb{E}_{t+1}\left[\max_{\{w_s\}_{s=t+1}^{s=T-1}} \mathbb{E}_{t+1}\left[\frac{W_T^{1-\gamma}}{1-\gamma}\right]\right] && (1.23)\\
&= \max_{w_t} \mathbb{E}_{t+1}\left[\max_{\{w_s\}_{s=t+1}^{s=T-1}} \mathbb{E}_{t+1}\left[\frac{\left(W_t \prod_{s=t}^{T-1}(1 + w_s'R_{s+1} + (1 - w_t'\mathbb{1})R_f)\right)^{1-\gamma}}{1-\gamma}\right]\right] && (1.24)\\
&= \max_{w_t} \mathbb{E}_{t+1}\left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma} \max_{\{w_s\}_{s=t+1}^{s=T-1}} \mathbb{E}_{t+1}\left[\left(\prod_{s=t+1}^{T-1}(1 + w_s'R_{s+1} + (1 - w_t'\mathbb{1})R_f)\right)^{1-\gamma}\right]\right] && (1.25)\\
&= \max_{w_t} \mathbb{E}_t\Big[U(W_{t+1})\Phi(t+1, z_{t+1})\Big] && (1.26)
\end{aligned}
$$

It shows that the value function is the product of the investor's next time-step utility $U(W_{t+1})$ and $\Phi(t+1, z_{t+1})$, which only depends on the remaining horizon and future state variables. This can be further expressed as

$$
V(t, W_t, z_t) = \frac{W_t^{1-\gamma}}{1-\gamma}\phi(t, z_t), \tag{1.27}
$$

with the terminal period $\Phi(T, \cdot) = 1$. Note that if the risky asset returns are independent and identically distributed throughout time, a multiperiod investor who has power utility will behave myopically. However, when the investment opportunity set varies over time, the above problem can only be solved using a dynamic programming approach. In both cases, continuously rebalancing the portfolio weights requires the payment of high *transaction costs*.

1.4 Transaction Costs

In general, transaction cost is a cost incurred in making an transaction. Incorporating transaction costs into the portfolio optimization model helps us converge towards the optimal portfolio in the most profitable way. For each stage, the transaction costs can be expressed in terms of trading vector Δx. Typically, it is a convex function of Δx that allows us to capture the cost associated with both large and small trades depending on the size of the transaction.

For small trades which do not distort market prices, the transaction costs generally come from the bid-ask spreads and other brokerage fees, which are simply modelled as a function that is proportional to the amount traded. There are several ways to characterize proportional transaction costs in the portfolio optimization literature. For example, Balduzzi and Lynch [1999] model the transaction cost such that it is paid from the investor's wealth, while Lynch and Tan [2010] assume that it is paid by costlessly liquidating the risk-free asset.

For large trades, on the other hand, there is market price impact which is very difficult to measure. In general, the market impact refers to the difference between the transaction price and what the market price would have been in the absence of the transaction. Moreover, the impact on market price can be both temporary and permanent. Market price impact is temporary when it affects a single transaction, and permanent when it affects every future transaction. The corresponding transaction costs caused by large trades are generally modelled in a strictly convex form. For example, Gârleanu and Pedersen [2013] consider transaction costs in a quadratic form, which is a special case for capturing temporary impact on market prices. In this dissertation, we consider transaction costs caused by both small and large trades in Chapter 2 and proportional transaction costs in Chapter 3.

1.5 Contributions

Ignoring transaction costs will induce an unrealistic level of trading activity. However, taking these costs into consideration will require a higher dimension in dynamic programming, and thus even more computation time. In this dissertation, we investigate the discrete-time portfolio rebalancing problem where transaction costs are charged when the holding in risky assets changes. We contribute to the existing literature in two main aspects: first, we focus on the multiperiod mean-variance problem and we analyze the optimal portfolio policy for an investor facing a large number of risky assets in the presence of different types of transaction costs. Second, we propose several feasible policies that are based on optimizing simple quadratic programs to solve a portfolio optimization problem for a multiperiod CRRA investor who incurs proportional transaction costs.

Specifically, in Chapter 2, we carry out a comprehensive analysis on the optimal portfolio policies for an investor whose objective is to maximize her expected portfolio return penalized by the portfolio variability. We consider a constant investment opportunity set and various types of transaction costs. We make several contributions in this field. First, we analytically characterize the optimal portfolio policy for the case with many risky assets and proportional transaction costs by providing a closed-form expression for the no-trade region. This analysis is related to Dybvig [2005], who considers the same problem but for a single-period setting. He shows graphically without an analytical proof that the optimal portfolio policy is characterized by a no-trade region shaped as a parallelogram. Second, we analytically study the optimal portfolio policy in the presence of market impact costs, and we show that the optimal policies are confined by a rebalancing region at each period. Third, we numerically study the utility losses associated with ignoring transaction costs and investing myopically depending on relevant parameters. We show that the resultant utility losses can be substantial. Finally, we gauge the robustness of our results obtained based on mean-variance utility to the use of a framework with CRRA utility, and we show that the associated certainty equivalent losses from following the mean-variance portfolio policy are considerably small.

In Chapter 3, we propose some approximate trading strategies for the multiperiod portfolio selection problem for an investor who maximizes her expected CRRA utility of intermediate consumption and suffers from proportional transaction costs. These feasible trading strategies are proposed based on an approximation for the CRRA utility framework with a framework where an investor has a mean-variance utility. We first propose some approximate trading strategies for the mean-variance investor in the presence of predictability. We then show how to adapt the proposed feasible strategies based on the mean-variance framework to that based on CRRA power utility. We also show that the multiperiod portfolio selection problem with multiple risky assets can be tackled through the use of the duality method developed in Brown et al. [2010]. Finally, to evaluate the performance of the proposed trading policies, we numerically computed the mean-variance utility losses of the associated feasible strategies as well as the corresponding certainty equivalent losses in the CRRA power utility framework. This chapter is related to Brown and Smith [2011], who propose some heuristic trading strategies for an investor with power utility of terminal wealth. However, we use a different objective function, namely the one with power utility of intermediate consumption. Moreover, our proposal of approximate trading strategies is based on quadratic programs which are easy to solve whereas theirs rely on approximations in the dynamic recursion of a primal problem.

We conclude and summarize the main findings of this dissertation in Chapter 4, and we also give a short description of possible research lines in Chapter 4.2.

Chapter 2

Multiperiod Portfolio Optimization with Many Risky Assets and General Transaction Costs

2.1 Overview

Merton [1971] showed that an investor who wishes to maximize her utility of consumption should hold a fixed proportion of her wealth on each of the risky assets, and consume at a rate proportional to her wealth.[1] Merton's seminal work relies on the assumptions that the investor has constant relative risk aversion (CRRA) utility, faces an infinite horizon, can trade continuously and (crucially) costlessly. Implementing Merton's policy, however, requires one to rebalance the portfolio weights continuously, and in practice this may result in high or even infinite transaction costs. Ever since Merton's breakthrough, researchers have tried to address this issue by characterizing the optimal portfolio policy in the presence of transaction costs.

Researchers focused first on the case with a single-risky asset. Magill and Constantinides [1976] consider a finite-horizon continuous-time investor subject to proportional transaction costs and for the first time *conjecture* that the optimal policy is characterized by a *no-trade interval*: if the portfolio weight on the risky-asset is inside this interval, then it is optimal not to trade, and if it is outside, then it is optimal to trade to the boundary of this interval. Constantinides [1979] demonstrates the optimality of the no-trade interval policy in a finite-horizon discrete-time setting. Constantinides [1986] considers the Merton framework with a single risky asset and proportional transaction costs, and computes *approximately-optimal* no-trade interval policies by requiring the investor's consumption rate to be a fixed proportion of her wealth, a condition that is not satisfied in general. Davis and Norman [1990] consider the same framework, show that the *optimal* no-trade interval policy exists, and propose a numerical method to compute

[1] Meton's result holds for either an investor facing a constant investment opportunity set, or an investor with logarithmic utility; see also Mossin [1968], Samuelson [1969], and Merton [1969, 1973].

it. Dumas and Luciano [1991] consider a continuous-time investor who maximizes utility of terminal wealth, and show how to calculate the boundaries of the no-trade interval for the limiting case when the terminal period goes to infinity.

The case with multiple risky assets is less tractable, and the bulk of the existing literature relies on numerical results for the case with only two risky assets. Akian et al. [1996] consider a multiple risky-asset version of the framework in Davis and Norman [1990], and for the restrictive case where the investor has power utility with relative risk aversion between zero and one[2] and risky-asset returns are uncorrelated, they show that there exists a unique optimal portfolio policy. They also compute numerically the no-trade region for the case with *two* uncorrelated stocks. Leland [2000] considers the tracking portfolio problem subject to proportional transaction costs and capital gains tax, and proposes a numerical approach to approximate the no-trade region.

Muthuraman and Kumar [2006] consider an infinite-horizon continuous-time investor and propose an efficient numerical approach to compute the no-trade region. Their numerical results show that the no-trade region for the case with two risky assets is characterized by four corner points, but these four corner points are not joined by straight lines, although their numerical experiments show that a quadrilateral no-trade region does provide a very close approximation. Lynch and Tan [2010] consider a finite-horizon discrete-time investor facing proportional and fixed transaction costs, and two risky assets with predictable returns. Using numerical dynamic programming, they show that for the case *without* predictability the no-trade region is closely approximated by a parallelogram, whereas for the case with predictability the no-trade region is closely approximated by a convex quadrilateral.[3]

Most of the aforementioned papers assume an investor with CRRA utility of consumption who faces borrowing constraints. These assumptions render the problem untractable analytically, and hence they generally rely on numerical analysis for the case with two risky assets. A notable exception is the work by Liu [2004], who obtains an analytically tractable framework by making several restrictive assumptions.[4] Specifically, he considers an investor with constant *absolute* risk aversion (CARA) and access to unconstrained borrowing[5], who can invest in multiple *uncorrelated* risky assets. For this framework, Liu shows *analytically* that there exists a box-shaped no-trade region.

Recently, Gârleanu and Pedersen [2013], herein G&P, consider a more tractable framework that allows them to provide closed-form expressions for the optimal portfolio policy in the

[2] Janeček and Shreve [2004] show that relative risk aversion parameters between one and zero lead to intolerably risky behavior.

[3] Brown and Smith [2011] also consider the case with proportional transaction costs and return predictability. Specifically, they propose several heuristic trading strategies for a finite-horizon discrete-time investor facing proportional transaction costs and multiple assets with predictable returns, and use upper bounds based on duality theory to evaluate the optimality of the proposed heuristics.

[4] Another important exception is Muthuraman and Zha [2008] who use a simulation-based numerical optimization to approximate the optimal portfolio policy of a continuous-time investor who maximizes her long-term expected growth rate for cases with up to seven risky assets. Also, in their early paper Magill and Constantinides [1976] conjecture the existence of a box-shaped no-trade region for the case where the portfolio weights are small.

[5] He does impose constraints to preclude arbitrage portfolio policies.

presence of *quadratic* transaction costs. Their investor maximizes the present value of the mean-variance utility of her *wealth changes* at multiple time periods, she has access to unconstrained borrowing, and she faces multiple risky assets with predictable *price changes*. Several features of this framework make it tractable. First, the focus on utility of wealth changes (rather than consumption) plus the access to unconstrained borrowing imply that there is no need to track the investor's total wealth evolution, and instead it is sufficient to track wealth change at each period. Second, the focus on price changes (rather than returns) implies that there is no need to track the risky-asset price evolution, and instead it is sufficient to account for price changes. Finally, the aforementioned features, combined with the use of mean-variance utility and quadratic transaction costs places the problem in the category of linear quadratic control problems, which are tractable.

In this section, we use the tractable formulation of G&P to study analytically the optimal portfolio policies for general transaction costs. Our portfolio selection framework is both more general and more specific than that considered by G&P. It is more general because we consider a broader class of transaction costs that includes not only quadratic transaction costs, but also the less tractable proportional and market impact costs. It is more specific because, consistent with most of the literature on proportional transaction costs, we consider the case with constant investment opportunity set, whereas G&P focus on the case with predictability. Specifically, we consider independent and identically distributed (iid) price changes, whereas G&P consider predictable price changes.

We make four contributions. Our first contribution is to characterize analytically the optimal portfolio policy for the case with many risky assets and *proportional* transaction costs. Specifically, we provide a closed-form expression for the no-trade region, which is shaped as a multi-dimensional parallelogram, and use the closed-form expressions to show how the size of the no-trade region shrinks with the investment horizon and the risk-aversion parameter, and grows with the level of proportional transaction costs and the discount factor.

Moreover, we show how the optimal portfolio policy can be computed by solving a quadratic program—a class of optimization problems that can be efficiently solved for cases with up to thousands of risky assets.

Our second contribution is to study analytically the optimal portfolio policy in the presence of market impact costs, which arise when the investor makes large trades that distort market prices.[6] Traditionally, researchers have assumed that the market *price* impact is linear on the amount traded, see Kyle [1985], and thus that market impact costs are quadratic. Under this assumption, Gârleanu and Pedersen [2013] derive closed-form expressions for the optimal portfolio policy within their multiperiod setting. However, Torre and Ferrari [1997], Grinold and Kahn [2000], and Almgren et al. [2005] show that the square root function is more appropriate for modeling market price impact, thus suggesting market impact costs grow at a rate slower than quadratic. Our contribution is to extend the analysis by G&P to a *general case* where we

[6] This is particularly relevant for optimal execution, where institutional investors have to execute an investment decision within a fixed time interval; see Bertsimas and Lo [1998] and Engle et al. [2012]

are able to capture the distortions on market price through a a power function with an exponent between one and two. For this general formulation, we show *analytically* that there exists a state-dependent *rebalancing region* for every time period, such that the optimal policy at each period is to trade to the boundary of the corresponding rebalancing region.

Our third contribution is to use an empirical dataset with the prices of 15 commodity futures to evaluate the losses associated with ignoring transaction costs and investing myopically, as well as identifying how these losses depend on relevant parameters. We find that the losses associated with either ignoring transaction costs or behaving myopically can be large. Moreover, the losses from ignoring transaction costs increase in the level of transaction costs, and decrease with the investment horizon, whereas the losses from behaving myopically increase with the investment horizon and are unimodal on the level of transaction costs.

Finally, our analysis relies on certain assumptions underlying the G&P framework such as focus on price changes and multiperiod mean-variance utility. Our fourth contribution is to gauge the robustness of our results to the use of a framework with iid returns and CRRA utility. To do this, we consider an investor who maximizes her CRRA utility of terminal wealth by investing in a risk-free and a risky asset with iid returns, in the presence of proportional transaction costs. We compute the investor's optimal portfolio policy numerically, and show that the certainty equivalent loss from following the G&P-type portfolio policy is typically smaller than 0.5%.

Our work is related to Dybvig [2005], who considers a single-period investor with mean-variance utility and proportional transaction costs. For the case with multiple risky assets, he shows that the optimal portfolio policy is characterized by a no-trade region shaped as a parallelogram, but the manuscript does not provide a detailed analytical proof. Like Dybvig [2005], we consider proportional transaction costs and mean-variance utility, but we extend the results to a multi-period setting, and show how the results can be rigorously proven analytically. In addition, we consider the case with market impact costs.

This section is organized as follows. Section 2.2 describes the multiperiod framework under general transaction costs. Section 2.3 studies the case with proportional transaction costs, Section 2.4 the case with market impact costs, and Section 2.5 the case with quadratic transaction costs. Section 2.6 evaluates the utility loss associated with ignoring transaction costs and with behaving myopically for an empirical dataset on 15 commodity futures. Section 2.7 studies numerically the robustness of our modelling framework, and Section 2.8 concludes. Appendix A.1 contains the figures, Appendix A.2 contains the tables, and Appendix A.3 contains the proofs for all results in the chapter.

2.2 General Framework

Our framework is closely related to the one proposed by G&P. Like G&P, we consider an investor who maximizes the present value of the mean-variance utility of excess wealth changes (net of transaction costs), by investing in multiple risky assets and for multiple periods, and who has access to unconstrained borrowing. Moreover, like G&P, we make assumptions on the distribution of risky-asset price changes rather than returns. One difference with the framework proposed by G&P is that, consistent with most of the existing transaction cost literature, we focus on the case with constant investment opportunity set; that is, we consider the case with iid price changes.

As discussed in previous section, these assumptions render the model analytically tractable. The assumptions seem reasonable in the context of institutional investors who typically operate many different and relatively unrelated investment strategies. Each of these investment strategies represents only a fraction of the institutional investor's portfolio, and thus focusing on excess wealth changes and assuming unconstrained borrowing is a good approximation. The stationarity of price changes is also a reasonable assumption for institutional investors who often have shorter investment horizons—because they operate each investment strategy only for a few months or at most a small number of years, and discontinue the investment strategy once its performance deteriorates.

The assumptions underlying our framework, however, do not seem suitable (a priori) to model individual investors who finance their lifetime consumption from the proceeds of their investments. These investors face constraints on borrowing, and have long investment horizons (their lifetime), during which one would expect returns, rather than price changes, to be stationary. To gauge the severity of our assumptions for the case of an individual investor, in Section 2.7 we compute numerically the optimal portfolio policy of an investor who faces a risk-free and a risky asset with iid returns, and who maximizes her CRRA utility of terminal wealth with a ten year horizon. We find that the certainty equivalent loss from using the G&P-type portfolio policy derived from our framework is less than 0.5% for a broad range of problem parameters.

Finally, there are three main differences between our model and the model by G&P. First, we consider a more general class of transaction costs that includes not only quadratic transaction costs, but also proportional and market impact costs. Second, we consider both finite and infinite investment horizons, whereas G&P focus on the infinite horizon case. Finally, we assume iid price changes, whereas G&P focus on the case with predictable price changes. We now rigorously state this assumption. Let r_{t+1} denote the vector of price changes (in excess of the risk-free asset price) between times t and $t+1$ for N risky assets.

Assumption 1. Price changes, r_{t+1}, are independently and identically distributed (iid) with mean vector $\mu \in \mathbb{R}^N$ and covariance matrix $\Sigma \in \mathbb{R}^{N \times N}$.

The investor's decision in our framework can be written as:

$$\max_{\{x_t\}_{t=1}^T} \quad \sum_{t=1}^{T} \left[(1-\rho)^t (x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t) - (1-\rho)^{t-1} \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right], \tag{2.1}$$

where $x_t \in \mathbb{R}^N$ contains the number of shares of each of the N risky assets held in period t, T is the investment horizon, ρ is the discount factor, and γ is the *absolute* risk-aversion parameter.[7]

The term $\kappa\|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p$ is the transaction cost for the tth period,[8] where $\kappa \in \mathbb{R}$ is the transaction cost parameter, $\Lambda \in \mathbb{R}^{N\times N}$ is the symmetric positive semidefinite transaction cost matrix, and $\|s\|_p$ is the p-norm of vector s; that is, $\|s\|_p^p = \sum_{i=1}^N |s_i|^p$. This term allows us to capture the transaction costs associated with both small and large trades. Small trades typically do not impact market prices, and thus their transaction costs come from the bid-ask spread and other brokerage fees, which are modeled as proportional to the amount traded. Our transaction cost term captures proportional transaction costs for the case with $p = 1$ and $\Lambda = I$, where I is the identity matrix.

Large trades can have both temporary as well as permanent impact on market prices. Market price impact is temporary when it affects a single transaction, and permanent when it affects every future transaction. For simplicity of exposition, we focus on the case with temporary market impact costs, but our analysis can be extended to the case with permanent impact costs following an approach similar to that in Section 4 of G&P. For market impact costs, Almgren et al. [2005] suggest that transaction costs grow as a power function with an exponent between one and two, and hence we consider in our analysis values of $p \in (1, 2]$. The transaction cost matrix Λ captures the distortions to market prices generated by the interaction between the multiple assets. G&P argue that it can be viewed as a multi-dimensional version of Kyle's lambda, see Kyle [1985], and they argue that a sensible choice for the transaction cost matrix is $\Lambda = \Sigma$. We consider this case as well as the case with $\Lambda = I$ to facilitate the comparison with the case with proportional transaction costs.

Note that, like G&P, we also assume the investor can costlessly short the different assets. The advantage of allowing costless shortsales is that it allows us to obtain analytical expressions for the no-trade region for the case with proportional transaction costs, and for the rebalancing regions for the case with market impact costs. We state this assumption explicitly now.

Assumption 2. The investor can costlessly short the risky assets.

Finally, the multiperiod mean-variance framework proposed by G&P and the closely related framework described in Equation (2.1) differ from the traditional dynamic mean-variance approach, which attempts to maximize the mean-variance utility of *terminal wealth*. Part 1 of Proposition 2.1 below, however, shows that the utility given in Equation (2.1) is equal to the mean-variance utility of the change in excess of *terminal* wealth for the case where the discount factor $\rho = 0$. This shows that the framework we consider is not too different from the traditional dynamic mean-variance approach. Also, a worrying feature of multiperiod mean-variance frameworks is that as demonstrated by Basak and Chabakauri [2010] they are often

[7] Because the investment problem is formulated in terms of wealth changes, the mean-variance utility is defined in terms of the *absolute* risk aversion parameter, rather than the *relative* risk aversion parameter. Note that the relative risk-aversion parameter equals the absolute risk-aversion parameter times the wealth.

[8] The pth root of the positive definite matrix Λ can be defined as $\Lambda^{1/p} = Q^T D^{1/p} Q$, where $\Lambda = Q^T D Q$ is the spectral decomposition of a symmetric positive definite matrix.

time-inconsistent: the investor may find it optimal to deviate from the ex-ante optimal policy as time goes by. Part 2 of Proposition 2.1 below, however, shows that the framework we consider is time consistent.

Proposition 2.1. *Let Assumption 1 hold, then the multiperiod mean-variance framework described in Equation (2.1) satisfies the following properties:*

1. *The utility given in Equation (2.1) is equivalent to the mean-variance utility of the change in excess terminal wealth for the case where the discount factor $\rho = 0$.*
2. *The optimal portfolio policy for the multiperiod mean-variance framework described in Equation (2.1) is time consistent.*

2.3 Proportional Transaction Costs

We now study the case where transaction costs are proportional to the amount traded. This type of transaction cost is appropriate to model small trades, where the transaction cost originates from the bid-ask spread and other brokerage commissions. Section 2.3.1 characterizes analytically the no-trade region and the optimal portfolio policy, and Section 2.3.2 shows how the no-trade region depends on the level of proportional transaction costs, the risk-aversion parameter, the discount factor, the investment horizon, and the correlation and variance of asset price changes.

2.3.1 The no-trade region

The investor's decision for this case can be written as:

$$\max_{\{x_t\}_{t=1}^T} \left\{ \sum_{t=1}^{T} \left[(1-\rho)^t \left(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t \right) - (1-\rho)^{t-1} \kappa \|x_t - x_{t-1}\|_1 \right] \right\}. \tag{2.2}$$

The following theorem characterizes the optimal portfolio policy.

Theorem 2.2. *Let Assumption 1 hold, then:*

1. *It is optimal not to trade at any period other than the first period; that is,*

$$x_1 = x_2 = \cdots = x_T. \tag{2.3}$$

2. *The investor's optimal portfolio for the first period x_1 (and thus for all subsequent periods) is the solution to the following quadratic programming problem:*

$$\min_{x_1} \quad (x_1 - x_0)^\top \Sigma (x_1 - x_0), \tag{2.4}$$

$$s.t \quad \|\Sigma(x_1 - x^*)\|_\infty \leq \frac{\kappa}{(1-\rho)\gamma} \frac{\rho}{1-(1-\rho)^T}. \tag{2.5}$$

where x_0 is the starting portfolio, and $x^* = \Sigma^{-1}\mu/\gamma$ is the optimal portfolio in the absence of transaction costs (the Markowitz or target portfolio).

3. *Constraint (2.5) defines a no-trade region shaped as a parallelogram centered at the target portfolio x^*, such that if the starting portfolio x_0 is inside this region, then it is optimal not to trade at any period, and if the starting portfolio is outside this no-trade region, then it is optimal to trade at the first period to the point in the boundary of the no-trade region that minimizes the objective function in (2.4), and not to trade thereafter.*

A few comments are in order. First, a counterintuitive feature of our optimal portfolio policy is that it only involves trading in the first period. A related property, however, holds for most of the policies in the literature. Liu [2004], for instance, explains that: "the optimal trading policy involves possibly an initial discrete change (jump) in the dollar amount invested in the asset, followed by trades in the minimal amount necessary to maintain the dollar amount within a constant interval." The "jump" in Liu's policy, is equivalent to the first-period investment in our policy. The reason why our policy does not require any rebalancing after the first period is that it relies on the assumption that *prices changes* are iid. As a result, the portfolio and no-trade region in our framework are defined in terms of *number of shares*, and thus no rebalancing is required after the first period because realized price changes do not alter the number of shares held by the investor. Finally, our numerical results in Section 2.7 show that the certainty equivalent loss from using the G&P-type policy given in Theorem 2.2 for the case where *returns* are iid is small.

Second, inequality (2.5) provides a closed-form expression for the no-trade region. This expression shows that it is optimal to trade only if the marginal increment in utility from trading in one of the assets is larger than the transaction cost parameter κ. To see this, note that inequality (2.5) can be rewritten as

$$-\kappa e \leq (\gamma(1-\rho)(1-(1-\rho)^T)/\rho)\Sigma(x_1 - x^*) \leq \kappa e, \tag{2.6}$$

where e is the N-dimensional vector of ones. Moreover, because Part 1 of Theorem 2.2 shows that it is optimal to trade only at the first period, it is easy to show that the term in the middle of (2.6) is the gradient (first derivative) of the discounted multiperiod mean variance utility $\sum_{t=1}^{T}(1-\rho)^t \left(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t\right)$ with respect to x_1. Consequently, it is optimal to trade at the first period only if the *marginal* increase in the present value of the multiperiod mean-variance utility is larger than the transaction cost parameter κ.

Finally, inequality (2.5) shows that the no-trade region is a multi-dimensional parallelogram centered around the target portfolio. Mathematically, this property follows from the *linearity* of the first derivative of the multiperiod mean-variance utility with respect to the portfolio x_1; that is, the linearity of the middle term in (2.6). Our result contrasts with the findings of Muthuraman and Kumar [2006], who show (numerically for the case with two risky assets) that the no-trade region is a convex quadrilateral (rather than a parallelogram), and it is *not* centered around the target portfolio. Mathematically, the reason for this is that the CRRA utility they consider results in a value function whose first derivative is *not* linear. In addition, they impose constraints on borrowing, which adds to the nonlinearity of the boundary conditions defining the no-trade region.[9] Economically, Muthuraman and Kumar [2006] consider an investor who maximizes CRRA utility of intermediate consumption and with constraints on borrowing. Consequently, the investor in Kumar and Muthuraman's framework is more willing to trade (and thus incur higher transaction costs) when she holds *large* positions on the risky assets, in order to guarantee a more stable level of wealth to finance her ongoing consumption. As a result, the no-trade region in Kumar and Muthuraman's framework is not centered around the target portfolio, and instead it is biased towards the risk-free asset.

Third, the optimal portfolio policy can be conveniently computed by solving the quadratic program (2.4)–(2.5). This class of optimization problems can be efficiently solved for cases with up to thousands of risky assets using widely available optimization software. As mentioned in the introduction, most of the existing results for the case with transaction costs rely on numerical analysis for the case with two risky assets. Our framework can be used to deal with cases with proportional transaction costs and hundreds or even thousands of risky assets. To gain understanding about the quadratic program (2.4)–(2.5), Figure A.1 depicts the no-trade region defined by inequality (2.5) and the level sets for the objective function given by (2.4) for a case with two assets with mean and covariance matrix equal to the sample estimators for two commodity futures on gasoil and sugar, which are part of the full dataset of 15 commodities described in Section 2.6. The figure shows that the optimal portfolio policy is to trade to the intersection between the no-trade region and the tangent level set, at which the marginal utility from trading equals the transaction cost parameter κ.

2.3.2 Comparative statics

The following corollary establishes how the no-trade region depends on the level of proportional transaction costs, the risk-aversion parameter, the discount factor, and the investment horizon.

Corollary 2.3. *The no-trade region for the multiperiod investor satisfies the following properties:*

1. *The no-trade region expands as the proportional transaction parameter κ increases.*
2. *The no-trade region shrinks as the risk-aversion parameter γ increases.*

[9] For the case where the investor maximizes her long-term expected growth rate, Muthuraman and Zha [2008] also find that the no-trade region is a quadrilateral that is not centered around the Merton portfolio. Essentially, maximizing the long-term growth rate is similar to maximizing a logarithmic utility function, which again results in a value function that has a nonlinear derivative. For the case with a single risky asset, Constantinides [1986] considers CRRA utility and constrained borrowing and also finds that the no-trade interval is not centered around the Merton portfolio. Dumas and Luciano [1991], on the other hand, consider the case with a single risky asset and an investor with constrained borrowing and CRRA utility of terminal wealth. For the case where the investor's horizon goes to infinity, they find that the no-trade interval *is* centered. Although Dumas and Luciano [1991] consider a nonlinear utility function and constrained borrowing, the focus on terminal wealth when the investment horizon goes to infinity results in a centered no-trade interval.

3. *The no-trade region expands as the discount factor parameter ρ increases.*

4. *The no-trade region shrinks as the investment horizon T increases.*

Part 1 of Corollary 2.3 shows that, not surprisingly, the size of the no-trade region grows with the transaction cost parameter κ. The reason for this is that the larger the transaction costs, the less willing the investor is to trade in order to diversify. This is illustrated in Panel (A) of Figure A.2, which depicts the no-trade regions for different values of the transaction cost parameter κ for the two commodity futures on gasoil and sugar.[10] Note also that (as discussed in Section 2.3.1) the no-trade regions for different values of the transaction cost parameter are all centered around the target portfolio.

Part 2 of Corollary 2.3 shows that the size of no-trade region decreases with the risk aversion parameter γ. Intuitively, as the investor becomes more risk averse, the optimal policy is to move closer to the diversified (safe) position x^*, despite the transaction costs associated with this. This is illustrated in Figure A.2, Panel (B), which also shows that, not surprisingly, the target portfolio shifts towards the risk-free asset as the risk-aversion parameter increases.

Part 3 of Corollary 2.3 shows that the size of the no-trade region increases with the discount factor ρ. This makes sense intuitively because the larger the discount factor, the less important the utility for future periods and thus the smaller the incentive to trade today. This is illustrated in Figure A.2, Panel (C).

Finally, Part 4 of Corollary 2.3 shows that the size of the no-trade region decreases with the investment horizon T. To see this intuitively, note that we have shown that the optimal policy is to trade at the first period and hold this position thereafter. Then, a multiperiod investor with shorter investment horizon will be more concerned about the transaction costs incurred at the first stage, compared with the investor who has a longer investment horizon. Finally, when $T \to \infty$, the no-trade region shrinks to the parallelogram bounded by $\kappa\rho/((1-\rho)\gamma)$, which is much closer to the center x^*. When $T = 1$, the multiperiod problem reduces to the single-period problem studied by Dybvig [2005]. This is illustrated in Figure A.3, Panel (A).

The no-trade region also depends on the correlation between assets. Figure A.3, Panel (B) shows the no-trade regions for different correlations[11]. When the two assets are positively correlated, the parallelogram leans to the left, reflecting the substitutability of the two risky assets, whereas with negative correlation it leans to the right. In the absence of correlations the no-trade region becomes a rectangle.

Finally, the impact of variance on the no-trade region is shown in Panel (C) of Figure A.3, where for expositional clarity we have considered the case with two uncorrelated symmetric

[10] Although we illustrate Corollary 2.3 using two commodity futures, the results apply to the general case with N risky assets.

[11] Because change in correlation also makes the target shift, in order to emphasize how correlation affects the shape of the region, we change the covariance matrix Σ in a manner so as to keep the Markowitz portfolio (the target) fixed, similar to the analysis in Muthuraman and Kumar [2006].

risky assets. Like Muthuraman and Kumar [2006], we find that as variance increases, the no-trade region moves towards the risk-free asset because the investor is less willing to hold the risky assets. Also the size of no-trade region shrinks as the variance increases because the investor is more willing to incur transaction costs in order to diversify her portfolio.

2.4 Market Impact Costs

We now consider the case of large trades that may impact market prices. As discussed in Section 2.2, to simplify the exposition we focus on the case with temporary market impact costs, but the analysis can be extended to the case with permanent impact costs following an approach similar to that in Section 4 of G&P. Almgren et al. [2005] suggest that market impact costs grow as a power function with an exponent between one and two, and hence we consider a general case, where the transaction costs are given by the p-norm with $p \in (1, 2)$, and where we capture the distortions on market price through the transaction cost matrix Λ. For exposition purposes, we first study the single-period case.

2.4.1 The Single-Period Case

For the single-period case, the investor's decision is:

$$\max_{x} \quad (1-\rho)(x^\top \mu - \frac{\gamma}{2} x^\top \Sigma x) - \kappa \|\Lambda^{1/p}(x - x_0)\|_p^p, \tag{2.7}$$

where $1 < p < 2$. Problem (2.7) can be solved numerically, but unfortunately it is not possible to obtain closed-form expressions for the optimal portfolio policy. The following proposition, however, shows that the optimal portfolio policy is to trade to the boundary of a *rebalancing region* that depends on the starting portfolio and contains the target or Markowitz portfolio.

Proposition 2.4. *Let Assumption 1 hold, then if the starting portfolio x_0 is equal to the target or Markowitz portfolio x^*, the optimal policy is not to trade. Otherwise, it is optimal to trade to the boundary of the following rebalancing region:*

$$\frac{\|\Lambda^{-1/p}\Sigma(x - x^*)\|_q}{p\|\Lambda^{1/p}(x - x_0)\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}, \tag{2.8}$$

where q is such that $\frac{1}{p} + \frac{1}{q} = 1$.

Comparing Proposition 2.4 with Theorem 2.2 we observe that there are three main differences between the cases with proportional and market impact costs. First, for the case with market impact costs it is always optimal to trade (except in the trivial case where the starting portfolio coincides with the target or Markowitz portfolio), whereas for the case with proportional transaction costs it may be optimal not to trade if the starting portfolio is inside the no-trade region. Second, the rebalancing region depends on the starting portfolio x_0, whereas the no-trade region is independent of it. Third, the rebalancing region contains the target or Markowitz portfolio, but it is not centered around it, whereas the no-trade region is centered around the Markowitz portfolio.

Note that, as in the case with proportional transaction costs, the size of the rebalancing region increases with the transaction cost parameter κ, and decreases with the risk-aversion parameter. Intuitively, the more risk averse the investor, the larger her incentives to trade and diversify her portfolio. Also, the rebalancing region grows with κ because the larger the transaction cost parameter, the less attractive to the investor is to trade to move closer to the target portfolio.

The following corollary gives the rebalancing region for two important particular cases. First, the case where the transaction cost matrix $\Lambda = I$, which is a realistic assumption when the amount traded is small, and thus the interaction between different assets, in terms of market impact, is small. This case also facilitates the comparison with the optimal portfolio policy for the case with proportional transaction costs. The second case corresponds to the transaction cost matrix $\Lambda = \Sigma$, which G&P argue is realistic in the context of quadratic transaction costs.

Corollary 2.5. *For the single-period investor defined in* (2.7)*:*

1. *When the transaction cost matrix is $\Lambda = I$, then the rebalancing region is*

$$\frac{\|\Sigma(x - x^*)\|_q}{p\|x - x_0\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{2.9}$$

2. *When the transaction cost matrix is $\Lambda = \Sigma$, then the rebalancing region is*

$$\frac{\|\Sigma^{1/q}(x - x^*)\|_q}{p\|\Sigma^{1/p}(x - x_0)\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{2.10}$$

Note that, in both particular cases, the Markowitz strategy x^* is contained in the rebalancing region.

To gain intuition about the form of the rebalancing regions characterized in (2.9) and (2.10), Panel (A) in Figure A.4 depicts the rebalancing region and the optimal portfolio policy for a two-asset example when $\Lambda = I$, while Panel (B) depicts the corresponding rebalancing region and optimal portfolio policy when $\Lambda = \Sigma$. The figure shows that, in both cases, the rebalancing region is a convex region containing the Markowitz portfolio. Moreover, it shows how the optimal trading strategy moves to the boundary of the rebalancing region.

2.4.2 The Multiperiod Case

The investor's decision for this case can be written as:

$$\max_{\{x_t\}_{t=1}^T} \quad \sum_{t=1}^{T} \left[(1-\rho)^t \left(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t\right) - (1-\rho)^{t-1} \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p\right]. \tag{2.11}$$

As in the single-period case, it is not possible to provide closed-form expressions for the optimal portfolio policy, but the following theorem illustrates the analytical properties of the optimal portfolio policy.

Theorem 2.6. *Let Assumption 1 hold, then:*

1. *If the starting portfolio x_0 is equal to the target or Markowitz portfolio x^*, then the optimal policy is not to trade at any period.*
2. *Otherwise it is optimal to trade at every period. Moreover, at the tth period it is optimal to trade to the boundary of the following rebalancing region:*

$$\frac{\| \sum_{s=t}^{T} (1-\rho)^{s-t} \Lambda^{-1/p} \Sigma (x_s - x^*) \|_q}{p \| \Lambda^{1/p} (x_t - x_{t-1}) \|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}, \tag{2.12}$$

where q is such that $\frac{1}{p} + \frac{1}{q} = 1$.

Theorem 2.6 shows that for the multiperiod case with market impact costs it is optimal to trade at every period (except in the trivial case where the starting portfolio coincides with the Markowitz portfolio). This is in contrast to the case with proportional transaction costs, where it is not optimal to trade at any period other than the first. The reason for this is that with strictly convex transaction costs, $p > 1$, the transaction cost associated with a trade can be reduced by breaking it into several smaller transactions. For the case with proportional transaction costs, on the other hand, the cost of a transaction is the same whether executed at once or broken into several smaller trades, and thus it makes economic sense to carry out all the trading in the first period to take advantage of the utility benefits from the beginning.

Note also that for the case with market impact costs at every period it is optimal to trade to the boundary of a *state-dependent* rebalancing region that depends not only on the starting portfolio, but also on the portfolio for every subsequent period.

Finally, note that the size of the rebalancing region for period t, assuming the portfolios for the rest of the periods are fixed, increases with the transaction cost parameter κ and decreases with the discount factor ρ and the risk-aversion parameter γ.

The following proposition shows that the rebalancing region for period t contains the rebalancing region for every subsequent period. Moreover, the rebalancing region converges to the Markowitz portfolio as the investment horizon grows, and thus the optimal portfolio x_T converges to the target portfolio x^* in the limit when T goes to infinity. The intuition behind this result is again that the strict convexity of the transaction cost function for the case with $p > 1$ implies that it is optimal to trade every period, but it is never optimal to trade all the way to the Markowitz portfolio, and thus the investor's portfolio converges to the Markowitz portfolio only as time goes to infinity.

Proposition 2.7. *Let Assumption 1 hold, then:*

1. *The rebalancing region for the t-th period contains the rebalancing region for every subsequent period,*

2. *Every rebalancing region contains the Markowitz portfolio,*

3. *The rebalancing region converges to the Markowitz portfolio in the limit when the investment horizon goes to infinity.*

The next corollary gives the rebalancing region for the two particular cases of transaction cost matrix we consider.

Corollary 2.8. *For the multiperiod investor defined in (2.11):*

1. *When the transaction cost matrix is $\Lambda = I$, then the rebalancing region is*

$$\frac{\| \sum_{s=t}^{T}(1-\rho)^{s-t}\Sigma(x_s - x^*)\|_q}{p\|x_t - x_{t-1}\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{2.13}$$

2. *When the transaction cost matrix is $\Lambda = \Sigma$, then the rebalancing region is*

$$\frac{\| \sum_{s=t}^{T}(1-\rho)^{s-t}\Sigma^{1/q}(x_s - x^*)\|_q}{p\|\Sigma^{1/p}(x_t - x_{t-1})\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{2.14}$$

To gain intuition about the shape of the rebalancing regions characterized in (2.13) and (2.14), Panel (A) in Figure A.5 shows the optimal portfolio policy and the rebalancing regions for the two commodity futures on gasoil and sugar with an investment horizon $T = 3$ when $\Lambda = I$, whereas Panel (B) depicts the corresponding optimal portfolio policy and rebalancing regions when $\Lambda = \Sigma$. The figure shows, in both cases, how the rebalancing region for each period contains the rebalancing region for subsequent periods. Moreover, every rebalancing region contains, but is not centered at, the Markowitz portfolio x^*. In particular, for each stage, any trade is to the boundary of the rebalancing region and the rebalancing is towards the Markowitz strategy x^*.

Finally, we study numerically the impact of the market impact cost growth rate p on the optimal portfolio policy. Figure A.6 shows the rebalancing regions and trading trajectories for investors with different transaction growth rates $p = 1, 1.25, 1.5, 1.75, 2$. When the transaction cost matrix $\Lambda = I$, Panel (A) shows how the rebalancing region depends on p. In particular, for $p = 1$ we recover the case with proportional transaction costs, and hence the rebalancing region becomes a parallelogram. For $p = 2$, the rebalancing region becomes an ellipse. And for values of p between 1 and 2, the shape of the rebalancing regions are similar to superellipses[12] but not centered at the target portfolio x^*. On the other hand, Panel (B) in Figure A.6 shows how the trading trajectories depend on p for a particular investment horizon of $T = 10$ days. We observe that, as p grows, the trading trajectories become more curved and the investor converges

[12] The general expression for a superellipse is $\left|\frac{x}{a}\right|^m + \left|\frac{y}{b}\right|^n = 1$ with $m, n > 0$.

towards the target portfolio at a slower rate. To conserve space, we do not provide the figure for the case $\Lambda = \Sigma$, but we find that for this case the trajectories are less curved as p grows, and becomes a straight line for $p = 2$.

2.5 Quadratic Transaction Costs

We now consider the case with quadratic transaction costs. The investor's decision is:

$$\max_{\{x_t\}_{t=1}^T} \quad \sum_{t=1}^{T} \left[(1-\rho)^t (x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t) - (1-\rho)^{t-1} \kappa \|\Lambda^{1/2}(x_t - x_{t-1})\|_2^2 \right]. \tag{2.15}$$

For the case with quadratic transaction costs, our framework differs from that in G&P in two respects only. First, G&P's work focuses on impact of predictability, whereas consistent with most of the existing literature on transaction costs we assume price changes are iid. Second, G&P consider an infinite horizon, whereas we allow for a finite investment horizon. The next theorem adapts the results of G&P to obtain an explicit characterization of the optimal portfolio policy.

Theorem 2.9. *Let Assumption 1 hold, then:*

1. *The optimal portfolio $x_t, x_{t+1}, \ldots, x_{t+T-1}$ satisfies the following equations:*

$$x_t = A_1 x^* + A_2 x_{t-1} + A_3 x_{t+1}, \qquad \textit{for} \quad t = 1, 2, \ldots, T-1 \tag{2.16}$$

$$x_t = B_1 x^* + B_2 x_{t-1}, \qquad \textit{for} \quad t = T. \tag{2.17}$$

where

$$\begin{aligned} A_1 &= (1-\rho)\gamma \left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2(1-\rho)\kappa\Lambda\right]^{-1} \Sigma, \\ A_2 &= 2\kappa \left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2(1-\rho)\kappa\Lambda\right]^{-1} \Lambda, \\ A_3 &= 2(1-\rho)\kappa \left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2(1-\rho)\kappa\Lambda\right]^{-1} \Lambda, \end{aligned}$$

with $A_1 + A_2 + A_3 = I$, and

$$\begin{aligned} B_1 &= (1-\rho)\gamma[(1-\rho)\gamma\Sigma + 2\kappa\Lambda]^{-1}\Sigma, \\ B_2 &= 2\kappa[(1-\rho)\gamma\Sigma + 2\kappa\Lambda]^{-1}\Lambda, \end{aligned}$$

with $B_1 + B_2 = I$.

2. *The optimal portfolio converges to the Markowitz portfolio as the investment horizon T goes to infinity.*

Theorem 2.9 shows that the optimal portfolio for each stage is a combination of the Markowitz strategy (the target portfolio), the previous period portfolio, and the next period portfolio.

The next corollary shows the specific optimal portfolios for two particular cases of transaction cost matrix. We consider the case where the transaction costs matrix is proportional to the covariance matrix, which G&P argue is realistic.[13] In addition, we also consider the case where the transaction costs matrix is proportional to the identity matrix; that is $\Lambda = I$.

Corollary 2.10. *For a multiperiod investor with objective function* (2.15)*:*

1. *When the transaction cost matrix is $\Lambda = I$, then the optimal trading strategy satisfies*

$$x_t = A_1 x^* + A_2 x_{t-1} + A_3 x_{t+1}, \qquad \text{for} \quad t = 1, 2, \ldots, T-1 \tag{2.18}$$
$$x_t = B_1 x^* + B_2 x_{t-1}, \qquad \text{for} \quad t = T \tag{2.19}$$

where

$$\begin{aligned} A_1 &= (1-\rho)\gamma\left[(1-\rho)\gamma\Sigma + 2\kappa I + 2(1-\rho)\kappa I\right]^{-1}\Sigma, \\ A_2 &= 2\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa I + 2(1-\rho)\kappa I\right]^{-1}, \\ A_3 &= 2(1-\rho)\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa I + 2(1-\rho)\kappa I\right]^{-1}, \end{aligned}$$

with $A_1 + A_2 + A_3 = I$, and

$$\begin{aligned} B_1 &= (1-\rho)\gamma[(1-\rho)\gamma\Sigma + 2\kappa I]^{-1}\Sigma, \\ B_2 &= 2\kappa[(1-\rho)\gamma\Sigma + 2\kappa I]^{-1}. \end{aligned}$$

with $B_1 + B_2 = I$.

2. *When the transaction cost matrix is $\Lambda = \Sigma$, then the optimal trading strategy satisfies*

$$x_t = \alpha_1 x^* + \alpha_2 x_{t-1} + \alpha_3 x_{t+1}, \qquad \text{for} \quad t = 1, 2, \ldots, T-1 \tag{2.20}$$
$$x_t = \beta_1 x^* + \beta_2 x_{t-1}, \qquad \text{for} \quad t = T. \tag{2.21}$$

where $\alpha_1 = (1-\rho)\gamma/((1-\rho)\gamma + 2\kappa + 2(1-\rho)\kappa)$, $\alpha_2 = 2\kappa/((1-\rho)\gamma + 2\kappa + 2(1-\rho)\kappa)$, $\alpha_3 = 2(1-\rho)\kappa/((1-\rho)\gamma + 2\kappa + 2(1-\rho)\kappa)$ with $\alpha_1 + \alpha_2 + \alpha_3 = 1$, and $\beta_1 = (1-\rho)\gamma/((1-\rho)\gamma + 2\kappa)$ $\beta_2 = 2\kappa/((1-\rho)\gamma + 2\kappa)$ with $\beta_1 + \beta_2 = 1$.

3. *When the transaction cost matrix is $\Lambda = \Sigma$, then the optimal portfolios for periods $t = 1, 2, \ldots, T$ lay on a straight line.*

Corollary 2.10 shows that, when $\Lambda = \Sigma$, the solution becomes simpler and easier to interpret than when $\Lambda = I$. Note that when $\Lambda = \Sigma$, matrices A and B in Theorem 2.9 become scalars α and β, respectively, and hence the optimal portfolio at period t can be expressed as a linear combination of the Markowitz portfolio, the previous period portfolio and the next period

[13]Note that although G&P argue that the case $\Lambda = \Sigma$ is realistic, they also solve explicitly the case with general transaction cost matrix Λ.

portfolio. For this reason, it is intuitive to observe that the optimal trading strategies for all the periods must lay on a straight line.

To conclude this section, Figure A.7 provides a comparison of the optimal portfolio policy for the case with quadratic transaction costs (when $\Lambda = \Sigma$), with those for the cases with proportional and market impact costs (when $\Lambda = I$), for a multiperiod investor with $T = 3$. We have also considered other transaction cost matrices, but the insights are similar. The figure confirms that, for the case with quadratic transaction costs, the optimal portfolio policy is to trade at every period along a straight line that converges to the Markowitz portfolio. It can also be appreciated that the investor trades more aggressively at the first periods compared to the final periods. For the case with proportional transaction costs, it is optimal to trade to the boundary of the no-trade region shaped as a parallelogram in the first period and not to trade thereafter. Finally, for the case with market impact costs, the investor trades at every period to the boundary of the corresponding rebalancing region. The resulting trajectory is not a straight line.

2.6 Empirical Analysis

In this section, we study empirically the losses associated with ignoring transaction costs and investing myopically, as well as how these losses depend on the transaction cost parameter, the investment horizon, the risk-aversion parameter, and the discount factor. We first consider the case with proportional transaction costs, and then study how the monotonicity properties of the losses change when transaction costs are quadratic. We have also considered the case with market impact costs ($p = 1.5$), but the monotonicity properties for this case are in the middle of those for the cases with $p = 1$ and $p = 2$ and thus we do not report the results to conserve space.

For each type of transaction cost (proportional or quadratic), we consider three different portfolio policies. First, we consider the *target* portfolio policy, which consists of trading to the target or Markowitz portfolio in the first period and not trading thereafter. This is the optimal portfolio policy for an investor in the absence of transaction costs. Second, the *static* portfolio policy, which consists of trading at each period to the solution to the single-period problem subject to transaction costs. This is the optimal portfolio policy for a myopic investor who takes into account transaction costs. Third, we consider the *multiperiod* portfolio policy, which is the optimal portfolio policy for a multiperiod investor who takes into account transaction costs.

Finally, we evaluate the utility of each of the three portfolio policies using the appropriate multiperiod framework; that is, when considering proportional transaction costs, we evaluate the investor's utility from each portfolio with the objective function in equation (2.2); and when considering quadratic transaction costs, we evaluate the investor's utility using the objective function (2.15). These utilities are in units of wealth, and to facilitate the comparison of the different policies, we report the percentage utility loss of the suboptimal portfolio policies with respect to the utility of the optimal multiperiod portfolio policy. Note also that, as argued by [DeMiguel et al., 2009, Footnote17], it can be shown that the mean-variance utility is approximately equal to the certainty equivalent of an investor with quadratic utility, and thus although we report utility losses in this section, they can be interpreted equivalently as certainty equivalent losses.

We consider an empirical dataset similar to the one used by Gârleanu and Pedersen [2013].[14] In particular, the dataset is constructed with 15 commodity futures: Aluminum, Copper, Nickel, Zinc, Lead, and Tin from the London Metal Exchange (LME), Gasoil from the Intercontinental Exchange (ICE), WTI crude, RBOB Unleaded gasoline, and Natural Gas from the New York Mercantile Exchange (NYMEX), Gold and Silver from the New York Commodities Exchange (COMEX), and Coffee, Cocoa, and Sugar from the New York Broad of Trade (NYBOT). The dataset contains daily data from July 7th, 2004 until September 19th, 2012. For our evaluation, we replace the mean and covariance matrix of price changes with their sample estimators.

2.6.1 Proportional Transaction Costs

Base Case

For our base case, we adapt the parameters used by G&P in their empirical analysis to the case with proportional transaction costs. We assume proportional transaction costs of 50 basis points ($\kappa = 0.005$), absolute risk-aversion parameter $\gamma = 10^{-6}$, which corresponds to a relative risk aversion of one for a small investor managing one million dollars[15], annual discount factor $\rho = 2\%$, and an investment horizon of $T = 22$ days (one month). For all the cases, the investor's initial portfolio is the equally weighted portfolio; that is, the investor splits her one million dollars equally among the 15 assets.

For our base case, we observe that the utility loss associated with investing myopically (that is, the relative difference between the utility of the *multiperiod* portfolio policy and the *static* portfolio policy) is 60.46%. This utility loss is large because the no-trade region corresponding to the static portfolio policy contains the equally-weighted portfolio, and thus the static portfolio policy is to remain at the starting equally-weighted portfolio, which attains a much lower multiperiod mean-variance utility than the multiperiod portfolio policy. The utility loss associated with ignoring transaction costs altogether (that is, the relative difference between the utility of the *multiperiod* portfolio policy and the *target* portfolio policy) is 49.33%. This loss is large because the target portfolio policy requires a large amount of trading in the first period that results in large transaction costs. Summarizing, we find that the loss associated

[14] We thank Alberto Martin-Utrera for making this dataset available to us.

[15] Gârleanu and Pedersen [2013] consider a smaller absolute risk aversion $\gamma = 10^{-9}$, which corresponds to a larger investor managing $M = 10^9$ dollars. It makes sense, however, to consider a smaller investor (and thus a larger absolute risk-aversion parameter) in the context of proportional transaction costs because these are usually associated with small trades.

with either ignoring transaction costs or behaving myopically can be substantial.[16] The next section confirms this is also true when we change relevant model parameters.

Comparative Statics

We study numerically how the utility losses associated with ignoring transaction costs (i.e., with the static portfolio), and investing myopically (i.e., with the target portfolio) depend on the transaction cost parameter, the investment horizon, the risk-aversion parameter, and the discount factor.

Panel (a) in Figure A.8 depicts the utility loss associated with the target and static portfolios for values of the proportional transaction cost parameter κ ranging from 0 basis point to 460 basis points (which is the value of κ for which the optimal multiperiod policy is not to trade). As expected, the utility loss associated with ignoring transaction costs is zero in the absence of transaction costs and increases monotonically with transaction costs. Moreover, for large transaction costs parameters, the utility loss associated with ignoring transaction costs grows linearly with κ and can be very large.[17] The utility losses associated with behaving myopically are unimodal (first increasing and then decreasing) in the transaction cost parameter, being zero for the case with zero transaction costs (because both the single-period and multiperiod portfolio policies coincide with the target or Markowitz portfolio), and for the case with large transaction costs (because both the single-period and multiperiod portfolio policies result in little or no trading). The utility loss of behaving myopically reaches a maximum of 80% for a level of transaction costs of around 5 basis points.

Panel (b) in Figure A.8 depicts the utility loss associated with investing myopically and ignoring transaction costs for investment horizons ranging from $T = 5$ (one week) to $T = 260$ (over one year). Not surprisingly, the utility loss associated with behaving myopically grows with the investment horizon. Also, the utility loss associated with ignoring transaction costs is very large for short-term investors, and decreases monotonically with the investment horizon. The reason for this is that the size of the no-trade region for the multiperiod portfolio policy decreases monotonically with the investment horizon, and thus the target and multiperiod policies become similar for long investment horizons. This makes sense intuitively: by adopting the Markowitz portfolio, a multiperiod investor incurs transaction cost losses at the first period, but makes mean-variance utility gains for the rest of the investment horizon. Hence, when the investment horizon is long, the transaction losses are negligible compared with the utility gains.

[16] We find that the multiperiod portfolio policy contains short positions of around 30% of the total wealth invested. We acknowledge that it would be more realistic to assume that there is a cost associated with holding short positions. However, the Assumption 2 allows us to provide closed-form expressions for the no-trade region, and helps to keep our exposition simple. We have computed numerically the no-trade region for the case with shortsale constraints and find that the qualitative findings about the shape and size of the no-trade region continue to hold in the presence of shortsale constraints.

[17] In fact, the utility of the target portfolio policy is negative (and thus the utility loss is larger than 100%) for κ larger than 69 basis points because of the high transaction costs associated with this policy.

2.6.2 Quadratic Transaction Costs

In this section we study whether and how the presence of quadratic transaction costs (as opposed to proportional transaction costs) affects the utility losses of the static and target portfolios.

The Base Case

Our base case parameters are similar to those adopted in Gârleanu and Pedersen [2013]. We assume that the matrix $\Lambda = \Sigma$ and set the absolute risk aversion parameter $\gamma = 10^{-8}$, which corresponds to an investor with relative-risk aversion of one who manages 100 million dollars[18], discount factor $\rho = 2\%$ annually, transaction costs parameter $\kappa = 1.5 \times 10^{-7}$ (which corresponds to $\lambda = 3 \times 10^{-7}$ in G&P's formulation), investment horizon $T = 22$ days (one month), and equal-weighted initial portfolio.

Similar to the case with proportional transaction costs, we find that the losses associated with either ignoring transaction costs or behaving myopically are substantial. For instance, for the base case with find that the utility loss associated with investing myopically is 28.98%, whereas the utility loss associated with ignoring transaction costs is 109.14%. Moreover, we find that the utility losses associated with the target portfolio are relatively larger, compared to those of the static portfolio, for the case with quadratic transaction costs. The explanation for this is that the target portfolio requires large trades in the first period, which are penalized heavily in the context of *quadratic* transaction costs. The static portfolio, on the other hand, results in smaller trades over successive periods and this will result in overall smaller *quadratic* transaction costs.

Finally, we have carried out comparative static analysis similar to those for proportional transaction costs, but the qualitative insights we have obtained are similar to those for proportional transaction costs, and thus we do not report the results to conserve space.

2.7 Model robustness

To gauge the robustness of our results to the use of a framework with iid returns and CRRA utility, we consider an investor who maximizes her CRRA utility (of either terminal wealth or intermediate consumption) by investing in a risk-free asset and a risky asset with iid returns, and who is subject to proportional transaction costs. We compute the investor's optimal portfolio policy using a numerical procedure similar to that used by DeMiguel and Uppal [2005], and find that the certainty equivalent loss from following the G&P-type portfolio policy given in Theorem 2.2 is typically smaller than 0.5%.

[18] Gârleanu and Pedersen [2013] choose a smaller absolute risk aversion parameter $\gamma = 10^{-9}$, which corresponds to an investor with relative-risk aversion of one who manages one billion dollars. The insights from our analysis are robust to the use of $\gamma = 10^{-9}$, but we choose $\gamma = 10^{-8}$ because this results in figures that are easier to interpret.

We consider a base case similar to that considered by DeMiguel and Uppal [2005]: the investor has a risk-aversion parameter $\gamma = 3$, faces a risk-free asset with annual rate of return of 6% and a risky asset with mean return of 10% and return volatility of 20%. The investor has an initial endowment of \$1 invested in the risk-free asset, faces proportional transaction costs of 50 basis points, rebalances her portfolio once a year, and has an investment horizon of ten years.

For the base case where the investor maximizes her CRRA utility of terminal wealth, we find that the certainty equivalent loss from using the G&P portfolio policy is only 0.20%.[19] To understand why this loss is so small, note that the main difference between the optimal portfolio policy and the G&P policy is that the optimal portfolio policy requires rebalancing, whereas the G&P policy is a buy-and-hold policy. Nonetheless, we find that, because of the presence of transaction costs, even the optimal portfolio policy requires very little rebalancing (on average the investor trades only 2% of her wealth per year), and thus the optimal and G&P policies are quite similar.

Table A.1 shows how the certainty equivalent loss of the G&P-type portfolio policy depends on the risk-aversion parameter, the stock return volatility, and the transaction cost parameter for the case where the investor maximizes her CRRA utility of terminal wealth. We find that the certainty equivalent loss decreases with risk aversion (because the amount invested in the risky asset, and thus the optimal amount of rebalancing, decreases), increases with the stock return volatility (because the optimal amount of rebalancing increases), and decreases with transaction costs (because the optimal amount of rebalancing decreases). For all values of these parameters that we try, we find that the certainty equivalent loss is below 0.50%.

We then consider an investor who maximizes her CRRA utility of intermediate consumption. To be able to make a sensible comparison, we augment the portfolio policy given by Theorem 2.2 by assuming that the investor's consumption to wealth ratio is equal to that of an investor who maximizes her CRRA utility of intermediate consumption, but ignores the presence of transaction costs.[20] We find that the certainty equivalent loss for our base case with intermediate consumption is 0.14%, and thus even smaller than that for the case where the investor maximizes utility of terminal wealth. The reason for this is that the investor progressively liquidates her wealth to finance consumption, both for the optimal and G&P policies, and thus the amount of trading required to rebalance the risky asset position is smaller.

Table A.2 shows how the certainty equivalent loss for the case with intermediate consumption depends on the risk-aversion, volatility, and transaction costs. As in the case without intermediate consumption, we find that the certainty equivalent loss decreases with risk aversion

[19] Note that the G&P-type portfolio policy given by Theorem 2.2 is a buy-and-hold policy. We compute the number of shares to be held by this policy by solving problem (2.4)–(2.5). Because we assume the starting price of the risky asset is \$1, the mean and covariance matrix of excess price changes is equal to the mean and covariance matrix of excess returns. It is straightforward to generalize this procedure to the case where the risky asset starting price is different from \$1, and it is easy to show that the certainty equivalent loss from using the G&P-type portfolio policy does not depend on the risky asset starting price.

[20] In particular, we use the investor's consumption to wealth ratio given by [Ingersoll, 1987, p. 243], after fixing the the well-known typo in the book. This is a conservative choice because this consumption to wealth ratio is not optimal in the presence of transaction costs.

because the amount invested in the risky asset and thus the optimal amount of rebalancing decreases. We find that the certainty equivalent loss is not very sensitive to the stock return volatility because volatility has two opposite effects on the amount of rebalancing. On the one hand, higher volatility implies smaller allocation to the risky asset, and thus less rebalancing. On the other hand, higher volatility implies higher fluctuations on portfolio weights, and thus more rebalancing. These two opposite effects seem to be of roughly similar strength and thus the certainty equivalent loss is not very sensitive to volatility. Finally, we find that the certainty equivalent loss increases with transaction costs (unlike in the case without intermediate consumption). The reason for this is that the particular procedure we have used to augment the G&P policy to finance intermediate consumption requires a large amount of trading on the risky asset, and thus large transaction costs. Nevertheless, we find that the certainty equivalent loss incurred from using the G&P-type policy is always below 0.50%.

Finally, we consider a case with two risky stocks with return correlation of 50%, and an investment horizon of 8 years, and find that the certainty equivalent loss of using the G&P policy is 0.11%. Overall, our results show that although the G&P-type policy relies on the assumptions of iid price changes and multiperiod mean-variance utility, it provides a reasonable approximation for the optimal portfolio policy for the case with iid returns and CRRA utility, and thus we think our analysis provides insights that are interesting beyond the G&P framework.

2.8 Summary

We study the optimal portfolio policy for a multiperiod mean-variance investor facing multiple risky assets subject to proportional, market impact, or quadratic transaction costs. For the case with proportional transaction costs, we provide a closed-form expression for a no-trade region shaped as a parallelogram, and use these closed-form expressions to show how the no-trade region shrinks with the investment horizon and the risk-aversion parameter, and grows with the level of proportional transaction costs and the discount factor. Moreover, we show that the optimal portfolio policy can be conveniently computed by solving a single quadratic program for problems with up to thousands of risky assets. For the case with market impact costs, the optimal portfolio policy is to trade to the boundary of a state-dependent rebalancing region. In addition, the rebalancing region converges to the Markowitz portfolio as the investment horizon grows large. We also show numerically that the losses associated with ignoring transaction costs or investing myopically may be large, and study how they depend on the relevant parameters. Finally, we compute the optimal portfolio policy for a case with CRRA utility and iid returns, and show that the G&P-type policy that we study provides a good approximation to this case.

Chapter 3

Portfolio Selection with Transaction Costs and Predictability

3.1 Overview

Dynamic portfolio choice is one of the most important practical problems in finance since the work of Merton [1971], which examines an investor who wishes to maximize her utility of consumption, and has access to multiple risky assets with a constant investment opportunity set.[1] Merton's policy indicates that an investor should continuously rebalance her portfolio weights in order to hold a fixed proportion of her wealth in each of the risky assets. However, continuously portfolio rebalancing requires high transaction costs. Since Merton's seminal work, researchers have tried to characterize the optimal portfolio policy in the presence of transaction costs.

The case with a single risky asset and proportional transaction costs is now well understood. Magill and Constantinides [1976] first consider proportional transaction costs and conjecture that for a finite-horizon continuous-time investor, the optimal trading policy can be characterized by a *no-trade interval*: if the portfolio weight on the risky asset is inside this interval, then it is optimal not to trade, and if the portfolio weight is outside, then it is optimal to trade to the boundary of this interval. Constantinides [1979] studies a general discrete-time model and demonstrates the optimality of no-trade interval policy with CRRA power utility of intermediate consumption and a single risky asset. Constantinides [1986] considers an infinite horizon problem with proportional transaction costs, and computes *approximately-optimal* no-trade intervals by assuming the investor's consumption rate is a fixed proportion of her wealth, a condition that is not satisfied in general. Davis and Norman [1990] address the same problem, establish analogous results on a no-trade interval, and provide a numerical method to calculate the optimal policy. Muthuraman [2007] develops an efficient computational scheme for the same problem.

[1] Merton also studies the case where the investor has logarithmic utility in the presence of predictability.

The case with multiple risky assets and proportional transaction costs is generally intractable analytically. With a constant opportunity set, Akian et al. [1996] prove the existence and uniqueness of the optimal portfolio policy for a CRRA investor who has power utility with relative risk aversion between zero and one and uncorrelated risky asset returns . They also present some numerical results for the case of two uncorrelated risky assets. Liu [2004] considers a constant absolute risk aversion (CARA) investor who has access to unconstrained borrowing and faces uncorrelated risky asset returns. He analytically shows that there exists a box-shaped no-trade region and numerically solves the case of two risky assets with a small correlation value. Muthuraman and Kumar [2006] propose an efficient numerical approach to compute the no-trade region for an infinite-horizon CRRA investor who makes decisions continuously.

In the presence of predictability, the case with multiple risky assets and proportional transaction costs is much more difficult to solve, and a small number of papers deal with this problem. Balduzzi and Lynch [1999] study the impact of return predictability on the utility costs and the optimal rebalancing rules for a single risky asset case. They show the costs of ignoring predictability can be substantial for a CRRA investor with a finite life. Lynch and Tan [2010] numerically investigate the model with two risky assets and predictable returns for a multiperiod CRRA investor who maximizes her power utility of intermediate consumption. Using numerical dynamic programming, they show that for each state, there is a quadrilateral-shaped no-trade region that confines the transaction. The numerical methods employed in their paper are based on a grid discretization of the state space; as such, their approach would run into the curse of dimensionality with more risky assets. Brown and Smith [2011] provide several heuristic trading strategies for a finite-horizon discrete-time investor facing proportional transaction costs and multiple risky assets in the presence of return predictability. They evaluate the optimality of the proposed heuristics based on upper bounds obtained through a dual approach. The dual method based on information-relaxation is initially developed in Brown et al. [2010] and provides a technique to construct valid dual bounds for any approximated solution.

The aforementioned papers show that, for a CRRA power utility investor facing a small number of risky assets (up to two), the model that incorporates return predictability with transaction costs generally admits only a numerical solution. With more risky assets, only an approximate solution can be obtained due to the curse of dimensionality. Gârleanu and Pedersen [2013], hereafter G&P, consider a more analytically tractable framework that allows them to achieve a closed-form solution for the optimal portfolio policy in the presence of *quadratic* transaction costs. Specifically, their investor maximizes the present value of the mean-variance utility of her wealth changes at multiple time periods, has access to unconstrained borrowing, and faces multiple risky assets with predictable price changes. With quadratic utility and quadratic transaction costs and no portfolio constraints, the model is formulated as a linear quadratic control problem which is straightforward to solve.

In this chapter, we consider the problem of dynamic portfolio selection in a discrete-time, finite-horizon setting. In our model, the investor maximizes her expected CRRA utility of intermediate consumption. We further assume that she faces multiple risky assets with predictable returns and constraints on borrowing, and incurs *proportional* transaction costs. We propose several approximate trading strategies that are based on solving simple quadratic programs and evaluate the sub-optimality of these strategies through the dual approach proposed by Brown et al. [2010]. In order to propose these approximate strategies, we first approximate our model for a CRRA power utility investor with the mean-variance problem considered in G&P. But instead of a model with infinite investment horizon and quadratic transaction costs, we considered a more realistic framework with finite investment horizon and proportional transaction costs. We then find some approximate solutions that induce low utility loss for the mean-variance problem. Finally, we adapt these approximate solutions for the mean-variance problem to the CRRA problem. Our numerical experiment suggests that these adapted approximate strategies perform reasonably well.

We make three contributions to the dynamic portfolio choice and transaction cost literature. Our first contribution is to provide several approximate trading strategies for a mean-variance utility investor who faces proportional transaction costs and predictability. Specifically, these approximate trading strategies are proposed using G&P's tractable mean-variance framework and can be conveniently computed by solving simple quadratic programs.

Our second contribution is to show how to adapt the strategies based on the mean-variance framework to a strategy based on a CRRA power utility. To do this, we consider an investor who wishes to maximize her CRRA utility of intermediate consumption with predictable returns, in the presence of proportional transaction costs. We numerically compute the corresponding upper bounds to the certainty equivalent of the investor and show that the certainty equivalent losses from using these approximate policies are reasonable.

Finally, in our third contribution we show that the multiperiod portfolio selection problem with multiple risky assets in the presence of predictability and proportional transaction costs can be tackled through the use of a duality method developed in Brown et al. [2010] based on information relaxation. The dual methods can be used to compute dual bounds on the optimal value function of dynamic portfolio selection problem through introducing proper penalty functions. We show that these dual bounds can significantly improve the bounds computed when no penalty function is considered.

Our work is related to Brown and Smith [2011] and DeMiguel et al. [2014]. Like Brown and Smith [2011], we propose some approximate trading strategies for a multiperiod investor with CRRA power utility, but instead of approximating the dynamic programming recursion (the continuation value functions) of the primal problem, we approximate the primal problem for each period with a quadratic program that can handle problems with more risky assets. In addition, we consider an investor who maximizes her power utility of intermediate consumption while Brown and Smith [2011] consider an investor who maximizes her utility of terminal wealth. DeMiguel et al. [2014] consider a mean-variance investor who faces general transaction costs and a constant investment opportunity set. For the case with proportional transaction costs, they give closed-form expressions for the no-trade region. In this chapter, we propose the approximate trading strategies based on their analysis on no-trade regions but we consider a more realistic case where there is predictability.

The remainder of this chapter is organized as follows. In Section 3.2, we introduce the dynamic portfolio selection problem in the presence of proportional transaction costs and predictability. In Section 3.3, we describe our approximate trading policies for a mean-variance investor and evaluate these approximate strategies under the mean-variance framework. Section 3.4 describes how to adapt these approximate strategies to a CRRA power utility framework and numerically evaluates these strategies. The evaluation is based on information relaxations that allow us to obtain dual bounds. Section 3.5 concludes. Appendix B.2 contains the tables. Appendix B.1 contains the derivation of the aim portfolio of linear policy while Appendix B.3 contains the derivation of the penalty function. In Appendix B.4 we describe how to approximate the consumption for each period for the model with transaction costs.

3.2 General Framework

We now describe the basic portfolio selection problem of an investor who needs to decide the portfolio weights for N risky assets. Assume time is discrete and indexed as $t = 1, \cdots, T$ with $t = 1$ being the current period and $t = T$ being the terminal period. There is also a risk-free asset being traded in the market and the risk-free rate R_f is assumed to be constant over time. From time $t-1$ to t, the risky asset returns are stochastic and denoted by $R_t = [R_{t,1}, \cdots, R_{t,N}]$, where $R_{t,i} \geq 0$ is the gross return on asset i. Based on the asset returns up to t, the investor then determines the decision vector $x_t = [x_{t,1}, \cdots, x_{t,N}]$, where $x_{t,i}$ is the weight of the ith asset hold in period t. Throughout this manuscript, we will use $\mathbf{x_t}$ to denote the $N \times t$ vector of decision variable, $[x_1, \cdots, x_t]$ and we always use $\mathbf{x}$ to denote $[x_1, \cdots, x_T]$.

Trading costs are imposed in the problem, and we assume that short selling is not allowed for risky assets. In this chapter, we will focus on the special case where the transaction cost for each period is proportional to the amount of trade. This type of transaction cost is realistic to model small trades, where the transaction costs come from the bid-ask spread and other brokerage fees. Let $x_{t,+}$ be the vector of allocation to the risky assets inherited from period t, that is

$$x_{t,+} = \frac{x_t \cdot R_{t+1}}{R_{p,t}}, \tag{3.1}$$

where $\cdot$ denotes component-wise multiplication and $R_{p,t}$ is the portfolio return which is defined as

$$R_{p,t+1} = x_t^\top R_{t+1} + [1 - x_t^\top e - \|K(x_t - x_{t-1,+})\|_1] R_f, \tag{3.2}$$

with e a vector of ones of length N. The term $\|K(x_t - x_{t-1,+})\|_1$ is the proportional transaction costs that the investor incurs. Here K is a $N \times N$ diagonal matrix with elements in the diagonal

$\mathrm{diag}(K) = [\kappa_1, \cdots, \kappa_N]$. Each κ_i denotes the proportional transaction-cost rate parameter for asset i. Taking transaction costs into account, the law of motion for the investor's wealth is given by

$$W_{t+1} = W_t(1 - c_t)R_{p,t}, \tag{3.3}$$

where W_t is the investor's wealth at t, and c_t is the fraction of wealth consumed at period t. The above law of motion assumes that the transaction costs are paid by costlessly liquidating the risk-free asset.

Let C_t be the total consumption, that is $C_t = c_t W_t$. The investor's objective is to maximize the expected utility of the intermediate consumption over all the periods:

$$\max_{\{c_t, x_t\}_{t=1}^T} \quad \mathbb{E}_1\Big[\sum_{t=1}^T \rho^t U_t(C_t)\Big], \tag{3.4}$$

where $\rho \in (0, 1)$ is the discount factor and U_t is the power utility function

$$U_t = \frac{C_t^{1-\gamma} - 1}{1 - \gamma}, \tag{3.5}$$

with relative risk aversion parameter $\gamma \geq 1$. In (3.4), $\mathbb{E}_1(\cdot)$ denotes the expectation conditioned on the information at the beginning of initial period $t = 1$.

Let $\{\mathcal{F}_t\}_{t=1}^T$ denote the filtration generated by the risky asset returns as well as other state variables in the model. This filtration is an indexed set that describes the investor's state of information that evolves over time. Each $\mathcal{F}_t$ represents the set of events that describes the investor's state of information at the beginning of period t and we require $\mathcal{F}_t \subseteq \mathcal{F}_{t+1}$ for all $t < T$ so the investor does not forget the past.

In our model, the investor must choose an x_t from a set P_t at each period t. Let $P(u) \subseteq P_1 \times \cdots \times P_T$ denote the set of all feasible action sequences $\mathbf{x} = [x_1, \cdots, x_T]$ for any given scenario of state variable u. A *feasible* policy is the one where each individual action x_t depends on previous actions $[x_1, \cdots, x_{t-1}]$ for all t. Let $\mathcal{P}$ denote the set of such policies. Besides feasibility, in our model, the investor's choice is also required to be nonanticipative such that the decision choice x_t for each period t must depend only on the information known at the beginning of t. To make the problem realistic, we assume action x_t satisfies the *nonanticipativity* constraints. Let $\mathcal{P}_\mathcal{F}$ be the set of nonanticipative feasible policies. With all of these definitions, we can now introduce the considered dynamic portfolio choice problem:

$$\max_{(\mathbf{c}, \mathbf{x}) \in \mathcal{P}_\mathcal{F}} \quad \mathbb{E}_1\Big[\sum_{t=1}^T \rho^t \frac{C_t^{1-\gamma}}{1 - \gamma}\Big] \tag{3.6}$$

$$\text{s.t.} \quad W_{t+1} = W_t(1 - c_t)R_{p,t}, \tag{3.7}$$

$$x_t \geq 0, c_t \geq 0. \tag{3.8}$$

Several comments are in order here. When we allow the presence of predictability, this portfolio optimization problem can be formulated as a stochastic dynamic program where the state variables include the current wealth level and portfolio position as well as the market state variables. Note that for the power utility function, the wealth level W_t can be factored out for each period; if we let $\mathbf{S}_t$ denote the vector of state variables, the Bellman equation for each period faced by the investor is,

$$\frac{\phi_t(\mathbf{S}_t, x_{t-1,+})}{1-\gamma} = \max_{(c_t,x_t)\in P_t(\mathbf{x_{t-1}})} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \frac{(1-c_t)^{1-\gamma}}{1-\gamma} \mathbb{E}\Big[\phi_{t+1}\big(\mathbf{S}_{t+1}, x_{t,+}\big)^{1-\gamma} R_{p,t+1}^{1-\gamma}\Big] \right\}, \quad (3.9)$$

for $t = 1, \cdots, T-1$. Note also that the above equation is solved by backward iteration, starting with $t = T-1$ and $\phi_T = 1$. In Section 3.2.1, we will introduce the model for the market state variables $\mathbf{S}_t$. Note that solving the dynamic program requires discretizing the state variables and the iteration in (3.9) involves the expectation of ϕ_{t+1}, while the dimension of the state space leads to the curse of dimensionality especially when there are more than two risky assets. In Section 3.3.1, we introduce the G&P framework that allows us to deal with many risky assets.

3.2.1 Predictability Model

As in Campbell and Viceira [1999], we assume the dynamics of asset returns and state variables follow a restricted first-order vector auto-regression model (VAR). With this model, the risky asset returns can be predicted by the log of the dividend-price ratio, which is the only state variable needed to forecast the risky asset dynamics. The restricted VAR model is also used in Balduzzi and Lynch [1999], Lynch and Tan [2010], Brown and Smith [2011] and Gârleanu and Pedersen [2013].

Specifically, let r_t be a vector consisting of the log of risky asset returns, $r_t = \log(R_t)$. Denote D_t the predictive variable (dividend yield) and let $d_t = \log(1 + D_t)$. We assume that r_t and d_t follow the given VAR model (expressed in terms of percentages):

$$r_{t+1} = A_r + B_r d_t + e_{t+1}, \quad (3.10)$$

$$d_{t+1} = a_d + b_d d_t + \epsilon_{t+1}. \quad (3.11)$$

Here, A_r is a $N \times 1$ vector, a_d is a scalar, B is a $N \times 1$ vector and b_d is a scalar. Moreover, $[e_{t+1}\ ;\ \epsilon_{t+1}]$ is a i.i.d. vector of mean-zero disturbances with constant covariance matrix $\Sigma_{e\epsilon}$. Without loss of generality, the mean of $\{d\}_{t=1}^{T}$ can be normalized to 0 and the variance to 1.

3.3 The Mean-Variance Approximation

In this section we propose several methods for constructing feasible sub-optimal trading strategies for a CRRA investor with objective (3.6). To avoid solving the dynamic program in (3.9) which may result in the curse of dimensionality, we will base our feasible strategies on the mean-variance framework adapted from G&P. Note that the performance estimates of these strategies can be obtained through simulation and these estimates provide lower bounds to the mean-variance utility. To test the sub-optimality of the approximate strategies, we also compute valid upper bounds for the utility of the mean-variance model by relaxing future information in Section 3.3.3.

3.3.1 Mean-variance Framework

As discussed in Section 3.2, problem (3.6)-(3.8) is difficult to solve even numerically when there are more than two risky assets. Heuristic trading strategies are proposed in Brown and Smith [2011] based on solving simpler optimization problems. These heuristic strategies help the investor to solve the dynamic program more efficiently based on an approximation in the continuation value function. But still, a significant amount of time is required to evaluate the quality of the heuristic strategies.

Compared with the power utility framework, the framework proposed by G&P is more analytically tractable. With quadratic transaction costs, the closed-form expressions for the optimal number of shares can be obtained based on their framework. With proportional transaction costs and a constant investment opportunity set, DeMiguel et al. [2014] analytically study the properties of optimal trading strategies and provide a closed-form expression for the no-trade regions based on the G&P framework. They also show that the certainty equivalent loss incurred from using a mean-variance utility instead of a CRRA utility of intermediate consumption is small. This implies that the optimal policy based on a quadratic utility specified in Gârleanu and Pedersen [2013] provides a reasonable approximation for the optimal policy implied by a CRRA utility in the presence of transaction costs. With the presence of predictability in price changes, the objective function for an investor with quadratic utility is

$$\max_{\mathbf{x}\in\mathcal{P}_{\mathcal{F}}} \quad \mathbb{E}_1\left\{ \sum_{t=1}^{T} \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \| K(x_t - x_{t-1}) \|_1 \right] \right\}, \tag{3.12}$$

where x_t denotes the number of shares, μ_t is the conditional expectation of price change, Σ is the covariance matrix of price changes, assumed to be constant, and γ is the *absolute* risk-aversion parameter.

With proportional transaction costs, it is still impossible to obtain a closed-form solution. But we can formulate the problem as a stochastic dynamic program with state variables consisting of the current number of shares in risky assets and the expected price changes conditional at the current period. To do that, note that the value function for the last period is:

$$V_T(x_{T-1}, \mu_T) = \rho^T (x_T^\top \mu_T - \frac{\gamma}{2} x_T^\top \Sigma x_T) - \rho^{T-1} \| K(x_T - x_{T-1}) \|_1, \tag{3.13}$$

and from (3.13), we can define the value functions for previous periods recursively using the Bellman equation

$$V_t(x_{t-1}, \mu_t) = \max_{x_t \in P_t(\mathbf{x_{t-1}})} \quad \rho^t(x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 + \mathbb{E}_t \left[V_{t+1}(x_t, \mu_{t+1}) \right], \tag{3.14}$$

for $t = 1, \cdots, T-1$. Still, the numerical solution is difficult to obtain when there are more than two risky assets. However, under the G&P framework, we only need to track the wealth change at each period instead of tracking the evolution of total wealth. Besides, unlike the model with power utility, the focus on price changes implies that there is no need to track the risky-asset price evolution. Hence, instead of considering the power utility framework, in this section, we are going to propose trading strategies based on the G&P framework. Later, in Section 3.4, we are going to adapt them to the CRRA utility framework.

3.3.2 Approximate Strategies

To avoid the difficulties we may be confronted with when solving the portfolio optimization model with predictability and transaction costs, we will propose several trading strategies to approximate the optimal trading strategy for (3.12). As in the G&P framework, we assume the dynamic of price changes follows the model specified in (3.10)-(3.11).

Simple Policy

First, we consider a *deterministic approximation* that ignores model predictability and simply follows the optimal trading strategy recommended by a deterministic model. In this model, the investor ignores the innovations in the predictability model. The resultant model becomes a deterministic problem which can be solved based on quadratic programming:

$$\max_{\{x_t\}_{t=1}^T} \quad \sum_{t=1}^T \left[\rho^t (x_t^\top \tilde{\mu}_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 \right], \tag{3.15}$$

where $\tilde{\mu}_t$ is the expectation of price changes conditional at the initial stage, that is $\tilde{\mu}_t = \mathbb{E}_1(\mu_t)$. The corresponding *simple policy* is defined as the solution to the deterministic portfolio problem (3.15). It is intuitive that the simple policy will perform well in practice when the volatility in the predictability model is small.

Linear Policy

Second, we consider the *linear approximation* where the investor trades linearly towards a next-period target. Gârleanu and Pedersen [2013] show that for an infinite-horizon investor who faces quadratic transaction costs and predictability, the next-period optimal number of shares is a linear combination of the existing position and the next-period target when transaction costs are quadratic. Moreover, the aim portfolio is a linear combination of the current optimal portfolio in the absence of transaction costs and the expected future aim portfolios. With a finite horizon and a constant opportunity set, DeMiguel et al. [2014] show that the optimal policy is a linear combination of the Markowitz portfolio, the previous period portfolio and the next period portfolio. The combination of the Markowitz strategy and the next period portfolio can be considered as the investor's target for the next period.

In this policy, we assume the investor has an aim portfolio at each period, and she chooses to trade partially to this aim portfolio based on the procedure that is optimal for the model with quadratic transaction costs. Specifically, following Gârleanu and Pedersen [2013], define the aim portfolio for period $t < T$ as

$$\text{aim}_t = z \text{ Markowitz}_t + (1-z)\mathbb{E}_t(\text{aim}_{t+1}), \tag{3.16}$$

where $z = \frac{c}{\gamma+c}$ and $c = \frac{-\gamma\rho-\delta(1-\rho)+\sqrt{(\gamma\rho+\delta(1-\rho))^2+4\delta\gamma\rho^2}}{2\rho}$. For $t = T$, define the aim portfolio as the optimal portfolio in the absence of transaction costs; that is,

$$\text{aim}_T = (\gamma\Sigma)^{-1}\mu_T. \tag{3.17}$$

We define the *linear policy* as

$$x_t = (1 - \frac{c}{\delta})x_{t-1} + \frac{c}{\delta}\text{aim}_t. \tag{3.18}$$

As in Gârleanu and Pedersen [2013], δ denotes the quadratic transaction cost parameter. Since we use (3.18) as an approximation to the model with proportional transaction costs, we propose to calibrate δ such that the realized utility is maximized. With the specified dynamics of price changes in (3.10)-(3.11), the aim portfolio for each period is

$$\text{aim}_t = (\gamma\Sigma)^{-1}(A_r + b_d^{T-t}B_r d_t) + z(\gamma\Sigma)^{-1}B_r d_t f_t(b_d), \tag{3.19}$$

where $f_t(b_d) = (1-z)b_d f_{t+1}(b_d) + 1 - b_d^{T-t}$ and $f_T = 0$.

No-trade Region Policy

Third, we consider the *no-trade region approximation* where, in each period, the investor's portfolio choice is confined by a no-trade region. DeMiguel et al. [2014] show analytically that the optimal trading strategies are confined by a no-trade region centered at a target portfolio in the presence of proportional transaction costs and they give a close-form expression for the no-trade region when there is no predictability. With predictability, Lynch and Tan [2010] numerically find the optimal rebalancing rule for each period to be a no-trade region with rebalancing to the boundary.

In the same spirit, we consider the investor's target position as the center of the no-trade region. But instead of trading linearly towards the aim portfolio, the investor will trade to the boundary of the no-trade region centered at the aim portfolio. Besides, following DeMiguel et al. [2014] and Lynch and Tan [2010], we assume the size of the no-trade region shrinks when the number of remaining periods increases. For each period, we define the *no-trade region policy* as the solution to the following optimization problem

$$\min_{x_t} \quad (x_t - x_{t-1})^\top \Sigma (x_t - x_{t-1}) \tag{3.20}$$

$$\text{s.t.} \quad \|K^{-1}\Sigma(x_t - \text{aim}_t)\|_\infty \le \frac{1}{\rho\gamma}\frac{1-\rho}{1-\rho^{T-t+1}}, \tag{3.21}$$

where aim_t denotes the aim portfolio for each period which is defined in (3.16), x_{t-1} is the position from the previous period. Note that the aim portfolio aim_t is specified in (3.19) as the dynamic of price changes follows (3.10)-(3.11).

Rolling Optimize-and-Hold Policy

Finally, we consider an approximate policy that assumes the investor has a buy-and-hold strategy at each stage. In Section 3.3.2, we assume the investor has an aim portfolio for each period and she trades to the boundary of the no-trade region centered at the aim portfolio. Taking into account that the aim portfolio is defined as the one under quadratic transaction costs and the optimal trading strategy for the model without predictability is a buy-and-hold strategy, we define the *rolling optimize-and-hold approximation* by assuming that the investor can trade only in the next period, but subsequently she will not be allowed to trade until time T.

Starting with the period before the last $t = T-1$, and assuming the investor does not trade at $t = T$ (which means $x_T = x_{T-1}$), the value function for the last period is

$$V_T^*(x_{T-1}, \mu_T) = \rho^T (x_{T-1}^\top \mu_T - \frac{\gamma}{2} x_{T-1}^\top \Sigma x_{T-1}). \tag{3.22}$$

The optimal strategy at $t = T-1$ is the solution to the following problem

$$\begin{aligned} V_{T-1}(x_{T-2}, \mu_{T-1}) = \max_{x_{T-1}} \quad & \rho^{T-1}(x_{T-1}^\top \mu_{T-1} - \frac{\gamma}{2} x_{T-1}^\top \Sigma x_{T-1}) - \rho^{T-2}\|K(x_{T-1} - x_{T-2})\|_1 \\ & + \mathbb{E}_{T-1}[V_T^*(\mu_T, x_{T-1})] \\ \equiv \max_{x_{T-1}} \; & \rho^{T-1}(x_{T-1}^\top[\mu_{T-1} + \rho\mu_{T\to T-1}] - \frac{\gamma(1+\rho)}{2} x_{T-1}^\top \Sigma x_{T-1}\Big\} \\ & - \rho^{T-2}\|K(x_{T-1} - x_{T-2})\|_1. \end{aligned} \tag{3.23}$$

In DeMiguel et al. [2014], they show that problem (3.23) is equivalent to the following constrained optimization problem

$$\min_{x_{T-1}} \quad (x_{T-1} - x_{T-2})^\top \Sigma (x_{T-1} - x_{T-2}) \tag{3.24}$$

$$\text{s.t.} \quad \|K^{-1}\Sigma(x_{T-1} - x^C_{T-1})\|_\infty \le \frac{1}{\rho\gamma_{T-1}}, \tag{3.25}$$

where $x^C_{T-1} = \frac{1}{\gamma_{T-1}}\Sigma^{-1}\mu'_{T-1}$, $\gamma_{T-1} = (1+\rho)\gamma$ and $\mu'_{T-1} = \mu_{T-1} + (1-\rho)\mu_{T|T-1}$. Analogously, assume trading only occurs in period t. In this case, in each period, the investor selects her portfolio by solving the following optimization problem

$$\min_{x_t} \quad (x_t - x_{t-1})^\top \Sigma (x_t - x_{t-1}) \tag{3.26}$$

$$\text{s.t.} \quad \|K^{-1}\Sigma(x_t - x^C_t)\|_\infty \le \frac{1}{\rho\gamma_t}, \tag{3.27}$$

where $x^C_t = (\gamma_t\Sigma)^{-1}\mu'_t$ with

$$\gamma_t = (1 + \rho + \cdots + \rho^{T-t})\gamma = \frac{1-\rho^{T-t+1}}{1-\rho}\gamma, \tag{3.28}$$

$$\mu'_t = \mu_t + \rho\mu_{t+1|t} + \cdots + \rho^{T-t}\mu_{T|t}. \tag{3.29}$$

In (3.29), each $\mu_{t+j|t}$ refers to the mean price changes for period $t+j$ conditioned on the information at period t. For each period, we define the *rolling optimize-and-hold policy* as the solution to problem (3.26)-(3.27) and the investor's transaction is confined by the no-trade region defined in (3.27). Also note that the center of the no-trade region x^C_t is a linear combination of all future optimal portfolios conditional at the current period in the absence of transaction costs.

3.3.3 Evaluation

In this section, we numerically study the performance of the proposed feasible policies. For each policy, we compare the realized utility with the upper bounds obtained from the *perfect hindsight solution*. Specifically, given the convexity of the objective function in (3.12), it is straightforward that

$$\begin{aligned} &\max_{\mathbf{x}\in\mathcal{P}_\mathcal{F}} \quad \mathbb{E}_1\left\{\sum_{t=1}^{T}\left[\rho^t(x_t^\top\mu_t - \frac{\gamma}{2}x_t^\top\Sigma x_t) - \rho^{t-1}\|K(x_t - x_{t-1})\|_1\right]\right\} \\ \le\quad &\mathbb{E}_1\left\{\max_{\mathbf{x}}\sum_{t=1}^{T}\left[\rho^t(x_t^\top\mu_t - \frac{\gamma}{2}x_t^\top\Sigma x_t) - \rho^{t-1}\|K(x_t - x_{t-1})\|_1\right]\right\}. \end{aligned} \tag{3.30}$$

In deriving the upper bounds for problem (3.12), we will focus on the perfect information relaxation that assumes the investor knows all market states and price changes before making any investment decisions. We obtain an estimate of the upper bound given on the right of (3.30) using simulation. In each trial of the simulation, we solve the following deterministic problem

$$\max_{\mathbf{x}} \sum_{t=1}^{T} \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \| K(x_t - x_{t-1}) \|_1 \right], \tag{3.31}$$

where we do not require $\mathbf{x}$ to be nonanticipative. The estimate of the upper bound is then obtained by averaging the optimal values from the above problems across all the sample paths.

We first generate a scenario tree for dividend yield and mean price changes. Consider the model with an investment horizon of T periods. Given any value of dividend yield d_t, n_d different values of dividend yields for the next period are generated according to the specified predictability model. Besides, the values of mean price changes are calculated at each node based on the specific price changes model. When evaluating the utility of each of the feasible policies, we calculate the utility for each period at each node and then average over all the branches for a given period. The realized utility is calculated by adding up the utilities for each period. Notice that the scenario tree defines n_d^{T-1} different sample paths for dividend yield and mean price changes, and the upper bound is then obtained by averaging the realized utilities over all the sample path.

We assume the predictability models for dividend yield and price changes are the ones given in (3.10)-(3.11) and that the initial dividend yield is neutral (i.e., $d_1 = 0$). We further assume that the investor has an initial wealth of \$1. With these assumptions, the absolute risk-aversion parameter under the mean-variance framework described in Section 3.3.1 is equivalent to the relative risk-aversion parameter under the power utility framework.

As an illustrative example, we consider the model with two risky assets.[2] Following the same example in Lynch and Tan [2010], we consider a model with two risky assets where we take a 12-month dividend yield on the value-weighted New York Stock Exchange (NYSE) index as a proxy for the predictive variable D. For these two risky assets, the first is the monthly rate of return on the value-weighted NYSE index while the second is the high BM portfolio which is formed from the six value-weighted portfolios SL, SM, SH, BL, BM and BH.[3] The parameters for the predictability model in (3.10)-(3.11) are estimated using ordinary least squares (OLS) with $A_r = [0.83; 0.54]$, $B_r = [0.47; 0.30]$, $a_d = 0$, $b_d = 0.98$ and stable state covariance matrix for both risky assets $\Sigma = [0.0054\ 0.0037; 0.0037\ 0.0030]$.

In our numerical experiment, we consider the model for annual price changes by annualizing the parameters for the model of monthly price changes. Specifically, let $A_r^Y = 12A_r$, $B_r^Y = 12B_r$ and the slope for the annualized d_t model be $b_d = 0.75$.

[2] Our approximate trading strategies can also be applied to the case with many risky assets. But the evaluation of the approximate strategies is computationally demanding since it requires us to discretize the state variable space for each of the assets. The resultant sample path is large especially when we use a scenario tree to capture the return dynamics.

[3] The notation S(B) indicates that the firms in the portfolio are smaller (larger) than 50% of NYSE stocks. The notation L indicates that the firms in the portfolio have BM ratios that place them in the bottom three deciles for all stocks; analogously, M indicates the middle four deciles and H indicates the top three deciles. The high BM portfolio is an equal-weighted portfolio of SH and BH.

The Base Case

For our base case, we consider a time horizon of six years, that is $T = 7$. We assume absolute risk aversion parameter $\gamma = 5$, matrix of proportional transaction cost $K = [0.0050\ 0; 0\ 0.0050]$, annual discount rate $\rho = \frac{1}{R_f}$, and the investor starts with zero shares in both assets. Let the number of branches at each node $n_d = 4$. With $T = 7$, we have $4^6 = 4096$ different sample paths.

For the *simple policy*, the expectation $\tilde{\mu}_t$ is $\tilde{\mu}_t = \mathbb{E}_1(\mu_t) = A_r^Y + B_r^Y b_d^{t-1} d_1$. Given the value of b_d, the innovation term ϵ_{t+1} is with zero mean and variance $1 - b_d^2$. Notice that with higher b_d, there is lower volatility in d_t.

For the *linear policy*, the value of δ is calibrated such that it provides maximum value of realized utility. Figure 3.1 depicts the utility of linear policy for values of δ. For the *no-trade region policy* and the *rolling optimize-and-hold policy*, the center of the no-trade region for each period can be calculated based on the conditional mean at each node.

For our base case, the realized utilities for the feasible policies and the associated upper bound are reported in Table 3.1. We observe that the utility loss associated with adopting the *simple policy* (that is, the relative difference between the utility of the *simple policy* and the upper bound obtained from perfect hindsight solution (3.30)) is as much as 18.27%. This utility loss is relatively large because the simple policy ignores the existence of volatility in the predictability model. The utility loss associated with trading linearly (that is, the relative difference between the utility of the *linear policy* and the upper bound obtained from perfect hindsight solution (3.30)) is 3.89%. To understand this result, it is important to note that this policy combines the previous stage position with the current stage target, and at each period it trades at a lower rate towards the target. The *no-trade region policy*, on the other hand, outperforms the other proposed strategies with the associated utility loss less than 0.91%. This result can be explained as follows: the approximate policy takes the target portfolio in linear policy as the center of the no-trade region and the existence of the no-trade region leads to a higher trading rate towards the target compared with linear policy. Finally, the utility loss associated with using *rolling optimize-and-hold policy* is 2.17%. This utility loss is relatively larger than that associated with the no-trade region policy. This is because compared with the target portfolio in linear policy, the center of the no-trade region may deviate from the true target for each period with the assumption that trade only occurs in the next period.

Table 3.1: Realized Utilities - The Base Case

This table shows the realized utilities for an investor with objective (3.12) based on the approximate policies proposed in Section 3.3.2. Columns 2-5 show the utilities obtained based on the *simple policy(S-P)*, the *linear policy(L-P)*, the *no-trade region policy(NTR-P)* and the *rolling optimize-and-hold policy(ROH-P)* respectively. The last column denotes the upper bounds to the utilities obtained based on a perfect hindsight solution. The values of the realized utilities are reported in the first row and the second row shows the corresponding gaps with respect to the upper bounds (in %). The gaps are computed by $(U_{bounds} - U_{heuristic})/U_{bounds}$.

Policies	S-P	L-P	NTR-P	ROH-P	Bounds
Utilities	0.0893	0.1050	0.1082	0.1068	0.1092
Utility losses (in %)	18.27	3.89	0.91	2.17	

Figure 3.1: Utility of Linear Policy Depending on δ

This figure plots the realized utilities for the linear policy depending on the value of δ while the other model parameters are fixed as $\kappa = [0.0050; 0.0050]$, discount factor $\frac{1}{R_f}$, risk-aversion parameter $\gamma = 3$ and the investment horizon T is fixed at $T = 8$.

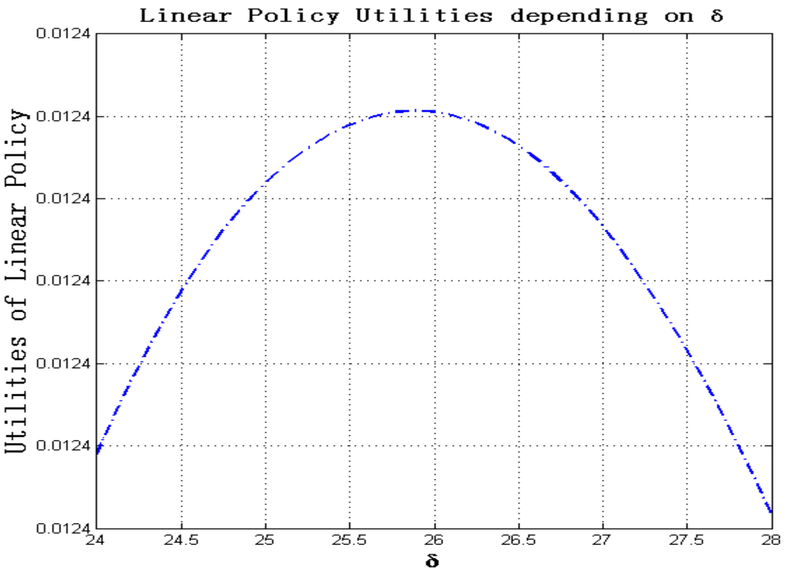

Comparative Statics

We study numerically how the utility loss associated with ignoring predictability (i.e., with the simple policy), trading linearly (i.e, with the linear policy), the no-trade region policy and the rolling optimize-and-hold policy depends on the transaction costs parameter, the risk-aversion parameter and the slope of predictability model for dividend yield.

Table B.1 shows how the utility loss of proposed approximate policies depends on different values of the above mentioned parameter. We find that regardless of the value of b_d, the realized utilities decrease monotonically for all policies as κ increases. The utility loss associated with ignoring uncertainty in the predictability model (i.e., the simple policy) is high for all the parameters we try especially when b_d is small (i.e., high volatility in dividend yield). Taking into account that bigger b_d indicates lower volatility in dividend yield, we observe that the utility loss in simple policy decreases dramatically when we increase b_d to 0.98. The reason is that the conditional mean $\tilde{\mu}_t$ in simple policy can approximate future means better when there is lower volatility in dividend yield. We also find that the utility loss associated with simple policy decreases as κ increases. This occurs because higher κ leads to less trade in risky assets, and thus predictability plays a less important role in the model.

When there is high volatility in d_t, the utility loss associated with trading linearly decreases with κ. As we explained in our discussion of the base case, linear policy indicates a slower trading rate compared with the other policies. The slower trading rate leads to partial loss in utility. As κ increases, there is less trade in risky assets. A slow trading rate in linear policy can compensate for the trade amount in subsequent periods. When there is low volatility in the predictability model, the utility loss first decreases and then increases slightly as κ increases. To understand this result, it is important to recognize that there is less trade involved in the optimal strategy when κ keeps increasing. The linear policy, however, incurs extra transaction costs because it trades at every period, which leads to the increment in utility loss.

The utility loss associated with the no-trade region policy increases as κ increases for both values of b_d. After looking into the solution, we find that for each sample path, the no-trade region policy encourages trading more frequently compared with the perfect information relaxation which incurs extra transaction costs. It is worth noting that the utility loss is low in both cases, with the highest loss no more than 1.5%.

When there is high volatility in dividend yield, the utility loss associated with the rolling optimize-and-hold policy is higher than that associated with no-trade region policy. Note that in this policy, the center of no-trade region depends directly on the assumption that there is no trade in future periods. With high volatility, the center of no-trade region defined in this policy cannot approximate accurately the true center. We also find that the trading rate in rolling optimize-and-hold policy decreases slower than that in the upper bound as κ increases. When there is low volatility, the center of no-trade region is more accurate, and the utility loss comes from the extra transaction costs incurred from extra trading, which leads to increment in utility loss as κ increases.

Finally, we find that the relative utility losses associated with the approximate policies do not depend on the risk-aversion parameter.

To conclude, the no-trade region policy and rolling optimize-and-hold policy perform quite well for all the values of κ (the utility loss is below 3% for all the values of κ when $b_d = 0.75$ and below 1% when $b_d = 0.98$). For the simple policy, it results in higher utility loss compared with the other policies. Based on the above robustness check, the benefits of the proposed approximation can be summarized as follows:

1. For the simple policy, if there is low volatility in dividend yield (hence the mean price changes conditional at the first period can predict the real conditional price changes at each period very well), it can perform well.

2. For the linear policy, it has lower utility loss than the simple policy but higher loss than the no-trade region policy and the rolling optimize-and-hold policy regardless of the volatility in dividend yield. The loss is always below 5%. So it can still be used as an approximation to the optimal solution.

3. For the no-trade region policy, it performs constantly well for all the cases. The utility loss for all cases is low. It can be used as a robust approximation for the optimal solution for the mean-variance portfolio model in the presence of proportional transaction costs and predictability. The gap between its realized utility and upper bound is relatively small for all cases.

4. For the rolling optimize-and-hold policy, just like the no-trade region policy, it can be a robust approximation for the optimal solution for the mean-variance portfolio model in the presence of proportional transaction costs and predictability. The gap between its realized utility and upper bound is relatively small for all cases.

3.4 Moving to the CRRA Framework

We now adapt each of the approximate trading strategies that are proposed for the mean-variance problem to the framework with CRRA power utility. Campbell and Viceira [2003] show that when the risky asset returns are lognormal, the portfolio choices resulting from the power utility and mean-variance frameworks are consistent. In the absence of transaction costs, the investor trades off mean against variance for a single period in both cases. In the presence of transaction costs, DeMiguel et al. [2014] show that the certainty equivalent loss from adapting the mean-variance framework is typically smaller than 0.5% for the case with a constant investment opportunity set.

In this section, we consider an investor who maximizes her CRRA utility of intermediate consumption by investing in a risk-free asset and N risky assets in the presence of predictability, and who is subject to proportional transaction costs (i.e, with preferences (3.6)-(3.8)). In Section 3.4.1, we obtain several approximate solutions for portfolio optimization problem (3.6)-(3.8) by adapting the feasible policies proposed based on the G&P framework. We check the robustness of these approximate policies by evaluating the certainty equivalent losses in Section 3.4.3.

3.4.1 Adapting the Mean-variance Framework to the CRRA Framework

Under the G&P framework, the proposed feasible policies in Section 3.3.2 provide a sub-optimal number of shares that the investor needs to hold for each period. Besides, there is no risk-free asset and consumption. In order to assess the robustness of these policies properly, we make several assumptions to adapt the feasible policies to the power utility framework.

First, we assume that the investor's consumption to wealth ratio for each period, c_t, is given by the model without transaction costs. Given the risky asset return dynamics in (3.10)-(3.11), the optimal solution to model (3.6)-(3.8) in the absence of transaction costs can easily be computed numerically using dynamic programming by discretizing market state variables. Note that this is a conservative choice since the consumption to wealth ratio given by this model is not optimal for the model in the presence of transaction costs. Second, we assume that the investor has an initial wealth of \$1 invested in the risk-free asset. With this assumption, the absolute risk aversion parameter in model (3.12) under the mean-variance framework equals the relative risk aversion parameter in (3.6)-(3.8) under the power utility framework. Consequently, for each period, the amount of money invested in the risk-free asset is

$$W_t^f = W_t(1-\hat{c}_t) - \hat{W}_t(1-\hat{c}_t)\hat{x}_t^\top P_t - \hat{W}_t(1-\hat{c}_t)\|K(\hat{x}_t - \hat{x}_{t-1,+})\|_1, \tag{3.32}$$

where $\hat{c}_t$ is the consumption to wealth ratio in the absence of transaction costs and $\hat{x}_t = x_t./P_t$ is the number of shares that the investor can hold when the price is P_t instead of \$1. Here ./ refers to the component-wise division of two vectors. The investor's wealth W_{t+1} in each period is the sum of the total holdings across the risky assets and risk-free asset, i.e.,

$$\hat{W}_{t+1} = W_t^f R_f + \hat{W}_t(1-\hat{c}_t)\hat{x}_t^\top P_{t+1}. \tag{3.33}$$

For each approximate policy, the realized power utility is

$$U^h = \mathbb{E}_1\Big[\sum_{t=1}^{T} \rho^t \frac{\hat{c}_t^{1-\gamma}\hat{W}_t^{1-\gamma}}{1-\gamma}\Big]. \tag{3.34}$$

To evaluate the expectation, we first generate a scenario tree with M sample paths for risky asset returns. Starting with any given initial dividend yield d_1, we generate n_d different values for d_2 based on (3.11) and save the realizations for the values of ϵ_2. The first period risky asset returns R_2 can be generated based on (3.10).[4] For each subsequent time period at each node, we repeat the process until we arrive at the final period $t = T$.

Assume that the initial price for each risky asset is \$1. Taking into account $P_{t+1} = P_t R_{t+1}$ for each period, the scenario tree for the price P_t can also be generated based on the scenario tree of risky asset returns. In each sample path, we compute the realized power utility for each feasible policy as

$$U_{(m)}^h = \sum_{t=1}^{T} \rho^t \frac{\hat{c}_{t,(m)}^{1-\gamma}\hat{W}_{t,(m)}^{1-\gamma}}{1-\gamma}, \tag{3.35}$$

[4]Following Balduzzi and Lynch [1999], in order to make sure d_t is the only state variable to predict risky asset returns, a regression between e_{t+1} and ϵ_{t+1} is designed when estimating the parameters in model (3.10)-(3.11).

for $m = 1, 2, \cdots, M$. Note that in each sample path, for any value of state variables that are not on the grid, the consumption can be computed based on interpolation. With M sample paths, an unbiased estimate for the utility given in (3.34) can be obtained by averaging the values of $U^h_{(m)}$. Let $\overline{U}_h$ be

$$\overline{U}_h = \frac{\sum_{m=1}^{M} U^h_{(m)}}{M}. \tag{3.36}$$

The certainty equivalent for each feasible policy is defined as

$$CE(\overline{U}_h) = ((1-\gamma)\overline{U}_h)^{\frac{1}{1-\gamma}}. \tag{3.37}$$

Given that each feasible policy provides an approximation to the optimal solution for the portfolio optimization problem with the CRRA power utility, the corresponding certainty equivalent is clearly a lower bound to that of the true model. Taking into account that it is very difficult to numerically solve the original model, we complement these proposed feasible policies with upper bounds on the certainty equivalent consumptions. This is based on the dual approach which will be explained in the next section.

3.4.2 Upper Bounds

Given any feasible sub-optimal policy to the portfolio selection problem (3.6)-(3.8), we can obtain an unbiased lower bound to the optimal power utility U^* by simulating different sample paths and taking the average of the realized utilities. However, since the optimal U^* is not computable, we cannot evaluate the optimality of the feasible policy. In order to assess the quality of a given feasible policy, it would be helpful if we could obtain a valid upper bound to the optimal power utility.

Brown et al. [2010] show that the valid upper bound can be constructed by relaxing future information and imposing a penalty function. Specifically, the upper bounds on the optimal value function can be obtained by relaxing the nonanticipativity constraints that require decisions to depend only on information available at the time when the decision is made and imposing a penalty function that punishes the violations of nonanticipativity. For each realized scenario, they show that an ideal penalty function $\pi(\mathbf{x})$ follows

$$\pi^*(\mathbf{x}) = \sum_{t=1}^{T} \Big\{ \mathbb{E}\big[V^*_{t+1}(\mathbf{x_t})|\mathcal{G}_t\big] - \mathbb{E}\big[V^*_{t+1}(\mathbf{x}_t)|\mathcal{F}_t\big] \Big\}, \tag{3.38}$$

where $\mathcal{G}_t$ is a relaxation of filtration $\mathcal{F}_t$ such that $\mathcal{F}_t \subseteq \mathcal{G}_t$. If we consider $\mathcal{G}_t$ to be the perfect information relaxation of $\mathcal{F}_t$, for each period the investor determines the actions with full knowledge of future scenarios. Obtaining the real optimal value function for each period is not practical. However, we can find an approximation of the optimal penalty function by approximating the optimal value function V_t^* with $\hat{V}_t$:

$$\hat{\pi}(\mathbf{x}) = \sum_{t=1}^{T} \{ \hat{V}_{t+1}(\mathbf{x_t}) - \mathbb{E}\left[\hat{V}_{t+1}(\mathbf{x_t}) | \mathcal{F}_t \right] \}, \tag{3.39}$$

where in (3.39), we drop the conditional expectation in $\mathcal{G}_t$ because when we assume perfect information relaxation, the term inside expectation is a constant. As mentioned in Brown and Smith [2014], weak duality implies that the upper bounds to the optimal value function $V_1^*(x_1)$ are given by

$$V_{ub} \;:=\; \mathbb{E}_1\Big[\max_{\mathbf{x} \in P(u)} \{ \sum_{t=1}^{T} I_t(\mathbf{x}_t) - \hat{\pi}(\mathbf{x}) \} \Big] \tag{3.40}$$

where $I_t(\mathbf{x}_t)$ denotes the reward function for each period. To obtain an unbiased estimate of the expectation, we simply simulate M sample paths of the state variables and disturbances, and solve the following maximization associated with penalty function (3.39) in each sample path,

$$\max_{\mathbf{c},\mathbf{x}} \sum_{t=1}^{T} \Big\{ \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + \mathbb{E}_t \left[\hat{V}_{t+1}(\mathbf{x_t}) \right] - \hat{V}_{t+1}(\mathbf{x_t}) \Big\}, \tag{3.41}$$

$$\text{s.t.} \quad W_{t+1} = W_t(1-c_t) R_{p,t}, \tag{3.42}$$

$$x_t \geq 0, c_t \geq 0. \tag{3.43}$$

The expected value in (3.40) is estimated by taking the average of the optimal objective function value of problem (3.41) over all sample paths.

In (3.40), each reward function $I_t(\mathbf{x}_t)$ is concave in decision variables $\mathbf{x}$, which leads to convexity in the original DP defined by (3.9). When taking the penalty function into account, the objective function (3.41) may not be concave and consequently, it cannot be solved computationally. Brown and Smith [2014] study DPs that have a convex structure and consider penalties based on a first-order approximation of the approximate value function $\hat{V}_t(\mathbf{x}_{t-1})$.

For the model with power utility and intermediate consumption, we consider $\tilde{I}_t(\mathbf{x}_t)$ to be a relaxation of the reward function of the original model,

$$\tilde{I}_t(c_t, \mathbf{x}_t) = \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} \tag{3.44}$$

subject to the law of motion for the investor's wealth which is defined in (3.3) but without transaction costs. With this relaxation, the realized utility for the model with reward function $\tilde{I}_t(\mathbf{c_t}, \mathbf{x_t})$ is greater than that for the original model. Let $\hat{V}_t(\mathbf{c_{t-1}}, \mathbf{x_{t-1}})$ denote the time t optimal value function corresponding to the above reward functions. It is then an approximation to the time t optimal function $V_t^*(\mathbf{c_{t-1}}, \mathbf{x_{t-1}})$ for the original model. Define the penalty function as

$$\pi(\mathbf{c}, \mathbf{x}) = \sum_{t=1}^{T} \left\{ \hat{V}_{t+1}(\mathbf{c_t}, \mathbf{x_t}) - \mathbb{E}_t \left[\hat{V}_{t+1}(\mathbf{c_t}, \mathbf{x_t}) \right] \right\}. \tag{3.45}$$

Following Brown and Smith [2014], weak duality implies that

$$\max_{(\mathbf{c},\mathbf{x}) \in \mathcal{P}_{\mathcal{F}}} \mathbb{E}_1 \Big[\sum_{t=1}^{T} \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} \Big] \;\leq\; \max_{(\mathbf{c},\mathbf{x}) \in \mathcal{P}_{\mathcal{G}}} \mathbb{E}_1 \Big[\sum_{t=1}^{T} \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} - \pi(\mathbf{c}, \mathbf{x}) \Big], \tag{3.46}$$

where $\mathcal{P}_{\mathcal{G}}$ denotes the set of feasible policies that are adapted to $\mathcal{G}$. Like the case with mean-variance utility, note that an unbiased estimate of the expectation on the right-hand side of the above inequality can be obtained via Monte Carlo simulation: we randomly generate M different sample paths of risky asset returns and in each sample path, solve the inner problem on the right-hand side of (3.46).

To guarantee the convex structure of the problem on the right-hand side of (3.46), we take the first-order linear approximation of $\hat{V}_t$ so that the resultant penalty function is linear in decision variables. Let $\tilde{y}_t^* = (\tilde{c}_t^*, \tilde{x}_t^*)$ denote the optimal policy for the no-transaction costs problem with reward function (3.44) and $\tilde{W}_t^*$ the corresponding wealth for each period. We define the gradient penalty as follows,[5]

$$\pi(\mathbf{c}, \mathbf{x}) = \sum_{t=1}^{T} \nabla \tilde{I}_t(\tilde{\mathbf{y}}_t^*)(\mathbf{y}_t - \tilde{\mathbf{y}}_t^*), \tag{3.47}$$

where $\nabla \tilde{I}_t(\tilde{\mathbf{y}}_t^*)$ denotes the gradient of the reward function for each period with respect to the decision vectors $\mathbf{c_t}$ and $\mathbf{x_t}$. To obtain the upper bound in each sample path, we solve the following deterministic problem

$$\max_{\mathbf{c},\mathbf{x}} \sum_{t=1}^{T} \Big\{ \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} - \nabla \tilde{I}_t(\tilde{\mathbf{y}}_t^*)(\mathbf{y}_t - \tilde{\mathbf{y}}_t^*) \Big\}, \tag{3.48}$$

$$\text{s.t.} \quad W_{t+1} = W_t(1 - c_t) R_{p,t}, \tag{3.49}$$

$$x_t \geq 0, c_t \geq 0. \tag{3.50}$$

Let $U_{(m)}^p$ be the corresponding optimal objective function from the above deterministic optimization for $m = 1, 2, \cdots, M$ and $\overline{U}_p = \frac{\sum_{m=1}^M U_{(m)}^p}{M}$. Then, $\overline{U}_p$ is clearly an upper bound to the utility of the heuristic policy $\overline{U}_h$.

[5]When we take first-order approximation of the approximate value function $\hat{V}_{t+1}(\mathbf{c_t}, \mathbf{x_t})$, the term which is constant in actions $\sum_{t=1}^T \hat{V}_{t+1}(\tilde{y}_t^*) - \mathbb{E}_t\big[\hat{V}_{t+1}(\tilde{y}_t^*)\big]$ is omitted from the penalty function $\pi(\mathbf{c}, \mathbf{x})$. When we calculate the upper bound for each scenario, we add the realized values for this term after (3.47), which serves as the control variate; see Brown and Smith [2014].

Given the fact that each reward function (3.44) is differentiable in decision variables, we can explicitly derive the penalty function in (3.47). Let the certainty equivalent be

$$CE(\overline{U}_p) = ((1-\gamma)\overline{U}_p)^{\frac{1}{1-\gamma}}. \tag{3.51}$$

$CE(\overline{U}_p)$ is clearly an upper bounds to the certainty equivalent provided by the approximate policies.

3.4.3 Numerical Results

In this section, we empirically study the certainty equivalent losses associated with adopting proposed feasible policies, as well as how those losses depend on the model parameters. Starting with an initial value of dividend yield $d_1 = 0$, we first generate a scenario tree of risky asset returns and dividend yield and then evaluate the realized utilities and dual bounds in each sample path. In the numerical experiments, we consider two risky assets with the same model parameters that we use in Section 3.3.3 for annual returns. We also consider a risk-free asset with an annual rate of return of 6%.

The Base Case

For our base case, we consider the same parameters that are used in Section 3.3.3. That is, the investor has a relative risk-aversion parameter $\gamma = 5$, has an initial wealth of \$1 invested in a risk-free asset, faces proportional transaction costs of 50 basis points for both assets, rebalances her portfolio once per year, and has an investment horizon of $T = 7$. Note that with $c_T = 1$, we only need to solve a six-period problem for the feasible policies to evaluate a $T = 7$ period problem based on the CRRA power utility.

For our base case, we observe that the certainty equivalent losses associated with adopting the simple policy (S-P), the linear policy (L-P), the no-trade region policy (NTR-P) and the rolling optimize-and-hold policy (ROH-P) are 7.08%, 6.84%, 5.40% and 5.70% respectively. The insights are similar to those in Section 3.3.3: the no-trade region policy is the best approximate feasible solution to the original model among all the feasible policies, followed by the rolling optimize-and-hold policy. Compared with the utility loss under the mean-variance framework, the certainty equivalent loss associated with adopting the simple policy is small, indicating a flatter utility than that of the mean-variance framework. This shows a proposed feasible policy which has large (small) utility loss under the mean-variance framework can have a small (large) certainty equivalent loss under the power utility framework.

Comparative Statics

In this section, we numerically study how the certainty equivalent loss associated with ignoring volatility in predictability (i.e., with the *simple policy*), trading linearly (i.e, with the *linear* *policy*), the *no-trade region policy* and the *rolling optimize-and-hold policy* depends on the transaction costs parameter, the risk-aversion parameter and the slope of predictability model for dividend yield.

Note that the loss in certainty equivalent can be decomposed into three parts that are not additive. The first part is the loss associated with adopting feasible policy instead of optimal policy. The utility losses associated with employing these policies are presented in Section 3.3.3. The second part is the loss from following portfolio policies that are derived based on the mean-variance framework. DeMiguel et al. [2014] show that the certainty equivalent loss from following such a policy is typically smaller than 0.5% when risky asset returns are i.i.d. Finally, the third part is the loss from comparing the lower bounds on certainty equivalent with an upper bound which is typically greater than the true optimal certainty equivalent.

Table B.2 shows how the certainty equivalent loss associated with adopting proposed approximate policies depends on different values of the above-mentioned parameters. It shows that there is significant improvement on the bounds provided based on perfect information relaxation with the penalty function.

In line with the conclusion in DeMiguel et al. [2014], the certainty equivalent loss associated with employing the feasible policies decreases as risk aversion parameter γ increases. Given that the risk-aversion parameter has no influence on the utility loss of the proposed approximate policies under the mean-variance framework, it can be understood that as γ increases, the amount invested in the risky assets decreases, and thus the optimal amount of rebalancing decreases. Moreover, we also observe that, the higher κ is always associated with higher certainty equivalent loss for NTR-P and ROH-P. This can be explained in two ways. First, the loss in mean-variance utility increases with κ, and second, the procedure we have used to augment the proposed policy to finance intermediate consumption requires a large amount of trading on risky assets, and thus large transaction costs.

Consistent with the conclusion in Section 3.3.3, the approximate policy that ignores uncertainty in predictability does not perform as well as the other heuristic policies, although there is less loss in certainty equivalent compared with the loss in mean-variance utility. Moreover, when there is high volatility in dividend yield, NTR-P performs better than ROH-P. When there is low volatility in dividend yield, it performs as well as ROH-P.

Interestingly, different from the case for NTR-P and ROH-P, the certainty equivalent loss associated with adopting simple policy and linear policy decreases as κ increases. This may be explained by the fact that the loss in mean-variance utility for these policies decreases as κ increases.

Overall, our results show that the certainty equivalent losses associated with the proposed feasible policies are quite acceptable with the no-trade region policy outperforming the other approximate strategies for most of the cases. Moreover, once given the approximate strategies for the mean-variance problem, the adapted strategies allow us to handle many risky assets simultaneously and the portfolio weights recommended by these adapted strategies can be determined in a very short time.

3.5 Summary

We consider a multiperiod CRRA individual who faces transaction costs and who has access to multiple risky assets in the presence of predictability. We propose some feasible trading strategies for the individual's multiperiod portfolio selection problem with proportional transaction costs, and construct lower and upper bounds on the certainty equivalent consumptions of these policies. In particular, we propose these feasible strategies based on the G&P framework which allows us to obtain the recommended portfolio weights through a quadratic problem with constraints. Our numerical experiments show that there is very little mean-variance utility loss, which indicates that some of the proposed strategies are nearly optimal in the G&P framework. With power utility, we find that the certainty equivalent losses are reasonable, and in addition, we can deal with the problem with many risky assets. For both cases, we have performed some comparative statics to better understand the losses associated with adopting the proposed approximate strategies. Moreover, we have shown how the upper bounds to the certainty equivalent consumption can be significantly improved through duality methods based on information relaxation.

Chapter 4

Conclusions and Future Research

4.1 Conclusion

Transaction cost plays an important role in optimal portfolio decision making. Ignoring transaction costs in the portfolio selection model clearly will lead to bankruptcy. Therefore, investors must consider transaction costs to obtain an optimal policy that trades at a carefully chosen discrete sequence of instances.

In this dissertation, we study the portfolio selection problem with multiple risky assets and different types of transaction costs in a multiperiod setting. We first consider the model for a multiperiod mean-variance investor facing general transaction costs in Chapter 2. With a constant investment opportunity set and proportional transaction costs, we simplify the multiperiod model into a single quadratic program that can deal with many risky assets simultaneously, and we provide the explicit expression for the no-trade region and analyze the property of the no-trade region based on relevant model parameters. For large trades that may impact market price, we show that the optimal policy for each period is confined by a state-dependent rebalancing region, and the rebalancing region converges to the Markowitz portfolio as the investment horizon grows large. Finally, we numerically study the utility loss associated with ignoring transaction costs, and we find that it is very important to take transaction costs into account.

We then consider the model for a multiperiod CRRA investor facing proportional transaction costs in Chapter 3. Instead of a constant investment opportunity set, we allow predictability in risky asset returns. In this chapter, we first approximate the model for a CRRA investor using a more tractable mean-variance framework, and then propose several approximate optimal solutions for such a framework based on solving simple quadratic programs. Numerical results shows that there is very small mean-variance utility loss associated with adopting these approximate strategies. In addition, we obtain several approximate solutions for the model with CRRA utility by adapting the feasible policies proposed based on the mean-variance framework. Finally, we complement the proposal of feasible policies with upper bounds to the certainty equivalent consumptions with duality methods based on information relaxation, and we show numerically that these adapted policies result in reasonable certainty equivalent losses.

4.2 Future Research Lines

Overall, we considered the mean-variance portfolio selection problems in the presence of different types of transaction costs in Chapter 2. However, for tractability reasons we make some assumptions in order to characterize the optimal trading strategies. Specifically, our results rely mainly on the assumption that the risky asset price changes are independently and identically distributed (iid). As a natural extension, we can relax this assumption and consider a more general model that takes into account the predictability in the risky assets price changes. Since Gârleanu and Pedersen [2013] already deal with predictability for the model with quadratic transaction costs, one of the future research lines to focus on is the market impact costs case. The same idea can apply to the extension for Chapter 3.

The above-mentioned research lines are natural extensions of this thesis and deal with the portfolio optimization problem such that an investor does not take into account any constraints on the transaction limit. An interesting future research line would be considering portfolio liquidation in a limit order book (LOB) model. In this model, the investor needs to trade a large block of orders of shares, which takes up a significant percentage of the daily traded volume of shares, and she needs to decide the best allocation strategy for each individual placement so as to minimize the overall price impact.

Moreover, modeling risky asset returns is attracting a great deal of interest in the literature as well. Among all the predictability models, the factor model is the most widely used. With a big data set, the selection of factors that can account for most of the variability of risky asset returns is a question for future research to explore.

Appendix A

Appendix to Chapter 2

A.1 Figures

Figure A.1: No-trade region and level sets for proportional transaction costs.

This figure depicts the no-trade region and the level sets when the investment horizon $T = 5$ for an investor facing proportional transaction costs with $\kappa = 0.005$, annual discount factor $\rho = 2\%$, absolute risk-aversion parameter $\gamma = 10^{-4}$, and mean and covariance matrix of price changes equal to the sample estimators for the commodity futures on gasoil and sugar.

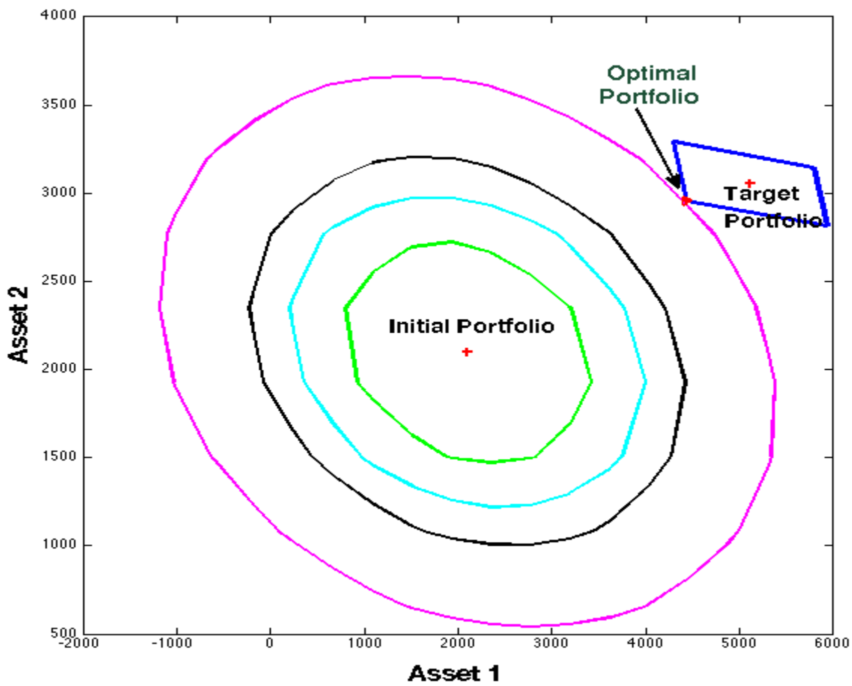

Figure A.2: No-trade region: comparative statics.

This figure shows how the no-trade region for a multiperiod investor subject to proportional transaction costs depends on relevant parameters. For the base case, we consider a proportional transaction cost parameter $\kappa = 0.005$, annual discount factor $\rho = 2\%$, absolute risk-aversion parameter $\gamma = 10^{-6}$, and mean and covariance matrix of price changes equal to the sample estimators for the commodity futures on gasoil and sugar.

(A) No-trade regions for different κ

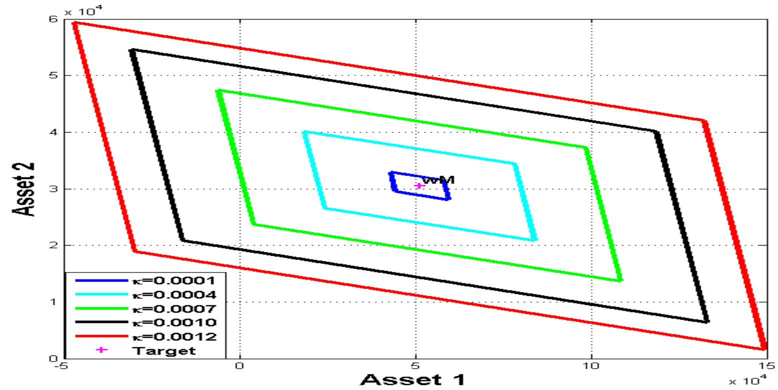

(B) No-trade regions for different γ

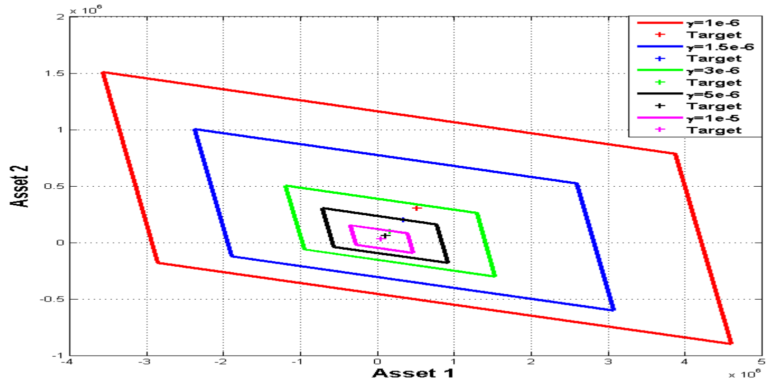

(C) No-trade regions for different ρ

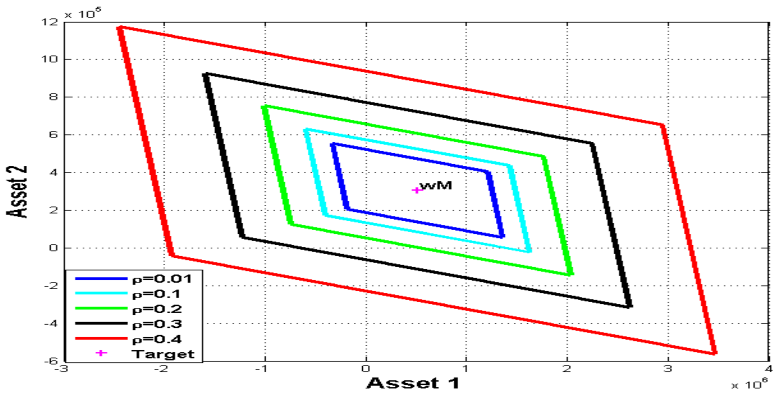

Figure A.3: No-trade region: comparative statics.

This figure shows how the no-trade region for a multiperiod investor subject to proportional transaction costs depends on relevant parameters. For the base case, we consider a proportional transaction cost parameter $\kappa = 0.005$, annual discount factor $\rho = 2\%$, absolute risk-aversion parameter $\gamma = 10^{-6}$, and mean and covariance matrix of price changes equal to the sample estimators for the commodity futures on gasoil and sugar.

(A) No-trade regions for different T

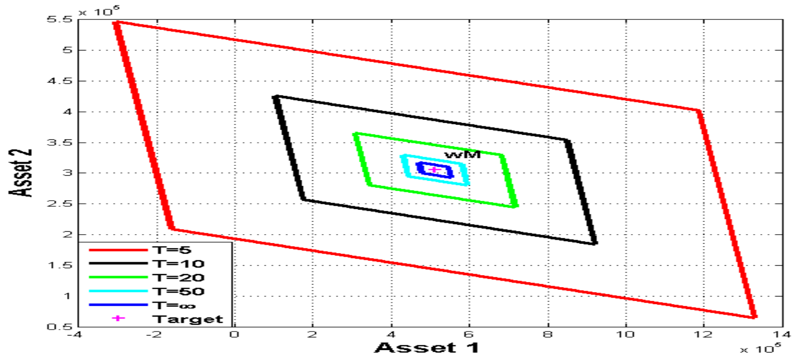

(B) No-trade regions for different correlations

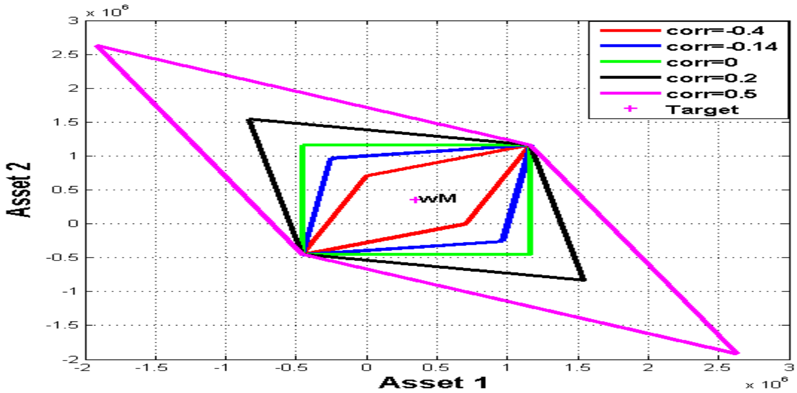

(C) No-trade regions for different σ^2

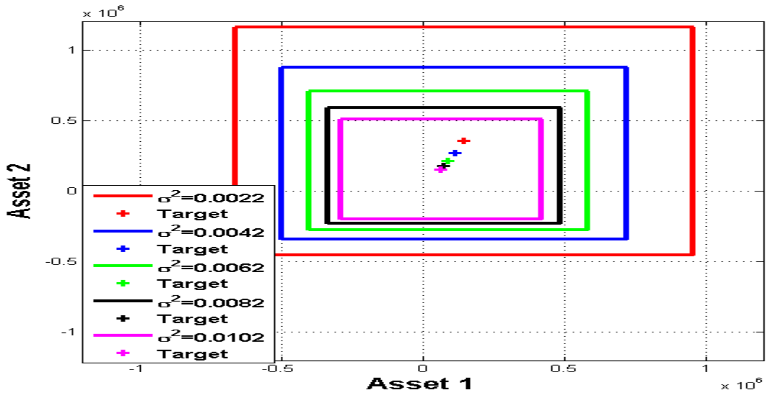

Figure A.4: Rebalancing region for single-period investor

This figure depicts the rebalancing region for a single-period investor subject to market impact costs. We consider an absolute risk aversion parameter $\gamma = 10^{-7}$, annual discount factor $\rho = 2\%$, and mean and covariance matrix of price changes equal to the sample estimators for the commodity futures on gasoil and sugar. Panel (A) depicts the rebalancing region when the transaction cost matrix $\Lambda = I$, the exponent of the power function is $p = 1.5$, and the transaction cost parameter $\kappa = 1.5 \times 10^{-8}$, and Panel (B) when $\Lambda = \Sigma$, $p = 1.5$, and $\kappa = 5 \times 10^{-6}$.

(A) Rebalancing region when $\Lambda = I$

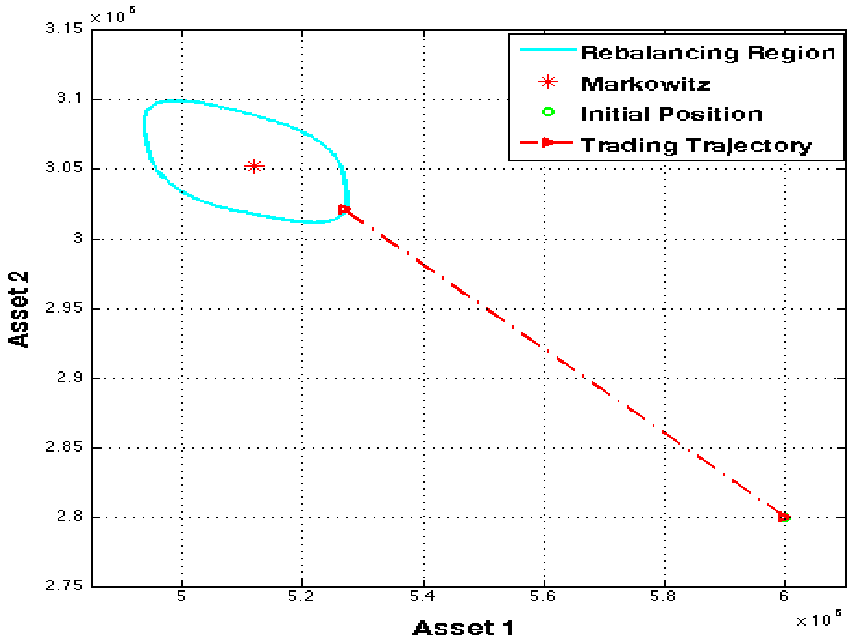

(B) Rebalancing region when $\Lambda = \Sigma$

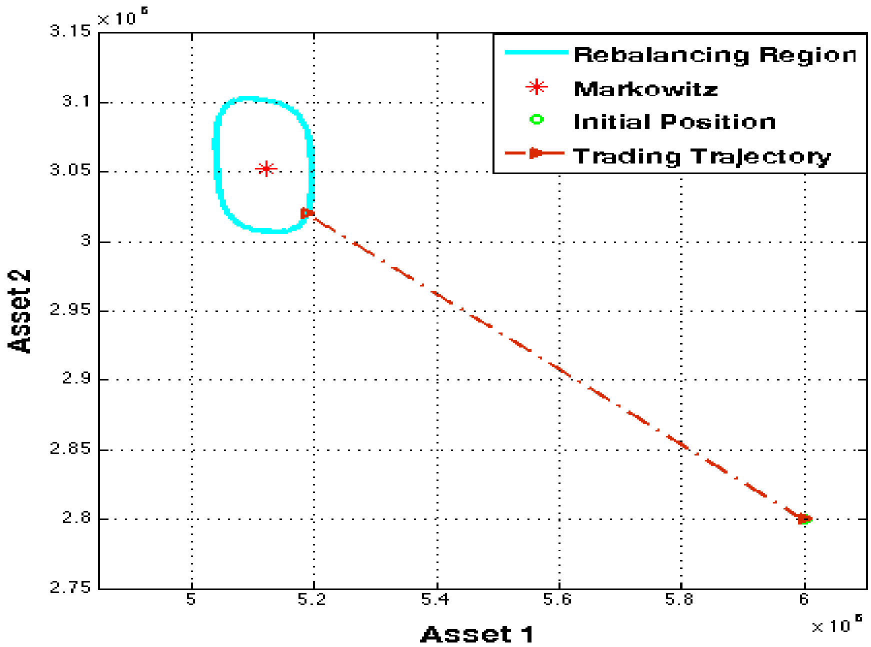

Figure A.5: Rebalancing region for multiperiod investor

This figure depicts the rebalancing region for a multiperiod investor subject to market impact costs. We consider an absolute risk aversion parameter $\gamma = 10^{-7}$, annual discount factor $\rho = 2\%$, and mean and covariance matrix of price changes equal to the sample estimators for the commodity futures on gasoil and sugar. Panel (A) depicts the rebalancing region when the transaction cost matrix $\Lambda = I$, the exponent of the power function is $p = 1.5$, and the transaction cost parameter $\kappa = 1.5 \times 10^{-8}$, and Panel (B) when $\Lambda = \Sigma$, $p = 1.5$, and $\kappa = 5 \times 10^{-6}$.

(A) Rebalancing region when $\Lambda = I$, $T = 3$

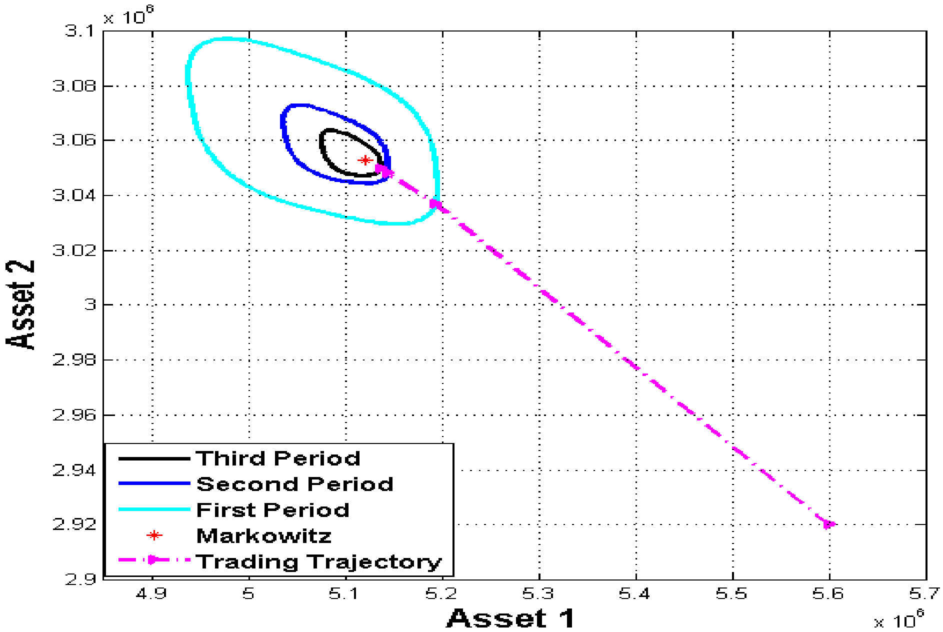

(B) Rebalancing region when $\Lambda = \Sigma$, $T = 3$

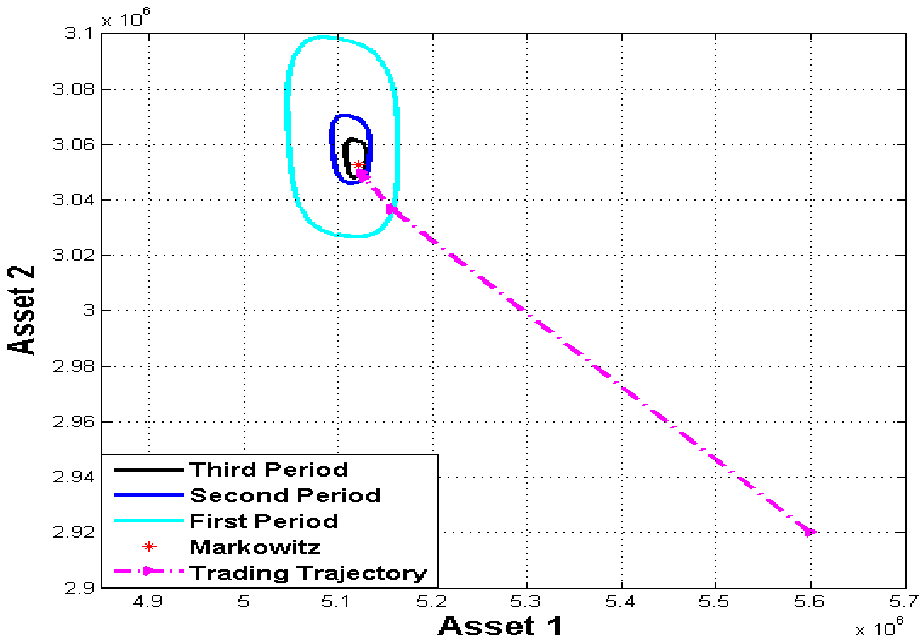

Figure A.6: Rebalancing regions and trading trajectories for different exponents p.

This figure shows how the rebalancing regions and trading trajectories for the market impact costs model change with the exponent of the transaction cost function p. Panel (A) depicts the rebalancing regions for the single-period investor, with transaction cost parameter $\kappa = 1.5 \times 10^{-8}$, annual discount factor $\rho = 50\%$. Panel (B) depicts the multiperiod optimal trading trajectories when the investment horizon $T = 10$, with transaction cost parameter $\kappa = 5 \times 10^{-6}$, and annual discount factor $\rho = 5\%$. In both cases, we consider transaction costs matrix $\Lambda = I$, the risk-aversion parameter $\gamma = 10^{-4}$, and mean and covariance matrix of price changes equal to the sample estimators for the commodity futures on gasoil and sugar.

(A) Rebalancing regions depending on exponent p.

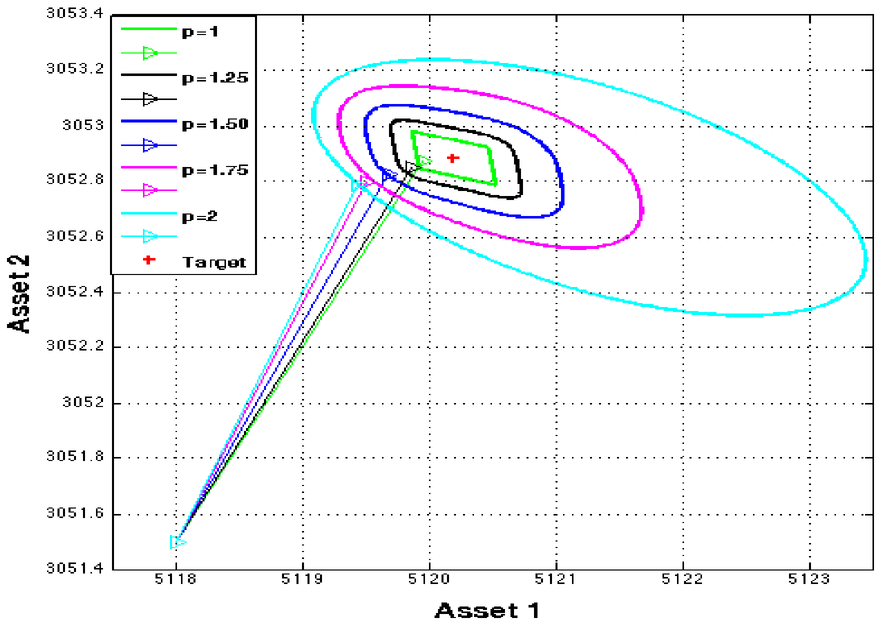

(B) Trading trajectories depending on exponent p.

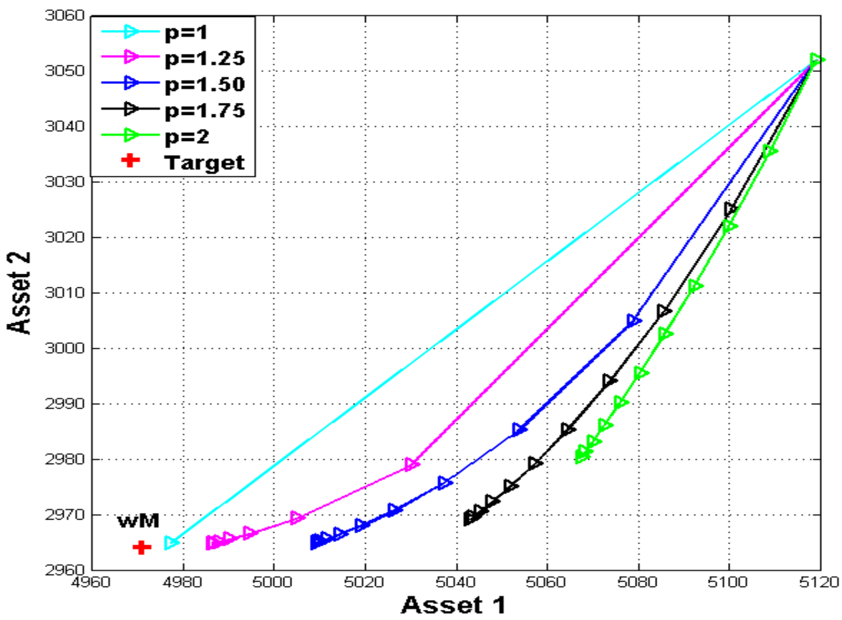

Figure A.7: Trading trajectories for different transaction costs.

This figure depicts the trading trajectories for a multiperiod investor facing different types of transaction cost. We consider an investment horizon $T = 3$, a risk-aversion parameter $\gamma = 10^{-4}$, and annual discount factor $\rho = 20\%$. In proportional transaction cost case, the transactions costs parameter $\kappa = 1.5 \times 10^{-8}$. In market impact cost case, $\kappa = 2\times10^{-8}$, $\Lambda = I$, and $p = 1.5$. In the quadratic transaction cost case, $\kappa = 2\times10^{-4}$ and $\Lambda = \Sigma$.

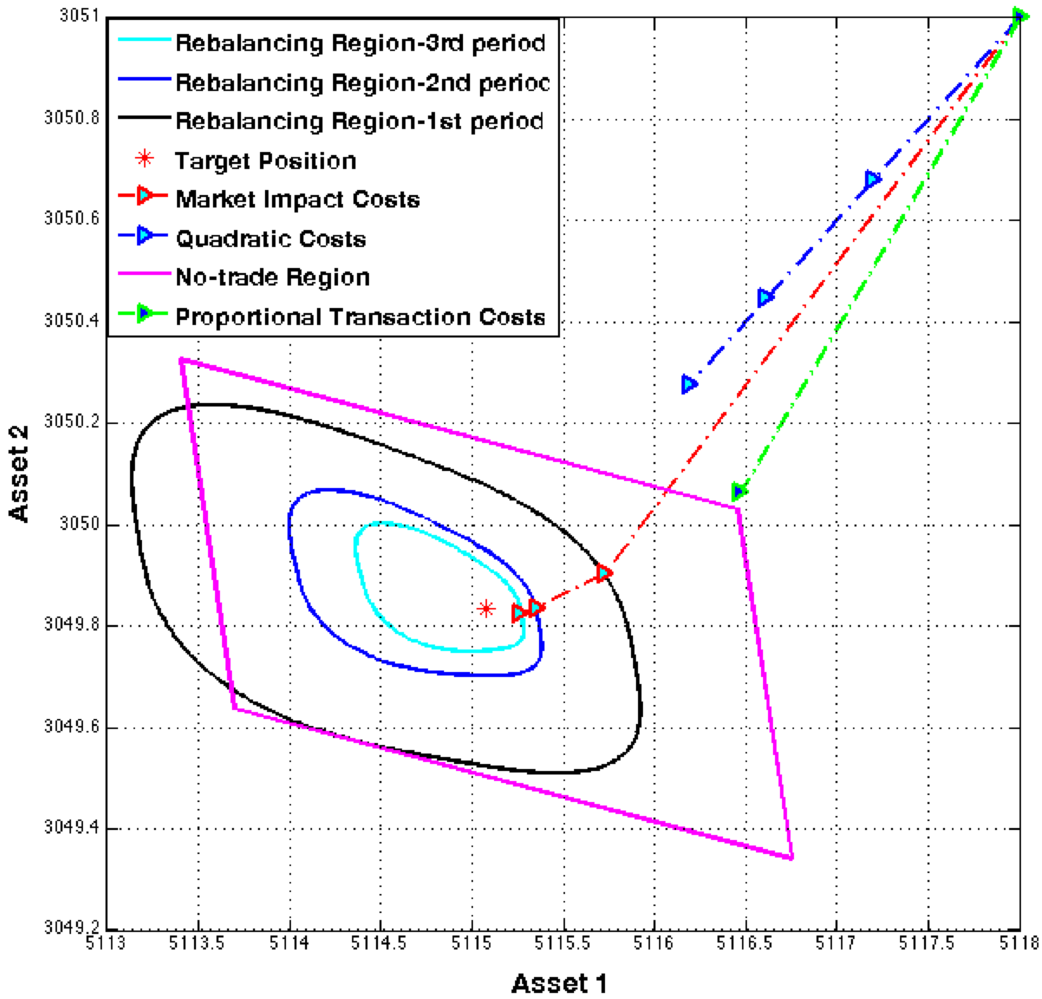

Figure A.8: Utility losses with proportional transaction costs.

This figure depicts the utility loss of the static and target portfolios for the dataset with 15 commodity futures as a function of the transaction cost parameter κ (Panel (A)), and the investment horizon T (Panel (B)). In the base case, we consider proportional transaction costs parameter $\kappa = 0.0050$, risk-aversion parameter $\gamma = 1e-6$, annual discount factor $\rho = 2\%$ and investment horizon $T = 22$. The price-change mean and covariance matrix are set equal to the sample estimators for the dataset that contains 15 commodity prices changes.

(A) Utility losses depending on κ.

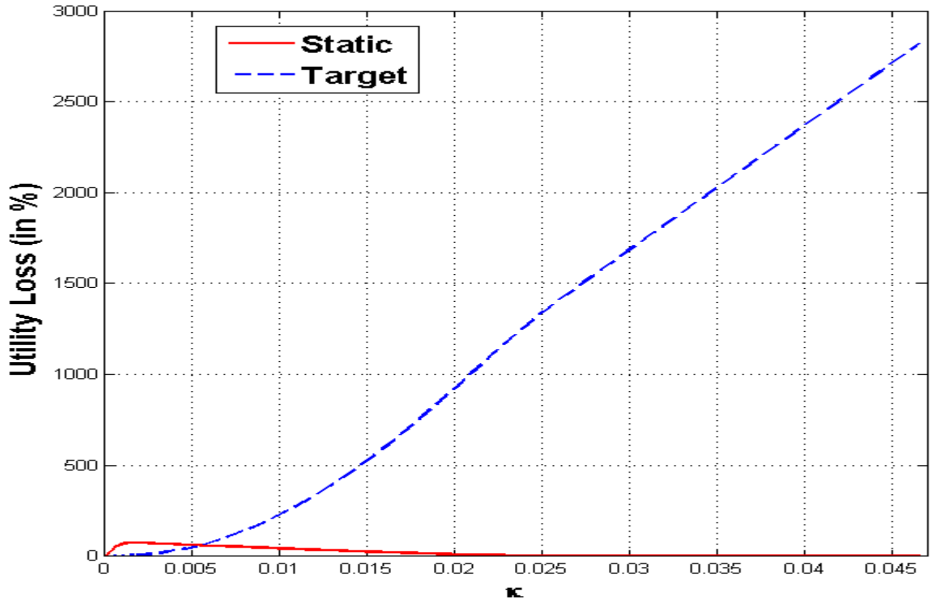

(B) Utility losses depending on investment horizon T.

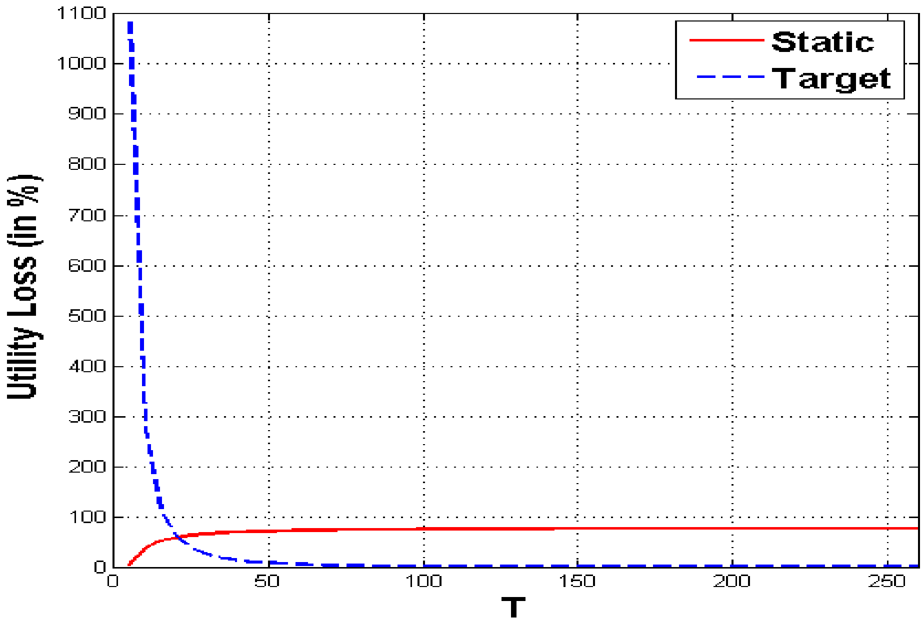

A.2 Tables

Table A.1: Certainty equivalent loss: CRRA utility of terminal wealth

This table shows the certainty equivalent (CEQ) wealth for an investor with CRRA utility of terminal wealth and for different values of the investor's risk aversion parameter (γ), the stock return volatility (σ), and the proportional transaction cost rate, which are given in the first three columns. The fourth column gives the certainty equivalent of the optimal portfolio policy, and the last column gives the percentage loss in certainty equivalent incurred by using the G&P-type portfolio policy.

Parameters			**CEQ**	**CEQ Loss**
γ	σ	κ	$	%
2	0.15	0.000	2.8756	0.19%
		0.005	2.8475	0.13%
		0.010	2.8206	0.08%
	0.20	0.000	2.6730	0.28%
		0.005	2.6565	0.18%
		0.010	2.6417	0.12%
	0.25	0.000	2.5857	0.45%
		0.005	2.5739	0.32%
		0.010	2.5640	0.25%
3	0.15	0.000	2.2643	0.19%
		0.005	2.2487	0.11%
		0.010	2.2342	0.07%
	0.20	0.000	2.1574	0.30%
		0.005	2.1479	0.20%
		0.010	2.1396	0.15%
	0.25	0.000	2.1103	0.40%
		0.005	2.1035	0.29%
		0.010	2.0980	0.24%
4	0.15	0.000	2.0720	0.19%
		0.005	2.0609	0.12%
		0.010	2.0508	0.07%
	0.20	0.000	1.9985	0.28%
		0.005	1.9917	0.19%
		0.010	1.9859	0.14%
	0.25	0.000	1.9658	0.34%
		0.005	1.9609	0.25%
		0.010	1.9570	0.21%

Table A.2: Certainty equivalent loss: CRRA utility of intermediate consumption

This table shows the certainty equivalent (CEQ) consumption for an investor with CRRA utility of intermediate consumption and for different values of the investor's risk aversion parameter (γ), the stock return volatility (σ), and the proportional transaction cost rate, which are given in the first three columns. The fourth column gives the certainty equivalent consumption of the optimal portfolio policy, and the last column gives the percentage loss in certainty equivalent consumption incurred by using the G&P-type portfolio policy.

Parameters			**CEQ**	**CEQ Loss**
γ	σ	κ	\$	%
2	0.15	0.000	0.0165	0.08%
		0.005	0.0164	0.22%
		0.010	0.0163	0.45%
	0.20	0.000	0.0160	0.09%
		0.005	0.0159	0.17%
		0.010	0.0158	0.29%
	0.25	0.000	0.0157	0.14%
		0.005	0.0157	0.18%
		0.010	0.0156	0.25%
3	0.15	0.000	0.0465	0.07%
		0.005	0.0463	0.16%
		0.010	0.0461	0.32%
	0.20	0.000	0.0455	0.09%
		0.005	0.0454	0.14%
		0.010	0.0453	0.22%
	0.25	0.000	0.0451	0.12%
		0.005	0.0450	0.14%
		0.010	0.0449	0.19%
4	0.15	0.000	0.0655	0.07%
		0.005	0.0652	0.13%
		0.010	0.0650	0.25%
	0.20	0.000	0.0644	0.09%
		0.005	0.0643	0.12%
		0.010	0.0642	0.18%
	0.25	0.000	0.0640	0.10%
		0.005	0.0639	0.12%
		0.010	0.0638	0.15%

A.3 Proofs

Proof of Proposition 2.1

Part 1. When $\rho = 0$, the change in excess terminal wealth net of transaction costs for a multiperiod investor is

$$W_T = \sum_{t=1}^{T} \left[x_t^\top r_{t+1} - \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right]. \tag{A.1}$$

From Assumption 1, it is straightforward that the expected change in terminal wealth is

$$E_0(W_T) = \sum_{t=1}^{T} \left(x_t^\top \mu - \kappa \|\Lambda^{1/p}\| x_t - x_{t-1}\|_p^p \right). \tag{A.2}$$

Using the law of total variance, the variance of change in terminal wealth can be decomposed as

$$var_0(W_T) = E_0\left[var_s(W_T)\right] + var_0\left[E_s(W_T)\right]. \tag{A.3}$$

Taking into account that

$$\begin{aligned}
E_0\left[var_s(W_T)\right] &= E_0 \left\{ var_s \left[\sum_{t=1}^{T} \left[x_t^\top r_{t+1} - \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right] \right] \right\} \\
&= E_0 \left[\sum_{t=s}^{T} x_t^\top \Sigma x_t \right] \\
&= \sum_{t=s}^{T} x_t^\top \Sigma x_t,
\end{aligned} \tag{A.4}$$

and

$$\begin{aligned}
var_0\left[E_s(W_T)\right] &= var_0 \left\{ E_s \left[\sum_{t=1}^{T} \left[x_t^\top r_{t+1} - \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right] \right] \right\} \\
&= var_0 \left\{ \sum_{t=1}^{s-1} \left[x_t^\top r_{t+1} - \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right] + \sum_{t=s}^{T} \left[x_t^\top \mu - \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right] \right\} \\
&= \sum_{t=1}^{s-1} x_t^\top \Sigma x_t,
\end{aligned} \tag{A.5}$$

the variance of the change in excess terminal wealth can be rewritten as

$$var_0(W_T) = \sum_{t=s}^{T} x_t^\top \Sigma x_t + \sum_{t=1}^{s-1} x_t^\top \Sigma x_t = \sum_{t=1}^{T} x_t^\top \Sigma x_t. \tag{A.6}$$

Consequently, the mean-variance objective of the change in excess terminal wealth for a multi-period investor is

$$
\begin{aligned}
& \max_{\{x_t\}_{t=1}^T} \quad E(W_T) - \frac{\gamma}{2} var(W_T) \\
\equiv & \max_{\{x_t\}_{t=1}^T} \quad \sum_{t=1}^T \left(x_t^\top \mu - \kappa \|\Lambda^{1/p}\|x_t - x_{t-1}\|_p^p \right) - \frac{\gamma}{2} \sum_{t=1}^T x_t^\top \Sigma x_t \\
\equiv & \max_{\{x_t\}_{t=1}^T} \quad \sum_{t=1}^T \left[(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \kappa \|\Lambda^{1/p}(x_t - x_{t-1})\|_p^p \right].
\end{aligned} \tag{A.7}
$$

Which is exactly objective function (2.1) when the value of discount factor $\rho = 0$.

Part 2. For the model with proportional transaction costs, the optimal policy is to trade at the first period to the boundary of the no-trade region given by (2.5), and not to trade for periods $t = 2, 3, \cdots, T$. For an investor who is now sitting at period j with $j > 0$, the no-trade region for the remaining periods $t = j+1, j+2, \cdots, T$ is given by

$$
\|\Sigma(x - x^*)\|_\infty \leq \frac{\kappa}{(1-\rho)\gamma} \frac{\rho}{1 - (1-\rho)^{T-j}}, \tag{A.8}
$$

which defines a region that contains the region defined by constraint (2.5). we could infer that the "initial position" for stage s, which is on the boundary of no-trade region defined by (2.5), is inside the no-trade region defined in (A.8). Hence for any $\tau > j$, the optimal strategy for an investor who is sitting at j is to stay at the boundary of no-trade region defined in (2.5), which is consistent with the optimal policy obtained at time $t = 0$.

For the model with temporary market impact costs, for simplicity of exposition we consider $\Lambda = \Sigma$. At $t = 0$, the optimal trading strategy for period τ is on the boundary of the following rebalancing region

$$
\frac{\| \sum_{s=\tau}^T (1-\rho)^{s-\tau} \Sigma^{1/q} (x_{s|0} - x^*) \|_q}{p \|\Sigma^{1/p}(x_{t|0} - x_{t-1|0})\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{A.9}
$$

For the investor who is now at $t = j$ for $j < \tau$, the optimal trading strategy is at the boundary of rebalancing region given by

$$
\frac{\| \sum_{s=\tau}^T (1-\rho)^{s-t} \Sigma^{1/q} (x_{s|j} - x^*) \|_q}{p \|\Sigma^{1/p}(x_{t|j} - x_{t-1|j})\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{A.10}
$$

Because $x_{\tau|j} = x_{\tau|0}$ for $\tau \leq j$ and taking into account that $x_{j|j} = x_{j|0}$, we can infer that

$$
\frac{\| \sum_{s=j}^T (1-\rho)^{s-t} \Sigma^{1/q} (x_{s|j} - x^*) \|_q}{p \|\Sigma^{1/p}(x_{j|j} - x_{j-1|j})\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma} \tag{A.11}
$$

defines the same rebalancing region for period j as the following one

$$\frac{\|\sum_{s=j}^{T}(1-\rho)^{s-t}\Sigma^{1/q}(x_{s|0}-x^*)\|_q}{p\|\Sigma^{1/p}(x_{j|0}-x_{j-1|0})\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{A.12}$$

That is, $x_{\tau|j} = x_{\tau|0}$ for $j < \tau$, and hence the optimal policies for the model with temporary market impact is time-consistent.

For the model with quadratic transaction costs, for simplicity of exposition we consider $\Lambda = \Sigma$. For an investor sitting at period $t = 0$, the optimal trading strategy for a future period $t = \tau$ is given by (2.20) (when $t = T$, simply let $\alpha_1 = \beta_1$, $\alpha_2 = \beta_2$, $\alpha_3 = 0$.), that is,

$$x_{\tau|0} = \alpha_1 x^* + \alpha_2 x_{\tau-1|0} + \alpha_3 x_{\tau+1|0}. \tag{A.13}$$

For the investor who is at $t = j$, for $j < \tau$, the optimal trading strategy is

$$x_{\tau|j} = \alpha_1 x^* + \alpha_2 x_{\tau-1|j} + \alpha_3 x_{\tau+1|j}. \tag{A.14}$$

Because $x_{\tau|j} = x_{\tau|0}$ for $\tau \leq j$ and taking into account that

$$\begin{aligned} &x_{j|0} = \alpha_1 x^* + \alpha_2 x_{j-1|0} + \alpha_3 x_{j+1|0} \\ =&x_{j|j} = \alpha_1 x^* + \alpha_2 x_{j-1|j} + \alpha_3 x_{j+1|j} \\ =&\alpha_1 x^* + \alpha_2 x_{j-1|0} + \alpha_3 x_{j+1|j}, \end{aligned} \tag{A.15}$$

which gives $x_{j+1|0} = x_{j+1|j}$. By using the relation recursively, we can show that for all $\tau > j$, it holds $x_{\tau|j} = x_{\tau|0}$. □

Proof of Theorem 2.2

Part 1. Define Ω_t as the subdifferential of $\kappa\|x_t - x_{t-1}\|_1$

$$s_t \in \Omega_t = \left\{ u_t \mid u_t^\top (x_t - x_{t-1}) = \kappa\|x_t - x_{t-1}\|_1, \|u_t\|_\infty \leq \kappa \right\}, \tag{A.16}$$

where s_t denotes a subgradient of $\kappa\|x_t - x_{t-1}\|_1$, $t = 1, 2, \cdots, T$. If we write $\kappa\|x_t - x_{t-1}\|_1 = \max_{\|s_t\|_\infty \leq \kappa} s_t^T (x_t - x_{t-1})$, objective function (2.2) can be sequentially rewritten as

$$\begin{aligned} &\max_{\{x_t\}_{t=1}^T} \sum_{t=1}^{T} \left[(1-\rho)^t \left(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t \right) - (1-\rho)^{t-1} \kappa \|x_t - x_{t-1}\|_1 \right] \\ =& \max_{\{x_t\}_{t=1}^T} \min_{\|s_t\|_\infty \leq \kappa} \sum_{t=1}^{T} \left[(1-\rho)^t \left(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t \right) - (1-\rho)^{t-1} s_t^\top (x_t - x_{t-1}) \right] \\ =& \min_{\|s_t\|_\infty \leq \kappa} \max_{\{x_t\}_{t=1}^T} \sum_{t=1}^{T} \left[(1-\rho)^t \left(x_t^\top \mu - \frac{\gamma}{2} x_t^\top \Sigma x_t \right) - (1-\rho)^{t-1} s_t^\top (x_t - x_{t-1}) \right]. \end{aligned} \tag{A.17}$$

The first order condition for the inner objective function of (A.17) with respect to x_t is

$$0 = (1-\rho)(\mu - \gamma\Sigma x_t) - s_t + (1-\rho)s_{t+1}, \tag{A.18}$$

and hence

$$x_t = \frac{1}{\gamma}\Sigma^{-1}(\mu + s_{t+1}) - \frac{1}{(1-\rho)\gamma}\Sigma^{-1}s_t, \qquad \mathit{for} \quad s_t \in \Omega_t,\ s_{t+1} \in \Omega_{t+1}. \tag{A.19}$$

Denote x_t^* as the optimal solution for stage t, there exists s_t^* and s_{t+1}^* such that

$$x_t^* = \frac{1}{\gamma}\Sigma^{-1}(\mu + s_{t+1}^*) - \frac{1}{(1-\rho)\gamma}\Sigma^{-1}s_t^*, \quad \forall\ t. \tag{A.20}$$

We now let $s_t^* = \frac{1-(1-\rho)^{T-t+2}}{\rho}s_T^*$, for $t = 1, 2, \cdots, T-1$ and $s_T^* = (1-\rho)(\mu - \gamma\Sigma x_T^*)$. Rewrite x_t^* as

$$\begin{aligned} x_t^* &= \frac{1}{\gamma}\Sigma^{-1}(\mu + s_{t+1}^*) - \frac{1}{(1-\rho)\gamma}\Sigma^{-1}s_t^* \\ &= \frac{1}{\gamma}\Sigma^{-1}(\mu + s_{r+1}^*) - \frac{1}{(1-\rho)\gamma}\Sigma^{-1}s_r^* = x_r^*, \quad \forall \quad t, r, \end{aligned} \tag{A.21}$$

where $\|s_t\|_\infty \le \kappa$. By this means, we find the value of s_t^* such that $x_t^* = x_r^*$ for all $t \neq r$. We conclude that $x_1 = x_2 = \cdots = x_T$ satisfies the optimality conditions.

Part 2. Because $x_1 = x_2 = \cdots = x_T$, one can rewrite the objective function (2.2) as

$$\begin{aligned} &\max_{\{x_t\}_{t=1}^T} \left\{ \sum_{t=1}^T \left[(1-\rho)^t \left(x_t^\top\mu - \frac{\gamma}{2}x_t^\top\Sigma x_t \right) - (1-\rho)^{t-1}\kappa\|x_t - x_{t-1}\|_1 \right] \right\} \\ =&\max_{x_1} \left\{ \sum_{t=1}^T \left[(1-\rho)^t \left(x_1^\top\mu - \frac{\gamma}{2}x_1^\top\Sigma x_1 \right) \right] - \kappa\|x_1 - x_0\|_1 \right\} \\ =&\max_{x_1} \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}\left(x_1^\top\mu - \frac{\gamma}{2}x_1^\top\Sigma x_1 \right) - \kappa\|x_1 - x_0\|_1. \end{aligned} \tag{A.22}$$

Let s be the subgradient of $\kappa\|x_1 - x_0\|_1$ and let Ω be the subdifferential

$$s \in \Omega = \left\{ u \mid u^\top(x_1 - x_0) = \kappa\|x_1 - x_0\|_1, \|u\|_\infty \le \kappa \right\}. \tag{A.23}$$

If we write $\kappa\|x - x_0\|_1 = \max_{\|s\|_\infty \le \kappa} s^\top(x - x_0)$, objective function (A.22) can be sequentially rewritten as:

$$\begin{aligned} &\max_{x_1} \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}(x_1^\top\mu - \frac{\gamma}{2}x_1^\top\Sigma x_1) - \kappa\|x_1 - x_0\|_1 \\ =&\max_{x_1}\min_{\|s\|_\infty \le \kappa} \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}(x_1^\top\mu - \frac{\gamma}{2}x_1^\top\Sigma x_1) - s^\top(x_1 - x_0) \\ =&\min_{\|s\|_\infty \le \kappa}\max_{x_1} \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}(x_1^\top\mu - \frac{\gamma}{2}x_1^\top\Sigma x_1) - s^\top(x_1 - x_0). \end{aligned} \tag{A.24}$$

The first order condition for the inner objective function in (A.24) is

$$0 = \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}(\mu - \gamma\Sigma x_1) - s, \tag{A.25}$$

and hence $x_1 = \frac{1}{\gamma}\Sigma^{-1}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s)$ for $s \in \Omega$. Then we plug x_1 into (A.24),

$$\begin{aligned}
\min_{\|s\|_\infty \le \kappa} \quad & \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}\Bigg\{\left[\frac{1}{\gamma}\Sigma^{-1}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s)\right]^\top \mu && \text{(A.26)}\\
& - \frac{\gamma}{2}\left[\frac{1}{\gamma}\Sigma^{-1}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s)\right]^\top \Sigma \left[\frac{1}{\gamma}\Sigma^{-1}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s)\right]\Bigg\}\\
& - s^\top\left[\frac{1}{\gamma}\Sigma^{-1}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s) - x_0\right] =\\
\min_{\|s\|_\infty \le \kappa} \quad & \frac{(1-\rho)-(1-\rho)^{T+1}}{2\rho\gamma}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s)^\top + \ldots\\
& \Sigma^{-1}(\mu - \frac{\rho}{(1-\rho)-(1-\rho)^{T+1}}s) + s^\top x_0. && \text{(A.27)}
\end{aligned}$$

Note that from (A.25), we have $s = \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}(\mu - \gamma\Sigma x_1) = \frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}[\gamma\Sigma(x^* - x_1)]$ as well as $\|s\|_\infty \le \kappa$, where $x^* = \frac{1}{\gamma}\Sigma^{-1}\mu$ is the Markowitz portfolio. Replacing s in (A.27), we conclude that problem (A.27) is equivalent to the following

$$\begin{aligned}
\min_{x_1} \quad & \frac{\gamma}{2}x_1^\top \Sigma x_1 - \gamma x_1^\top \Sigma x_0 && \text{(A.28)}\\
s.t. \quad & \|\frac{(1-\rho)-(1-\rho)^{T+1}}{\rho}\gamma\Sigma(x_1 - x^*)\|_\infty \le \kappa. && \text{(A.29)}
\end{aligned}$$

Be aware of that the term $\frac{\gamma}{2}x_0^\top \Sigma x_0$ is a constant term for objective function (A.28), it can be then rewritten as:

$$\begin{aligned}
& \max_{x_1} \quad \frac{\gamma}{2}x_1^\top \Sigma x_1 - \gamma x_1^\top \Sigma x_0\\
\equiv & \max_{x_1} \quad \frac{\gamma}{2}x_1^\top \Sigma x_1 - \gamma x_1^\top \Sigma x_0 + \frac{\gamma}{2}x_0^\top \Sigma x_0. && \text{(A.30)}
\end{aligned}$$

It follows immediately that objective function (B.3) together with constraint (B.2) is equivalent to

$$\begin{aligned}
\min_{x_1} \quad & (x_1 - x_0)^\top \Sigma (x_1 - x_0), && \text{(A.31)}\\
s.t \quad & \|\Sigma(x_1 - x^*)\|_\infty \le \frac{\kappa}{(1-\rho)\gamma}\frac{\rho}{1-(1-\rho)^T}. && \text{(A.32)}
\end{aligned}$$

Part 3. Note that constraint (2.5) is equivalent to

$$-\frac{\kappa}{(1-\rho)\gamma}\frac{\rho}{1-(1-\rho)^T}e \le \Sigma(x_1 - x^*) \le \frac{\kappa}{(1-\rho)\gamma}\frac{\rho}{1-(1-\rho)^T}e,$$

which is a parallelogram centered at x^*. To show that constraint (2.5) defines a no-trade region, note that when the starting portfolio x_0 satisfies constraint (2.5), $x_1 = x_0$ is a minimizer of objective function (2.4) and is feasible with respect to the constraint. On the other hand, when x_0 is not inside the region defined by (2.5), the optimal solution x_1 must be the point on the boundary of the feasible region that minimizes the objective. We then conclude that constraint (2.5) defines a no-trade region. □

Proof of Proposition 2.4

Differentiating objective function (2.7) with respect to x_t gives

$$(1-\rho)(\mu - \gamma\Sigma x) - \kappa p \Lambda^{1/p} |\Lambda^{1/p}(x - x_0)|^{p-1} \cdot \mathrm{sign}(\Lambda^{1/p}(x - x_0)) = 0, \tag{A.33}$$

where $|a|^{p-1}$ denotes the absolute value to the power of $p-1$ for each component

$$|a|^{p-1} = (|a_1|^{p-1}, |a_2|^{p-1}, \cdots, |a_N|^{p-1}),$$

and $\mathrm{sign}(\Lambda^{1/p}(x - x_0))$ is a vector containing the sign of each component for $\Lambda^{1/p}(x - x_0)$. Given that Λ is symmetric, rearranging (A.33) we have

$$(1-\rho)\Lambda^{-1/p}\gamma\Sigma(x^* - x) = \kappa p |\Lambda^{1/p}(x - x_0)|^{p-1} \cdot sign(\Lambda^{1/p}(x - x_0)). \tag{A.34}$$

Note that $x = x_0$ cannot be the optimal solution unless the initial position x_0 satisfies $x_0 = x^*$. Otherwise, take q-norm on both sides of (A.34):

$$\|\Lambda^{-1/p}\Sigma(x - x^*)\|_q = \frac{\kappa}{(1-\rho)\gamma} p \| |\Lambda^{1/p}(x - x_0)|^{p-1} \cdot \mathrm{sign}(\Lambda^{1/p}(x - x_0))\|_q, \tag{A.35}$$

where q is the value such that $\frac{1}{p} + \frac{1}{q} = 1$. Note that $\| |\Lambda^{1/p}(x - x_0)|^{p-1} \cdot \mathrm{sign}(\Lambda^{1/p}(x - x_0))\|_q = \|x - x_0\|_p^{p-1}$, we conclude that the optimal trading strategy satisfies

$$\frac{\|\Lambda^{-1/p}\Sigma(x - x^*)\|_q}{p\|\Lambda^{1/p}(x - x_0)\|_p^{p-1}} = \frac{\kappa}{(1-\rho)\gamma}. \tag{A.36}$$

□

Proof of Theorem 2.6

Differentiating objective function (2.11) with respect to x_t gives

$$\begin{aligned}(1-\rho)^t(\mu - \gamma\Sigma x_t) - (1-\rho)^{t-1} p\kappa\Lambda^{1/p}|\Lambda^{1/p}(x_t - x_{t-1})|^{p-1} \cdot \mathrm{sign}(\Lambda^{1/p}(x_t - x_{t-1})) \\ + (1-\rho)^t p\kappa\Lambda^{1/p}|\Lambda^{1/p}(x_{t+1} - x_t)|^{p-1} \cdot \mathrm{sign}(\Lambda^{1/p}(x_{t+1} - x_t)) = 0.\end{aligned} \tag{A.37}$$

Specifically, given that Λ is symmetric, the optimality condition for stage T reduces to

$$(1-\rho)\Lambda^{-1/p}(\mu - \gamma\Sigma x_T) = p\kappa|\Lambda^{1/p}(x_T - x_{T-1})|^{p-1} \cdot sign(\Lambda^{1/p}(x_T - x_{T-1})). \tag{A.38}$$

Note that $x_T = x_{T-1}$ cannot be the optimal solution unless $x_{T-1} = x^*$. Otherwise, take q-norm on both sides of (A.38) and rearrange terms,

$$\frac{\|\Lambda^{-1/p}\Sigma(x_T - x^*)\|_q}{p\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}} \leq \frac{\kappa}{(1-\rho)\gamma}. \tag{A.39}$$

Summing up the optimal conditions recursively gives

$$p\kappa|x_t - x_{t-1}|^{p-1} \cdot \mathrm{sign}(x_t - x_{t-1}) = \sum_{s=t}^{T}(1-\rho)^{s-t+1}\gamma\Lambda^{-1/p}\Sigma(x^* - x_s). \tag{A.40}$$

Note that $x_t = x_{t-1}$ cannot be the optimal solution unless $x_{t-1} = x^*$. Otherwise, take q-norm on both sides, it follows straightforwardly that

$$\frac{\|\sum_{s=t}^{T}(1-\rho)^{s-t}\Lambda^{-1/p}\Sigma(x_s - x^*)\|_q}{p\|\Lambda^{1/p}(x_t - x_{t-1})\|_p^{p-1}} = \frac{\kappa}{(1-\rho)\gamma}. \tag{A.41}$$

We conclude that the optimal trading strategy for period t satisfies (A.41) whenever the initial portfolio is not x^*. $\square$

Proof of Proposition 2.7

Part 1. We first define the function $g(x) = (1-\rho)\Lambda^{-1/p}\Sigma(x - x^*)$. For period T it holds that $\|g(x_T)\|_q \leq p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}$. Moreover, for the following last period it holds $\|g(x_{T-1}) + (1-\rho)g(x_T)\|_q \leq p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_{T-1} - x_{T-2})\|_p^{p-1}$. Noting that $\|A+B\|_q \geq \|A\|_q - \|B\|_q$, it follows immediately that

$$\begin{aligned}
p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_{T-1} - x_{T-2})\|_p^{p-1} &\geq \|g(x_{T-1}) + (1-\rho)g(x_T)\|_q \\
&\geq \|g(x_{T-1})\|_q - (1-\rho)\|g(x_T)\|_q \\
&\geq \|g(x_{T-1})\|_q - (1-\rho)p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1},
\end{aligned}$$

where the last inequality holds based on the fact $\|g(x_T)\|_q \leq p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}$. Rearranging terms,

$$\|g(x_{T-1})\|_q \leq p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_{T-1} - x_{T-2})\|_p^{p-1} + (1-\rho)p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1},$$

which implies that

$$\frac{\|g(x_{T-1})\|_q}{p\|\Lambda^{1/p}(x_{T-1} - x_{T-2})\|_p^{p-1}} \leq \frac{\kappa}{\gamma} + (1-\rho)\frac{\kappa}{\gamma}\frac{\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}}{\|\Lambda^{1/p}(x_{T-1} - x_{T-2})\|_p^{p-1}},$$

which gives a wider area than the rebalancing region defined for x_T: $\frac{\|g(x_T)\|_q}{p\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}} \le \frac{\kappa}{\gamma}$.
Similarly, $\|g(x_{T-2}) + (1-\rho)g(x_{T-1}) + (1-\rho)^2 g(x_T)\|_q \le \frac{\kappa}{\gamma} p\|\Lambda^{1/p}(x_{T-2} - x_{T-3})\|_p^{p-1}$, it follows

$$\begin{aligned}
\frac{\kappa}{\gamma} p\|\Lambda^{1/p}(x_{T-2} - x_{T-3})\|_p^{p-1} &\ge \|g(x_{T-2}) + (1-\rho)g(x_{T-1}) + (1-\rho)^2 g(x_T)\|_q \\
&\ge \|g(x_{T-2}) + (1-\rho)\|g(x_{T-1})\|_p - (1-\rho)^2\|g(x_T)\|_q, \\
&\ge \|g(x_{T-2}) + (1-\rho)\|g(x_{T-1})\|_p - (1-\rho)^2 p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1},
\end{aligned}$$

where the last inequality holds because $\|g(x_T)\|_q \le p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}$.
Rearranging terms

$$\|g(x_{T-2}) + (1-\rho)\|g(x_{T-1})\|_p \le \frac{\kappa}{\gamma} p\|\Lambda^{1/p}(x_{T-2} - x_{T-3})\|_p^{p-1} + (1-\rho)^2 p\frac{\kappa}{\gamma}\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1},$$

which implies that,

$$\frac{\|g(x_{T-2}) + (1-\rho)\|g(x_{T-1})\|_p}{p\|x_{T-2} - x_{T-3}\|_p^{p-1}} \le \frac{\kappa}{\gamma} + (1-\rho)^2 \frac{\kappa}{\gamma} \frac{\|\Lambda^{1/p}(x_T - x_{T-1})\|_p^{p-1}}{\|\Lambda^{1/p}(x_{T-2} - x_{T-3})\|_p^{p-1}}.$$

The above inequality defines a region which is wider than the rebalancing region defined by $\frac{\|g(x_{T-1}) + (1-\rho)\|g(x_T)\|_p}{p\|\Lambda^{1/p}(x_{T-1} - x_{T-2})\|_p^{p-1}} \le \frac{\kappa}{\gamma}$ for x_{T-1}.
Recursively, we can deduce the rebalancing region corresponding to each period shrinks along t.

Part 2. Note that the rebalancing region for period t relates with the trading strategies thereafter. Moreover, the condition $x_1 = x_2 = \cdots = x_T = x^*$ satisfies inequality (2.12), we then conclude that the rebalancing region for stage t contains Markowitz strategy x^*.

Part 3. The optimality condition for period T satisfies

$$(1-\rho)(\mu - \gamma\Sigma x_T) - p\kappa\Lambda^{1/p}|\Lambda^{1/p}(x_T - x_{T-1})|^{p-1} \cdot sign(\Lambda^{1/p}(x_T - x_{T-1})) = 0. \tag{A.42}$$

Let ω to be the vector such that $\lim_{T\to\infty} x_T = \omega$. Taking limit on both sides of (A.42)

$$(1-\rho)(\mu - \gamma\Sigma\omega) - p\kappa\Lambda^{1/p}|\Lambda^{1/p}(\omega - \omega)|^{p-1} \cdot sign(\Lambda^{1/p}(\omega - \omega)) = 0. \tag{A.43}$$

Noting that $\lim_{T\to\infty} x_T = \lim_{T\to\infty} x_{T-1} = \omega$, it follows

$$(1-\rho)(\mu - \gamma\Sigma\omega) = 0,$$

which gives that $\omega = \frac{1}{\gamma}\Sigma^{-1}\mu = x^*$. We conclude that the investor will move to Markowitz strategy x^* in the limit case. □

Proof of Theorem 2.9

Part 1. For $t = 1, 2, \cdots, T-1$, differentiating objective function (2.15) with respect to x_t gives

$$(1-\rho)(\mu - \gamma\Sigma x_t) - \kappa(2\Lambda x_t - 2\Lambda x_{t-1}) - \kappa(1-\rho)(2\Lambda x_t - 2\Lambda x_{t+1}) = 0, \tag{A.44}$$

rearranging terms

$$[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2\kappa(1-\rho)\Lambda]\, x_t = (1-\rho)\mu + 2\kappa\Lambda x_{t-1} + 2(1-\rho)\kappa\Lambda x_{t+1}. \tag{A.45}$$

The solution can be written explicitly as following

$$\begin{aligned} x_t =&(1-\rho)\gamma\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2\kappa(1-\rho)\Lambda\right]^{-1}\Sigma x^* \\ &+ 2\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2\kappa(1-\rho)\Lambda\right]^{-1}\Lambda x_{t-1} \\ &+ 2(1-\rho)\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2\kappa(1-\rho)\Lambda\right]^{-1}\Lambda x_{t+1}. \end{aligned} \tag{A.46}$$

Define

$$\begin{aligned} A_1 &= (1-\rho)\gamma\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2(1-\rho)\kappa\Lambda\right]^{-1}\Sigma, \\ A_2 &= 2\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2(1-\rho)\kappa\Lambda\right]^{-1}\Lambda, \\ A_3 &= 2(1-\rho)\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda + 2(1-\rho)\kappa\Lambda\right]^{-1}\Lambda, \end{aligned}$$

where $A_1 + A_2 + A_3 = I$.
For $t = T-1$, the optimality condition is

$$(1-\rho)(\mu - \gamma\Sigma x_T) - \kappa(2\Lambda x_T - 2\Lambda x_{T-1}) = 0, \tag{A.47}$$

the explicit solution is

$$x_T = (1-\rho)\gamma\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda\right]^{-1}\Sigma x^* + 2\kappa\left[(1-\rho)\gamma\Sigma + 2\kappa\Lambda\right]^{-1}\Lambda x_{T-1}. \tag{A.48}$$

Define

$$\begin{aligned} B_1 &= (1-\rho)\gamma[(1-\rho)\gamma\Sigma + 2\kappa\Lambda]^{-1}\Sigma, \\ B_2 &= 2\kappa[(1-\rho)\gamma\Sigma + 2\kappa\Lambda]^{-1}\Lambda, \end{aligned}$$

where $B_1 + B_2 = I$.

Part 2. As $T \to \infty$, taking limit on both sides of (A.48) we then conclude that $\lim_{T\to\infty} x_T = x^*$.
□

Proof of Corollary 2.10

Substituting Λ with $\Lambda = I$ and $\Lambda = \Sigma$ respectively we obtain the optimal trading strategy.

To show that the trading trajectory of the case $\Lambda = \Sigma$ follows a straight line, noting that x_T is a linear combination of x_{T-1} and x^*, which indicates that x_T, x_{T-1} and x^* are on a straight line. On the contrary, if we assume that x_{T-2} is not on the same line, then x_{T-2} cannot be expressed as linear combination of x_T, x_{T-1} and x^*. Recall from equation (2.20) when $t = T - 1$ that

$$x_{T-1} = \alpha_1 x^* + \alpha_2 x_{T-2} + \alpha_3 x_T, \tag{A.49}$$

rearranging terms

$$x_{T-2} = \frac{1}{\alpha_2} x_{T-1} - \frac{\alpha_1}{\alpha_2} x^* - \frac{\alpha_3}{\alpha_2} x_T.$$

Noting that $\frac{1}{\alpha_2} - \frac{\alpha_1}{\alpha_2} - \frac{\alpha_3}{\alpha_2} = 1$, it is contradictory with the assumption that x_{T-2} is not a linear combination of x_T, x_{T-1} and x^*. By this means, we can show recursively that all the policies corresponding to the model when $\Lambda = \Sigma$ lay on the same straight line. $\square$

Appendix B

Appendix to Chapter 3

B.1 Aim Portfolio of Linear Policy

Starting from the last period, we know the aim portfolio for an investor is just Markowitz strategy

$$Markowitz_T = (\gamma\Sigma)^{-1}\mu_T.$$

Given the model dynamics specified in (3.10)-(3.11), the conditional mean for each period is $\mu_t = A_r + B_r d_t$. The aim portfolio for the second last period is

$$\begin{aligned} aim_{T-1} &= zMarkowitz_{T-1} + (1-z)\mathbb{E}_{T-1}(aim_T) \\ &= zMarkowitz_{T-1} + (1-z)\mathbb{E}_{T-1}(Markowitz_T) \\ &= \frac{1}{\gamma}\Sigma^{-1}(A_r + B_r b_d d_{T-1}) + \frac{z}{\gamma}\Sigma^{-1}B_r d_{T-1}(1-b_d). \end{aligned} \tag{B.1}$$

Analogously, the aim portfolio for period t can be derived as follows

$$\begin{aligned} aim_t &= zMarkowitz_t + (1-z)\mathbb{E}_t(aim_{t+1}) \\ &= \frac{1}{\gamma}\Sigma^{-1}(A_r + b_d^t B_r d_t) + \frac{z}{\gamma}\Sigma^{-1}B_r d_t f_t(b_d), \end{aligned} \tag{B.2}$$

where $f_t(b_d)$ is a polynomial of b_d changing along period with and $f_T(b_d) = 0$ and

$$f_t(b_d) = (1-z)b_d f_{t+1}(b_d) + 1 - b_d^{T-t}. \tag{B.3}$$

B.2 Tables

Table B.1: Utilities Depending on Different Parameters

This table shows the realized utilities for an investor with objective (3.12) based on the heuristic policies proposed in Section 3.3.2. The first three columns give the values of b_d, γ and κ. Column 4-7 show the utilities obtained based on the simple policy(S-P), linear policy(L-P), the no-trade region policy(NTR-P) and rolling optimize-and-hold policy(ROH-P). The last column shows the upper bounds to the utilities obtained based on perfect hindsight solution. For each row, the realized utilities and the gaps which are computed by $(U_{bounds} - U_{heuristic})/U_{bounds}$ are given.

Parameters			**Utilities**				**Upper Bounds**
b_d	γ	κ	S-P	L-P	NTR-P	ROH-P	Bounds
0.75	3	0.0020	0.1504 19.29%	0.1786 4.16%	0.1849 0.80%	0.1823 2.20%	0.1864
0.75	3	0.0050	0.1488 18.27%	0.1749 3.89%	0.1804 0.91%	0.1781 2.17%	0.1820
0.75	3	0.0100	0.1460 17.03%	0.1698 3.50%	0.1739 1.14%	0.1722 2.12%	0.1759
0.75	5	0.0020	0.0903 19.29%	0.1072 4.17%	0.1109 0.80%	0.1094 2.20%	0.1118
0.75	5	0.0050	0.0893 18.27%	0.1050 3.89%	0.1082 0.91%	0.1068 2.17%	0.1092
0.75	5	0.0100	0.0876 17.03%	0.1019 3.50%	0.1044 1.14%	0.1033 2.12%	0.1056
0.75	7	0.0020	0.0645 19.29%	0.0766 4.17%	0.0792 0.80%	0.0781 2.20%	0.0799
0.75	7	0.0050	0.0638 18.27%	0.0750 3.89%	0.0773 0.91%	0.0763 2.17%	0.0780
0.75	7	0.0100	0.0626 17.03%	0.0728 3.50%	0.0745 1.14%	0.0738 2.12%	0.0754
0.98	3	0.0020	0.1425 4.83%	0.1441 3.72%	0.1490 0.46%	0.1489 0.47%	0.1497
0.98	3	0.0050	0.1408 4.13%	0.1419 3.41%	0.1461 0.50%	0.1461 0.51%	0.1469
0.98	3	0.0100	0.1381 3.42%	0.1381 3.42%	0.1422 0.53%	0.1422 0.54%	0.1430
0.98	5	0.0020	0.0855 4.83%	0.0865 3.72%	0.0894 0.46%	0.0894 0.47%	0.0898
0.98	5	0.0050	0.0845 4.13%	0.0851 3.41%	0.0877 0.50%	0.0877 0.51%	0.0881
0.98	5	0.0100	0.0829 3.42%	0.0829 3.42%	0.0853 0.53%	0.0853 0.54%	0.0858
0.98	7	0.0020	0.0611 4.83%	0.0618 3.72%	0.0639 0.46%	0.0638 0.47%	0.0642
0.98	7	0.0050	0.0603 4.13%	0.0608 3.41%	0.0626 0.50%	0.0626 0.51%	0.0630
0.98	7	0.0100	0.0592 3.42%	0.0592 3.42%	0.0610 0.53%	0.0610 0.54%	0.0613

Table B.2: Certainty Equivalent Depending on Different Parameters

This table shows the certainty equivalents for an investor with objective (3.6)-(3.8) based on the heuristic policies proposed in Section 3.3.2. The first three columns give the values of b_d, the risk-aversion parameter γ as well as the proportional transaction costs parameter κ. Column 4-7 show the certainty equivalent based on the simple policy(S-P), linear policy(L-P), the no-trade region policy(NTR-P) and rolling optimize-and-hold policy(ROH-P) respectively. The last two columns show the upper bounds to the certainty equivalent obtained based on perfect information relaxation with and without penalty function. For each row, the values of the certainty equivalent and the corresponding gaps respect to the upper bounds with penalty function (in %) are reported. The gaps are computed by $(CE_{bounds} - CE_{heuristic})/CE_{bounds}$. The number of investment periods $T = 7$.

Parameters			**CEQ**				**Upper Bounds**	
b_d	γ	κ	S-P	L-P	NTR-P	ROH-P	Penalty	No-penalty
0.75	3	0.0020	0.0691 11.17%	0.0695 10.70%	0.0716 8.02%	0.0712 8.53%	0.0778	0.0943 17.54%
0.75	3	0.0050	0.0690 10.98%	0.0694 10.55%	0.0712 8.23%	0.0708 8.75%	0.0776	0.0938 17.32%
0.75	3	0.0100	0.0693 10.42%	0.0697 9.92%	0.0705 8.87%	0.0703 9.11%	0.0774	0.0930 16.85%
0.75	5	0.0020	0.1075 7.15%	0.1078 6.89%	0.1097 5.22%	0.1094 5.52%	0.1158	0.1407 17.75%
0.75	5	0.0050	0.1074 7.08%	0.1077 6.85%	0.1094 5.40%	0.1090 5.70%	0.1156	0.1401 17.47%
0.75	5	0.0100	0.1078 6.56%	0.1082 6.27%	0.1090 5.61%	0.1088 5.75%	0.1155	0.1392 17.03%
0.75	7	0.0020	0.1248 5.25%	0.1250 5.10%	0.1266 3.88%	0.1263 4.08%	0.1317	0.1588 17.10%
0.75	7	0.0050	0.1247 5.21%	0.1249 5.08%	0.1262 4.02%	0.1260 4.24%	0.1316	0.1582 16.83%
0.75	7	0.0100	0.1251 4.81%	0.1254 4.59%	0.1261 4.12%	0.1259 4.23%	0.1315	0.1572 16.36%
0.98	3	0.0020	0.0662 14.25%	0.0680 11.96%	0.0690 10.73%	0.0690 10.73%	0.0773	0.0931 17.05%
0.98	3	0.0050	0.0662 14.02%	0.0679 11.86%	0.0685 11.12%	0.0685 11.12%	0.0771	0.0926 16.79%
0.98	3	0.0100	0.0662 13.78%	0.0678 11.70%	0.0679 11.64%	0.0679 11.64%	0.0768	0.0918 16.35%
0.98	5	0.0020	0.1051 8.83%	0.1067 7.44%	0.1076 6.69%	0.1077 6.68%	0.1153	0.1393 17.29%
0.98	5	0.0050	0.1051 8.68%	0.1066 7.38%	0.1071 6.93%	0.1071 6.93%	0.1151	0.1386 16.99%
0.98	5	0.0100	0.1051 8.57%	0.1066 7.32%	0.1066 7.28%	0.1066 7.28%	0.1150	0.1377 16.47%
0.98	7	0.0020	0.1229 6.41%	0.1242 5.41%	0.1249 4.86%	0.1249 4.86%	0.1314	0.1577 16.78%
0.98	7	0.0050	0.1229 6.31%	0.1241 5.37%	0.1246 5.04%	0.1246 5.04%	0.1312	0.1570 16.46%
0.98	7	0.0100	0.1229 6.22%	0.1241 5.33%	0.1241 5.29%	0.1241 5.29%	0.1311	0.1559 15.95%

B.3 Derivation of Penalty Function

A brief review about duality based on information relaxation has been made in Haugh and Wang [2014]. This approach has been explained in detail in Brown et al. [2010] as well as in Brown and Smith [2014]. The following we are going to derive the expression for penalty function based on approximate reward functions (3.44). Given $\tilde{y}_t^*$ to be the optimal solution to the problem without transaction costs, for each period the derivative for reward function is

$$\nabla \tilde{I}_t(\tilde{y}_t^*) = (\frac{\partial \tilde{I}_t}{\partial c_1}, \frac{\partial \tilde{I}_t}{\partial x_1}, \cdots, \frac{\partial \tilde{I}_t}{\partial c_t}, \frac{\partial \tilde{I}_t}{\partial x_t})|_{y=\tilde{y}_t^*}, \tag{B.4}$$

where

$$\begin{aligned}
\frac{\partial \tilde{I}_t}{\partial c_j} &= -\rho^t c_t^{1-\gamma} W_t^{1-\gamma} \frac{R_{p,j+1}}{W_{j+1}}, \text{for } j = 1, \cdots, t-1 \\
\frac{\partial \tilde{I}_t}{\partial x_j} &= \rho^t c_t^{1-\gamma} W_t^{1-\gamma} \frac{(1-c_j) R_{j+1}^e}{W_{j+1}}, \text{for } j = 1, \cdots, t-1 \\
\frac{\partial \tilde{I}_t}{\partial c_t} &= \rho^t c_t^{-\gamma} W_{t-1}^{-\gamma} W_t \\
\frac{\partial \tilde{I}_t}{\partial x_t} &= \mathbf{0}
\end{aligned} \tag{B.5}$$

with R_t^e the risky asset returns excess of risk-free asset. When we solve problem (3.48) for each sample path, $\tilde{y}_t^*$ is computed through dynamic programming by discretizing the state variable space. Taking into account the fact that $\tilde{y}_t^*$ does not involve transaction costs, it can be solved in short time.

B.4 Approximate Consumption for the Model with Transaction Costs

When we adapt the heuristic policies to power utility framework, we let the consumption for each period to be the one for the model without transaction costs. We discretize first the state variable space in order to implement dynamic programming procedure.

The VAR model discretized using a variation of Tauchen and Hussey [1991] Gaussian quadrature method, which has been described in Balduzzi and Lynch [1999]. The variation is designed to ensure that d is the only state variable to predict the risky assets returns. Specifically, let η be a $N \times 1$ vector of coefficients of the regression model such

$$e_{t+1} = \eta \epsilon_{t+1} + u_{t+1}, \tag{B.6}$$

with u_{t+1} $N \times 1$ i.i.d. normally distributed vector which is uncorrelated with ϵ_{t+1}.

Following Lynch [2001] and Lynch and Tan [2010], the discretization is implemented so as to match both the conditional mean vector and the covariance matrix for log returns at all grid points of the predictive variable as well as the unconditional volatilities of the predictive variables. We choose 19 grid points for the dividend yield and 3 grid points for each of the stock-return innovations since Balduzzi and Lynch [1999] find that the resulting approximation is able to capture important dimensions of the return predictability in the data.

In Table B.3, we reproduce the quadrature method based on the estimated VAR model given in Lynch [2001]. The information indicated under the table 'DATA' is the estimation from the sample, while the values under the table 'QUADRATURE' are the results obtained from quadrature approximations.

Besides, the Bellman equation corresponding to problem (3.4) with utility function (3.5), without transaction costs is

$$\frac{V_t(d_t)}{1-\gamma} = \max_{c_t, x_t} \left[\frac{c_t^{1-\gamma}}{1-\gamma} + \rho \frac{(1-c_t)^{1-\gamma}}{1-\gamma} \mathbb{E}\Big[\big(x_t^\top R_{t+1} + (1 - x_t^\top e) R_f \big)^{1-\gamma} V_{t+1}(d_{t+1}) \Big] \right]. \tag{B.7}$$

Notice the above equation does not depend on previous stage position. Starting from the last period, it costs no effort to solve for the optimal consumption of each period at each node point by backward iteration. In numerical experiments, the corresponding consumption which is not on the grid can be computed by linear interpolation.

Table B.3: Sample Statistics, VAR coefficients and Quadrature Approximation: High and Low Book-to-Market Portfolios

This table reports the moments and parameters for the high and low book-to-market portfolios adapted from Lynch and Tan (2010) and calculated for the quadrature approximation based on the VAR model that uses log dividend yield as the only state variable. High and low book-to-market portfolios are denoted as High BM and Low BM respectively. Panel A reports unconditional means, VAR slopes and the corresponding coefficients of determination for the data and the quadrature approximation. Panel B reported the unconditional covariance matrix for the data and for the quadrature approximation. All results are for continuously compounded returns. Returns are expressed per month and in percentages. The coefficients of determination are expressed in percentages.

Panel A: Unconditional Sample Moments and VAR coefficients

DATA

Asset/Variable	**Uncond. Mean**	**Slope**	R^2
High BM	0.80	0.49	0.45
Low BM	0.53	0.39	0.38
Dividend yield	0.00	0.98	96.03

QUADRATURE

Asset/Variable	**Uncond. Mean**	**Slope**	R^2
High BM	0.83	0.47	0.40
Low BM	0.51	0.39	0.37
Dividend yield	0.00	0.98	95.98

Panel B: Unconditional Standard Deviations, Covariance(above diagonal) and Correlations(below)

DATA

Asset/Variable	**High BM**	**Low BM**	**Div.Yield**
High BM	7.33	41.70	-0.82
Low BM	0.89	6.41	-0.70
Dividend yield	-0.11	-0.11	1.00

QUADRATURE

Asset/Variable	**High BM**	**Low BM**	**Div.Yield**
High BM	7.37	42.12	-0.83
Low BM	0.88	6.44	-0.70
Dividend yield	-0.11	-0.11	1.00

Panel C: Unconditional Standard Deviations, Covariance and Correlations for Residuals

DATA

Asset/Variable	**High BM**	**Low BM**	**Div.Yield**
High BM	7.32	41.52	-1.27
Low BM	0.89	6.40	-1.08
Dividend yield	-0.88	-0.85	0.20

QUADRATURE

Asset/Variable	**High BM**	**Low BM**	**Div.Yield**
High BM	7.35	41.94	-1.29
Low BM	0.89	6.43	-1.09
Dividend yield	-0.88	-0.84	0.20

Bibliography

Akian, M., J. L. Menaldi, and A. Sulem (1996). On an investment-consumption model with transaction costs. *SIAM Journal on Control and Optimization 34*(1), 329–364.

Almgren, R., C. Thum, E. Hauptmann, and H. Li (2005). Direct estimation of equity market impact. *Risk 18*(7), 58–62.

Balduzzi, P. and A. W. Lynch (1999). Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics 52*(1), 47–78.

Basak, S. and G. Chabakauri (2010). Dynamic mean-variance asset allocation. *Review of Financial Studies 23*(8), 2970–3016.

Bertsimas, D. and A. W. Lo (1998). Optimal control of execution costs. *Journal of Financial Markets 1*(1), 1–50.

Brown, D. B. and J. E. Smith (2011). Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science 57*(10), 1752–1770.

Brown, D. B. and J. E. Smith (2014). Information relaxations, duality, and convex stochastic dynamic programs. *Operations Research 62*(6), 1394–1415.

Brown, D. B., J. E. Smith, and P. Sun (2010). Information relaxations and duality in stochastic dynamic programs. *Operations research 58*(4-part-1), 785–801.

Campbell, J. Y. and L. M. Viceira (1999). Consumption and portfolio decisions when expected returns are time varying. *The Quarterly Journal of Economics 114*(2), 433–495.

Campbell, J. Y. and L. M. Viceira (2003). *Strategic asset allocation: portfolio choice for long-term investors.* Oxford University Press.

Constantinides, G. M. (1979). Multiperiod consumption and investment behavior with convex transactions costs. *Management Science 25*(11), 1127–1137.

Constantinides, G. M. (1986). Capital market equilibrium with transaction costs. *The Journal of Political Economy 94*(4), 842–862.

Davis, M. H. and A. R. Norman (1990). Portfolio selection with transaction costs. *Mathematics of Operations Research 15*(4), 676–713.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *Review of Financial Studies 22*(5), 1915–1953.

DeMiguel, V., X. Mei, and F. J. Nogales (2014). Multiperiod portfolio optimization with many risky assets and general transaction costs. *Available at SSRN 2295345*.

DeMiguel, V. and R. Uppal (2005). Portfolio investment with the exact tax basis via nonlinear programming. *Management Science 51*(2), 277–290.

Dumas, B. and E. Luciano (1991). An exact solution to a dynamic portfolio choice problem under transactions costs. *The Journal of Finance 46*(2), 577–595.

Dybvig, P. H. (2005). Mean-variance portfolio rebalancing with transaction costs. *Manuscript Washington University, St. Louis*.

Engle, R. F., R. Ferstenberg, and J. R. Russell (2012). Measuring and modeling execution cost and risk. *The Journal of Portfolio Management 38*(2), 14–28.

Fama, E. F. and K. R. French (1989). Business conditions and expected returns on stocks and bonds. *Journal of financial economics 25*(1), 23–49.

Gârleanu, N. and L. H. Pedersen (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance 68*(6), 2309–2340.

Grinold, R. C. and R. N. Kahn (2000). *Active portfolio management* (2 ed.). McGraw-Hill New York.

Haugh, M. and C. Wang (2014). Dynamic portfolio execution and information relaxations. *SIAM Journal on Financial Mathematics 5*(1), 316–359.

Ingersoll, J. E. (1987). *Theory of financial decision making*, Volume 3. Rowman & Littlefield.

Janeček, K. and S. E. Shreve (2004). Asymptotic analysis for optimal investment and consumption with transaction costs. *Finance and Stochastics 8*(2), 181–206.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society 53*(6), 1315–1335.

Leland, H. E. (2000). Optimal portfolio implementation with transactions costs and capital gains taxes. *Haas School of Business Technical Report*.

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

Lynch, A. W. (2001). Portfolio choice and equity characteristics: Characterizing the hedging demands induced by return predictability. *Journal of Financial Economics 62*(1), 67–130.

Lynch, A. W. and S. Tan (2010). Multiple risky assets, transaction costs, and return predictability: Allocation rules and implications for us investors. *Journal of Financial and Quantitative Analysis 45*(4), 1015.

Magill, M. J. and G. M. Constantinides (1976). Portfolio selection with transactions costs. *Journal of Economic Theory 13*(2), 245–263.

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance 7*(1), 77–91.

Markowitz, H. (1959). Portfolio selection: Efficient diversification of investments.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics 51*(3), 247–257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory 3*(4), 373–413.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, 867–887.

Mossin, J. (1968). Optimal multiperiod portfolio policies. *The Journal of Business 41*(2), 215–229.

Muthuraman, K. (2007). A computational scheme for optimal investment–consumption with proportional transaction costs. *Journal of Economic Dynamics and Control 31*(4), 1132–1159.

Muthuraman, K. and S. Kumar (2006). Multidimensional portfolio optimization with proportional transaction costs. *Mathematical Finance 16*(2), 301–335.

Muthuraman, K. and H. Zha (2008). Simulation-based portfolio optimization for large portfolios with transaction costs. *Mathematical Finance 18*(1), 115–134.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *The Review of Economics and Statistics 51*(3), 239–246.

Tauchen, G. and R. Hussey (1991). Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica: Journal of the Econometric Society*, 371–396.

Torre, N. G. and M. Ferrari (1997). Market impact model handbook. *BARRA Inc, Berkeley*.