



UC3M Working Papers
Statistics and Econometrics
15-21
ISSN 2387-0303
November 2015

Departamento de Estadística
Universidad Carlos III de Madrid
Calle Madrid, 126
28903 Getafe (Spain)
Fax (34) 91 624-98-48

Portfolio Selection with Proportional Transaction Costs and Predictability

Xiaoling Mei^a, Francisco J. Nogales^a

Abstract

We consider the portfolio optimization problem for a multiperiod investor who seeks to maximize her utility of consumption facing multiple risky assets and proportional transaction costs in the presence of return predictability. Due to the curse of dimensionality, this problem is very difficult to solve even numerically. In this paper, we propose several feasible policies that are based on optimizing quadratic programs. These proposed feasible policies can be easily computed even for many risky assets. We show how to compute upper bounds and use them to study how the losses associated with using the approximate policies depend on different problem parameters.

Keywords: *Portfolio Optimization, Dynamic Portfolio Choice, Information Relaxations*

^a Department of Statistics, Universidad Carlos III de Madrid.

Acknowledgements: the authors acknowledge financial support from the Spanish Government Project MTM2013-44902

Portfolio Selection with Proportional Transaction Costs and Predictability

Xiaoling Mei

Javier Nogales

Abstract

We consider the portfolio optimization problem for a multiperiod investor who seeks to maximize her utility of consumption facing multiple risky assets and proportional transaction costs in the presence of return predictability. Due to the curse of dimensionality, this problem is very difficult to solve even numerically. In this paper, we propose several feasible policies that are based on optimizing quadratic programs. These proposed feasible policies can be easily computed even for many risky assets. We show how to compute upper bounds and use them to study how the losses associated with using the approximate policies depend on different problem parameters.

Keywords: portfolio optimization; dynamic portfolio choice; information relaxations.

1 Introduction

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 on each of the risky assets. However, continuously portfolio rebalancing requires the payment of 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 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 no-trade interval, and provide a numerical method to calculate the optimal policy. Muthuraman (2007) develops an efficient computational scheme for the same problem.

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 risky-asset returns uncorrelated. They also present some numerical results for the two uncorrelated risky assets case. Liu (2004) considers a constant absolute risk aversion (CARA) investor who has access to unconstrained borrowing and faces uncorrelated risky asset returns. He shows analytically 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.

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

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) investigate numerically 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 and then 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 it 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 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 paper, 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 the tractable mean-variance framework by G&P 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 that 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 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 these dual bounds can improve significantly 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 that maximizes her power utility of intermediate consumption while Brown and Smith (2011) consider an investor maximizes her utility of terminal wealth. DeMiguel et al. (2014) consider a mean-variance investor who faces general transaction costs and constant opportunity set. For the case with proportional transaction costs, they give closed-form expressions for the no-trade region. In this paper, 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 paper is organized as follows. In Section 2, we introduce the dynamic portfolio selection problem in the presence of proportional transaction costs and predictability. In Section 3, we describe our approximate trading policies for a mean-variance investor and evaluate these approximate strategies under the mean-variance framework. Section 4 describes how to adapt these approximate strategies to a CRRA power utility framework and evaluates numerically these strategies. The evaluation is based on information relaxations that allow to obtain dual bounds. Section 5 concludes. Appendix A contains the tables. Appendix B contains the derivation of the aim portfolio of linear policy while Appendix C contains the derivation of the penalty function. In Appendix D we describe how to approximate the consumption for each period for the model with transaction costs.

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 $t = 1, \dots, T$ 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 $R_t = [R_{t,1}, \dots, R_{t,N}]$, 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 $x_t = [x_{t,1}, \dots, x_{t,N}]$, where $x_{t,i}$ is the weight of the i th asset hold in period t . Throughout this manuscript, we will use \mathbf{x}_t to denote the $N \times t$ vector of decision variable, $[x_1, \dots, x_t]$ and we always use \mathbf{x} to denote $[x_1, \dots, x_T]$.

Trading costs are imposed in the problem, and we assume short selling is not allowed for risky assets. In this paper, we will focus on the special case where 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

$$x_{t,+} = \frac{x_t \cdot R_{t+1}}{R_{p,t}}, \quad (1)$$

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

$$R_{p,t+1} = x_t^\top R_{t+1} + [1 - x_t^\top e - \|K(x_t - x_{t-1,+})\|_1] R_f, \quad (2)$$

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

$$W_{t+1} = W_t(1 - c_t)R_{p,t}, \quad (3)$$

where W_t is 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:

$$\max_{\{c_t, x_t\}_{t=1}^T} \mathbb{E}_1 \left[\sum_{t=1}^T \rho^t U_t(C_t) \right], \quad (4)$$

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

$$U_t = \frac{C_t^{1-\gamma} - 1}{1 - \gamma}, \quad (5)$$

with relative risk aversion parameter $\gamma \geq 1$. In (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 $P(u) \subseteq P_1 \times \dots \times P_T$ denote the set of all feasible action sequences $\mathbf{x} = [x_1, \dots, x_T]$ for any given scenario of state variable u . A *feasible* policy is the one where each individual action x_t depends on previous actions $[x_1, \dots, x_{t-1}]$ 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:

$$\max_{(\mathbf{c}, \mathbf{x}) \in \mathcal{P}_{\mathcal{F}}} \mathbb{E}_1 \left[\sum_{t=1}^T \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} \right] \quad (6)$$

$$\text{s.t. } W_{t+1} = W_t(1 - c_t)R_{p,t}, \quad (7)$$

$$x_t \geq 0, c_t \geq 0. \quad (8)$$

Several comments are in order. 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 power utility function, the wealth level W_t can be factored out for each period and let \mathbf{S}_t denote the vector of state variables, the Bellman equation for each period faced by the investor is,

$$\frac{\phi_t(\mathbf{S}_t, x_{t-1,+})}{1-\gamma} = \max_{(c_t, x_t) \in P_t(\mathbf{x}_{t-1})} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} + \frac{(1-c_t)^{1-\gamma}}{1-\gamma} \mathbb{E} \left[\phi_{t+1}(\mathbf{S}_{t+1}, x_{t,+})^{1-\gamma} R_{p,t+1}^{1-\gamma} \right] \right\}, \quad (9)$$

for $t = 1, \dots, T-1$. Note the above equation is solved by backward iteration, starting with $t = T-1$ and $\phi_T = 1$. In Section 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 (9) involves the expectation of ϕ_{t+1} , while the dimension of the state space leads to the problem of curse of dimensionality especially when there are more than two risky assets. In Section 3.1, we introduce the G&P framework that allows us to deal with many risky assets.

2.1 Predictability Model

Like 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 log of 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 log 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):

$$r_{t+1} = A_r + B_r d_t + e_{t+1}, \quad (10)$$

$$d_{t+1} = a_d + b_d d_t + \epsilon_{t+1}. \quad (11)$$

Here, A_r is $N \times 1$ 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 Σ_{ee} . Without loss of generality, the mean of $\{d\}_{t=1}^T$ can be normalized to 0 and the variance to 1.

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 (6). To avoid solving the dynamic program in (9) which may result in 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.

3.1 Mean-variance Framework

As discussed in Section 2, problem (6)-(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 facilitate the investor to solve the dynamic program more efficiently based on an approximation in continuation value function. But still, it requires much time 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 constant opportunity set, DeMiguel et al. (2014) study analytically the properties of optimal trading strategies and provide 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

$$\max_{\mathbf{x} \in \mathcal{P}_{\mathcal{F}}} \mathbb{E}_1 \left\{ \sum_{t=1}^T \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 \right] \right\}, \quad (12)$$

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, a closed-form solution still cannot be obtained. 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 current period. To do that, note the value function for the last period is:

$$V_T(x_{T-1}, \mu_T) = \rho^T (x_T^\top \mu_T - \frac{\gamma}{2} x_T^\top \Sigma x_T) - \rho^{T-1} \|K(x_T - x_{T-1})\|_1, \quad (13)$$

and from (13), we can define the value functions for previous periods recursively using the Bellman equation

$$V_t(x_{t-1}, \mu_t) = \max_{x_t \in P_t(\mathbf{x}_{t-1})} \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 + \mathbb{E}_t [V_{t+1}(x_t, \mu_{t+1})] \right], \quad (14)$$

for $t = 1, \dots, T-1$. 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, and later, in Section 4 we are going to adapt them to the CRRA utility framework.

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 (12). As in the G&P framework, we assume the dynamic of price changes follows the model specified in (10)-(11).

3.2.1 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:

$$\max_{\{x_t\}_{t=1}^T} \sum_{t=1}^T \left[\rho^t (x_t^\top \tilde{\mu}_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 \right], \quad (15)$$

where $\tilde{\mu}_t$ is the expectation of price changes $\tilde{\mu}$ conditional at the initial stage, that is $\tilde{\mu}_t = E_1(\mu_t)$. The corresponding *simple policy* is defined as the solution to the deterministic portfolio problem (15). It is intuitive that the simple policy will perform well in practice when the volatility in the predictability model is small.

3.2.2 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 target portfolio is a linear combination of the current optimal portfolio in the absence of transaction costs and the expected future target portfolios. With a finite horizon and a constant opportunity set, DeMiguel et al. (2014) show 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 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

$$\text{aim}_t = z \text{Markowitz}_t + (1 - z) E_t(\text{aim}_{t+1}), \quad (16)$$

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,

$$\text{aim}_T = (\gamma\Sigma)^{-1} \mu_T. \quad (17)$$

We define the *linear policy* as

$$x_t = \left(1 - \frac{c}{\delta}\right)x_{t-1} + \frac{c}{\delta}\text{aim}_t. \quad (18)$$

As in Gârleanu and Pedersen (2013), δ denotes the quadratic transaction cost parameter. Since we use (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 (10)-(11), the aim portfolio for each period is

$$\text{aim}_t = (\gamma\Sigma)^{-1}(A_r + b_d^{T-t}B_r d_t) + z(\gamma\Sigma)^{-1}B_r d_t f_t(b_d), \quad (19)$$

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

3.2.3 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 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.

Following 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 no-trade region centred 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

$$\min_{x_t} (x_t - x_{t-1})^\top \Sigma (x_t - x_{t-1}) \quad (20)$$

$$\text{s.t.} \quad \|K^{-1}\Sigma(x_t - \text{aim}_t)\|_\infty \leq \frac{1}{\rho\gamma} \frac{1 - \rho}{1 - \rho^{T-t+1}}, \quad (21)$$

where aim_t denotes the aim portfolio for each period which is defined in (16), x_{t-1} is the position from previous period. Note the aim portfolio aim_t is specified in (19) as the dynamic of price changes follows (10)-(11).

3.2.4 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 3.2.3, we assume the investor has an aim portfolio for each period and she trades to the boundary of 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

$$V_T^*(x_{T-1}, \mu_T) = \rho^T (x_{T-1}^\top \mu_T - \frac{\gamma}{2} x_{T-1}^\top \Sigma x_{T-1}). \quad (22)$$

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

$$\begin{aligned} V_{T-1}(x_{T-2}, \mu_{T-1}) &= \max_{x_{T-1}} \rho^{T-1} (x_{T-1}^\top \mu_{T-1} - \frac{\gamma}{2} x_{T-1}^\top \Sigma x_{T-1}) - \rho^{T-2} \|K(x_{T-1} - x_{T-2})\|_1 \\ &\quad + \mathbb{E}_{T-1}[V_T^*(\mu_T, x_{T-1})] \\ &\equiv \max_{x_{T-1}} \rho^{T-1} (x_{T-1}^\top [\mu_{T-1} + \rho \mu_{T \rightarrow T-1}] - \frac{\gamma(1+\rho)}{2} x_{T-1}^\top \Sigma x_{T-1}) \\ &\quad - \rho^{T-2} \|K(x_{T-1} - x_{T-2})\|_1. \end{aligned} \quad (23)$$

In DeMiguel et al. (2014), they show problem (23) is equivalent to the following constrained optimization problem

$$\min_{x_{T-1}} (x_{T-1} - x_{T-2})^\top \Sigma (x_{T-1} - x_{T-2}) \quad (24)$$

$$\text{s.t. } \|K^{-1} \Sigma (x_{T-1} - x_{T-1}^C)\|_\infty \leq \frac{1}{\rho \gamma_{T-1}}, \quad (25)$$

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

$$\min_{x_t} (x_t - x_{t-1})^\top \Sigma (x_t - x_{t-1}) \quad (26)$$

$$\text{s.t. } \|K^{-1} \Sigma (x_t - x_t^C)\|_\infty \leq \frac{1}{\rho \gamma_t}, \quad (27)$$

where $x_t^C = (\gamma_t \Sigma)^{-1} \mu_t'$ with

$$\gamma_t = (1 + \rho + \dots + \rho^{T-t})\gamma = \frac{1 - \rho^{T-t+1}}{1 - \rho}\gamma, \quad (28)$$

$$\mu_t' = \mu_t + \rho\mu_{t+1|t} + \dots + \rho^{T-t}\mu_{T|t}. \quad (29)$$

In (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 (26)-(27) and the investor's transaction is confined by the no-trade region defined in (27). Also note that the center of the no-trade region x_t^C is a linear combination of all future optimal portfolios conditional at current period in the absence of transaction costs.

3.3 Evaluation

In this section, we study numerically the performance of the proposed feasible policies. For each policy, we compare the realized utility with the upper bounds obtained from *perfect hindsight solution*. Specifically, given the convexity of the objective function in (12), it is straightforward that

$$\begin{aligned} & \max_{\mathbf{x} \in \mathcal{P}_{\mathcal{F}}} \mathbb{E}_1 \left\{ \sum_{t=1}^T \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 \right] \right\} \\ & \leq \mathbb{E}_1 \left\{ \max_{\mathbf{x}} \sum_{t=1}^T \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 \right] \right\}. \end{aligned} \quad (30)$$

In deriving the upper bounds for problem (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 (30) using simulation. In each trial of the simulation, we solve the following deterministic problem

$$\max_{\mathbf{x}} \sum_{t=1}^T \left[\rho^t (x_t^\top \mu_t - \frac{\gamma}{2} x_t^\top \Sigma x_t) - \rho^{t-1} \|K(x_t - x_{t-1})\|_1 \right], \quad (31)$$

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 T periods investment horizon, given any value of dividend yield d_t , n_d different values of dividend yields for 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 (10)-(11) and that the initial dividend yield is neutral (i.e., $d_1 = 0$). Further assume that the investor has an initial wealth of 1 dollar. With these assumptions, the absolute risk-aversion parameter under the mean-variance framework described in Section 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 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 6 value-weighted portfolios SL,SM,SH,BL,BM and BH³. The parameters for the predictability model in (10)-(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 $b_d = 0.75$.

3.3.1 The Base Case

For our base case, we consider a time horizon of 6 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 holding zero share 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.

²Our 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 scenario tree to capture the return dynamics.

³The 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 3 deciles for all stocks; analogously, M indicates the middle 4 deciles and H indicates the top 3 deciles. The high BM portfolio is an equal-weighted portfolio of SH and BH.

For the *simple policy*, the expectation $\tilde{\mu}_t$ is $\tilde{\mu}_t = E_1(\mu_t) = A_r^Y + B_r^Y b_d^{t-1} d_1$. 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 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 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 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 (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 (30)) is 3.89%. To understand this result, it is important to note that this policy combines the previous stage position with 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 no-trade region and the existence of no-trade region leads to 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 no-trade region policy. This is because compared with the target portfolio in linear policy, the center of no-trade region may deviate from the true target for each period with the assumption that trade only occurs in next period.


Table 1: Realized Utilities - The Base Case

This table shows the realized utilities for an investor with objective (12) based on the approximate policies proposed in Section 3.2. Column 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 perfect perfect hindsight solution. The values of the realized utilities are reported in the first row and the second row shows the corresponding gaps 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 1: Utility of Linear Policy depending on δ

This figure plots the realized utilities for 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 $T = 8$.



3.3.2 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), no-trade region policy and 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 2 shows how the utility loss of proposed approximate policies depends on different values of 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 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 the utility loss associated with simple policy decreases as κ increases. This occurs because higher κ leads to less trade in risky assets, 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 it has been explained in the base case, linear policy indicates 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. Slow trading rate in linear policy can compensate the trade amount in subsequent periods. When there is low volatility in 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 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 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 highest loss no more than 1.5%.

When there is high volatility in dividend yield, the utility loss associated with 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 $b_d = 0.98$) while simple policy results in higher utility loss compared with the other policies, followed by linear policy. Based on the above robustness check, the benefits of the proposed approximation can be summarized as follows:

1. For simple policy, if there is low volatility in dividend yield (hence the mean price changes conditional at first period can predict the real conditional price changes at each period very well), it can perform well.
2. For linear policy, it has lower utility loss than the simple policy but higher loss than no-trade region policy and 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 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 cost and predictability. The gap between its realized utility and upper bound is relatively small for all the cases.
4. For rolling optimize-and-hold policy, just like 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 cost and predictability. The gap between its realized utility and upper bound is relatively small for all the cases.

4 Moving to 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 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 (6)-(8)). In Section 4.1, we obtain several approximate solutions for portfolio optimization problem (6)-(8) by adapting the feasible policies proposed based on G&P framework. We check the robustness of these approximate policies by evaluating the certainty equivalent losses in Section 4.3.

4.1 Adapting Mean-variance Framework to CRRA Framework

Under G&P framework, the proposed feasible policies in Section 3.2 provide 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 risky asset return dynamics in (10)-

(11), the optimal solution to model (6)-(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 (12) under mean-variance framework equals the relative risk aversion parameter for in (6)-(8) under power utility framework. Consequently, for each period, the amount of money invested in risk-free asset is

$$W_t^f = W_t(1 - \hat{c}_t) - \hat{W}_t(1 - \hat{c}_t)\hat{x}_t^\top P_t - \hat{W}_t(1 - \hat{c}_t)\|K(\hat{x}_t - \hat{x}_{t-1,+})\|_1, \quad (32)$$

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 holding across the risky assets and risk-free asset, i.e.,

$$\hat{W}_{t+1} = W_t^f R_f + \hat{W}_t(1 - \hat{c}_t)\hat{x}_t^\top P_{t+1}. \quad (33)$$

For each approximate policy, the realized power utility is

$$U^h = \mathbb{E}_1 \left[\sum_{t=1}^T \rho^t \frac{\hat{c}_t^{1-\gamma} \hat{W}_t^{1-\gamma}}{1-\gamma} \right]. \quad (34)$$

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 (11) and save the realizations for the values of ϵ_2 . The first period risky asset returns R_2 can be generated based on (10)⁴. For each subsequent time period at each node, we repeat the process until we arrive the final period $t = T$.

Assume 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

$$U_{(m)}^h = \sum_{t=1}^T \rho^t \frac{\hat{c}_{t,(m)}^{1-\gamma} \hat{W}_{t,(m)}^{1-\gamma}}{1-\gamma}, \quad (35)$$

⁴Following 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 (10)-(11).

for $m = 1, 2, \dots, M$. 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 (34) can be obtained by averaging the values of $U_{(m)}^h$. Let \bar{U}_h be

$$\bar{U}_h = \frac{\sum_{m=1}^M U_{(m)}^h}{M}. \quad (36)$$

The certainty equivalent for each feasible policy is defined as

$$CE(\bar{U}_h) = ((1 - \gamma)\bar{U}_h)^{\frac{1}{1-\gamma}}. \quad (37)$$

Given that each feasible policy provides an approximation to the optimal solution for the portfolio optimization problem with 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 solve numerically the original model, we complement these proposed feasible policies with upper bounds on the certainty equivalent based on the dual approach which will be explained in the following section.

4.2 Upper Bounds

Given any feasible sub-optimal policy to the portfolio selection problem (6)-(8), we can obtain an unbiased lower bound to the optimal power utility U^* by simulating different sample paths and taking 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 can obtain a valid upper bound to the optimal power utility.

In Brown et al. (2010), they show that the valid upper bound can be constructed by relaxing future information and meanwhile 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 meanwhile imposing penalty function that punishes the violations of nonanticipativity. For each realized scenario, they show that an ideal penalty function $\pi(\mathbf{x})$ follows

$$\pi^*(\mathbf{x}) = \sum_{t=1}^T \left\{ \mathbb{E}[V_{t+1}^*(\mathbf{x}_t) | \mathcal{G}_t] - \mathbb{E}[V_{t+1}^*(\mathbf{x}_t) | \mathcal{F}_t] \right\}, \quad (38)$$

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. However, it is not practical to have the real optimal value function for each period. A natural idea is to find an approximation to the optimal penalty function by approximating the optimal value function V_t^* with \hat{V}_t :

$$\hat{\pi}(\mathbf{x}) = \sum_{t=1}^T \left\{ \hat{V}_{t+1}(\mathbf{x}_t) - \mathbb{E} \left[\hat{V}_{t+1}(\mathbf{x}_t) | \mathcal{F}_t \right] \right\}, \quad (39)$$

where in (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 $V_1^*(x_1)$ is given by

$$V_{ub} := \mathbb{E}_1 \left[\max_{\mathbf{x} \in P(u)} \left\{ \sum_{t=1}^T I_t(\mathbf{x}_t) - \hat{\pi}(\mathbf{x}) \right\} \right]. \quad (40)$$

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 (39) in each sample path,

$$\max_{\mathbf{c}, \mathbf{x}} \sum_{t=1}^T \left\{ \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} + \mathbb{E}_t \left[\hat{V}_{t+1}(\mathbf{x}_t) \right] - \hat{V}_{t+1}(\mathbf{x}_t) \right\}, \quad (41)$$

$$\text{s.t. } W_{t+1} = W_t(1 - c_t)R_{p,t}, \quad (42)$$

$$x_t \geq 0, c_t \geq 0. \quad (43)$$

The expected value in (40) is estimated by taking the average of the optimal objective function value of problem (41) over all sample paths.

In (40), each reward function $I_t(\mathbf{x}_t)$ is concave in decision variables \mathbf{x} , and it leads to convexity in the original DP defined by (9). When taking penalty function into account, the objective function (41) may not be concave and consequently, it cannot be solved computationally. In Brown and Smith (2014), they study DPs that have a convex structure and consider penalties based on 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 reward function of the original model,

$$\tilde{I}_t(c_t, \mathbf{x}_t) = \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} \quad (44)$$

subject to the law of motion for investor's wealth which is defined in (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 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

$$\pi(\mathbf{c}, \mathbf{x}) = \sum_{t=1}^T \{ \hat{V}_{t+1}(\mathbf{c}_t, \mathbf{x}_t) - \mathbb{E}_t [\hat{V}_{t+1}(\mathbf{c}_t, \mathbf{x}_t)] \}. \quad (45)$$

Following Brown and Smith (2014), weak duality implies that

$$\max_{(\mathbf{c}, \mathbf{x}) \in \mathcal{P}_{\mathcal{F}}} \mathbb{E}_1 \left[\sum_{t=1}^T \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} \right] \leq \max_{(\mathbf{c}, \mathbf{x}) \in \mathcal{P}_{\mathcal{G}}} \mathbb{E}_1 \left[\sum_{t=1}^T \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} - \pi(\mathbf{c}, \mathbf{x}) \right], \quad (46)$$

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 above inequality can be obtained via Monte Carlo simulation: we generate randomly M different sample paths of risky asset returns and in each sample path, solve the inner problem on the right-hand side of (46).

To guarantee the convex structure of the problem on the right-hand side of (46), we take the first order linear approximation of \hat{V}_t so that the resultant penalty function is linear in decision variables. Let $\tilde{\mathbf{y}}_t^* = (\tilde{c}_t^*, \tilde{x}_t^*)$ denote the optimal policy for the no-transaction costs problem with reward function (44) and \tilde{W}_t^* the corresponding wealth for each period. We define the gradient penalty as follows⁵,

$$\pi(\mathbf{c}, \mathbf{x}) = \sum_{t=1}^T \nabla \tilde{I}_t(\tilde{\mathbf{y}}_t^*)(\mathbf{y}_t - \tilde{\mathbf{y}}_t^*), \quad (47)$$

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

⁵When 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 $\sum_{t=1}^T \hat{V}_{t+1}(\tilde{\mathbf{y}}_t^*) - \mathbb{E}_t [\hat{V}_{t+1}(\tilde{\mathbf{y}}_t^*)]$ 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 (47), which serves as control variate, see Brown and Smith (2014).

the following deterministic problem

$$\max_{\mathbf{c}, \mathbf{x}} \sum_{t=1}^T \left\{ \rho^t \frac{C_t^{1-\gamma}}{1-\gamma} - \nabla \tilde{I}_t(\tilde{\mathbf{y}}_t^*)(\mathbf{y}_t - \tilde{\mathbf{y}}_t^*) \right\}, \quad (48)$$

$$\text{s.t. } W_{t+1} = W_t(1 - c_t)R_{p,t}, \quad (49)$$

$$x_t \geq 0, c_t \geq 0. \quad (50)$$

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

Given the fact that each reward function (44) is differentiable in decision variables, we can derive explicitly the penalty function in (47). Let the certainty equivalent be

$$CE(\bar{U}_p) = ((1 - \gamma)\bar{U}_p)^{\frac{1}{1-\gamma}}. \quad (51)$$

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

4.3 Numerical Results

In this section, we study empirically 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 for annual returns. We also consider a risk-free asset with annually rate of return of 6%.

4.3.1 The Base Case

For our base case, we consider the same parameters that are used in Section 3.3.1. That is, the investor has a relative risk-aversion parameter $\gamma = 5$, has an initial wealth of \$1 invested in 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 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: the no-trade region policy is the best approximate feasible solution to the original model among all the feasible policies, followed by rolling optimize-and-hold policy. Compared with the utility loss under mean-variance framework, the certainty equivalent loss associated with adopting simple policy is small, indicating a flatter utility than that of the mean-variance framework. This shows that, for a proposed feasible policy which has large (small) utility loss under mean-variance framework, it can have small (large) certainty equivalent loss under power utility framework.

4.3.2 Comparative Statics

In this section, we study numerically 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 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. The second part is the loss from following portfolio policies that are derived based on mean-variance framework. In DeMiguel et al. (2014), they show that the certainty equivalent loss from following such 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 3 shows how the certainty equivalent loss associated with adopting proposed approximate policies depends on different values of above mentioned parameters. It shows that there is significant improvement on the bounds provided based on perfect information relaxation with 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 mean-variance framework, it can be understood in the way 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 pro-

cedure 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, 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 meanwhile the portfolio weights recommended by these adapted strategies can be determined in very short time.

5 Conclusion

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 G&P framework which facilitates 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 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 the 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.

References

- 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.
- Balduzzi, P. and A. W. Lynch (1999). Transaction costs and predictability: Some utility cost calculations. *Journal of Financial Economics* 52(1), 47–78.
- 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*, 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., X. Mei, and F. J. Nogales (2014). Multiperiod portfolio optimization with many risky assets and general transaction costs. *Available at SSRN 2295345*.
- Gârleanu, N. and L. H. Pedersen (2013). Dynamic trading with predictable returns and transaction costs. *The Journal of Finance* 68(6), 2309–2340.
- Haugh, M. and C. Wang (2014). Dynamic portfolio execution and information relaxations. *SIAM Journal on Financial Mathematics* 5(1), 316–359.
- 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.
- Magill, M. J. and G. M. Constantinides (1976). Portfolio selection with transactions costs. *Journal of Economic Theory* 13(2), 245–263.

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

Appendices

A Tables

Table 2: Utilities Depending on Different Parameters

This table shows the realized utilities for an investor with objective (12) based on the heuristic policies proposed in Section 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 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) respectively. The last two columns show the upper bounds to the utilities obtained based on perfect information relaxation with and without penalty function. For each row, the values of the realized utilities and the corresponding gaps respect to the upper bounds with penalty function (in %) are reported. The gaps are computed by $(U_{bounds} - U_{heuristic})/U_{bounds}$. The number of investment periods $T = 7$.

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

Table 3: Certainty Equivalent Depending on Different Parameters

This table shows the certainty equivalents for an investor with objective (6)-(8) based on the heuristic policies proposed in Section 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	0.0695	0.0716	0.0712	0.0778	0.0943
			11.17%	10.70%	8.02%	8.53%		17.54%
0.75	3	0.0050	0.0690	0.0694	0.0712	0.0708	0.0776	0.0938
			10.98%	10.55%	8.23%	8.75%		17.32%
0.75	3	0.0100	0.0693	0.0697	0.0705	0.0703	0.0774	0.0930
			10.42%	9.92%	8.87%	9.11%		16.85%
0.75	5	0.0020	0.1075	0.1078	0.1097	0.1094	0.1158	0.1407
			7.15%	6.89%	5.22%	5.52%		17.75%
0.75	5	0.0050	0.1074	0.1077	0.1094	0.1090	0.1156	0.1401
			7.08%	6.85%	5.40%	5.70%		17.47%
0.75	5	0.0100	0.1078	0.1082	0.1090	0.1088	0.1155	0.1392
			6.56%	6.27%	5.61%	5.75%		17.03%
0.75	7	0.0020	0.1248	0.1250	0.1266	0.1263	0.1317	0.1588
			5.25%	5.10%	3.88%	4.08%		17.10%
0.75	7	0.0050	0.1247	0.1249	0.1262	0.1260	0.1316	0.1582
			5.21%	5.08%	4.02%	4.24%		16.83%
0.75	7	0.0100	0.1251	0.1254	0.1261	0.1259	0.1315	0.1572
			4.81%	4.59%	4.12%	4.23%		16.36%
0.98	3	0.0020	0.0662	0.0680	0.0690	0.0690	0.0773	0.0931
			14.25%	11.96%	10.73%	10.73%		17.05%
0.98	3	0.0050	0.0662	0.0679	0.0685	0.0685	0.0771	0.0926
			14.02%	11.86%	11.12%	11.12%		16.79%
0.98	3	0.0100	0.0662	0.0678	0.0679	0.0679	0.0768	0.0918
			13.78%	11.70%	11.64%	11.64%		16.35%
0.98	5	0.0020	0.1051	0.1067	0.1076	0.1077	0.1153	0.1393
			8.83%	7.44%	6.69%	6.68%		17.29%
0.98	5	0.0050	0.1051	0.1066	0.1071	0.1071	0.1151	0.1386
			8.68%	7.38%	6.93%	6.93%		16.99%
0.98	5	0.0100	0.1051	0.1066	0.1066	0.1066	0.1150	0.1377
			8.57%	7.32%	7.28%	7.28%		16.47%
0.98	7	0.0020	0.1229	0.1242	0.1249	0.1249	0.1314	0.1577
			6.41%	5.41%	4.86%	4.86%		16.78%
0.98	7	0.0050	0.1229	0.1241	0.1246	0.1246	0.1312	0.1570
			6.31%	5.37%	5.04%	5.04%		16.46%
0.98	7	0.0100	0.1229	0.1241	0.1241	0.1241	0.1311	0.1559
			6.22%	5.33%	5.29%	5.29%		15.95%

B 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 (10)-(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

$$\begin{aligned} aim_{T-1} &= z Markowitz_{T-1} + (1-z) E_{T-1}(aim_T) \\ &= z Markowitz_{T-1} + (1-z) E_{T-1}(Markowitz_T) \\ &= \frac{1}{\gamma} \Sigma^{-1} (A_r + B_r b_d d_{T-1}) + \frac{z}{\gamma} \Sigma^{-1} B_r d_{T-1} (1 - b_d). \end{aligned} \quad (B1)$$

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

$$\begin{aligned} aim_t &= z Markowitz_t + (1-z) E_t(aim_{t+1}) \\ &= \frac{1}{\gamma} \Sigma^{-1} (A_r + b_d^t B_r d_t) + \frac{z}{\gamma} \Sigma^{-1} B_r d_t f_t(b_d), \end{aligned} \quad (B2)$$

where $f_t(b_d)$ is a polynomial of b_d changing along period with expression

$$f_t(b_d) = (1-z) b_d f_{t+1}(b_d) + 1 - b_d^{T-t}, \quad (B3)$$

and $f_T(b_d) = 0$.

C 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 (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

$$\nabla \tilde{I}_t(\tilde{y}_t^*) = \left(\frac{\partial \tilde{I}_t}{\partial c_1}, \frac{\partial \tilde{I}_t}{\partial x_1}, \dots, \frac{\partial \tilde{I}_t}{\partial c_t}, \frac{\partial \tilde{I}_t}{\partial x_t} \right) \Big|_{y=\tilde{y}_t^*}, \quad (C4)$$

where

$$\begin{aligned}
\frac{\partial \tilde{I}_t}{\partial c_j} &= -\rho^t c_t^{1-\gamma} W_t^{1-\gamma} \frac{R_{p,j+1}}{W_{j+1}}, \text{ for } j = 1, \dots, t-1 \\
\frac{\partial \tilde{I}_t}{\partial x_j} &= \rho^t c_t^{1-\gamma} W_t^{1-\gamma} \frac{(1-c_j)R_{j+1}^e}{W_{j+1}}, \text{ for } j = 1, \dots, t-1 \\
\frac{\partial \tilde{I}_t}{\partial c_t} &= \rho^t c_t^{-\gamma} W_{t-1}^{-\gamma} W_t \\
\frac{\partial \tilde{I}_t}{\partial x_t} &= \mathbf{0}
\end{aligned} \tag{C5}$$

with R_t^e the risky asset returns excess of risk-free asset. When we solve problem (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.

D 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

$$e_{t+1} = \eta \epsilon_{t+1} + u_{t+1}, \tag{D6}$$

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 4, 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 (4) with utility function (5), without transaction costs is

$$\frac{V_t(d_t)}{1-\gamma} = \max_{c_t, x_t} \left[\frac{c_t^{1-\gamma}}{1-\gamma} + \rho \frac{(1-c_t)^{1-\gamma}}{1-\gamma} \mathbf{E} \left[(x_t^\top R_{t+1} + (1-x_t^\top e)R_f)^{1-\gamma} V_{t+1}(d_{t+1}) \right] \right]. \quad (\text{D7})$$

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 4: 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
QUADRATURE

Asset/Variable	Uncond. Mean	Slope	R^2	Asset/Variable	Uncond. Mean	Slope	R^2
High BM	0.80	0.49	0.45	High BM	0.83	0.47	0.40
Low BM	0.53	0.39	0.38	Low BM	0.51	0.39	0.37
Dividend yield	0.00	0.98	96.03	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	Asset/Variable	High BM	Low BM	Div. Yield
High BM	7.33	41.70	-0.82	High BM	7.37	42.12	-0.83
Low BM	0.89	6.41	-0.70	Low BM	0.88	6.44	-0.70
Dividend yield	-0.11	-0.11	1.00	Dividend yield	-0.11	-0.11	1.00

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

Asset/Variable	High BM	Low BM	Div. Yield	Asset/Variable	High BM	Low BM	Div. Yield
High BM	7.32	41.52	-1.27	High BM	7.35	41.94	-1.29
Low BM	0.89	6.40	-1.08	Low BM	0.89	6.43	-1.09
Dividend yield	-0.88	-0.85	0.20	Dividend yield	-0.88	-0.84	0.20