

Dynamic Portfolio Choice with Transaction Costs and Return Predictability: Linear Rebalancing Rules

Ciamac C. Moallemi
Graduate School of Business
Columbia University
email: ciamac@gsb.columbia.edu

Mehmet Sağlam
Graduate School of Business
Columbia University
email: msaglam13@gsb.columbia.edu

September 5, 2011

Abstract

We provide a highly tractable rebalancing rule for dynamic portfolio choice problems with return predictability and transaction costs. Our rebalancing rule is a linear function of the return predicting factors and can be utilized in a wide spectrum of portfolio choice models with minimal assumptions. Linear rebalancing rules enable to compute exact and efficient formulations of portfolio choice models with linear constraints, proportional and nonlinear transaction costs, and terminal wealth objectives. We illustrate the implementation of the best linear rebalancing rule in the context of portfolio execution with positivity constraints in the presence of short-term predictability. We show that there exists a considerable performance gain in using linear rebalancing rules compared to various static and dynamic heuristics.

1. Introduction

Dynamic portfolio allocation has been a central objective for institutional investors in active asset management due to changes in the estimates of expected future returns. A typical asset manager observes predictions for the expected future returns using return predicting factors such as market capitalization, book-to-market ratio, lagged returns, dividend yields, gross industrial production and other security specific or macroeconomic variables (see for example, Chen et al. (1986), Fama and French (1996), Goetzmann and Jorion (1993)). With the new estimates for the future returns, the asset manager needs to update the holdings of the portfolio while aligning with the risk objectives of the fund and keeping trading costs to a minimum. On top of these tradeoffs, there are further regulatory enforcements and diversification needs such as short-sale restriction and market neutrality (or a specific industry neutrality) that complicate the portfolio allocation in a dynamic fashion. Characterizing an optimal rebalancing rule under these complex dynamics and restrictions is a daunting task if not impossible.

Many dynamic portfolio choice models need to impose restrictive assumptions, yet often unrealistic, about each ingredient of the model i.e., return predictability, risk measure, transaction

costs and constraints, in order to achieve a tractable solution. The analytical structure will be very sensitive to these assumptions and the tractability of the model will disappear easily as soon as some simple relaxations are needed in the model. In this paper, we will overcome this rigidity by proposing a simple rebalancing rule that is a linear function of the predicting factors. Then, we solve for the optimal parameters of the linear functional form. As long as the original optimization problem is a convex programming problem, the modified optimization problem that seeks to find the optimal parameters of the linear decision rule will be a convex programming problem.

We present a large class of dynamic portfolio choice models that differ in their modeling of risk measures, transaction costs and constraints which can be formulated efficiently using a linear decision rule. Specifically, we introduce a quadratic utility function on the terminal wealth, proportional and nonlinear transaction cost functions and finally linear equality and inequality constraints. In these cases, the portfolio optimization problem reduces to a deterministic convex optimization program and the optimal parameters of the linear policy can be solved with any standard off-the-shelf solver. In the remaining cases, the optimal parameters can be solved via the sampling techniques available from the sample average and stochastic approximation literature (see for example Shapiro (2003) and Nemirovski et al. (2009)).

Finally, we illustrate our methodology empirically in the context of portfolio execution, the execution of a large long position. In order to highlight the performance gain using linear decision rules, we use the identical discrete-time setup of Garleanu and Pedersen (2009) that follows the classical linear dynamics, quadratic cost model from the control literature. However, we now introduce linear inequality constraints on portfolio positions in order to enforce decreasing positions in the security during the execution horizon. We estimate the model parameters from transactions data and propose an analogue of value and momentum predictors in a high-frequency setting. The simulation that uses these factors and the estimated parameters show that the best linear policy performs better than the deterministic policy, model predictive control and a projected version of the optimal policy proposed by Garleanu and Pedersen (2009).

The rest of this paper is organized as follows: In Section 1.1, we review the related literature. In Section 2, we present the abstract form of a dynamic portfolio choice model and provide various specific problems that satisfy the assumptions of the abstract model. We formally describe the class of linear decision rules in Section 2 and provide solution techniques in order to find the optimal parameters of the linear policy. In Section 4, we provide efficient and exact formulations of dynamic portfolio choice models using linear decision models while incorporating linear equality and inequality constraints, proportional and nonlinear transaction costs and a measure of terminal wealth risk. In Section 5, we apply our methodology in an optimal execution problem and evaluate the performance of the best linear policy. Finally, in Section 6 we conclude and discuss some future directions.

1.1. Related literature

Our paper is related to two different strands of literature, dynamic portfolio choice with or without return predictability and transaction costs and the use of linear decision rules in the optimal control problems.

The vast literature on dynamic portfolio choice starts with the seminal paper by Merton (1971) which studies the optimal dynamic allocation of one risky asset and one bond in the portfolio in a continuous-time setting. Following this seminal paper, there has been a significant literature aiming to incorporate the impact of various frictions on the optimal portfolio choice. For a survey on this literature, see Cvitanic (2001).

Constantinides (1986) is an early example that studies the impact of proportional transaction costs on the optimal investment decision and the liquidity premium in the context of CAPM. Davis and Norman (1990), Dumas and Luciano (1991) and Shreve and Soner (1994) provide the exact solution for the optimal investment and consumption decision by formally characterizing the trade and no-trade regions. One drawback of all these papers is that the optimal solution is only computed in the case of a single stock and bond. Liu (2004) extends this result to multiple assets but assumes that asset returns are not correlated.

There is a growing literature on portfolio selection that incorporates return predictability with transaction costs. Balduzzi and Lynch (1999) and Lynch and Balduzzi (2000) illustrate the impact of return predictability and transaction costs on the utility costs and the optimal rebalancing rule by discretizing the state space of the dynamic program. With a similar state space discretization, Lynch and Tan (2010) model the dynamic portfolio decision with multiple risky assets under return predictability and transaction costs and provide numerical experiments with two risky assets. Recently, Brown and Smith (2010) provides heuristic trading strategies and dual bounds for a general dynamic portfolio optimization problem with transaction costs and return predictability. Our paper is closely related to these papers. However, our approach with linear decision rules scales better with multiple assets compared to that of discretization methods as grid approximations with multiple risky assets will suffer from the curse of dimensionality.

Brandt et al. (2009) parameterizes the rebalancing rule as a function of security characteristics and estimates the parameters of the rule from empirical data without modeling the distribution of the returns and the return predicting factors. Even though our approach is also a linear parametrization of return predicting factors, there are fundamental differences between our approach and that of Brandt et al. (2009). First, the class of linear policies we consider is much larger than the specific linear functional form in Brandt et al. (2009). In our approach the parameters are time-varying and cross-sectionally different for each security. Second, the extensions provided in Brandt et al. (2009) for imposing positivity constraints and transaction costs are ad-hoc and cannot be generalized to arbitrary convex constraints or transaction cost functions.

Garleanu and Pedersen (2009) achieve a closed-form solution for a model with linear dynamics in return predictors and quadratic function for transaction costs and quadratic penalty term for

risk. However, the analytic solution is highly sensitive to the quadratic cost structure with linear dynamics (see Bertsekas (2000)). This special case cannot handle any inequality constraints on portfolio positions, non-quadratic transactions costs, such as proportional transaction cost, or risk considerations on terminal wealth. On the other hand, our approach can be implemented efficiently in these realistic scenarios and provides more flexibility in the objective function of the investor.

The use of linear decision rules in optimal control problems has been abundant in the literature. This approximation technique has attracted considerable interest recently in robust and two-stage adaptive optimization context (see Ben-Tal et al. (2004), Ben-Tal et al. (2005), Chen et al. (2007), Chen et al. (2008), Bertsimas et al. (2010) and Bertsimas and Goyal (2011)). Shapiro and Nemirovski (2005) illustrate that linear decision rules can reduce the complexity of multistage stochastic programming problems. Kuhn et al. (2009) proposes an efficient method to estimate the loss of optimality incurred by linear decision rule approximation.

In this strand of literature, we believe the closest works to the methodology described in our paper are Calafiore (2009) and Skaf and Boyd (2010). Both of these papers use linear decision rules to address dynamic portfolio choice problems with proportional transaction costs without return predictability. Calafiore (2009) compute lower and upper bounds on the expected transaction costs and solves two convex optimization problems to get upper and lower bounds on the optimal value of the simplified dynamic optimization program with linear decision rules. On the other hand, Skaf and Boyd (2010) study the dynamic portfolio choice problem as an application to their general methodology of using affine controllers on convex stochastic programs. They first linearize the dynamics of the wealth process and then solve the resulting convex optimization via sampling techniques. The foremost difference between our approach and these papers is the modeling of return predictability. Hence, the optimal rebalancing rule in our model is a linear function of the predicting factors. Furthermore, we derive exact reductions to deterministic convex programs in the cases of proportional and nonlinear transaction costs.

2. Dynamic Portfolio Choice with Return Predictability and Transaction Costs

We consider a dynamic portfolio choice problem with return predictability in security returns and costly trading. The number of investable securities is N , time is discrete and indexed by $t = 1, \dots, T$, where T is the investment horizon. Each security i has a price change of $r_{i,t+1}$ from time t to $t + 1$. We collect these price changes in the return vector $r_{t+1} \triangleq (r_{1,t+1}, \dots, r_{N,t+1})^\top$. We assume that security returns can be predicted by a set of K factors. These factors could be security specific characteristics such as market capitalization of the stock, book-to-market ratio of the stock, lagged twelve month return of the stock or dividend yields (see Fama and French (1996) and Goetzmann and Jorion (1993)). Alternatively, they could be macroeconomic signals that affect the return of each security, e.g., inflation, treasury bill rate, industrial production (see Chen et al.

(1986) for a complete list of variables).

We let $x_{i,t}$ denote the number of shares that the investor holds in i th security at time t . We collect the portfolio holdings in each security in the vector $x_t \triangleq (x_{1,t}, \dots, x_{N,t})^\top$. Similarly, let the trade vector $u_t \triangleq (u_{1,t+1}, \dots, u_{N,t+1})^\top$ be the amount of shares that the investor wants to trade at the end of t th period when he observes f_t with a portfolio x_t . Consequently, we have the following linear dynamics for our position and trade vector:

$$x_{t+1} = x_t + u_{t+1}$$

We assume very general dynamics, possibly nonlinear, for the evolution of returns and factors.

Assumption 1 (Dynamics). *On a complete filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$, we assume that factors and returns evolve according to*

$$\begin{aligned} f_{t+1} &= G(f_t, \dots, f_1, \epsilon_{t+1}), \\ r_{t+1} &= H(f_t, \dots, f_1, \epsilon_{t+1}), \end{aligned}$$

where $H(\cdot)$ and $G(\cdot)$ are known functions and the filtration $\mathbb{F} \triangleq \{\mathcal{F}_t\}_{t \geq 0}$ is generated by the exogenous noise terms ϵ_{t+1} and ϵ_{t+1} .

Let the time series vector of portfolio positions, factors and returns be given by $\mathbf{x} \triangleq (x_0, x_1, \dots, x_T)$, $\mathbf{f} \triangleq (f_1, \dots, f_T)$, and $\mathbf{r} \triangleq (r_1, \dots, r_T)$ respectively. Let $U \subseteq \mathbb{R}^N \times \dots \times \mathbb{R}^N$ denote the set of all feasible trade sequences $\mathbf{u} = (u_1, \dots, u_T)$. The investor's trading decisions are determined by a policy π that selects a sequences of trades \mathbf{u} in U for each realization of \mathbf{r} and \mathbf{f} . We let \mathcal{U} be the set of all policies. We assume that the investor's trading decisions are nonanticipative in that the trade vector u_{t+1} in period t depends only on what is known at the beginning of period t . Formally, we require policies to be adapted to the filtration \mathbb{F} such that a policy's selection of the first t trades (u_1, \dots, u_t) must be measurable with respect to \mathcal{F}_t . Let $\mathcal{U}_{\mathbb{F}}$ be the set of all nonanticipative policies.

The objective of the investor is to select a nonanticipative policy π that maximizes the expected total payoff. Since r_{t+1} is a known function of f_t , the extended and real-valued reward function can purely be written by the portfolio position vector, \mathbf{x} and the factor vector, \mathbf{f} . The investor wants to solve the following maximization problem:

$$(1) \quad \sup_{\pi \in \mathcal{U}_{\mathbb{F}}} E_{\pi}[p(\mathbf{x}, \mathbf{f})].$$

Assumption 2 (Concave objective function). *We assume that the total reward function $p(\mathbf{x}, \mathbf{f})$ is jointly concave in \mathbf{x} and \mathbf{f} and the set of all nonanticipative policies $\mathcal{U}_{\mathbb{F}}$ is convex.*

In this paper, we will consider dynamic portfolio choice models that satisfy Assumptions 1 and 2. We will now provide specific examples of dynamic portfolio choice models that satisfy Assumptions 1 and 2.

Example 1 (Discrete-time model in Garleanu and Pedersen (2009)). *This model has the following dynamics which fit into our general framework.*

$$(2) \quad \begin{aligned} r_{t+1} &= \mu_t + Bf_t + \epsilon_{t+1} \\ f_{t+1} &= (I - \Phi) f_t + \varepsilon_{t+1} \end{aligned}$$

where μ_t is the deterministic “fair return,” e.g., from the CAPM, B is a constant matrix of factor loadings, Φ is a matrix of mean reversion coefficients for the factors. It is assumed that $\text{Var}_t(\epsilon_{t+1}) = \Sigma$ and $\text{Var}_t(\varepsilon_{t+1}) = \Psi$. Trading is costly and the transaction cost to execute $u_t = x_t - x_{t-1}$ shares, is given by

$$TC(u_t) \triangleq \frac{1}{2} u_t^\top \Lambda u_t,$$

where Λ measures the level of trading costs.

The investor’s objective function is to choose the trading strategy (u_1, \dots, u_T) to maximize discounted future expected excess return penalized for risks and transaction costs:

$$(3) \quad \begin{aligned} & \text{maximize}_{\pi \in \mathcal{U}_{\mathbb{F}}} \mathbb{E} \left[\sum_{t=1}^T \rho^t \left(x_t^\top (r_{t+1} - \mu_t) - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{1}{2} u_t^\top \Lambda u_t \right) \right] \\ & = \text{maximize}_{\pi \in \mathcal{U}_{\mathbb{F}}} \mathbb{E} \left[\sum_{t=1}^T \rho^t \left(x_t^\top (Bf_t) - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{1}{2} u_t^\top \Lambda u_t \right) \right] \end{aligned}$$

where ρ is the discount factor and γ is the coefficient of risk aversion. Garleanu and Pedersen (2009) notes that this objective function belongs to an investor who is compensated based on his performance relative to a benchmark. Hence, $x_t^\top (Bf_t)$ measures the excess return over the benchmark, and $x_t^\top \Sigma x_t$ measures the variance of the tracking error relative to the benchmark¹.

This problem can be solved fairly easily using the classical results from the linear-quadratic control (LQC) literature. On the other hand, the tractability of this model is very sensitive to quadratic cost structure with linear dynamics (see Bertsekas (2000)) in x_t and u_t . This special case cannot handle any inequality constraints on our decision variables. However, in many real world examples, the investor’s reward function is nonlinear and the investor needs to satisfy certain conditions related to his portfolio holdings at all times. Therefore, many examples of portfolio choice models will not have the simple analytical structure of Example 1. In the following examples, we will provide concrete examples of such models that will involve portfolio position or trade constraints, different forms of transaction costs and risk measures.

Example 2 (Constraints). *In practice, the most common constraint in constructing equity portfolios is the short-sale restriction. Most of the mutual funds are enforced not to have any short positions*

¹See Garleanu and Pedersen (2009) for other interpretations.

by law. This requires the portfolio optimization problem to include simple positivity constraint

$$x_t \geq \mathbf{0} \quad \forall t = 1, \dots, T.$$

We observe a similar restriction when an execution desk needs to sell or buy a large portfolio on behalf of an investor. Due to the regulatory rules in agency trading, the execution desk is only allowed to sell or buy during the trading horizon. In the “pure-sell” scenario, the execution desk needs to impose the following constraint

$$u_t \leq \mathbf{0} \quad \forall t = 1, \dots, T.$$

These simple positivity or negativity constraints may appear easy to handle but the portfolio optimization problem does no longer fall into LQC framework.

Example 3 (Non-quadratic transaction costs). In practice, trading costs can result from service fees such as bid-ask spread, commissions, exchange fees which are intrinsically proportional to the trade size. Furthermore, trading costs can also occur due to disadvantageous transaction price caused by the price impact of the trade. The case of proportional transaction cost is a classical cost structure which is well studied in the literature (see Constantinides (1986) and the references therein). Letting χ_i be the the proportional transaction cost rate (an aggregate sum of bid-ask cost and commission fees) for trading security i , the investor will incur a total cost of

$$TC(u_t) \triangleq \sum_{i=1}^N \chi_i |u_{t,i}|.$$

The management of the trading costs due to price impact has recently attracted considerable interest as well (see Obizhaeva and Wang (2005) and Almgren and Chriss (2000)). In this case, models of price impact imply a nonlinear relationship between trade size and the resulting transaction cost. The following general transaction cost function falls under this subcategory:

$$TC(u_t) = \sum_{i=1}^N \chi_i |u_{t,i}|^\beta,$$

where $\beta \geq 1$ and χ_i is the security specific proportionality constant. Example 1 only considers the case of $\beta = 2$ as any other value will cause the model not to be in the LQC framework.

Example 4 (Terminal wealth risk). The objective function of Example 1 includes a nonstandard penalty term to disallow taking excessive risk. A quadratic penalty term, $x_t^\top \Sigma x_t$, is used in order to satisfy the requirements of the LQC model. However, this measure of per-period risk is not very appropriate as the investor is more concerned about the volatility of his terminal wealth as opposed to the sum of per-period penalty terms which does not have much economic meaning. There are other appropriate risk measures which again break down the restrictive assumptions of the LQC

model.

Consider a quadratic utility function on the terminal wealth. Denoting the terminal wealth by W_T and the total trading cost by C_T , we have

$$W_T = \sum_{t=1}^T x_t^\top r_{t+1}$$

$$C_T = \frac{1}{2} \sum_{t=1}^T (u_t^\top \Lambda u_t)^\beta$$

and the investor's objective would be

$$\text{maximize}_{\pi \in \mathcal{U}_{\mathbb{F}}} \mathbb{E} \left[W_T - C_T - \frac{\gamma}{2} W_T^2 \right]$$

Example 5 (Maximum drawdown risk). *Instead of introducing per-period penalty terms to avoid excessive risk, a fund manager would be more interested in minimizing the probability of extreme successive losses. Since the asset management business is solely dependent on management and incentive fees proportional to the size of assets under management, successive losses will trigger the start of withdrawals from the accounts and hence, jeopardize the profitability of the business.*

We define maximum drawdown, MD , by the worst loss of the portfolio between any two points of time during the investment horizon. Formally,

$$MD \triangleq \text{maximize}_{t_1, t_2} \left(- \sum_{t=t_1}^{t_2} x_t^\top r_{t+1} \right).$$

Suppose that the fund manager wants to control the expected maximum drawdown by a certain threshold, C . Therefore, he imposes the following constraint

$$\mathbb{E} [MD] \leq C.$$

This constraint certainly violates the assumptions of the LQC framework but satisfies Assumption 2.

Example 6 (Complicated dynamics). *We can also relax the the dynamics in Example 1. Consider the following generalization of the factor dynamics:*

$$r_{t+1} = \mu_t + B_t(\lambda + \omega_{t+1}) + \epsilon_{t+1}$$

$$B_{t+1} = (I - \Phi) B_t + \varepsilon_{t+1},$$

where $\mathbb{E}_t[\omega_{t+1}] = 0$ and $\text{Var}_t[\omega_{t+1}] = \Upsilon$. In this model, the factor loading matrix, B_t , is the return predicting factor and follows a mean-reverting process. These dynamics are very sensible as the conditional variance of the return becomes dependent on the factor structure and time-varying, i.e.,

$\text{Var}_t[r_{t+1}] = B_t \Upsilon B_t^\top + \Sigma$. Thus, in this case, the penalty term for risk becomes $x_t^\top (B_t \Upsilon B_t^\top + \Sigma) x_t$ and the optimization problem stated as in Example 1 does no longer fall into the LQC framework.

The dynamics and the reward functions considered in these examples satisfy our basic requirements provided in Assumption 1 and Assumption 2. These examples illustrate that in many real-world considerations for return predictability, transaction costs, risk measures and constraints, the dynamic portfolio choice becomes difficult to solve analytically.

3. Best Linear Model

As emphasized in the examples provided in Section 2, many appropriate models for portfolio choice do not fall into the LQC framework. Without the special restrictions as imposed in the LQC framework, the optimal dynamic policy cannot be computed analytically. In this section, in order to find an approximate solution, we will restrict our feasible policies into a subset which is parsimoniously parameterized. Instead of solving for an optimal policy, we will solve for the optimal parameters of the restricted subset of policies.

Consider the following restricted set of policies, linear rebalancing policies, which is obtained by taking the affine combinations of the factors.

Definition 1. A linear policy π selecting the feasible trade sequences $\mathbf{u} = (u_1, \dots, u_T)$ can be represented in the following form,

$$u_t = c_t + \sum_{s=1}^t E_{s,t} f_s \quad t = 1, 2, \dots, T$$

parameterized by a collection of vectors $\{c_t \in \mathbb{R}^N, 1 \leq t \leq T\}$ and a collection of matrices $\{E_{s,t} \in \mathbb{R}^{N \times K}, 1 \leq s \leq t \leq T\}$.

There are some important properties of linear rebalancing policies. First, these policies allow recourse based on the evolution of the factors but in the limited linear functional form. Second, the set of affine policies are much more general than deterministic policies as setting $E_{s,t}$ to matrices of zeroes would yield the deterministic policy. Finally, if the optimization problem satisfies the requirements of the LQC model, i.e., linear dynamics, quadratic cost, equality constraints, affine policies will be optimal.

Note that the functional form provided in Definition 1 can be extended with basis functions as well:

$$u_t = \sum_{s=1}^t \varphi_{s,t}(f_s) \quad t = 1, 2, \dots, T.$$

where the collection of basis functions, $\{\varphi_{s,t}, 1 \leq s \leq t \leq T\}$, are real-valued functions on the state space.

We constrain our policies to be in the set of linear policies, denoted by $\mathcal{L} \subseteq \mathcal{U}_{\mathbb{F}}$. Instead of solving the original problem in (1), an approximate solution can be obtained by solving the following maximization problem:

$$(4) \quad \sup_{\pi \in \mathcal{L}} \mathbb{E}[p(\mathbf{x}, \mathbf{f})].$$

which is equivalent to

$$(5) \quad \sup_{\mathbf{E}, \mathbf{c}} \mathbb{E}[p(\mathbf{x}, \mathbf{f})]$$

where $\mathbf{E} \triangleq (E_{1,1}, E_{1,2}, \dots, E_{T-1,T}, E_{T,T})$ and $\mathbf{c} \triangleq (c_1, \dots, c_T)$.

Proposition 1. *The optimization problem given by*

$$(6) \quad \begin{aligned} & \underset{\mathbf{E}, \mathbf{c}}{\text{maximize}} \mathbb{E}[p(x_0, x_1, \dots, x_T, f_0, f_1, \dots, f_T)] \\ & \text{subject to } x_t = x_{t-1} + u_t && t = 1, 2, \dots, T, \\ & u_t = c_t + \sum_{s=1}^t E_{s,t} f_s && t = 1, 2, \dots, T, \end{aligned}$$

is a convex optimization problem.

Proof. Since $p(\mathbf{x}, \mathbf{f})$ is jointly concave in \mathbf{x} and \mathbf{f} , and \mathbf{x} can be written as an affine transformation of $E_{s,t}$ and c_t , $p(\cdot)$ is jointly concave in $E_{s,t}$, c_t and \mathbf{f} . Since expectation preserves concavity, the maximization problem in (6) is a convex optimization problem. ■

Note that this is a much simpler problem to solve as instead of solving for the optimal policy, we are now solving for the optimal parameters of a suboptimal policy. There exist three solution techniques to solve the stochastic optimization problem in (6).

- **Sample average approximation (SAA).**

Using Monte Carlo sampling techniques, we can generate a sample of M independent replications of the factor data, $\mathbf{f}^1, \dots, \mathbf{f}^M$, according to the functional form given by (1). Then, we can compute the sample average approximation of the original objective function, $h(\mathbf{E}, \mathbf{c}) \triangleq \mathbb{E}[p(\mathbf{x}, \mathbf{f})]$, by

$$\hat{h}(\mathbf{E}, \mathbf{c}) \triangleq \frac{1}{M} \sum_{i=1}^M p(\mathbf{x}, \mathbf{f}^i).$$

Then, we solve the deterministic optimization problem given by

$$(7) \quad \begin{aligned} & \underset{\mathbf{E}, \mathbf{c}}{\text{maximize}} \frac{1}{M} \sum_{i=1}^M p(\mathbf{x}, \mathbf{f}^i) \\ & \text{subject to } x_t = x_{t-1} + u_t && t = 1, \dots, T, \\ & u_t = c_t + \sum_{s=1}^t E_{s,t} f_s^i && t = 1, \dots, T. \end{aligned}$$

Let us denote V^* and S the optimal value and the set of optimal solutions of the original problem in (6), respectively, and analogously denote \hat{V}_M^* and \hat{S}_M the optimal value and the set of optimal solutions of the SAA problem in (7). Under appropriate regularity conditions (see Shapiro (2003)), we have $\hat{V}_M^* \rightarrow V^*$ and $\mathbb{D}(\hat{S}_M, S) \rightarrow 0$ with probability 1 as $M \rightarrow \infty^2$.

- **Stochastic approximation.**

Since we can stack all our decision variables, c_t and $E_{s,t}$, on top of each other in a single vector \mathbf{z} , we can rewrite the original problem in the abstract form,

$$\underset{\mathbf{z} \in \Theta}{\text{maximize}} \{h(\mathbf{z}) \triangleq \mathbb{E}[p(\mathbf{z}, \mathbf{f})]\}$$

where Θ is a closed and convex set. Since $p(\mathbf{z}, \mathbf{f})$ is convex for every realization of \mathbf{f} and $h(\mathbf{z})$ is finite valued (due to Assumption 2), then (cf. Strassen (1965))

$$\partial h(\mathbf{z}) = \mathbb{E}[\partial_{\mathbf{z}} p(\mathbf{z}, \mathbf{f})]$$

Using Monte Carlo sampling techniques, we can generate a sample of M independent replications of the factor data, $\mathbf{f}^1, \dots, \mathbf{f}^M$, according to the functional form given by (1). Then, we can use a subgradient of $P(\mathbf{z})$ denoted by $G(\mathbf{z}, \mathbf{f}) \in \partial_{\mathbf{z}} p(\mathbf{z}, \mathbf{f})$ as a stochastic subgradient. With these constructs at hand, we can employ the classical stochastic approximation algorithm, an application of the subgradient descent method. Starting from an initial point $\mathbf{z}_0 \in \Theta$, and using a sequence of appropriately chosen stepsizes (see Nemirovski et al. (2009) for the details), $\gamma_j > 0$, we update our approximate solution by the formula

$$\mathbf{z}_{j+1} = \Pi_{\Theta} \left(\mathbf{z}_j - \gamma_j G(\mathbf{z}_j, \mathbf{f}^j) \right)$$

where the operation $\Pi_{\Theta}(\mathbf{z})$ denotes $\text{argmin}_{y \in \Theta} \|\mathbf{z} - y\|_2$.

Suppose further that \mathbf{z}^* is the unique optimal solution of the original problem and V^* be the corresponding optimal value. Then, under certain regularity conditions, Nemirovski et al. (2009) show that after k iterations, the expected error of the current solution in terms of the

²The assertion, $\mathbb{D}(\hat{S}_M, S) \rightarrow 0$ with probability 1, means that for any selection of optimal solution, $(\mathbf{E}, \mathbf{c}) \in \hat{S}_M$ of the SAA problem, we have $\text{dist}((\mathbf{E}, \mathbf{c}), S) \rightarrow 0$ with probability 1.

distance to \mathbf{z}^* is of order $O(\frac{1}{\sqrt{k}})$ and the expected error in terms of the objective value is of order $O(\frac{1}{k})$.

- **Efficient exact formulation.** If the expectation in the objective of (6) can be computed analytically, the stochastic program will reduce to a deterministic convex optimization. The linear functional form helps in obtaining analytical tractability with further assumptions for the distribution of the noise terms in (1). In Section 4, we will provide the details of these exact formulations when the noise terms are serially independent and identically distributed with multivariate normal distribution.

4. Efficient exact formulations

In this section, we will provide efficient exact formulations of dynamic portfolio choice problems using the class of linear policies for our feasible set of policies. We will introduce various extensions with constraints on portfolio holdings, transaction costs, and risk measures and show that the stochastic program can be reduced to a deterministic convex program.

Throughout this section, we will assume that the vector constructed by stacking all factors on top of each other, $F_t \triangleq (f_1, \dots, f_t)^\top$, is jointly normal with mean θ_t and covariance matrix Ω_t . With this assumption, the distribution of any linear policy will also be jointly normal as each policy is affine transformations of the factors. In order to see this formally, let

$$(8) \quad M_t \triangleq \begin{bmatrix} E_{1,t} & E_{2,t} & \dots & E_{t,t} \end{bmatrix}$$

Then, we can write our trade vector as $u_t = c_t + M_t F_t$. With this representation, it is easy to compute the first and second moments of u_t :

$$(9) \quad \begin{aligned} \mu_t &\triangleq \mathbf{E}(u_t) = c_t + M_t \theta_t \\ V_t &\triangleq \mathbf{Var}(u_t) = M_t \Omega_t M_t^\top. \end{aligned}$$

Therefore, u_t is normally distributed with mean μ_t and covariance matrix V_t . Similarly, we can obtain the statistics for x_t . Using,

$$(10) \quad x_t = x_0 + \sum_{i=1}^t u_i = x_0 + \sum_{i=1}^t \left(c_i + \sum_{s=1}^i E_{s,i} f_s \right) = d_t + \sum_{s=1}^t J_{s,t} f_s$$

where $d_t \triangleq x_0 + \sum_{i=1}^t c_i$ and $J_{s,t} \triangleq \sum_{i=s}^t E_{s,i}$. Thus, with this representation, x_t is also normally distributed with mean vector κ_t and covariance matrix Y_t where

$$(11) \quad \begin{aligned} \kappa_t &\triangleq \mathbf{E}(x_t) = d_t + P_t \theta_t \\ Y_t &\triangleq \mathbf{Var}(x_t) = P_t \Omega_t P_t^\top. \end{aligned}$$

and

$$(12) \quad P_t \triangleq \begin{bmatrix} J_{1,t} & J_{2,t} & \dots & J_{t,t} \end{bmatrix}$$

4.1. Linear constraints

We will provide formulations for linear equality and inequality constraints. These type of constraint appear frequently in portfolio choice due to regulatory reasons such as short sale restriction, liquidation purposes or diversification needs such as keeping a specific industry exposure under a certain limit. We will provide the exact formulations of such constraints when we use the class of linear policies.

4.1.1. Equality constraints

Equality constraints in portfolio choice appear especially in portfolio execution problems when the investor needs to liquidate a certain portfolio or construct a certain target portfolio directed by his own investment research. Any type of equality constraints including these target portfolios can be implemented with linear policies.

Suppose that for some period t , we need $Au_t = b$ where $A \in \mathbb{R}^{M \times N}$ and $b \in \mathbb{R}^N$. Substituting the functional for the linear policy, the equality constraint can be written by

$$Ac_t + A \left(\sum_{s=1}^t E_{s,t} f_s \right) = b,$$

which is equivalent to

$$\begin{aligned} Ac_t &= b \\ AE_{s,t}^k &= \mathbf{0} \quad \forall (s, t, k) \text{ where } s = 1, \dots, t \text{ and } k = 1, \dots, K, \end{aligned}$$

and where $E_{s,t}^k$ denotes the k th column of $E_{s,t}$. Analogously, a similar computation can be used to include linear equality constraints on the portfolio positions, x_t .

4.1.2. Inequality constraints

In practice, the most common inequality constraint is the short-sale restriction. Most of the mutual funds are enforced not to have any short positions by law. Therefore, we need to impose positivity constraints on our portfolio positions, i.e., $x_t \geq \mathbf{0}$.

When the noise terms do not have a bounded support, such constraints can only be implemented with probabilistic constraints for the class of linear policies. In the case of bounded noise terms, such as the case with truncated multivariate normal distribution, the linear policies can also handle the linear inequality constraints in the robust sense.

Suppose that we have the following probabilistic constraint for an arbitrary time period t with the desired confidence level of η . Then, this constraint can be implemented exactly using a second-order cone constraint ³.

Fact 1. *The chance constraint $\mathbb{P}(a^\top u_t + b \geq 0) \leq \eta$ can be reduced to a second order cone constraint*

$$(13) \quad a^\top (c_t + M_t \theta_t) + b + \Phi^{-1}(1 - \eta) \left\| \left(M_t \Omega_t M_t^\top \right)^{1/2} a \right\|_2 \leq 0$$

where M_t is defined in (8).

Proof. Let μ_t and V_t be the mean and the variance of u_t as defined in Equation (9). Then,

$$\mathbb{P}(a^\top u_t + b \geq 0) = \mathbb{P}(\beta_t + \sigma_t Z \geq 0)$$

where $\beta_t \triangleq a^\top \mu_t + b$, $\sigma_t \triangleq \left\| \left(M_t \Omega_t M_t^\top \right)^{1/2} a \right\|_2$ and Z is a standard normal random variable. Then,

$$\mathbb{P}(\beta_t + \sigma_t Z \geq 0) = \mathbb{P}\left(Z \geq -\frac{\beta_t}{\sigma_t}\right)$$

Note that $\mathbb{P}\left(Z \geq -\frac{\beta_t}{\sigma_t}\right)$ is less than η if and only if $-\frac{\beta_t}{\sigma_t} \geq \Phi^{-1}(1 - \eta)$ where $\Phi^{-1}(y)$ is the inverse function of the cumulative distribution function of a standard normal random variable. Then, our probabilistic constraint is equivalent to

$$\beta_t + \Phi^{-1}(1 - \eta) \sigma_t \leq 0.$$

Substituting the values for β_t and σ_t , we finally obtain

$$a^\top (c_t + M_t \theta_t) + b + \Phi^{-1}(1 - \eta) \left\| \left(M_t \Omega_t M_t^\top \right)^{1/2} a \right\|_2 \leq 0$$

which is a second order cone constraint. ■

A similar approach can also be used to incorporate linear inequality constraints on the portfolio positions, x_t .

4.2. Transaction costs

In this section, we will provide formulations for proportional and nonlinear transaction cost functions with the class of linear policies. In practice, trading costs for small trades are mostly proportional as costs due from bid-ask spread, commissions, exchange fees are intrinsically proportional to the trade size. However, for large trades, trading costs do also include a price impact term which is

³This is a well known result in the literature, see Boyd and Vandenberghe (2004) for the derivation.

nonlinear with respect to the trade size. However, empirical studies studying price impact models favor nonlinear (not necessarily quadratic) transaction cost functions.

4.2.1. Proportional transaction cost

Proportional transaction costs are widely used in the portfolio choice literature. (see Constantinides (1986) for an early example). Letting χ_i (per share cost of total fees due the bid-ask spread, commission and exchange fees) be the proportional transaction cost rate for i th security, let

$$\text{TC}(u_t) = \sum_{i=1}^N \chi_i |u_{t,i}|,$$

where $u_{t,i}$ is the amount of shares traded in the i th security.

With this representation, one needs to compute $\mathbb{E}[\text{TC}(u_t)]$ analytically in order to solve the optimization program using deterministic techniques. Using the class of linear policies for u_t , we can compute this expectation in closed-form:

$$\begin{aligned} \mathbb{E}[\text{TC}(u_t)] &= \mathbb{E} \left[\sum_{i=1}^N \chi_i |u_{t,i}| \right] \\ &= \sum_{i=1}^N \chi_i \mathbb{E} \left[\left| \mu_{t,i} + \sqrt{V_{t,i}} Z \right| \right], \end{aligned}$$

where Z is a standard normal random variable and $\mu_{t,i}$ and $V_{t,i}$ are the i th rows of the mean vector, μ_t and the covariance matrix, V_t , respectively, which are the first and second moments of u_t as defined in (9). Using the properties of the folded normal distribution, and letting $\sigma_{t,i} \triangleq \sqrt{V_{t,i}}$ we obtain

$$\mathbb{E}[\text{TC}(u_t)] = \sum_{i=1}^N \chi_i \left(\sqrt{\frac{2}{\pi}} \sigma_{t,i} \exp \left\{ -\frac{\mu_{t,i}^2}{2\sigma_{t,i}^2} \right\} + \mu_{t,i} \left(1 - 2\Phi \left(-\frac{\mu_{t,i}}{\sigma_{t,i}} \right) \right) \right)$$

where $\Phi(y)$ is the cumulative distribution function of a standard normal random variable. This expression is necessarily convex in $c_{t,i}$ and $E_{s,t,i}$, the i th rows of c_t and $E_{s,t}$ as the expression

$$\mathbb{E}[\text{TC}(u_t)] = \sum_{i=1}^N \chi_i \mathbb{E} \left| \left(c_{t,i} + \sum_{s=1}^t E_{s,t,i} f_s \right) \right|$$

is the sum of the absolute values of linear functions in $c_{t,i}$ and $E_{s,t,i}$.

4.2.2. Nonlinear transaction cost

Nonlinear trading costs due to price impact have recently attracted considerable interest as well (see Obizhaeva and Wang (2005) and Almgren and Chriss (2000)). In these types of transaction

cost models, there is still a security specific proportionality constant χ_i , due to proportional nature of bid-ask spread, commission and exchange fees. Thus, in the presence of price impact

$$\text{TC}(u_t) = \sum_{i=1}^N \chi_i |u_{t,i}|^\beta,$$

where $u_{t,i}$ is the amount of shares traded in the i th security and $\beta \geq 1$ is the estimated coefficient of nonlinearity ⁴. This functional form again preserves the convexity of our optimization program when the class of linear policies is used. Formally,

$$\begin{aligned} \mathbb{E}[\text{TC}(u_t)] &= \mathbb{E} \left[\sum_{i=1}^N \chi_i |u_{t,i}|^\beta \right] \\ &= \sum_{i=1}^N \chi_i \mathbb{E} \left[|\mu_{t,i} + \sigma_{t,i} Z|^\beta \right] \\ &= \sum_{i=1}^N \chi_i \left((2\sigma_{t,i}^2)^{\frac{\beta}{2}} \frac{\Gamma\left(\frac{1+\beta}{2}\right)}{\sqrt{\pi}} {}_1F_1\left(-\frac{\beta}{2}; \frac{1}{2}; -\frac{\mu_{t,i}^2}{2\sigma_{t,i}^2}\right) \right) \end{aligned}$$

where $\Gamma(z)$ is the gamma function and ${}_1F_1(a; b; z)$ is confluent hypergeometric function of the first kind.

4.3. Risk measures

Investors are more concerned about fluctuations in their terminal wealth as opposed to per-period measures of risk. Consider an individual investor who would like to achieve a certain level of wealth for his expenses in the retirement period. Instead of penalizing per-period fluctuations, his main aversion would be the extreme fluctuations of his terminal wealth. Similarly, an institutional investor would also be more interested in minimizing the risk associated with the amount of assets under management (AUM) for next year as any significant drop in AUM will jeopardize the profitability of the business.

In this section, using the class of linear policies, we will provide an efficient formulation for a risk measure associated with the terminal wealth of the investor.

4.3.1. Terminal wealth risk

Suppose that the investor wants to maximize a quadratic-type of utility function on his terminal wealth. He wants to maximize the dollar value of his terminal holdings penalized by the square of his total wealth before transaction costs. Formally,

$$\underset{c_t, E_{s,t}}{\text{maximize}} \mathbb{E} \left[W_T - C_T - \frac{\gamma}{2} W_T^2 \right],$$

⁴This is typically reported to be around $\frac{3}{2}$.

where W_T denotes the total wealth before transaction costs and C_T denotes the total transaction costs paid during investment horizon. This objective aims to lower the risk associated with terminal wealth with the penalty term, $\frac{\gamma}{2}W_T^2$, where γ is the coefficient of investor's risk aversion.

Suppose the investor starts with an initial position, $x_0 = \mathbf{0}$. Then, we can write

$$W_T = \sum_{t=1}^T \left(x_t^\top r_{t+1} \right),$$

$$C_T = \sum_{t=1}^T \text{TC}(u_t).$$

Assume that the return dynamics and transaction cost function are given as in Example 1 (with $\mu_t = \mathbf{0}$ for simplicity), i.e.,

$$r_{t+1} = Bf_t + \epsilon_{t+1}$$

$$\text{TC}(u_t) = \frac{1}{2} u_t^\top \Lambda u_t,$$

where ϵ_{t+1} is serially independent zero-mean noise term with conditional covariance matrix, $\text{Var}_t(\epsilon_{t+1}) = \Sigma$ and is independent of all factor data, F_t . Total transaction cost, C_T by

$$C_T = \sum_{t=1}^T \frac{1}{2} u_t^\top \Lambda u_t.$$

Proposition 2. *Let the investor's objective be*

$$\underset{c_t, E_{s,t}}{\text{maximize}} \mathbf{E} \left[W_T - C_T - \frac{\gamma}{2} W_T^2 \right].$$

This program can be reduced to an exact deterministic convex program.

5. Application: Optimal execution in the presence of alpha

In this section, we will provide an empirical application to illustrate the implementation of the best-linear policy. As our empirical example, we consider a classical problem in equity agency trading.

Agency trading in equities has witnessed tremendous growth over the past quarter of century driven mostly by the increasing inflow of assets into the highly competitive market of institutional asset management such as mutual, pension, and hedge funds. Due to the constant changes in market variables, asset managers need to update the holdings of these large portfolios while minimizing trading costs, often labeled as “execution costs,” consisting of commissions, bid/ask spreads and more importantly, price impact from trading. For this purpose, they work closely with execution desks in investment banks and allow them to execute the trades on their behalf so that they can

achieve their desired portfolio with minimal transaction costs. The responsibility of the execution desk is to find a feasible execution schedule over the client-specified trading horizon while minimizing trading costs and aligning with the risk objectives of the client. This problem of finding an optimal execution schedule has received a lot of attention in the literature since the initial paper by Bertsimas and Lo (1998). In their model, when price impact is proportional to the number of shares traded, the optimal execution schedule is to trade equal number of shares at each trading time. There are number of papers that extend this model to incorporate the risk of the execution strategy. For example, Almgren and Chriss (2000) derive that risk averse agents need to liquidate their portfolio faster in order to reduce the uncertainty of the execution cost. With a different specification of preferences, Hora (2006) find that the execution schedule that trades rapidly at the beginning and the end of the execution horizon and slowly in-between is optimal.

As high-frequency data become more and more available to investors, there is a growing interest to model return predictability in intraday stock returns, often called as “short-term alpha models”, similar to those well-known factor models in the literature, e.g., Capital Asset Pricing Model (CAPM), and Fama-French Three Factor Model. For example, Heston et al. (2010) document that systematic trading as described in the examples above and institutional fund flows lead to predictable patterns in intraday returns of common stocks. Motivated by this result, we will consider an optimal execution problem in the presence of short-term predictability with a factor model.

If the optimal execution problem in this framework satisfies Assumption 1 and Assumption 2, then we can compute the best execution schedule in the space of linear execution schedules, i.e., the number of shares to trade at each time is a linear function of the previous return predicting factors.

In the next section, we will formulate the optimal execution model as outlined in Example 1. With this intentional choice, we have the advantage of comparing the performance of the best-linear policy with the dynamic policy provided in Garleanu and Pedersen (2009)⁵.

5.1. Formulation

Suppose that x_0 denotes the portfolio of N securities we hold in number of shares that we would like to sell before time T . We assume that trades can occur at discrete times, $t = 1, \dots, T$. We define an *execution schedule* to be a collection u_1, \dots, u_T , where u_t denotes the number of shares traded at time t . Note that a negative (positive) value of $u_{i,t}$ denotes a sell (buy) trade. Our total holdings of the security at time t , x_t , equal $x_t = x_0 + \sum_{s=1}^t u_s$ due to the linear relationship in the position and trade vector, i.e., $x_t = x_{t-1} + u_t$.

We follow the discrete time model of Garleanu and Pedersen (2009) closely for the dynamics of the return and factor realizations. We assume that the price change of the security from t to

⁵Note that Garleanu and Pedersen (2009) use an infinite-horizon model without any inequality constraints, thus we will need to solve for the dynamic policy in finite-horizon and project that into the feasible space.

$t + 1$, r_{t+1} is predicted by K factors collected in a vector f_t . Furthermore, the evolution of factor realizations follow a mean reverting process. Formally, we have the following dynamics for price changes and factor realizations:

$$\begin{aligned} r_{t+1} &= Bf_t + \epsilon_{t+1}, \\ f_{t+1} &= (I - \Phi) f_t + \varepsilon_{t+1}, \end{aligned}$$

where B ($N \times K$) is a constant matrix of factor loadings, Φ ($K \times K$) is a diagonal matrix of mean reversion coefficients for the factors. We assume that the noise terms are normally distributed with the conditional covariance matrices given by $\text{Var}_t(\epsilon_{t+1}) = \Sigma$ ($N \times N$) and $\text{Var}_t(\varepsilon_{t+1}) = \Psi$ ($K \times K$), respectively.

The investor would like to find a nonanticipative policy π in the feasible set $\mathcal{U}_{\mathbb{F}}$ that selects a sequence of trades u_1, \dots, u_T in order to maximize expected total profits after transaction costs and quadratic penalty term of risk:

$$(14) \quad \begin{aligned} & \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \left\{ \mathbb{E} \left[\sum_{t=1}^T \left(x_t^\top (Bf_t) - \frac{\gamma}{2} x_t^\top \Sigma x_t - \frac{1}{2} u_t^\top \Lambda u_t \right) \right] \right\} \\ & \text{subject to } u_t = x_t - x_{t-1} \text{ (a.s.)} && t = 1, \dots, T, \\ & \quad \quad u_t \leq \mathbf{0} \text{ (a.s.)} && t = 1, \dots, T, \\ & \quad \quad x_T = \mathbf{0} \text{ (a.s.)}, \end{aligned}$$

where γ denotes the coefficient of risk aversion. Note that compared to the optimization program in Example 1, this is only different with the negativity constraints on u_t .

5.2. Model calibration

In this section, we will estimate the parameters of the optimal execution problem formulated in Section 5.1 by using intraday transactions data from the NYSE TAQ database. We will assume that there is a single stock ($N = 1$) with two return predicting factors ($K = 2$) each with a different mean reversion speed. We chose one of the most liquid stocks, Apple, Inc. (NASDAQ: AAPL), for our empirical study. We set the execution horizon to be 1 hour and trade intervals to be 5 minutes. Thus, setting a trade interval to be a one unit of time, our execution period, T , equals 12. We assume that the trader is risk-neutral, i.e., $\gamma = 0$ and the initial position to be liquidated, x_0 , is 100,000 shares.

In trade execution problems, the time horizon is typically a day, thus we will construct a factor model in the same time-frequency. We will use the transaction prices of AAPL on the trading days of January 4, 2010 (day 0) and January 5, 2010 (day 1) to estimate our factor model. We first divide each trading day into 78 buckets, i.e., one bucket for each five minutes. For each 5-minute interval, we calculate the average transaction price from all the transactions happened in that interval. Let

$p_{d,t}$ be the average price for the t th interval ($t = 1, \dots, 78$) on day d ($d = 0, 1$). Let $f_{k,t}$ be the value of the k th factor value ($k = 1, 2$) for the t th interval ($t = 2, \dots, 78$). We define our factors as follows:

$$\begin{aligned} f_{1,t} &\triangleq p_{2,t} - p_{2,t-1} & t = 2, \dots, 78, \\ f_{2,t} &\triangleq p_{2,t} - p_{1,t} & t = 2, \dots, 78. \end{aligned}$$

Here, we can interpret the factors as the representations of value and momentum signals. Intuitively, the first factor can be considered as a ‘‘momentum’’ type signal with fast mean reversion and the second factor as a ‘‘value’’ type signal with slow mean reversion.

Using the price change of the security with $r_{t+1} \triangleq p_{2,t+1} - p_{2,t}$, we can compute the estimate of the factor loading matrix, B , using the following pooled regression,

$$\begin{aligned} r_{t+1} &= 0.0726 + 0.3375 f_{1,t} - 0.0720 f_{2,t} + \epsilon_{t+1} \\ &\quad (1.96) \quad (3.11) \quad (-2.2) \end{aligned}$$

where the OLS t-statistics are reported in brackets. Thus,

$$B = \begin{bmatrix} 0.3375 & -0.072 \end{bmatrix}.$$

Similarly, we obtain the mean reversion rates for the factors:

$$\begin{aligned} \Delta f_{1,t+1} &= -0.0353 f_{1,t} + \varepsilon_{1,t+1} \\ \Delta f_{2,t+1} &= -0.7146 f_{2,t} + \varepsilon_{2,t+1}. \end{aligned}$$

Thus,

$$\Phi = \begin{bmatrix} 0.0353 & 0 \\ 0 & 0.7146 \end{bmatrix}.$$

We compute the sample variance of the error terms in the returns and factors:

$$\begin{aligned} \text{Var}(\epsilon_t) &= \Sigma = 0.0428, \\ \text{Var}(\varepsilon_t) &= \Psi = \begin{bmatrix} 0.0378 & 0 \\ 0 & 0.0947 \end{bmatrix} \end{aligned}$$

We set the distribution of the initial factor realization, f_0 , to the stationary distribution of f_t if the process were to start from $f_0 = \mathbf{0}$. Therefore, f_0 is distributed with $\mathcal{N}(\mathbf{0}, \Omega_0)$ where Ω_0 is defined as

$$\Omega_0 \triangleq \sum_{t=1}^{\infty} (I - \Phi)^t \Psi (I - \Phi)^t = \begin{bmatrix} 0.0412 & 0 \\ 0 & 1.3655 \end{bmatrix}.$$

As assumed in Garleanu and Pedersen (2009), the transaction cost matrix, Λ , is taken to be proportional to the variance of the returns, Σ . Therefore, we assume that $\Lambda = \lambda\Sigma$ where $\lambda = 0.0005$.⁶

Using these parameters, we run a simulation with 50000 trials to estimate each policy's expected payoff and standard error. In each trial, we sample f_0 , solve for the execution schedules of each approximate method and compute its corresponding payoff.

5.3. Approximate Policies

In this section, we provide four approximate solution techniques to solve the optimal execution problem in (14) with the calibrated parameters.

- **Deterministic policy (DP).** Instead of solving for a dynamic policy, we can solve for the deterministic sequence of trades at the beginning of the trading horizon. For each simulation trial, we solve a deterministic quadratic program:

$$(15) \quad \begin{aligned} & \underset{u_1, \dots, u_T}{\text{maximize}} \left\{ \sum_{t=1}^T \left(x_t^\top (B\delta_t) - \frac{1}{2} u_t^\top \Lambda u_t \right) \right\} \\ & \text{subject to } u_t = x_t - x_{t-1} && t = 1, \dots, T, \\ & u_t \leq \mathbf{0} && t = 1, \dots, T, \\ & x_T = \mathbf{0} \end{aligned}$$

where $\delta_t \triangleq E[f_t|f_0]$ and equals $(I - \Phi)^t f_0$.

- **Model predictive control (MPC).** In this approximation, at each trading time, we solve for the deterministic sequence of trades and implement only the first trade. Thus, this policy can be considered as an immediate extension of the deterministic policy with the addition of resolving at each trading time. Formally, at time t , we solve the following deterministic quadratic program, DP_t :

$$(16) \quad \begin{aligned} & \underset{u_t, \dots, u_T}{\text{maximize}} \left\{ \sum_{s=t}^T x_s^\top \left(B(I - \Phi)^{(s-t)} f_t \right) - \frac{1}{2} u_s^\top \Lambda u_s \right\} \\ & \text{subject to } u_s = x_s - x_{s-1} && s = t, \dots, T, \\ & u_s \leq \mathbf{0} && s = t, \dots, T, \\ & x_T = \mathbf{0}. \end{aligned}$$

⁶We obtain a rough estimate of λ as follows. Suppose that you pay the bid-ask spread of the stock, say 1 cent, for each share traded. Then, for a typical trade of 1000 shares, your total transaction cost is \$10. Then, according to our transaction cost model, we need $\frac{1}{2}\lambda\Sigma(1000^2) = 10$, which requires λ to equal 0.000467.

If (u_t^*, \dots, u_T^*) is the corresponding optimal solution, then the investor trades u_t^* at time t and solves DP_{t+1} in the next period.

- **Projected dynamic policy (PDP).** If the inequality constraints in (14) are relaxed, the program would reduce to the classical linear quadratic control problem:

$$(17) \quad \begin{aligned} & \underset{\pi \in \mathcal{L}_{\mathbb{F}}}{\text{maximize}} \left\{ \mathbb{E} \left[\sum_{t=1}^T \left(x_t^\top (B f_t) - \frac{1}{2} u_t^\top \Lambda u_t \right) \right] \right\} \\ & \text{subject to } u_t = x_t - x_{t-1} \text{ (a.s.)} & t = 1, \dots, T, \\ & x_T = \mathbf{0} \text{ (a.s.)}. \end{aligned}$$

The optimal dynamic policy for the program in (17) is given by

$$(18) \quad u_t = \left(\Lambda + A_{xx}^{t+1} \right)^{-1} \left(\Lambda x_{t-1} + \left(B + A_{xf}^{t+1} (I - \Phi) \right) f_t \right) - x_{t-1}$$

where the matrices A_{xx}^t and A_{xf} are defined in the Appendix. The dynamic rule for u_t in (18) which may not be feasible for the original program in (14). Thus, the projected dynamic policy seeks a trade decision, \hat{u}_t , which is the projection of u_t onto the constraint space of (14), i.e.,

$$\hat{u}_t = \max \{ -x_{t-1}, \min \{ 0, u_t \} \}$$

for each time $t < T$, and sets $\hat{u}_T = -x_{T-1}$.

- **Best linear policy (BLP).** As formulated in Proposition 1, the best linear policy represents the optimal trade vector with $u_t = c_t + \sum_{s=1}^t E_{s,t} f_s$ for $t = 1, \dots, T$. Due to the linear relationship between position and trade vectors, we can represent the position vector in the similar form, i.e., $x_t = d_t + \sum_{s=1}^t J_{s,t} f_s$ where $d_t \triangleq x_0 + \sum_{i=1}^t c_i$ and $J_{s,t} \triangleq \sum_{i=s}^t E_{s,i}$. We replace the almost-sure constraints with their probabilistic counterparts. Using these representations, we compute the parameters of the best linear policy by solving the following

stochastic program:

$$\begin{aligned}
(19) \quad & \text{maximize}_{c_t, E_{s,t}} \left\{ \mathbb{E} \left[\sum_{t=1}^T \left(\left(d_t + \sum_{s=1}^t J_{s,t} \right)^\top (Bf_t) - \frac{1}{2} \left(c_t + \sum_{s=1}^t E_{s,t} f_s \right)^\top \Lambda \left(c_t + \sum_{s=1}^t E_{s,t} f_s \right) \right) \right] \right\} \\
& \text{subject to } d_t = x_0 + \sum_{i=1}^t c_i \quad t = 1, \dots, T, \\
& J_{s,t} = \sum_{i=s}^t E_{s,i} \quad 1 \leq s \leq t \leq T, \\
& \mathbb{P} \left(\left(d_t + \sum_{s=1}^t J_{s,t} \right) \leq 0 \right) \leq \delta \quad t = 1, \dots, T, \\
& \mathbb{P} \left(\left(c_t + \sum_{s=1}^t E_{s,t} f_s \right) \geq 0 \right) \leq \delta \quad t = 1, \dots, T, \\
& d_T = 0 \text{ and } J_{s,T} = \mathbf{0}.
\end{aligned}$$

Proposition 3. *Computing the expectation in the objective function and replacing probabilistic constraints with deterministic constraints using Fact 1, we obtain the deterministic version of the stochastic program in (19), a second-order cone program:*

$$\begin{aligned}
(20) \quad & \text{maximize}_{c_t, E_{s,t}} \sum_{t=1}^T \left\{ d_t^\top B \delta_t + \sum_{s=1}^t \left(\delta_s^\top (B(I - \Phi)^{t-s} J_{s,t}^\top) \delta_s + \text{tr} \left(B(I - \Phi)^{t-s} E_{s,t}^\top \Omega_t^s \right) \right) \right. \\
& \quad \left. + \frac{1}{2} \left((c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \text{tr} \left(M_t^\top \Lambda M_t \Omega_t \right) \right) \right\} \\
& \text{subject to } d_t = x_0 + \sum_{i=1}^t c_i \quad t = 1, \dots, T, \\
& J_{s,t} = \sum_{i=s}^t E_{s,i} \quad 1 \leq s \leq t \leq T, \\
& (-d_t - P_t \theta_t) + \Phi^{-1}(1 - \delta) \left\| \left(P_t \Omega_t P_t^\top \right)^{1/2} \right\|_2 \leq 0 \quad t = 1, \dots, T, \\
& (c_t + M_t \theta_t) + \Phi^{-1}(1 - \delta) \left\| \left(M_t \Omega_t M_t^\top \right)^{1/2} \right\|_2 \leq 0 \quad t = 1, \dots, T, \\
& d_T = 0 \text{ and } J_{s,T} = \mathbf{0}.
\end{aligned}$$

where M_t and P_t are defined in (8) and (12), δ_t , θ_t and Ω_t are defined in (24) and (25). Finally, Ω_t^s is the s th diagonal block in Ω_t .

Note that the number of decision variables is considerably greater than that of the original execution problem in (14). Total number of decision variables in a problem with N securities,

K factors and T periods equals $2NT + 2NK \sum_{t=1}^T \frac{t(t+1)}{2}$ which is on the order of $O(NKT^3)$. The solution of the deterministic program in (19) provides the desired linear policy, $u_t = c_t + \sum_{s=1}^t E_{s,t} f_s$, in the return predicting factors. However, due to the probabilistic constraints, u_t may not be feasible for the original program in (14). Thus, the projected best linear policy seeks a trade decision, \hat{u}_t , which is the projection of u_t onto the constraint space of (14), i.e.,

$$\hat{u}_t = \max \{-x_{t-1}, \min \{0, u_t\}\}$$

for each time $t < T$, and sets $\hat{u}_T = -x_{T-1}$.

5.4. Upper bounds

In this section, we provide upper bounds for the optimal value of the program in (14).

- **Perfect Hindsight (PH).** In this upper bound, we compute the optimal with the perfect knowledge of the factors. For each simulation trial, we sample sequences of $\hat{f}_1, \dots, \hat{f}_T$, and remove all the stochasticity from the model. Then, we solve the following deterministic quadratic program with linear constraints:

$$(21) \quad \begin{aligned} & \underset{u_1, \dots, u_T}{\text{maximize}} \left\{ \sum_{t=1}^T \left(x_t^\top (B \hat{f}_t) - \frac{1}{2} u_t^\top \Lambda u_t \right) \right\} \\ & \text{subject to } u_t = x_t - x_{t-1} && t = 1, \dots, T, \\ & u_t \leq \mathbf{0} && t = 1, \dots, T, \\ & x_T = \mathbf{0} \end{aligned}$$

The solution of this program gives the optimal policy given full knowledge of future factor realizations for a given simulation path. By computing the average optimal value across all simulation paths, we obtain an upper bound for the original problem in (14).

- **Unprojected dynamic policy (UDP).** If the inequality constraints are relaxed in the original problem in (14), we obtain a linear quadratic control problem which we can solve exactly:

$$(22) \quad \begin{aligned} & \underset{\pi \in \mathcal{U}_E}{\text{maximize}} \left\{ \mathbb{E} \left[\sum_{t=1}^T \left(x_t^\top (B f_t) - \frac{1}{2} u_t^\top \Lambda u_t \right) \right] \right\} \\ & \text{subject to } u_t = x_t - x_{t-1} \text{ (a.s.)} && t = 1, \dots, T, \\ & x_T = \mathbf{0} \text{ (a.s.)} \end{aligned}$$

Since this is a relaxed version of the original program in (14), we obtain an upper bound for the original problem. In this case, we do not need to simulate sequences of factor realizations

to obtain the value of the upper bound. The value of the dynamic program in (22) equals $-\frac{1}{2}x_0^\top A_{xx}^0 x_0 + \frac{1}{2} \left(\text{tr}(\Omega_0 A_{ff}^0) + m_0 \right)$ where the recursions for A_{xx}^t , A_{xf}^t and m_t are defined in the appendix.

- **Pathwise optimization (PO).** Desai et al. (2011) proposes a pathwise optimization technique that can be applied for this problem as well.

5.5. Numerical results

Using the calibrated parameters from Section 5.2, we run a simulation with 50000 trials to estimate each policy’s average payoff and standard error. In each trial, we sample f_0 , solve for the execution schedules of each approximate method and compute its corresponding payoff. In order to evaluate the performance of each policy effectively, we use the same set of simulation paths in each policy’s computation of average payoff. We used CVX (see Grant and Boyd (2011)), a package for solving convex optimization problems, to solve the optimization problems in the computation of deterministic policy, model predictive control and best linear policy.

Table 1 summarizes the performance of each policy. For each policy, we divide the total payoff into two categories, alpha gains ($\sum_{t=1}^T x_t^\top B f_t$) and transaction costs ($\sum_{t=1}^T -u_t^\top \Lambda u_t$). For both categories, we report the mean value, (denoted by Avg.) and the associated standard error (denoted as S.E.). Finally, we report the average computation time (in seconds) per simulation trial for each policy evaluation.

We observe that the performance gain by the best linear policy is around 7% over the projected dynamic policy. Model predictive control achieves a closer payoff to that of the projected dynamic policy whereas deterministic policy achieves a significantly lower payoff compared to other payoffs. Since the projected dynamic policy has a closed form expression, its computation time is much smaller than the other policies. The remaining policies have roughly the same order of magnitudes in computation time with model predictive control having the longest running time.

	Best Linear			Projected Dynamic			Deterministic			Model Predictive		
	Alpha	TC	Total	Alpha	TC	Total	Alpha	TC	Total	Alpha	TC	Total
Avg.	23.24	-17.11	6.13	25.13	-19.40	5.73	19.34	-15.81	3.53	21.25	-16.54	4.71
S.E.	0.233	0.025	0.224	0.227	0.039	0.229	0.229	0.025	0.224	0.231	0.024	0.224
Time		5.40			0.0001			1.01			8.09	

Table 1: Summary of the performance statistics of each policy. For each approximate policy, we divide the total payoff into two categories, alpha gains and transaction costs. For both categories, we report the mean value, (denoted by Avg.) and the associated standard error (denoted as S.E.) which are reported in thousands of dollars. Finally, we report the average computation time (in seconds) per simulation trial for each policy evaluation. We observe that the best linear policy achieves a higher payoff than projected dynamic policy and model predictive control.

Despite the higher average payoff, relatively high standard errors impedes the immediate conclusion that best linear achieves a statistically significant higher payoff. Thus, in order to compare

the performances of the best linear policy and projected dynamic policy in detail, for each simulation trial, we compute the differences in alpha gains and transaction costs between these policies. Table 2 illustrates that best linear policy performs statistically better than the projected dynamic policy. We observe that the standard error for the difference in total profits is much smaller than the reported standard errors of the total profits in Table 1. In conclusion, this results emphasizes further that the performance gain with using the best linear policy is statistically significant.

	BLP-PDP		
	Alpha	TC	Total
Avg.	-1.89	2.29	0.40
S.E.	0.0137	0.0196	0.0095

Table 2: The detailed comparison between the alpha gains and transaction costs of the best linear policy and projected dynamic policy. We observe that the standard error for the difference in total profits is very small, thus, the performance gain by employing the best linear policy is statistically significant.

Finally, Table 3 summarizes the performance statistics for each upper bound policy. We observe that perfect hindsight and pathwise optimization provide much tighter bounds than the unprojected dynamic policy with pathwise optimization yielding the tightest bound. Comparing the tightest bound to the performance of the best linear policy, we conclude that the optimality gap by employing the best linear policy is fairly small, at most 5% of the optimal value of the original program in (14), which indicates that the best linear policy is nearly optimal.

	Perfect Hindsight			Unprojected Dynamic	Pathwise Optimization
	Alpha	TC	Total	Total	Total
Avg.	26.28	-17.70	8.57	12.58	6.46
S.E.	0.231	0.0271	0.223	0.00	0.04
Time		0.87		0.0001	N/A

Table 3: Summary of the performance statistics for each upper bound policy. We observe that pathwise optimization provides the tightest bound which is very close to the payoff of the best linear policy. This indicates that the best linear policy is nearly optimal.

6. Conclusion

This paper provides a highly tractable rebalancing rule for dynamic portfolio choice problems with return predictability and transaction costs. Our rebalancing rule is a linear function of return predicting factors and can be utilized in a wide spectrum of portfolio choice models with realistic considerations for risk measures, transaction costs and constraints. As long as the starting dynamic portfolio optimization problem is a convex programming problem, the modified optimization problem seeking the optimal parameters of the linear decision rule will be a convex programming problem.

We provide a large class of dynamic portfolio choice models that differ in their modeling of risk measures, transaction costs and constraints which can be formulated as deterministic convex optimization problems. Specifically, we compute the analytic expression of the objective function in the cases with quadratic utility function on the terminal wealth or proportional and nonlinear transaction cost functions. Finally we derive efficient formulations for incorporating linear equality and inequality constraints. If there does not exist an analytic expression for the objective, the optimal parameters can be solved via the sampling techniques available from the sample average and stochastic approximation literature.

Finally, we implement the computation of the best linear policy in the context of portfolio execution, the execution of a large long position in a single security. For this purpose, we need positivity constraints on portfolio positions and the amount of shares sold in each period in order to achieve a feasible execution. In order to compare the performance of the best linear rebalancing rule, we use the identical discrete-time setup of Garleanu and Pedersen (2009) for which a closed-form solution is available in the lack of constraints. We calibrate the model parameters using two-days of transactions data on a liquid stock and construct two predictors in a high-frequency setting with different mean reversion speeds. The simulation implemented with these predictors and calibrated parameters reveal that the best linear policy performs better than the deterministic policy, model predictive control and a projected version of the optimal policy proposed by Garleanu and Pedersen (2009).

References

- R. F. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2): 5–39, 2000.
- Pierluigi Balduzzi and Anthony W. Lynch. Transaction costs and predictability: some utility cost calculations. *Journal of Financial Economics*, 52(1):47–78, apr 1999.
- A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.
- Aharon Ben-Tal, Boaz Golany, Arkadi Nemirovski, and Jean-Philippe Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7:248–271, July 2005.
- D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2nd edition, 2000. ISBN 1886529094.
- D. Bertsimas and A. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1(1): 1–50, 1998. ISSN 1386-4181.
- Dimitris Bertsimas and Vineet Goyal. On the power and limitations of affine policies in two-stage adaptive optimization. *Mathematical Programming*, pages 1–41, 2011.
- Dimitris Bertsimas, Dan A. Iancu, and Pablo A. Parrilo. Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.*, 35, May 2010.
- Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, March 2004. ISBN 0521833787.
- Michael W. Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22(9): 3411–3447, September 2009.
- David B. Brown and James E. Smith. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. Working paper, 2010.
- Giuseppe Carlo Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control Optim.*, 48:2254–2274, June 2009.
- N. Chen, R. Roll, and S. A. Ross. Economic forces and the stock market. *The Journal of Business*, 59(3):pp. 383–403, 1986.
- X. Chen, M. Sim, and P. Sun. A robust optimization perspective on stochastic programming. *Operations Research*, 55:1058–1071, November 2007.

- X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56:344–357, March 2008.
- G.M. Constantinides. Capital market equilibrium with transaction costs. *The Journal of Political Economy*, 94(4):pp. 842–862, August 1986.
- J. Cvitanic. Theory of portfolio optimization in markets with frictions. In E. Jouini, J. Cvitanic, and Marek Musiela, editors, *Handbooks in Mathematical Finance*. Cambridge University Press, 2001.
- M. H. A. Davis and A. R. Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, 1990.
- V.V. Desai, V.F. Farias, and C.C. Moallemi. Pathwise optimization for linear convex systems. *Working paper*, 2011.
- Bernard Dumas and Elisa Luciano. An exact solution to a dynamic portfolio choice problem under transactions costs. *The Journal of Finance*, 46(2):pp. 577–595, 1991.
- E. F. Fama and K. R. French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1):pp. 55–84, 1996.
- N. Garleanu and L.H. Pedersen. Dynamic trading with predictable returns and transaction costs. *Working paper*, July 2009.
- W. N. Goetzmann and P. Jorion. Testing the predictive power of dividend yields. *Journal of Finance*, 48(2):663–679, 1993.
- M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. <http://cvxr.com/cvx>, April 2011.
- S.L. Heston, R.A. Korajczyk, and R. Sadka. Intraday patterns in the cross-section of stock returns. *Journal of Finance*, 65(4):1369–1407, 2010.
- Merrell Hora. Tactical liquidity trading and intraday volume. *Working Paper*, 2006.
- D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, pages 1–33, 2009.
- Hong Liu. Optimal consumption and investment with transaction costs and multiple risky assets. *Journal of Finance*, 59(1):289–338, 02 2004.
- Anthony W. Lynch and Pierluigi Balduzzi. Predictability and transaction costs: The impact on rebalancing rules and behavior. *Journal of Finance*, 55(5):2285–2309, October 2000.

- Anthony W. Lynch and Sinan Tan. Multiple risky assets, transaction costs, and return predictability: Allocation rules and implications for u.s. investors. *Journal of Financial and Quantitative Analysis*, 45(04):1015–1053, August 2010.
- Robert C. Merton. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4):373–413, December 1971.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- A. A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. Working paper, 2005.
- A. Shapiro. Monte carlo sampling methods. In A. Ruszczyński and A. Shapiro, editors, *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*, pages 353 – 425. Elsevier, 2003.
- A. Shapiro and A. Nemirovski. On complexity of stochastic programming problems. In Panos M. Pardalos, Donald Hearn, Vaithilingam Jeyakumar, and Alexander Rubinov, editors, *Continuous Optimization*, volume 99 of *Applied Optimization*, pages 111–146. 2005.
- S. E. Shreve and H. M. Soner. Optimal investment and consumption with transaction costs. *The Annals of Applied Probability*, 4(3):pp. 609–692, 1994.
- J. Skaf and S.P. Boyd. Design of affine controllers via convex optimization. *Automatic Control, IEEE Transactions on*, 55(11):2476 –2487, November 2010.
- V. Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.

A. Proofs

Proof of Proposition 2 Since $x_t = d_t + \sum_{s=1}^t J_{s,t} f_s$,

$$\begin{aligned}
E[W_T] &= \sum_{t=1}^T \left(\kappa_t \mu_t + \mathbb{E} \left[d_t^\top B f_t + \sum_{s=1}^t f_s^\top J_{s,t}^\top B f_t \right] \right) \\
&= \sum_{t=1}^T \left(\kappa_t \mu_t + d_t^\top B \delta_t + \sum_{s=1}^t \mathbb{E} \left[f_s^\top J_{s,t}^\top B \mathbb{E} [f_t | f_s] \right] \right) \\
&= \sum_{t=1}^T \left(\kappa_t \mu_t + d_t^\top B \delta_t + \sum_{s=1}^t \mathbb{E} \left[f_s^\top J_{s,t}^\top B (I - \Phi)^{t-s} f_s \right] \right) \\
&= \sum_{t=1}^T \left(\kappa_t \mu_t + d_t^\top B \delta_t + \sum_{s=1}^t \left(\delta_s^\top (B(I - \Phi)^{t-s} J_{s,t}^\top) \delta_s + \text{tr} \left(B(I - \Phi)^{t-s} J_{s,t}^\top \Omega_t^s \right) \right) \right),
\end{aligned}$$

where the mean vector of x_t , κ_t , is defined in (11). For the transaction cost terms, we use the representation, $u_t = d_t + M_t F_t$, and then,

$$\begin{aligned}
E[C_T] &= \frac{1}{2} \sum_{t=1}^T \mathbb{E} [u_t^\top \Lambda u_t] = \frac{1}{2} \sum_{t=1}^T \mathbb{E} [(c_t + M_t F_t)^\top \Lambda (c_t + M_t F_t)] \\
&= \frac{1}{2} \sum_{t=1}^T (c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \text{tr} (\Lambda M_t \Omega_t M_t^\top) \\
&= \frac{1}{2} \sum_{t=1}^T (c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \left\| \left(\Lambda^{\frac{1}{2}} M_t \Omega_t^{\frac{1}{2}} \right) \right\|_F^2.
\end{aligned}$$

For the penalty term

$$\begin{aligned}
E[W_T^2] &= \sum_{t=1}^T \sum_{k=1}^T \mathbb{E} [(x_t^\top r_{t+1})(x_k^\top r_{k+1})] \\
&= \sum_{t=1}^T \mathbb{E} [x_t^\top \epsilon_{t+1} \epsilon_{t+1}^\top x_t] + \sum_{t=1}^T \sum_{k=1}^T \mathbb{E} [x_t^\top B f_t f_k^\top B^\top x_k] \\
&= \sum_{t=1}^T \mathbb{E} [x_t^\top \epsilon_{t+1} \epsilon_{t+1}^\top x_t + x_t^\top B f_t f_t^\top B^\top x_t] + 2 \sum_{t=1}^T \sum_{k=1}^t \mathbb{E} [x_t^\top B f_t f_k^\top B^\top x_k]
\end{aligned}$$

First, note

$$\begin{aligned}
\mathbb{E} [x_t^\top \epsilon_{t+1} \epsilon_{t+1}^\top x_t] &= \mathbb{E} [x_t^\top \mathbb{E} [\epsilon_{t+1} \epsilon_{t+1}^\top | \mathcal{F}_t] x_t] \\
&= \mathbb{E} [x_t^\top \Sigma x_t] \\
&= (d_t + P_t \theta_t)^\top \Sigma (d_t + P_t \theta_t) + \text{tr} (\Sigma P_t \Omega_t P_t^\top).
\end{aligned}$$

Finally,

$$\begin{aligned}
\mathbb{E} \left[x_t^\top B f_t f_t^\top B^\top x_t \right] &= \mathbb{E} \left[\left(d_t + \sum_{s=1}^t J_{s,t} f_s \right)^\top B f_t f_t^\top B^\top \left(d_t + \sum_{s=1}^t J_{s,t} f_s \right) \right] \\
&= \mathbb{E} \left\{ \left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} ((I - \Phi) f_{t-1} + \varepsilon_t) \right)^\top B ((I - \Phi) f_{t-1} + \varepsilon_t) ((I - \Phi) f_{t-1} + \varepsilon_t)^\top B^\top \right. \\
&\quad \left. \left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} ((I - \Phi) f_{t-1} + \varepsilon_t) \right) \right\} \\
&= \mathbb{E} \left[\left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} (I - \Phi) f_{t-1} \right)^\top B (I - \Phi) f_{t-1} (B (I - \Phi) f_{t-1})^\top \left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} (I - \Phi) f_{t-1} \right) \right] \\
&\quad + \mathbb{E} \left[\left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} (I - \Phi) f_{t-1} \right)^\top \varepsilon_t \varepsilon_t^\top \left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} (I - \Phi) f_{t-1} \right) \right] \\
&\quad + \mathbb{E} \left[\varepsilon_t^\top B (I - \Phi) f_{t-1} (B (I - \Phi) f_{t-1})^\top \varepsilon_t \right] + \mathbb{E} \left[\varepsilon_t^\top \varepsilon_t \varepsilon_t^\top \varepsilon_t \right] \\
&= \mathbb{E} \left[\left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} (I - \Phi) f_{t-1} \right)^\top B (I - \Phi) f_{t-1} (B (I - \Phi) f_{t-1})^\top \left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t} (I - \Phi) f_{t-1} \right) \right] \\
&\quad + \left(d_t + \bar{P}_{t-1} \theta_{t-1} \right)^\top \Psi \left(d_t + \bar{P}_{t-1} \theta_{t-1} \right) + \text{tr} \left(\bar{P}_{t-1} \Omega_{t-1} \bar{P}_{t-1}^\top \right) \text{tr} \left(B (I - \Phi) \Omega_t^{t-1} (B (I - \Phi))^\top \Psi \right) + \psi^2
\end{aligned}$$

where in the last step we use iterated expectations by conditioning on the information up to $t-1$, \mathcal{F}_{t-1} , and we define

$$(23) \quad \bar{P}_{t-i} \triangleq \begin{bmatrix} J_{1,t} & J_{2,t} & \dots & J_{t-i-1,t} & \bar{J}_{t-i,t} \end{bmatrix},$$

with

$$\bar{J}_{t-i,t} \triangleq \left(\sum_{k=0}^i J_{t-k,t} (I - \Phi)^{i-k} \right),$$

and

$$\psi^2 \triangleq \sum_{i=1}^K 3\Psi_{ii} + 2 \sum_{i=1}^K \sum_{j=1}^i \Psi_{ij},$$

where Ψ_{ij} is the (i, j) th entry of the covariance matrix of the error terms for factor dynamics, Ψ . Thus, by conditioning recursively, we obtain,

$$\begin{aligned}
\mathbb{E} \left[x_t^\top B f_t f_t^\top B^\top x_t \right] &= \left(d_t + \sum_{s=1}^t J_{s,t} (I - \Phi)^s f_0 \right)^\top B (I - \Phi)^t f_0 (B (I - \Phi)^t f_0)^\top \left(d_t + \sum_{s=1}^t J_{s,t} (I - \Phi)^s f_0 \right) \\
&\quad + \sum_{i=1}^t \left(\left(d_t + \bar{P}_{t-i} \theta_{t-i} \right)^\top \Psi \left(d_t + \bar{P}_{t-i} \theta_{t-i} \right) + \text{tr} \left(\bar{P}_{t-i} \Omega_{t-i} \bar{P}_{t-i}^\top \right) \text{tr} \left(B (I - \Phi)^i \Omega_t^{t-i} (B (I - \Phi)^i)^\top \Psi \right) + \psi^2 \right)
\end{aligned}$$

The last term, $\mathbb{E} \left[x_t^\top B f_t f_k^\top B^\top x_k \right]$, can also be computed similarly.

Derivation of the Projected and Unprojected Dynamic Policies Using dynamic programming principle and $u_t = (x_t - x_{t-1})$, the value function $V_t(x_{t-1}, f_t)$ satisfies

$$V_{t-1}(x_{t-1}, f_t) = \underset{x_t}{\text{maximize}} \left(x_t^\top (B f_t) - \frac{1}{2} (x_t - x_{t-1})^\top \Lambda (x_t - x_{t-1}) + \mathbb{E}[V_t(x_t, f_{t+1})] \right).$$

We guess the following quadratic form for our value function:

$$V_t(x_t, f_{t+1}) = -\frac{1}{2} x_t^\top A_{xx}^t x_t + x_t^\top A_{xf}^t f_{t+1} + \frac{1}{2} f_{t+1}^\top A_{ff}^t f_{t+1} + \frac{1}{2} m_t.$$

Then,

$$\mathbb{E}[V_t(x_t, f_{t+1})] = -\frac{1}{2} x_t^\top A_{xx}^t x_t + x_t^\top A_{xf}^t (I - \Phi) f_t + \frac{1}{2} f_t^\top (I - \Phi) A_{ff}^t (I - \Phi) f_t + \frac{1}{2} \left(\text{tr}(\Psi A_{ff}^t) + m_t \right).$$

At the the last period, we need $x_T = 0$, and our value function equals

$$V_{T-1}(x_{T-1}, f_t) = -\frac{1}{2} x_{T-1}^\top \Lambda x_{T-1}$$

which satisfies our functional form with

$$A_{xx}^{T-1} = \Lambda \quad A_{xf}^{T-1} = \text{zero}(N, K) \quad A_{ff}^{T-1} = \text{zero}(K, K) \quad m_{T-1} = 0$$

where $\text{zero}(m, n)$ denotes a matrix of size $(m \times n)$ with each entry equaling zero.

For all $t < T - 1$, we maximize the quadratic objective $-\frac{1}{2} x_t^\top Q_t x_t + x_t^\top q_t + b_t$ where

$$\begin{aligned} Q_t &= \Lambda + A_{xx}^t \\ q_t &= \Lambda x_{t-1} + \left(B + A_{xf}^t (I - \Phi) \right) f_t \\ b_t &= -\frac{1}{2} x_{t-1}^\top \Lambda x_{t-1} + \frac{1}{2} f_t^\top (I - \Phi) A_{ff}^t (I - \Phi) f_t + \text{tr}(\Psi A_{ff}^t) + m_t \end{aligned}$$

Then, the optimal x_t is given by $Q_t^{-1} q_t$ and x_t and u_t are given by

$$\begin{aligned} x_t &= \left(\Lambda + A_{xx}^t \right)^{-1} \left(\Lambda x_{t-1} + \left(B + A_{xf}^t (I - \Phi) \right) f_t \right) \\ u_t &= \left(\Lambda + A_{xx}^t \right)^{-1} \left(\Lambda x_{t-1} + \left(B + A_{xf}^t (I - \Phi) \right) f_t \right) - x_{t-1} \end{aligned}$$

The maximum then occurs at $\frac{1}{2}q_t^\top Q_t^{-1}q_t + b_t$ and we obtain the following recursions:

$$\begin{aligned}
A_{xx}^{t-1} &= -\Lambda \left(\Lambda + A_{xx}^t \right)^{-1} \Lambda + \Lambda \\
A_{xf}^{t-1} &= \Lambda \left(\Lambda + A_{xx}^t \right)^{-1} \left(B + A_{xf}^{t-1} (I - \Phi) \right) \\
A_{ff}^{t-1} &= \left(B + A_{xf}^t (I - \Phi) \right)^\top \left(\Lambda + A_{xx}^{t-1} \right)^{-1} \left(B + A_{xf}^t (I - \Phi) \right) + (I - \Phi) A_{ff}^t (I - \Phi) \\
m_{t-1} &= \text{tr}(\Psi A_{ff}^t) + m_t
\end{aligned}$$

Using these recursions, we can compute the optimal expected payoff of the dynamic program. Using $f_0 = N(0, \Omega)$,

$$\begin{aligned}
\mathbb{E}[V_0(x_0, f_1)] &= \mathbb{E}[\mathbb{E}[V_0(x_0, f_1)|f_0]] \\
&= \mathbb{E} \left[-\frac{1}{2}x_0^\top A_{xx}^0 x_0 + x_0^\top A_{xf}^0 (I - \Phi) f_0 + \frac{1}{2}f_0^\top (I - \Phi) A_{ff}^0 (I - \Phi) f_0 + \frac{1}{2} \left(\text{tr}(\Omega A_{ff}^0) + m_0 \right) \right] \\
&= -\frac{1}{2}x_0^\top A_{xx}^0 x_0 + \frac{1}{2} \left(\text{tr}(\Omega A_{ff}^0 (I - \Phi)) + \text{tr}(\Psi A_{ff}^0) + m_0 \right).
\end{aligned}$$

Proof of Proposition 3 We first compute the objective function. For each t , we have to compute the expectation of the following two terms, $\mathbb{E} \left[x_t^\top (B f_t) \right]$, and $\mathbb{E} \left[u_t^\top \Lambda u_t \right]$. First, we derive the statistics for f_t , u_t and x_t . We first note that

$$f_t = (I - \Phi)^t f_0 + \sum_{s=1}^t (I - \Phi)^{t-s} \varepsilon_s.$$

Letting $F_t \triangleq (f_1, \dots, f_t)^\top$, Then, in vector form, we have the following representation

$$F_t = \begin{bmatrix} (I - \Phi)f_0 \\ (I - \Phi)^2 f_0 \\ \vdots \\ (I - \Phi)^{t-1} f_0 \\ (I - \Phi)^t f_0 \end{bmatrix} + \underbrace{\begin{bmatrix} I & 0 & \dots & 0 & 0 \\ (I - \Phi) & I & 0 & \dots & 0 \\ \vdots & (I - \Phi) & \ddots & \dots & 0 \\ (I - \Phi)^{t-1} & \dots & \ddots & I & 0 \\ (I - \Phi)^t & \dots & & (I - \Phi) & I \end{bmatrix}}_{\triangleq A_t} \begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_{t-1} \\ \varepsilon_t \end{bmatrix}.$$

Using this representation, we compute the mean

$$(24) \quad \theta_t \triangleq \mathbf{E}[F_t] = \begin{bmatrix} \delta_1 \\ \delta_2 \\ \vdots \\ \delta_{t-1} \\ \delta_t \end{bmatrix} \triangleq \begin{bmatrix} (I - \Phi)f_0 \\ (I - \Phi)^2 f_0 \\ \vdots \\ (I - \Phi)^{t-1} f_0 \\ (I - \Phi)^t f_0, \end{bmatrix}$$

and the covariance matrix

$$(25) \quad \Omega_t \triangleq \text{Var}[F_t] = A_t \begin{bmatrix} \Psi & 0 & \dots & 0 & 0 \\ 0 & \Psi & & \dots & 0 \\ \vdots & & \ddots & \dots & 0 \\ & \dots & & \Psi & 0 \\ 0 & \dots & & 0 & \Psi \end{bmatrix} A_t^\top.$$

Note that Ω_t is a block diagonal matrix with t blocks of size $K \times K$. Recall that in Section 4, we defined

$$(26) \quad M_t \triangleq \begin{bmatrix} E_{1,t} & E_{2,t} & \dots & E_{t,t} \end{bmatrix}$$

Then, $u_t = c_t + M_t F_t$ and we have the following moments for u_t :

$$(27) \quad \begin{aligned} \mu_t &\triangleq \mathbf{E}(u_t) = c_t + M_t \theta_t \\ V_t &\triangleq \text{Var}(u_t) = M_t \Omega_t M_t^\top. \end{aligned}$$

Therefore, u_t is normally distributed with mean μ_t and covariance matrix V_t . Similarly, we can obtain the statistics for x_t . Using (11),

$$\begin{aligned} \kappa_t &\triangleq \mathbf{E}(x_t) = d_t + P_t \theta_t \\ Y_t &\triangleq \text{Var}(x_t) = P_t \Omega_t P_t^\top. \end{aligned}$$

We note the following fact from multivariate statistics.

Fact 2. *If z is a random vector with mean μ and variance Σ , and Q is positive definite matrix, then*

$$\begin{aligned} \mathbf{E}(z' Q z) &= \text{tr}(Q \Sigma) + \mu' Q \mu \\ &= \|Q^{\frac{1}{2}} \Sigma^{\frac{1}{2}}\|_F^2 + \mu' Q \mu \end{aligned}$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

We can now compute each term in the objective function.

$$\begin{aligned}
\mathbb{E} \left[x_t^\top (Bf_t) \right] &= \mathbb{E} \left[d_t^\top Bf_t + \sum_{s=1}^t f_s^\top J_{s,t}^\top Bf_t \right] \\
&= d_t^\top B\delta_t + \sum_{s=1}^t \mathbb{E} \left[f_s^\top J_{s,t}^\top B \mathbb{E}[f_t | f_s] \right] \\
&= d_t^\top B\delta_t + \sum_{s=1}^t \mathbb{E} \left[f_s^\top J_{s,t}^\top B (I - \Phi)^{t-s} f_s \right] \\
&= d_t^\top B\delta_t + \sum_{s=1}^t \left(\delta_s^\top (B(I - \Phi)^{t-s} J_{s,t}^\top) \delta_s + \text{tr} \left(B(I - \Phi)^{t-s} J_{s,t}^\top \Omega_t^s \right) \right)
\end{aligned}$$

where Ω_t^s is the s th diagonal block in Ω_t . Finally, for the transaction cost terms,

$$\begin{aligned}
\mathbb{E} \left[u_t^\top \Lambda u_t \right] &= \mathbb{E} \left[(c_t + M_t F_t)^\top \Lambda (c_t + M_t F_t) \right] \\
&= (c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \text{tr} \left(\Lambda M_t \Omega_t M_t^\top \right) \\
&= (c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \left\| \left(\Lambda^{\frac{1}{2}} M_t \Omega_t^{\frac{1}{2}} \right) \right\|_F^2
\end{aligned}$$

Summing up all the terms, our final objective becomes

$$\begin{aligned}
\underset{c_t, E_{s,t}}{\text{maximize}} \sum_{t=1}^T \left\{ d_t^\top B\delta_t + \sum_{s=1}^t \left(\delta_s^\top (B(I - \Phi)^{t-s} J_{s,t}^\top) \delta_s + \text{tr} \left(B(I - \Phi)^{t-s} E_{s,t}^\top \Omega_t^s \right) \right) \right. \\
\left. + \frac{1}{2} \left((c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \left\| \left(\Lambda^{\frac{1}{2}} M_t \Omega_t^{\frac{1}{2}} \right) \right\|_F^2 \right) \right\}
\end{aligned}$$

In order to enforce $x_T = 0$, we need

$$\begin{aligned}
d_T &= 0 \\
J_{s,T,k}^\top &= \mathbf{0} \quad \forall s = 1, \dots, T
\end{aligned}$$

where $J_{s,T,k}$ is the k th row of $J_{s,T}$.

For the probabilistic constraints, we use Fact 1. Note that $\mathbb{P}(x_t \leq 0) \leq \delta$ can be written as $\mathbb{P}(-x_t \geq 0) \leq \delta$ which equals by Fact 1,

$$(-d_t - P_t \theta_t) + \Phi^{-1}(1 - \delta) \left\| \left(P_t \Omega_t P_t^\top \right)^{1/2} \right\|_2 \leq 0.$$

Similarly, we obtain that $\mathbb{P}(u_t \geq 0) \leq \delta$ can be written as

$$(c_t + M_t \theta_t) + \Phi^{-1}(1 - \delta) \left\| \left(M_t \Omega_t M_t^\top \right)^{1/2} \right\|_2 \leq 0.$$