JOURNAL OF FINANCIAL AND QUANTITATIVE ANALYSIS Vol. 52, No. 3, June 2017, pp. 1247–1278
COPYRIGHT 2017, MICHAEL G. FOSTER SCHOOL OF BUSINESS, UNIVERSITY OF WASHINGTON, SEATTLE, WA 98195
doi:10.1017/S0022109017000345

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

Dynamic Portfolio Choice with Linear Rebalancing Rules

Ciamac C. Moallemi and Mehmet Sağlam*

Abstract

We consider a broad class of dynamic portfolio optimization problems that allow for complex models of return predictability, transaction costs, trading constraints, and risk considerations. Determining an optimal policy in this general setting is almost always intractable. We propose a class of linear rebalancing rules and describe an efficient computational procedure to optimize with this class. We illustrate this method in the context of portfolio execution and show that it achieves near optimal performance. We consider another numerical example involving dynamic trading with mean-variance preferences and demonstrate that our method can result in economically large benefits.

I. Introduction

Dynamic portfolio optimization has been a central and essential objective for institutional investors in active asset management. Real-world portfolio allocation problems of practical interest have a number of common objectives. First, predicting future asset returns lies at the heart of active portfolio management. Such predictions are not limited to simple unconditional estimates of expected future returns but often involve predictions on short- and long-term expected returns using complex models based on observable return-predicting factors. Second, the fund manager aims to keep the transaction costs at a minimum. These costs can arise from sources ranging from the bid–offer spread or execution commissions to price impact, where the manager's own trading affects the subsequent evolution of prices. Third, managers oftentimes cannot make arbitrary investment decisions but, rather, face exogenous constraints on their trades or their resulting portfolios. Examples include short-sale constraints, leverage constraints, and

*Moallemi, ciamac@gsb.columbia.edu, Graduate School of Business, Columbia University; Sağlam (corresponding author), mehmet.saglam@uc.edu, Lindner College of Business, University of Cincinnati. We are grateful for helpful comments from an anonymous referee, David Brown, Stephen Brown (the editor), Sylvain Champonnois (discussant), Michael Sotiropoulos, and conference participants at the 2011 Annual Conference on Advances in the Analysis of Hedge Fund Strategies at Imperial College London. Moallemi acknowledges the support of National Science Foundation (NSF) grant CMMI-1235023. Sağlam acknowledges support from the Eugene M. Lang Doctoral Student Grant.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

restrictions requiring market neutrality (or specific-industry neutrality). Finally, portfolio managers seek to control the risk of their portfolios. In practical settings, risk aversion is not accomplished by the specification of an abstract utility function. Rather, managers specify limits or penalties for multiple summary statistics that capture aspects of portfolio risk that are easy to interpret and known to be important. For example, a manager could be interested both in the change in the risk of the portfolio value over various intervals of time, including, for example, short intervals (e.g., daily or weekly risk), and in the risk associated with the portfolio's terminal value. Such single-period risk can be measured in a number of ways (e.g., variance, value at risk). A manager could be further interested in multiperiod measures of portfolio risk, such as the portfolio's maximum drawdown.

The fact that the underlying problem is multiperiod, however, significantly complicates the analysis of portfolio choice. Generally, in this case, the decision made by a manager at a given instant of time could depend on all information realized up to that point. Traditional approaches to multiperiod portfolio choice, dating back at least to the work of Merton (1971), have focused on analytically determining the optimal dynamic policy. Although this work has brought forth important structural insights, it is fundamentally quite restrictive: Exact analytical solutions require very specific assumptions about investor objectives and market dynamics. These assumptions cannot accommodate flexibility in, for example, the return-generating process, trading frictions, and constraints and are often practically unrealistic. Absent such restrictive assumptions, analytical solutions are not possible. Motivated by this problem, much of the subsequent academic literature on portfolio choice seeks to develop modeling assumptions that allow for analytical solutions; however, the resulting formulations are often not representative of real-world problems of practical interest. Further, because of the "curse of dimensionality," exact numerical solutions are often intractable in cases of practical interest, where the universe of tradable assets is large.

In search of tractable alternatives, many practitioners eschew multiperiod formulations. Instead, they consider portfolio choice problems within a myopic, single-period setting when the underlying application is clearly multiperiod (e.g., Grinold and Kahn (1999)). Another tractable possibility is to consider portfolio choice problems that are multiperiod but without the possibility of recourse. In this case, a fixed set of deterministic decisions for the entire time horizon is made at the initial time. Both single-period and deterministic portfolio choice formulations are quite flexible and can accommodate many of the features just described. They are typically applied in a quasi-dynamic fashion through the method of model predictive control (MPC). Here, at each time period, the simplified portfolio choice problem is re-solved based on the latest available information.

Although these simplified approaches are extremely flexible and have been broadly adopted in practice, these methods have important flaws. Generally, such methods are heuristics; to achieve tractability, they neglect the explicit consideration of the possibility of future recourse. Hence, these methods could be significantly suboptimal. Moreover, single-period formulations, which are the most popular among practitioners, pose a number of additional challenges. Generally, they do not effectively manage transaction costs; re-solving a single-period model repeatedly causes portfolio churn. They are also difficult to apply in situations in

https://doi.org/10.1017/S0021109017000345 Published online by Cambridge University Press

which returns are predicted across multiple time horizons. Ideally, an investor should be very responsive to short-term predictions that will be realized quickly and respond less aggressively to long-term predictions where there is time to work into a position. It is not clear how to accommodate this scenario in a single-period setting that allows only a single choice of time horizon. Generally, practitioners adopt ad hoc heuristics to address these issues. For example, one can introduce artificial transaction costs to limit portfolio churn or artificially scale return predictors based on their relative horizons.

Another tractable alternative is the use of linear quadratic control in the formulation of portfolio choice problems (e.g., Hora (2006), Gârleanu and Pedersen (2013)). Linear quadratic control (LQC) problems have been an important class of tractable multiperiod optimal control problems since the 1950s. In the portfolio choice setting, these methods apply if the return dynamics are linear, transaction costs and risk-aversion penalties can be decomposed into per-period quadratic functions, and security holdings and trading decisions are unconstrained. However, many important problem cases simply do not fall into the linear quadratic framework.

In this paper, our central innovation is to propose a framework for multiperiod portfolio optimization that allows a broad class of problems, including many of the features described earlier. Our formulation maintains tractability by restricting the problem to determining the best policy out of a restricted class of linear rebalancing policies. Such policies allow planning for future recourse, but only of a form that can be parsimoniously parameterized in a specific affine fashion. In particular, the contributions of this paper are as follows: First, we define a flexible, general setting for portfolio optimization. Our setting allows for very general dynamics of asset prices, with arbitrary dependence on the history of "return-predictive factors." We allow for any convex constraints on trades and positions. Finally, the objective is allowed to be an arbitrary concave function of the sample path of positions. Our framework allows, for example, many complex models for transaction costs or risk aversion. We can consider both traditional problem formulations for portfolio optimization (e.g., maximization of the expected terminal utility of wealth) and formulations more popular among practitioners (e.g., maximization of expected wealth subject to risk constraints).

Second, our portfolio optimization problem is computationally tractable. In our setting, determination of the optimal linear rebalancing policy is a convex program. Convexity guarantees that the globally optimal policy can be generally tractably found. This is in contrast to nonconvex portfolio choice parameterizations (e.g., Brandt, Santa-Clara, and Valkanov (2009)), where only local optimality can be guaranteed.

In our case, numerical solutions can be obtained via, for example, sample average approximation or stochastic approximation methods (e.g., Shapiro (2003), Nemirovski, Juditsky, Lan, and Shapiro (2009)). These methods can be applied in a data-driven fashion, with access only to simulated trajectories and without an explicit model of system dynamics. In a number of instances in which the factor and return dynamics are driven by Gaussian uncertainty, we illustrate that our portfolio optimization problem can be reduced to a standard form of convex

optimization program, which can be solved with off-the-shelf commercial optimization solvers.

Third, our class of linear rebalancing policies subsumes many common heuristic portfolio policies. Both single-period and deterministic policies are special cases of linear rebalancing policies; however, linear rebalancing policies are a broader class. Hence, the optimal linear rebalancing policy will outperform policies from these more restricted classes. Further, our method can also be applied in the context of MPC. In addition, portfolio optimization problems that can be formulated as LQC problems also fit in our setting, and their optimal policies are linear rebalancing rules.

Finally, we demonstrate the practical benefits of our method with two examples: optimal execution with trading constraints and dynamic trading with mean-variance preferences. First, we consider an optimal execution problem where an investor seeks to liquidate a position over a fixed time horizon in the presence of transaction costs and a model for predicting returns. We further introduce linear inequality constraints that require the trading decisions to be only sales. Such sale-only constraints are common in agency algorithmic trading. The resulting optimal execution problem does not allow for an exact solution. Hence, we compare the best linear policy to a number of tractable alternative approximate policies, including a deterministic policy, MPC, and a projected variation of the LQC formulation of Gârleanu and Pedersen (2013). We demonstrate that the performance of the best linear policy is superior to that of the alternatives. Moreover, we compute a number of upper bounds on the performance of any policy for the problem at hand. Using these upper bounds, we see that the best linear policy is near optimal, with an optimality gap of 5% at most. Our sensitivity analysis shows that the improvement obtained using linear rebalancing rules can be as high as 18% compared with the best alternative policy. Second, we consider a dynamic trading problem where an investor with mean-variance preferences makes intraday trading decisions in the presence of return predictability. Using the same model calibration as that in the optimal execution example, we illustrate that the gains from using our best linear policy can be economically substantial when the model does not fall within the realm of a linear quadratic formulation. Moreover, our sensitivity analysis reveals that this outperformance is robust to different model calibrations and can provide a 72% improvement when benchmarked against a trading rule based on a linear quadratic formulation.

Literature Review

Our paper is related to two different strands of literature: the literature on dynamic portfolio choice with return predictability and transaction costs and the literature on the use of linear decision rules in optimal control problems.

First, we consider the literature on dynamic portfolio choice. This vast body of work begins with the seminal paper of Merton (1971). Following that, a significant literature has aimed to incorporate the impact of various frictions, such as transaction costs, on optimal portfolio choice.[1] Liu and Loewenstein (2002)

[1]The work of Constantinides (1986) is an early example that studies the impact of proportional transaction costs on optimal investment decisions and the liquidity premium in the context of the capital asset pricing model (CAPM). Davis and Norman (1990), Dumas and Luciano (1991), and Shreve

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

study the optimal trading strategy for an investor with constant relative risk aversion (CRRA) in the presence of transaction costs and obtain closed-form solutions when the finite horizon is uncertain. Detemple, Garcia, and Rindisbacher (2003) develop a simulation-based methodology for optimal portfolio choice in complete markets with complex state dynamics.

There is also a significant literature on portfolio optimization that incorporates return predictability (e.g., Campbell and Viceira (2002)). Balduzzi and Lynch (1999) and Lynch and Balduzzi (2000) illustrate the impact of return predictability and transaction costs on utility costs and the optimal rebalancing rule by discretizing the state space of the dynamic program. Using similar state-space discretization, Lynch and Tan (2010) model dynamic portfolio decisions with multiple risky assets under return predictability and transaction costs and conduct numerical experiments with two risky assets.

Much of the aforementioned literature seeks to find the best rebalancing policy out of the universe of all possible rebalancing policies. As discussed earlier, this leads to highly restrictive modeling primitives. Conversely, our work is in the spirit of Brandt et al. (2009), who allow for broader modeling flexibility at the expense of considering a restricted class of rebalancing policies. They parameterize the rebalancing rule as a function of security characteristics and estimate the parameters of the rule from empirical data without modeling the distribution of the returns or the return-predicting factors. Even though our approach is also a linear parameterization 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 their specific linear functional form. In our approach, the parameters are time varying and cross-sectionally different for each security. Second, the extensions Brandt et al. (2009) provide for imposing positivity constraints and transaction costs are ad hoc and cannot be generalized to arbitrary convex constraints or transaction cost functions. Finally, the authors' objective function is a nonconvex function of the policy parameters. Hence, it is not generally possible to obtain a globally optimal set of parameters. Our setting, in contrast, is convex, and hence globally optimal policies can be efficiently determined. Brandt and Santa-Clara (2006) use a different approximate policy for the optimal solution that invests in conditional portfolios, which invest in each asset an amount proportional to the conditioning variables. Furthermore, Brandt, Goyal, Santa-Clara, and Stroud (2005) compute approximate portfolio weights using a Taylor expansion of the value function and approximate conditional expected returns as affine parameterizations of nonlinear functions.

Gârleanu and Pedersen (2013) achieve a closed-form solution for a model with linear dynamics for return predictors, quadratic functions for transaction costs, and quadratic penalty terms for risk.[2] However, the analytic solution

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 these papers is that the optimal solution is only computed in the case of a single stock and bond. For a survey of this literature, see Cvitanic (2001). Liu (2004) extends these results to multiple assets with fixed and proportional transaction costs in the case of uncorrelated asset prices.

[2]Boyd, Mueller, O'Donoghue, and Wang (2014) consider an alternative generalization of the linear quadratic case, using ideas from approximate dynamic programming. Glasserman and Xu (2013)

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

is highly sensitive to the quadratic cost structure with linear dynamics (e.g., Bertsekas (2000)). This special case cannot handle inequality constraints on portfolio positions, nonquadratic transactions costs, or more complicated risk considerations. Conversely, our approach can be efficiently implemented in these realistic scenarios and provides more flexibility in the investor's objective function and the constraints the investor faces.

Second, there is also a strand of the literature that examines the use of linear decision rules in optimal control problems. This approximation technique has recently attracted considerable interest in the context of robust and 2-stage adaptive optimization.[3] In this strand of literature, we believe the works most closely related to the methodology described in our paper are those of 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) computes lower and upper bounds on the expected transaction costs and solves two convex optimization problems to obtain upper and lower bounds on the optimal value of the simplified dynamic optimization program with linear decision rules. In contrast, Skaf and Boyd (2010) study the dynamic portfolio choice problem as an application of 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 case of proportional and nonlinear transaction costs.

II. Dynamic Portfolio Choice with Return Predictability and Transaction Costs

We consider a problem of dynamic portfolio choice that allows general models for the predictability of security returns and for trading frictions. The number of investable securities is N, and time is discrete and indexed by $t = 1, \ldots, T$, 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 $r_{t+1} \triangleq (r_{1,t+1}, \ldots, r_{N,t+1})$. We assume that the investor has a predictive model of future security returns and that these predictions are made through a set of K return-predictive factors. These factors could be security-specific characteristics, such as the market capitalization of the stock, the book-to-market ratio of the stock, and the lagged 12-month

develop a linear quadratic formulation for portfolio optimization that offers robustness to modeling errors and misspecifications.

[3] See, for example, Ben-Tal, Goryashko, Guslitzer, and Nemirovski (2004), Ben-Tal, Golany, Nemirovski, and Vial (2005), Chen, Sim, and Sun (2007), Chen, Sim, Sun, and Zhang (2008), Bertsimas, Iancu, and Parrilo (2010), Bertsimas and Goyal (2012). Shapiro and Nemirovski (2005) show that linear decision rules can reduce the complexity of multistage stochastic programming problems. Kuhn, Wiesemann, and Georghiou (2011) propose an efficient method to estimate the loss of optimality incurred by linear decision rule approximation.

return of the stock (e.g., Goetzmann and Jorion (1993), Fama and French (1996)). Alternatively, they could be macroeconomic signals that affect the return of each security, such as inflation, the Treasury bill rate, and industrial production (e.g., Chen, Roll, and Ross (1986)). We denote by $f_t \in \mathbb{R}^K$ the vector of factor values at time t. Under the following assumption, we allow for very general dynamics, possibly nonlinear and with a general dependence on history, for the evolution of returns and factors.

Assumption 1. General Return and Factor Dynamics. Over a complete filtered probability space given by $\left(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathsf{P}\right)$, we assume that factors and returns evolve according to

$$f_{t+1} = G_{t+1}(f_t, \ldots, f_1, \epsilon_{t+1}), \qquad r_{t+1} = H_{t+1}(f_t, \epsilon_{t+1}),$$

for each time t, where $G_{t+1}(\cdot)$ and $H_{t+1}(\cdot)$ are known functions that describe the evolution of the factors and returns in terms of the history of factor values and exogenous independent and identically distributed (IID) disturbances ϵ_{t+1}. We assume that the filtration $\mathbb{F} \triangleq \{\mathcal{F}_t\}_{t\geq 0}$ is the natural filtration generated by the exogenous noise terms $\{\epsilon_t\}$.

Note that we choose to describe the evolution of asset prices in our framework in terms of absolute price changes; we will also refer to these as (absolute) returns. This choice is purely notational and without loss of generality. Because the return dynamics specified by Assumption 1 allow for arbitrary dependence on history, our framework also allows, for example, models that describe the percentage return of each security. Example 1 in Section II.B illustrates such a model.

Let $x_{i,t}$ denote the number of shares that the investor holds in the ith security over period t. We collect the portfolio holdings across all securities at time t in the vector $x_t \triangleq (x_{1,t}, \ldots, x_{N,t})$, and we denote the fixed initial portfolio of the investor by x_0. Similarly, let the trade vector $u_t \triangleq (u_{1,t}, \ldots, u_{N,t})$ denote the amount of shares that the investor wants to trade at the beginning of the tth period, when the investor inherits the portfolio x_{t-1} from the prior period and observes the latest realization of factor values f_t. Consequently, we have the following linear dynamics for our position and trade vector: $x_t = x_{t-1} + u_t$ for each t.

Let the entire sample path of portfolio positions, factor realizations, and security returns be denoted by $\mathbf{x} \triangleq (x_1, \ldots, x_T)$, $\mathbf{f} \triangleq (f_1, \ldots, f_T)$, and $\mathbf{r} \triangleq (r_2, \ldots, r_{T+1})$, respectively. Similarly, the sample path of trades over time is denoted by $\mathbf{u} = (u_1, \ldots, u_T)$. We make the following assumption on the feasible sample paths of trades:

Assumption 2. Convex Trading Constraints. The sample path of trades $\mathbf{u}$ is restricted to the nonempty, closed, and convex set $U \subseteq \mathbb{R}^N \times \cdots \times \mathbb{R}^N$.

The investor's trading decisions are determined by a policy π that selects a sample path 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 nonanticipating in that the trade vector u_t 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 trade vector u_t at time t must be measurable with respect to $\mathcal{F}_t$. Let $\mathcal{U}_{\mathbb{F}}$ be the set of all nonanticipating policies.

https://doi.org/10.1017/S0021109017000345 Published online by Cambridge University Press

The investor's objective is to select a policy $\pi \in \mathcal{U}_{\mathbb{F}}$ that maximizes the expected value of a total reward or payoff function $\mathrm{p}(\cdot)$. Formally, we consider the following optimization problem for the investor:

$$\sup_{\pi \in \mathcal{U}_{\mathbb{F}}} \mathsf{E}_{\pi}[\mathrm{p}(\mathbf{x}, \mathbf{f}, \mathbf{r})], \tag{1}$$

where the real-valued reward function $\mathrm{p}(\cdot)$ is a function of the entire sample path of portfolio positions $\mathbf{x}$, the factor realization $\mathbf{f}$, and security returns $\mathbf{r}$. For example, $\mathrm{p}(\cdot)$ may have the form

$$\mathrm{p}(\mathbf{x}, \mathbf{f}, \mathbf{r}) \triangleq W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) - \mathsf{RA}(\mathbf{x}, \mathbf{f}, \mathbf{r}), \tag{2}$$

where W denotes terminal wealth (total trading gains, ignoring transaction costs), that is,

$$W(\mathbf{x}, \mathbf{r}) \triangleq W_0 + \sum_{t=1}^{T} x_t^{\top} r_{t+1}, \tag{3}$$

with W_0 as initial wealth, $\mathsf{TC}(\cdot)$ captures the transaction costs associated with a set of trading decisions, and $\mathsf{RA}(\cdot)$ is the penalty term that incorporates risk aversion.

We make the following assumption about our objective function:

Assumption 3. Concave Objective Function. Given arbitrary, fixed sample paths of factor realizations $\mathbf{f}$ and security returns $\mathbf{r}$, assume that the reward function $\mathrm{p}(\mathbf{x}, \mathbf{f}, \mathbf{r})$ is a concave function of the sequence of positions $\mathbf{x}$.

If $\mathrm{p}(\cdot)$ has the specified form in equation (2), then Assumption 3 will be satisfied when the transaction cost term $\mathsf{TC}(\cdot)$ is a convex function of trades and the risk-aversion term $\mathsf{RA}(\cdot)$ is a convex function of portfolio positions.

A. Optimal Linear Model

Assumptions 1–3 can be satisfied by a broad range of portfolio optimization problems. However, without special restrictions, the optimal dynamic policy for such a broad set of problems cannot be computed either analytically or computationally. In this section, to obtain policies in a computationally tractable way, we consider a more modest goal. Instead of finding the optimal policy among all admissible dynamic policies, we restrict our search to a subset of policies that are parsimoniously parameterized. That is, instead of solving for a globally optimal policy, we instead find an approximately optimal policy by finding the best policy over the restricted subset of policies.

To simplify, we assume that the investor's reward function in equation (1) depends on only the sample path of portfolio positions $\mathbf{x}$ and factor realizations $\mathbf{f}$ and not on the security returns $\mathbf{r}$ explicitly. In other words, we assume that the reward function takes the form $\mathrm{p}(\mathbf{x}, \mathbf{f})$. This is without loss of generality. Given our general specification for factors under Assumption 1, we can simply include each security return as a factor. With this assumption, the investor's trading decisions will generally be a nonanticipating function of the sample path of factor realizations $\mathbf{f}$. However, consider the following restricted set of policies, linear rebalancing policies, which are obtained by taking the affine combinations of the factors:

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

Definition 1. Linear Rebalancing Policy. A linear rebalancing policy π is a nonanticipating policy parameterized by a collection of vectors $\mathbf{c} \triangleq \{c_t \in \mathbb{R}^N, 1 \leq t \leq T\}$ and a collection of matrices $\mathbf{E} \triangleq \{E_{s,t} \in \mathbb{R}^{N \times K}, 1 \leq s \leq t \leq T\}$ that generates a sample path of trades $\mathbf{u} \triangleq (u_1, \ldots, u_T)$ according to

$$u_t \triangleq c_t + \sum_{s=1}^{t} E_{s,t} f_s, \quad \text{for } t = 1, 2, \ldots, T. \tag{4}$$

Define $\mathcal{C}$ to be the set of parameters $(\mathbf{E}, \mathbf{c})$ such that the resulting sequence of trades $\mathbf{u}$ is contained in the constraint set U with probability 1; that is, $\mathbf{u}$ is feasible. Denote by $\mathcal{L} \subset \mathcal{U}_{\mathbb{F}}$ the corresponding set of feasible linear policies.

An alternative to solving the original optimal control problem in equation (1) is to consider the problem

$$\sup_{\pi \in \mathcal{L}} \mathsf{E}_{\pi}[\mathrm{p}(\mathbf{x}, \mathbf{f})], \tag{5}$$

which is restricted to linear rebalancing rules. Generally, solving the problem in equation (5) does not yield an optimal control for the main problem in equation (1). The exception is if the optimal control for the problem is indeed a linear rebalancing rule (e.g., in an LQC problem). However, solving the problem in equation (5) will yield the optimal linear rebalancing rule. Further, in contrast to the original optimal control problem, the program in equation (5) has the great advantage of being tractable, as suggested by the following result:

Proposition 1. The optimization problem given by

$$\begin{aligned} \underset{\mathbf{E},\mathbf{c}}{\text{maximize}} \ & \mathsf{E}\big[\mathrm{p}(\mathbf{x}, \mathbf{f})\big], \\ \text{subject to } & x_t = x_{t-1} + u_t, && \forall\, 1 \leq t \leq T, \\ & u_t = c_t + \sum_{s=1}^{t} E_{s,t} f_s, && \forall\, 1 \leq t \leq T, \\ & (\mathbf{E}, \mathbf{c}) \in \mathcal{C}, \end{aligned} \tag{6}$$

is a convex optimization problem; that is, it involves the maximization of a concave function subject to convex constraints.

Proof. Note that $\mathrm{p}(\cdot, \mathbf{f})$ is concave for a constant $\mathbf{f}$ by Assumption 3. Because $\mathbf{x}$ can be written as an affine transformation of $(\mathbf{E}, \mathbf{c})$, then, for each fixed $\mathbf{f}$, the objective function is concave in $(\mathbf{E}, \mathbf{c})$. Taking an expectation over realizations of $\mathbf{f}$ preserves this concavity. Finally, the convexity of the constraint set $\mathcal{C}$ follows from the convexity of U, under Assumption 2. □

B. Examples

We can compute optimal linear rebalancing rules in dynamic portfolio choice models that satisfy Assumptions 1–3. To illustrate the generality of this setting, we now provide a number of specific examples that satisfy these assumptions.

In many cases, it may be more natural to model the percentage returns associated with an asset rather than with nominal price changes. Our framework accommodates such models, as seen in the following example:

https://doi.org/10.1017/S0021109017000345 Published online by Cambridge University Press

Example 1. Models of Asset Returns. Consider an asset with price P_t with log-returns evolving according to

$$\log\left(\frac{P_{t+1}}{P_t}\right) \quad = \quad g(F_t, \epsilon_{t+1}^{(1)}),$$

where F_t is a vector of predictive variables and $\epsilon_{t+1}^{(1)}$ is an IID disturbance term. We assume that F_t is a Markov process; that is,

$$F_{t+1} \quad = \quad h(F_t, \epsilon_{t+1}^{(2)}),$$

where $\epsilon_{t+1}^{(2)}$ is another IID disturbance term.

In this setting, we can define the factor process $f_t \triangleq (P_t, P_{t-1}, F_t)$. This process evolves according to

$$f_{t+1} \quad = \quad G_{t+1}(f_t, \epsilon_{t+1}) \quad \triangleq \quad \left(P_t e^{g(F_t, \epsilon_{t+1}^{(1)})}, P_t, h\left(F_t, \epsilon_{t+1}^{(2)}\right)\right),$$

where $\epsilon_t \triangleq (\epsilon_t^{(1)}, \epsilon_t^{(2)})$. Similarly, we define the price change process as $r_t \triangleq P_t - P_{t-1}$. We obtain

$$r_{t+1} \quad = \quad H_{t+1}(f_t, \epsilon_{t+1}) \quad \triangleq \quad P_t e^{g(F_t, \epsilon_{t+1}^{(1)})} - P_t.$$

Then, the joint dynamics of (f_t, r_t) satisfy Assumption 1.

Note that the Markovian assumption on the predictive variables in Example 1 is just for notational convenience and is not strictly necessary; we can always augment the vector with sufficient history so that the process becomes Markovian. What is necessary is only that F_t be measurable with respect to the filtration generated by the disturbance processes. Indeed, the only real restriction that Assumption 1 imposes is that asset prices are exogenous and are not influenced by trades.

Example 2. Gârleanu and Pedersen (2013). This model has the following dynamics, where returns are driven by mean-reverting factors, that fit into our general framework:

$$f_{t+1} \quad = \quad (I - \Phi) f_t + \epsilon_{t+1}^{(1)}, \qquad r_{t+1} \quad = \quad \mu_t + B f_t + \epsilon_{t+1}^{(2)},$$

for each time $t \geq 0$, where μ_t is the deterministic "fair return," for example, derived from the CAPM, and $B \in \mathbb{R}^{N \times K}$ is a matrix of constant factor loadings. The factor process f_t is a vector mean-reverting process, with $\Phi \in \mathbb{R}^{K \times K}$ being a matrix of mean-reversion coefficients for the factors. It is assumed that the IID disturbances $\epsilon_{t+1} \triangleq (\epsilon_{t+1}^{(1)}, \epsilon_{t+1}^{(2)})$ have zero mean, with covariance given by $\text{var}(\epsilon_{t+1}^{(1)}) = \Psi$ and $\text{var}(\epsilon_{t+1}^{(2)}) = \Sigma$.

Trading is costly, and the transaction cost to execute $u_t = x_t - x_{t-1}$ shares is given by $\mathsf{TC}_t(u_t) \triangleq \frac{1}{2} u_t^\top \Lambda u_t$, where $\Lambda \in \mathbb{R}^{N \times N}$ is a positive semi-definite matrix that measures the level of trading costs. There are no trading constraints (i.e., $U \triangleq \mathbb{R}^{N \times T}$). The investor's objective function is to choose a trading strategy to

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

maximize discounted future expected excess returns while accounting for transaction costs and adding a per-period penalty for risk; that is,

$$\underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}}\ \mathsf{E}_{\pi}\left[\sum_{t=1}^{T}\left(x_t^{\top} B f_t - \mathsf{TC}_t(u_t) - \mathsf{RA}_t(x_t)\right)\right], \tag{7}$$

where $\mathsf{RA}_t(x_t) \triangleq \frac{\gamma}{2} x_t^{\top} \Sigma x_t$ is a per-period risk-aversion penalty, with γ being a coefficient of risk aversion. Gârleanu and Pedersen (2013) suggest this objective function for an investor who is compensated based on performance relative to a benchmark. Each $x_t^{\top} B f_t$ term measures the excess return over the benchmark, and each $\mathsf{RA}_t(x_t)$ term measures the variance of the tracking error relative to the benchmark.[4]

The problem in equation (7) clearly falls within our framework. The objective function is similar to that of equation (2), with the minor variation that expected excess return rather than expected wealth is considered. Further, the problem in equation (7) has the special property that the total transaction costs and penalty for risk aversion decompose over time:

$$\mathsf{RA}(\mathbf{x}, \mathbf{f}, \mathbf{r}) \triangleq \sum_{t=1}^{N} \mathsf{RA}_t(x_t), \qquad \mathsf{TC}(\mathbf{u}) \triangleq \sum_{t=1}^{N} \mathsf{TC}_t(u_t).$$

Note that this problem can be easily handled using classical theory from the LQC literature (e.g., Bertsekas (2000)). This theory provides an analytical characterization of the optimal solution, for example, that the value function at any time t is a quadratic function of the state (x_t, f_t) and that the optimal trade at each time is an affine function of the state. Moreover, efficient computational procedures are available to solve for the optimal policy.

However, the tractability of this model relies critically on three requirements:

- The state variables (x_t, f_t) at each time t must evolve as linear functions of the control u_t and the IID disturbances ϵ_t (i.e., linear dynamics).
- Each control decision u_t must be unconstrained.
- The objective function must decompose across time into a positive definite quadratic function of (x_t, u_t) at each time t.

These requirements are not satisfied by many real-world examples, which can involve portfolio position or trade constraints, different forms of transaction costs and risk measures, and more complicated return dynamics. In Example 3 and Internet Appendix I (available at www.jfqa.org), we provide concrete evidence of many such cases that do not satisfy these requirements but remain within our framework.

Example 3. Portfolio or Trade Constraints. In practice, short-sale restrictions are a common constraint in the construction of equity portfolios. Most mutual funds

[4] See Gârleanu and Pedersen (2013) for other interpretations.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

are prohibited by law from holding short positions. This requires the portfolio optimization problem to include the linear constraint

$$x_t \quad = \quad x_0 + \sum_{s=1}^{t} u_t \quad \geq \quad \mathbf{0},$$

for each t. This is clearly a convex constraint on the set of feasible trade sequences $\mathbf{u}$.

We observe a similar restriction when an execution desk needs to sell or buy a large portfolio on behalf of an investor. Because of regulatory rules in agency trading, the execution desk is allowed to sell or buy only during the trading horizon. In the "pure-sell" scenario, the execution desk needs to impose the negativity constraint $u_t \leq \mathbf{0}$ for each time t.

A third case arises in the context of insurance companies and banks, which often need to satisfy certain minimum capital requirements to reduce the risk of insolvency. Therefore, they need to choose a dynamic investment portfolio so that their total wealth net of transaction costs exceeds a certain threshold C at all times. In our framework, this translates into the constraint

$$W_0 + \sum_{s=1}^{t} \left(x_s^\top r_{s+1} - \mathsf{TC}_s(u_s) \right) \quad \geq \quad C,$$

for each time t and for each possible realization of returns $\mathbf{r}$. If each transaction cost function $\mathsf{TC}_s(\cdot)$ is a convex function, then this constraint is also convex.

Each of the well-known portfolio construction constraints just described fits easily in our framework.

Finally, the standard form for problems of dynamic portfolio choice is also accommodated in our framework, as shown in Example 4.

Example 4. Expected Utility of Terminal Wealth. Suppose that $U : \mathbb{R} \to \mathbb{R}$ is an increasing and concave utility function and consider the optimization problem

$$\underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \; \mathsf{E}_\pi \left[U\left(W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) \right) \right]. \tag{8}$$

In this case, the objective is to maximize the expected utility of terminal wealth net of transaction costs. If the transaction cost function TC(·) is convex, the objective in equation (8) is the composition of a concave and increasing function and a concave function of $\mathbf{x}$; it will be concave and satisfy Assumption 3.

The dynamics and reward functions considered in these examples satisfy the basic requirements of Assumptions 1–3. These examples illustrate that linear rebalancing rules can be used in many real-world problems with complex primitives for return predictability, transaction costs, risk measures, and constraints.

C. Policies within Linear Rebalancing Rules

Linear rebalancing rules allow for recourse, albeit in a restricted functional form. The affine specification in equation (4) includes several classes of policies of particular interest as special cases, as follows:

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

- *Deterministic Policies.* By taking $E_{s,t} \triangleq \mathbf{0}$ for all $1 \leq s \leq t \leq T$, it is easy to see that any deterministic policy is a linear rebalancing policy.

- *LQC Optimal Policies.* Optimal portfolios for the LQC framework of Example 2 take the form $x_t = \Gamma_{x,t} x_{t-1} + \Gamma_{f,t} f_t$, given matrices $\Gamma_{x,t} \in \mathbb{R}^{N \times N}$, $\Gamma_{f,t} \in \mathbb{R}^{N \times K}$, for all $1 \leq t \leq T$; that is, the optimal portfolio is linear in the previous position and the current factor values. Equivalently, by induction on t,

$$x_t \quad = \quad \left(\prod_{s=1}^{t} \Gamma_{x,s} \right) x_0 + \sum_{s=1}^{t} \left(\prod_{\ell=1}^{s-1} \Gamma_{x,\ell} \right) \Gamma_{f,s} f_s.$$

  Because $u_t = x_t - x_{t-1}$, it is clear that the optimal trade u_t is a linear function of the fixed initial position x_0 and the factor realizations $\{f_1, \ldots, f_t\}$ and is therefore of the form (4).

- *Linear Portfolio Policies.* Brandt et al. (2009) suggest a class of policies in which portfolios are determined by adjusting a deterministic benchmark portfolio according to a linear function of a vector of stochastic, time-varying firm characteristics. In our setting, the firm characteristics would be interpreted as stochastic return-predicting factors. An analogous rule would determine the positions at each time t via $x_t = \bar{x}_t + \Theta_t^\top (f_t - \bar{f}_t)$, where $\bar{f}_t$ is the expected factor realization at time t. The policy is parameterized by $\bar{x}_t$, the deterministic benchmark portfolio at time t, and the matrix $\Theta_t \in \mathbb{R}^{N \times K}$, which maps firm characteristics (standardized to have a mean of 0) to adjustments to the benchmark portfolio. Such a portfolio rule is clearly of the form in equation (4).

- *Policies Based on Basis Functions.* Instead of having policies that are direct affine functions of factor realizations, it is also possible to introduce basis functions (e.g., Skaf and Boyd (2009)). One could consider, for example, $\varphi : \mathbb{R}^K \to \mathbb{R}^D$, a collection of D (nonlinear) functions that capture particular features of the factor space that are important for good decision making. Consider a class of policies of the form

$$u_t \quad \triangleq \quad c_t + \sum_{s=1}^{t} E_{s,t} \varphi(f_s).$$

  Such policies belong to the linear rebalancing class if the factors are also augmented to include the value of the basis functions. This is easily accommodated in our framework, given the flexibility of Assumption 1. Similarly, policies that depend on past security returns (in addition to factor realizations) can be accommodated by augmenting the factors with past returns.

- *Policies Based on Other Policies.* One source of basis functions could be existing heuristic portfolio policies. For example, assume a collection of heuristic policies is available, each of which maps the history of factor realizations into a trading decision at each moment in time. Each such map can be used to define a set of basis functions, as described previously. The corresponding set of linear rebalancing policies would consist of all policies that are linear combinations of the heuristic policies.

https://doi.org/10.1017/S0021109017000345 Published online by Cambridge University Press

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

D. Solution Methodologies

Optimal linear rebalancing policy can be computed from the finite-dimensional convex optimization problem specified in equation (6). This problem is also a stochastic optimization problem, in the sense that the objective is the expectation of a random quantity. Generally, there are a number of effective numerical methods that can be applied to solve such problems, as follows:

- *Efficient Exact Formulation.* In many cases, with further assumptions on the problem primitives (the reward function $\mathrm{p}(\cdot)$, the dynamics of the factor realizations $\mathbf{f}$, and the trading constraint set U), the objective $\mathsf{E}\left[\mathrm{p}(\mathbf{x},\mathbf{f})\right]$ and the constraint set $\mathcal{C}$ of the program in equation (6) can be analytically expressed explicitly in terms of the decision variables $(\mathbf{E},\mathbf{c})$. In some of these cases, the program in equation (6) can be transformed into a standard form of convex optimization program, such as a quadratic program or a second-order cone program. In such cases, off-the-shelf solvers specialized for these standard forms (e.g., Grant and Boyd (2011)) can be used. Alternatively, generic methods for constrained convex optimization, such as interior-point methods (e.g., Boyd and Vandenberghe (2004)), can be applied to efficiently solve large-scale instances of equation (6). We explore this topic further, developing a number of efficient exact formulations in Internet Appendix II and providing numerical examples in Sections III–IV.

- *Sample Average Approximation (SAA).* In the absence of further structure on the problem primitives, the program in equation (6) can also be solved via Monte Carlo sampling. Specifically, suppose that $\mathbf{f}^{(1)},\ldots,\mathbf{f}^{(S)}$ are S independent sample paths of factor realization. The objective and constraints of equation (6) can be replaced with sampled versions, to obtain

$$
\begin{aligned}
(9)\qquad \underset{\mathbf{E},\mathbf{c}}{\text{maximize}}\ & \frac{1}{S}\sum_{\ell=1}^{S} p\left(\mathbf{x}^{(\ell)},\mathbf{f}^{(\ell)}\right), \\
\text{subject to}\ & x_t^{(\ell)} = x_{t-1}^{(\ell)} + u_t^{(\ell)}, && \forall\, 1\le t\le T, 1\le \ell\le S, \\
& u_t^{(\ell)} = c_t + \sum_{s=1}^{t} E_{s,t} f_s^{(\ell)}, && \forall\, 1\le t\le T, 1\le \ell\le S, \\
& \mathbf{u}^{(\ell)}\in U, && \forall\, 1\le \ell\le S.
\end{aligned}
$$

  The SAA in equation (9) can be solved via standard convex optimization methods (e.g., interior-point methods). Moreover, under appropriate regularity conditions, convergence of the SAA in equation (9) to the original program in equation (6) can be established as $S\to\infty$, along with guarantees on the rate of convergence (e.g., Shapiro (2003)).

- *Stochastic Approximation.* Denote the collection of decision variables in the problem in equation (6) by $\mathbf{z}\triangleq(\mathbf{E},\mathbf{c})$, and, allowing a minor abuse of notation, define $\mathrm{p}(\mathbf{z},\mathbf{f})$ as the reward when the sample path of factor realizations is given by $\mathbf{f}$ and the trading policy is determined by $\mathbf{z}$. Then, defining $h(\mathbf{z})\triangleq\mathrm{p}(\mathbf{z},\mathbf{f})$, the problem in equation (6) is simply to maximize $\mathsf{E}[h(\mathbf{z})]$

subject to the constraint $\mathbf{z} \in \mathcal{C}$. Under suitable technical conditions, gradients[5] of h and p are related according to $\nabla h(\mathbf{z}) = \mathsf{E}[\nabla_{\mathbf{z}} \mathrm{p}(\mathbf{z}, \mathbf{f})]$. Stochastic approximation methods are incremental methods that seek to estimate ascent directions for $h(\cdot)$ from sampled ascent directions for $\mathrm{p}(\cdot, \mathbf{f})$. For example, given a sequence of IID sample paths of factor realizations $\mathbf{f}^{(1)}, \mathbf{f}^{(2)}, \ldots$, a sequence of parameter estimates $\mathbf{z}^{(1)}, \mathbf{z}^{(2)}, \ldots$ can be constructed according to

$$\mathbf{z}^{(\ell+1)} \quad = \quad \Pi_{\mathcal{C}} \left(\mathbf{z}^{(\ell)} + \gamma_\ell \zeta_\ell \right),$$

where $\Pi_{\mathcal{C}}(\cdot)$ is the projection onto the feasible set $\mathcal{C}$, $\zeta_\ell = \nabla_{\mathbf{z}} p\left(\mathbf{z}^{(\ell)}, \mathbf{f}^{(\ell)}\right)$ is a gradient, and $\gamma_\ell > 0$ is a step size. Stochastic approximation methods have the advantage of being incremental, thus requiring minimal memory relative to the SAA, and are routinely applied in large-scale convex stochastic optimization (e.g., Nemirovski et al. (2009)).

One attractive feature of our framework is that it can often be applied in a data-driven fashion, without separately specifying and estimating an explicit functional form for the factor and return dynamics. For example, the SAA and stochastic approximation approaches require only access to simulated trajectories of factors and returns; they do not require explicit knowledge of the dynamics in Assumption 1 that drive these processes. It may be possible to use historical factor and return realizations (possibly in combination with nonparametric methods such as bootstrapping) to generate sample trajectories without an explicit model of the underlying dynamics. Similarly, in many of the exact formulations developed in Internet Appendix II, including the numerical examples of Sections III–IV, only moments of the factor realizations are necessary to find the optimal linear rebalancing policy. These can be estimated from historical data without an explicit, calibrated model.

Finally, observe that optimal linear policies can also be applied in concert with MPC. Here, at each time step t, the program in equation (6) is re-solved beginning from time t. This method determines the optimal linear rebalancing rule from time t onward, conditioned on the realized history up to time t. The resulting policy is only used to determine the current trading decision at time t, and the program in equation (6) is subsequently re-solved at each future time period. At the cost of additional computational burden, the use of optimal linear policies with MPC subsumes standard MPC approaches, such as re-solving a myopic variation of the portfolio optimization problem (ignoring the true multiperiod nature) or solving a deterministic variation of the portfolio optimization problem (and ignoring the possibility of future recourse).

III. Application: Equity Agency Trading

In this section, we provide an empirical application to illustrate the implementation and benefits of the optimal linear policy. Our example considers an important problem in equity agency trading. Equity agency trading seeks to address

[5] Note that differentiability of the functions h and p is not required for the application of stochastic approximation methods. Because these functions are concave, super-differentials, which generalize the concept of a gradient, can be used in place of gradients.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

the problem faced by large investors, such as pension funds, mutual funds, and hedge funds, that need to update large portfolio holdings. Here, the investor seeks to minimize the trading costs associated with a large portfolio adjustment. These costs, often labeled execution costs, consist of commissions, bid–ask spreads, and, most importantly in the case of large trades, price impacts from trading. Efficient execution of large trades is accomplished via algorithmic trading and requires significant technical expertise and infrastructure. Large investors therefore utilize algorithmic trading service providers, such as execution desks in investment banks. Such services are often provided on an agency basis, where the execution desk trades on behalf of the client in exchange for a fee. 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 client's risk objectives.

The problem of finding an optimal execution schedule has received much attention in the literature since the initial paper of Bertsimas and Lo (1998). In their model, when the price impact is proportional to the number of shares traded, the optimal execution schedule is to trade equal numbers of shares at each trading time. A number of papers extend this model to incorporate the risk of the execution strategy. For example, Almgren and Chriss (2000) show that risk-averse agents need to liquidate their portfolios faster to reduce the uncertainty of execution costs.

The models just described seek mainly to minimize execution costs by accounting for price impact and supply/demand imbalances caused by the investor's trading. In a complementary manner, an investor can also seek to exploit the short-term predictability of stock returns to inform the design of a trade schedule. As such, there is growing interest in modeling return predictability in intraday stock returns. Often called short-term alpha models, some of the predictive models are similar to well-known factor models for the study of long-term stock returns, such as the CAPM or the Fama and French (1996) 3-factor model. Alternatively, short-term predictions can be developed from microstructure effects, for example, the imbalance of orders in an electronic limit-order book. Heston, Korajczyk, and Sadka (2010) show that systematic trading as described in the previous examples and institutional fund flows lead to predictable patterns in the intraday returns of common stocks.

We consider an agency-trading optimal execution problem in the presence of short-term predictability. One issue that arises here is that because of regulatory rules in agency trading, the execution desk is only allowed to either sell or buy a particular security over the course of the trading horizon, depending on whether the ultimate position adjustment desired for that security is negative or positive. However, given a model for short-term predictability, an optimal trading policy that minimizes execution cost could result in both buy and sell trades for the same security because it seeks to exploit short-term signals. Hence, it is necessary to impose constraints on the signs of trades, as in Example 3.

If an agency-trading execution problem has price and factor dynamics that satisfy Assumption 1 and an objective (including transaction costs, price impact, and risk aversion) that satisfies Assumption 3, then we can compute the best execution schedule in the space of linear execution schedules; that is, the number of

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

https://doi.org/10.1017/S0021109017000345 Published online by Cambridge University Press

shares to trade at each time is a linear function of the previous return-predicting factors. We consider a particular formulation that involves linear price and factor dynamics and a quadratic objective function (as in Example 2). Note that this example does not highlight the full generality of our framework. More interesting cases would involve nonlinear factor dynamics (e.g., microstructure-based order imbalance signals) or a nonquadratic objective (e.g., transaction costs, as in Example 5). However, this example is intentionally chosen because in the absence of a trade sign constraint, the problem can be solved exactly with LQC methods. Hence, we are able to compare the optimal linear policy to policies derived from LQC methods applied to the unconstrained problem.

The rest of this section is organized as follows: We present our optimal execution problem formulation in Section III.A. An exact analytical solution to this problem is not available; hence, in Section III.B, we describe several approximate solution techniques, including how to find the best linear policy. To evaluate the quality of the approximate methods, in Section III.C, we describe several techniques for computing upper bounds on the performance of any policy for our execution problem. In Section III.D, we describe the empirical calibration of the parameters of our problem. Finally, in Section III.E, we present and discuss the numerical results.

A. Formulation

We follow the general framework of Section II. Suppose that $x_0 \in \mathbb{R}^N$ denotes the number of shares in each of N securities that we would like to sell before time T. We assume that trades can occur at discrete times, $t = 1, \ldots, T$. We define an execution schedule to be the collection $\mathbf{u} \triangleq (u_1, \ldots, u_T)$, where each $u_t \in \mathbb{R}^N$ denotes the number of shares traded at time t. Note that a negative (positive) value of $u_{i,t}$ denotes a sell (buy) trade of security i at time t. The total position at time t is given by $x_t = x_0 + \sum_{s=1}^{t} u_s$.

The formulation of the agency-trading optimal execution problem is as follows:

- *Constraints.* Without loss of generality, we assume that the initial position is positive; that is, $x_0 > \mathbf{0}$. The execution schedule must liquidate the entire initial position by the end of the time horizon; therefore:

$$x_T = x_0 + \sum_{t=1}^{T} u_t = \mathbf{0}. \tag{10}$$

Further, agency trading regulations allow only sell trades; therefore:

$$u_t \leq \mathbf{0}, \quad t = 1, \ldots, T. \tag{11}$$

Note that any schedule satisfying equations (10) and (11) will also satisfy

$$x_t = x_0 + \sum_{s=1}^{t} u_s \geq \mathbf{0}, \quad t = 1, \ldots, T. \tag{12}$$

We denote by $\mathcal{U}_{\mathbb{F}}^0$ the set of nonanticipating policies almost surely satisfying the constraint in equation (10) and by $\mathcal{U}_{\mathbb{F}}$ the set of nonanticipating policies almost surely satisfying the constraints in equations (10)–(12).

- *Return and Factor Dynamics.* We follow the discrete-time linear dynamics of Gârleanu and Pedersen (2013),[6] as described in Example 2. We assume that the price change of each security from t to $t+1$ is given by the vector r_{t+1} and is predicted by K factors collected in a vector f_t. Furthermore, the evolution of factor realizations follows a mean-reverting process. Formally, we have the following dynamics for price changes and factor realizations:

$$f_{t+1} = (I-\Phi)\, f_t + \epsilon^{(1)}_{t+1}, \qquad r_{t+1} = \mu + B f_t + \epsilon^{(2)}_{t+1},$$

where $B \in \mathbb{R}^{N\times K}$ is a constant matrix of factor loadings, $\Phi \in \mathbb{R}^{K\times K}$ is a diagonal matrix of the factors' mean-reversion coefficients, and $\mu \in \mathbb{R}^N$ is the mean return. We assume that the noise terms are IID and normally distributed, with zero mean and with covariance matrices given by $\mathrm{var}(\epsilon^{(1)}_{t+1}) = \Psi \in \mathbb{R}^{N\times N}$ and $\mathrm{var}(\epsilon^{(2)}_{t+1}) = \Sigma \in \mathbb{R}^{K\times K}$. We discuss the precise choice of return-predicting factors and the calibration of the dynamics in Section III.D.

- *Objective.* We assume that the investor is risk neutral and seeks to maximize total excess profits after quadratic transaction costs, that is:

$$\mathcal{V}^* \triangleq \underset{\pi\in\mathcal{U}_{\mathbb{F}}}{\text{maximize}}\ \mathsf{E}_\pi\left[\sum_{t=1}^{T}\left(x_t^\top B f_t - \tfrac{1}{2}u_t^\top \Lambda u_t\right)\right], \tag{13}$$

where $\Lambda \in \mathbb{R}^{N\times N}$ is a matrix parameterizing the quadratic transaction costs.

Note that the problem in equation (13) is a special case of the optimization program in Example 2, with the exception of the constraints in equations (10)–(12).

B. Approximate Policies

Because a tractable analytical or computational solution to the optimal execution problem in equation (13) is not available, we compare four approximate solution techniques, as follows:

- *Time-Weighted Average Price (TWAP) Policy.* A TWAP policy seeks to sell a fixed quantity $u_t = -x_0/T$ of shares in each of the T periods. This policy minimizes transaction costs and would be optimal in the absence of a predictive model for returns.

- *Deterministic Policy.* Instead of allowing for a nonanticipating dynamic policy, where the trade at each time t is allowed to depend on all events that have occurred before t, we can solve for an optimal static policy, that is, a deterministic sequence of trades over the entire time horizon that is decided at the beginning of the time horizon. Here, observe that at the beginning of the time horizon, the expected future factor vector is given

[6] Note that Gârleanu and Pedersen (2013) consider an infinite-horizon setting, whereas our setting involves a finite horizon. Further, the authors solve for dynamic policies in the absence of the constraints in equations (10)–(12).

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

by $\mathsf{E}[f_t \mid f_0] = (I - \Phi)^t f_0$. Therefore, to find the optimal deterministic policy, given f_0, we maximize the conditional expected value of the stochastic objective in the problem in equation (13) by solving the quadratic program

$$\begin{aligned} (14) \qquad \underset{\mathbf{u}}{\text{maximize}} \quad & \sum_{t=1}^{T} \left(x_t^\top B (I - \Phi)^t f_0 - \frac{1}{2} u_t^\top \Lambda u_t \right), \\ \text{subject to} \quad & u_t = x_t - x_{t-1}, \quad t = 1, \ldots, T, \\ & u_t \leq \mathbf{0},\ x_t \geq \mathbf{0}, \quad t = 1, \ldots, T, \\ & x_T = \mathbf{0}, \end{aligned}$$

to yield a deterministic sequence of trades $\mathbf{u}$.

- *MPC Policy.* In this approximation, at each trading time, we solve for the deterministic sequence of trades conditional on the available information and implement only the first trade. Thus, this policy is an immediate extension of the deterministic policy, with the addition of re-solving at each trading time. Formally, at time t, we solve the quadratic program

$$\begin{aligned} (15) \qquad \underset{u_t, \ldots, u_T}{\text{maximize}} \quad & \sum_{s=t}^{T} \left(x_s^\top B (I - \Phi)^{(s-t)} f_t - \frac{1}{2} u_s^\top \Lambda u_s \right), \\ \text{subject to} \quad & u_s = x_s - x_{s-1}, \quad s = t, \ldots, T, \\ & u_s \leq \mathbf{0},\ x_s \geq \mathbf{0}, \quad s = t, \ldots, T, \\ & x_T = \mathbf{0}. \end{aligned}$$

  If $(u_t^*, \ldots, u_T^*)$ is the optimal solution, then the investor trades u_t^* at time t.

- *Projected LQC Policy.* If the inequality constraints in equations (11) and (12) are eliminated, the program would reduce to the following classical LQC problem:

$$(16) \qquad \underset{\pi \in \mathcal{U}_{\mathbb{F}}^0}{\text{maximize}} \ \mathsf{E}_\pi \left[\sum_{t=1}^{T} \left(x_t^\top B f_t - \frac{1}{2} u_t^\top \Lambda u_t \right) \right].$$

  The optimal dynamic policy for the program in equation (16) yields the trade

$$(17) \qquad u_t = \left(\Lambda + A_{xx,t} \right)^{-1} \left(\Lambda x_{t-1} + \left(B + A_{xf,t} (I - \Phi) \right) f_t \right) - x_{t-1},$$

  at each time t as a function of the previous position x_{t-1} and the current factor values f_t. Here, the matrices $A_{xx,t}$ and $A_{xf,t}$ are derived in Internet Appendix III. The dynamic rule for u_t in equation (17) will, of course, not be feasible for the constrained program in equation (13), in general. Thus, the projected LQC policy seeks a trade decision, $\hat{u}_t$, that is the projection of u_t onto the constraint set in equations (11) and (12). In other words, given a trading decision $u_{i,t}$, we find the closest trade $\hat{u}_{i,t}$ among all trades satisfying equations (11) and (12), according to $\hat{u}_{i,t} = \max\left\{ -x_{i,t-1}, \min\left\{ 0, u_{i,t} \right\} \right\}$, for each time $t < T$ and for each security i.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

- *Optimal Linear Policy.* As formulated in Definition 1, a linear rebalancing policy specifies trades according to

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

for each time $t = 1, 2, \ldots, T$, given parameters $(\mathbf{E}, \mathbf{c})$. Because of the linear relationship between position and trade vectors, we can represent the position vector in a similar form; that is, $x_t = d_t + \sum_{s=1}^{t} J_{s,t} f_s$, where $d_t \triangleq x_0 + \sum_{i=1}^{t} c_i$ and $J_{s,t} \triangleq \sum_{i=s}^{t} E_{s,i}$. As shown in Internet Appendix II, we implement the almost-sure equality constraint in equation (10) via equality constraints on the policy parameters by setting $d_T = \mathbf{0}$ and $J_{t,T} = \mathbf{0}$ for all t. We replace the almost-sure inequality constraints in equations (11) and (12) with probabilistic relaxations, as discussed in Internet Appendix II. With these assumptions, we compute the parameters of the optimal linear policy by solving the following stochastic program:

$$\begin{aligned}
(18) \qquad \underset{(\mathbf{E},\mathbf{c})}{\text{maximize}} \; & \mathsf{E}\left[\sum_{t=1}^{T}\left(\left(d_t + \sum_{s=1}^{t} J_{s,t} f_s\right)^{\top} B f_t \right.\right. \\
& \left.\left. - \frac{1}{2}\left(c_t + \sum_{s=1}^{t} E_{s,t} f_s\right)^{\top} \Lambda \left(c_t + \sum_{s=1}^{t} E_{s,t} f_s\right)\right)\right], \\
\text{subject to} \quad & d_t = x_0 + \sum_{i=1}^{t} c_i, \quad 1 \le t \le T, \\
& J_{s,t} = \sum_{i=s}^{t} E_{s,i}, \quad 1 \le s \le t \le T, \\
& \mathsf{P}\left(d_t + \sum_{s=1}^{t} J_{s,t} f_s < \mathbf{0}\right) \le \eta, \quad 1 \le t \le T, \\
& \mathsf{P}\left(c_t + \sum_{s=1}^{t} E_{s,t} f_s > \mathbf{0}\right) \le \eta, \quad 1 \le t \le T, \\
& d_T = \mathbf{0}, \\
& J_{t,T} = \mathbf{0}, \quad 1 \le t \le T,
\end{aligned}$$

where the parameter $\eta \in (0, 1/2)$ controls the probability that the constraints in equations (11) and (12) are violated.[7] Using the facts that the objective is an expectation of a quadratic expression in Gaussian random variables and that the chance constraints can be handled using Lemma 1 in Internet Appendix II, the program in equation (18) can be explicitly written as a second-order cone program. This calculation is detailed in Internet Appendix IV. Then, the program in equation (18) can be solved using an off-the-shelf convex optimization solver.

[7] We used the value $\eta = 0.2$ in our simulation results.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

The solution of the program in equation (18) provides the desired linear policy, $u_t = c_t + \sum_{s=1}^{t} E_{s,t} f_s$, in the return-predicting factors. However, because some of the constraints of the original program in equation (13) are only probabilistically enforced, u_t may not be feasible for the original program. The projected optimal linear policy seeks a trade decision, $\hat{u}_t$, that is the projection of u_t onto the constraint set in equations (11) and (12). In other words, given a trading decision $u_{i,t}$, we find the closest trade $\hat{u}_{i,t}$ among all trades satisfying the constraints in equations (11) and (12), according to $\hat{u}_{i,t} = \max\left\{-x_{i,t-1}, \min\left\{0, u_{i,t}\right\}\right\}$, for each time $t < T$ and security i.

C. Upper Bounds

To evaluate the quality of the policies described in Section III.B, we compute a number of upper bounds on the performance of any policy for the program in equation (13), as follows:

- *Perfect Foresight.* In this upper bound, we compute the value of an optimal policy with perfect knowledge of future factor values. In particular, given a vector of factor realizations $\mathbf{f}$, consider the following optimization problem:

$$
\begin{aligned}
(19) \qquad \mathcal{V}^{\mathsf{PF}}(\mathbf{f}) \triangleq \underset{\mathbf{u}}{\text{maximize}} \quad & \sum_{t=1}^{T} \left(x_t^\top B f_t - \tfrac{1}{2} u_t^\top \Lambda u_t\right), \\
\text{subject to} \quad & u_t = x_t - x_{t-1}, \quad t = 1, \ldots, T, \\
& u_t \leq \mathbf{0},\ x_t \geq \mathbf{0}, \quad t = 1, \ldots, T, \\
& x_T = \mathbf{0}.
\end{aligned}
$$

  The value $\mathcal{V}^{\mathsf{PF}}(\mathbf{f})$ is the best that can be achieved with perfect foresight of a particular sample path of factors $\mathbf{f}$. Note that this can be readily computed by solving the quadratic program in equation (19). Because the nonanticipating policies of the original program in equation (13) are not able to utilize future factor information in making trading decisions, we obtain the upper bound $\mathcal{V}^* \leq \mathsf{E}[\mathcal{V}^{\mathsf{PF}}(\mathbf{f})]$. This upper bound can be computed via Monte Carlo simulation over sample paths of factor realizations.

- *Unconstrained LQC.* The value of the LQC problem in equation (16) in which the inequality constraints in equations (11) and (12) are relaxed also provides an upper bound to the problem in equation (13). The expected value of the relaxed program can be exactly computed and yields the following upper bound:

$$
(20) \quad \mathcal{V}^* \leq -\frac{1}{2} x_0^\top A_{xx,0} x_0 + \frac{1}{2} \left(\operatorname{tr}\left(\Omega_0 (I - \Phi)^\top A_{ff,0} (I - \Phi)\right) + \sum_{t=0}^{T-2} \operatorname{tr}(\Psi A_{ff,t}) \right),
$$

  where the matrices $A_{xx,0}$ and $A_{ff,t}$ are derived as shown in Internet Appendix III.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

- *Pathwise Optimization.* Given a sequence $\boldsymbol{\zeta} \triangleq (\zeta_1, \ldots, \zeta_T)$ of vectors $\zeta_t \in \mathbb{R}^K$ for each t, define the following process:

$$M_t(\boldsymbol{\zeta}) \quad \triangleq \quad \sum_{s=1}^{t} \zeta_s^\top \epsilon_s^{(1)}, \quad 1 \le t \le T.$$

Observe that for any choice of $\boldsymbol{\zeta}$, $M_t(\boldsymbol{\zeta})$ is a zero-mean martingale, so that

$$\begin{aligned}
\mathcal{V}^* &\triangleq \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}}\ \mathsf{E}_\pi \left[\sum_{t=1}^{T} \left(x_t^\top B f_t - \tfrac{1}{2} u_t^\top \Lambda u_t \right) \right] \\
&= \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}}\ \mathsf{E}_\pi \left[\sum_{t=1}^{T} \left(x_t^\top B f_t - \tfrac{1}{2} u_t^\top \Lambda u_t \right) - M_T(\boldsymbol{\zeta}) \right] \\
&\le \mathsf{E} \left[\underset{u \in U}{\text{maximize}} \sum_{t=1}^{T} \left(x_t^\top B f_t - \tfrac{1}{2} u_t^\top \Lambda u_t \right) - M_T(\boldsymbol{\zeta}) \right] \\
&\triangleq \mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta}).
\end{aligned}$$

Here, the equality follows from the martingale property, and the inequality follows from a relaxation of the nonanticipatory requirement for trading decisions. The upper-bounding quantity $\mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta})$ can be written as the following quadratic optimization program:

$$\begin{aligned}
(21)\ \mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta}) \quad \triangleq \quad & \underset{\mathbf{u}}{\text{maximize}} \sum_{t=1}^{T} \left(x_t^\top B f_t - \zeta_t^\top \epsilon_t^{(1)} - \frac{1}{2} u_t^\top \Lambda u_t \right), \\
& \text{subject to} \quad \epsilon_t^{(1)} = f_t - (I - \Phi) f_{t-1}, \quad t = 1, \ldots, T, \\
& \qquad\qquad\ \ u_t = x_t - x_{t-1}, \quad t = 1, \ldots, T, \\
& \qquad\qquad\ \ u_t \le \mathbf{0},\ x_t \ge \mathbf{0}, \quad t = 1, \ldots, T, \\
& \qquad\qquad\ \ x_T = \mathbf{0}.
\end{aligned}$$

Upper bounds of this nature are often known as information relaxations (Brown and Smith (2011), Desai, Farias, and Moallemi (2011)). Observe that the perfect-foresight upper bound is a special case of this when $\boldsymbol{\zeta}$ is 0. Roughly speaking, this upper bound corresponds to a relaxation of the nonanticipating policy requirement. Note that the difference between the objective in the program in equation (21) and the original problem has zero mean under any nonanticipating policy. $\boldsymbol{\zeta}$ corresponds to a choice of Lagrange multipliers or penalties for violating this requirement. Although any choice of $\boldsymbol{\zeta}$ results in an upper bound, the pathwise optimization upper bound corresponds to making a choice for $\boldsymbol{\zeta}$ that results in an optimal upper bound; that is, $\mathcal{V}^* \le \min_{\boldsymbol{\zeta}} \mathsf{E}[\mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta})]$. This minimization involves a convex objective function and can be computed efficiently (for details, see Desai et al. (2011)).

D. Model Calibration

In this section, we describe the calibration of the parameters of the optimal execution problem formulated in Section III.A. We choose one of the most liquid

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

stocks, Apple, Inc. (NASDAQ: AAPL), for our empirical study. We set the execution horizon at 1 hour and the trade intervals at 5 minutes. Thus, by setting a trade interval to be a single unit of time, we obtain a time horizon $T = 12$. We assume that the initial position to be liquidated is $x_0 = 100{,}000$ shares.

In trade execution problems, the time horizon is typically a day; we therefore construct a factor model in the same time frequency. We use the intraday transaction prices of AAPL from the New York Stock Exchange Trade and Quote database on the trading days of Jan. 4, 2010 (day 0) and Jan. 5, 2010 (day 1) to construct $K = 2$ return-predicting factors, each with a different mean-reversion speed. We first divide each trading day into 78 time intervals, each 5 minutes long. For each 5-minute interval, we calculate the average transaction price from all transactions in that interval. Let $p_t^{(d)}$ be the average price for interval $t = 1, \ldots, 78$ on day $d = 0, 1$. Let $f_{k,t}$ be the value of factor $k = 1, 2$ for interval $t = 2, \ldots, 78$, defined as follows:

$$f_{1,t} \triangleq p_t^{(1)} - p_{t-1}^{(1)}, \qquad f_{2,t} \triangleq p_t^{(1)} - p_t^{(0)}.$$

In other words, $f_{1,t}$ is the average price change over the previous 5-minute interval, and $f_{2,t}$ is the average price change relative to the previous day. Here, we can interpret the factors as the representations of value and momentum signals. Intuitively, the first factor can be considered a "momentum-type" signal, with fast mean reversion, and the second factor can be considered a "value-type" signal, with slow mean reversion.

Given the price change $r_{t+1} \triangleq p_{t+1}^{(1)} - p_t^{(1)}$ of the security, we can compute the estimate of the factor-loading matrix, B, using the following linear regression:

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

where the ordinary least squares t-statistics are reported in parentheses. Therefore,

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

Similarly, we obtain the mean-reversion rates for the factors as follows:

$$\begin{array}{rcccc} \Delta f_{1,t+1} & = & -0.7146 & f_{1,t} + \epsilon_{1,t+1}^{(1)}, & \qquad \Delta f_{2,t+1} = -0.0353 \quad f_{2,t} + \epsilon_{2,t+1}^{(1)}. \\ & & (-6.62) & & \qquad\qquad\quad (-1.16) \end{array}$$

Therefore,

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

The variance of the error terms is estimated to be

$$\Sigma \triangleq \operatorname{var}(\epsilon_t^{(2)}) = 0.0428, \qquad \Psi \triangleq \operatorname{var}(\epsilon_t^{(1)}) = \begin{bmatrix} 0.0378 & 0 \\ 0 & 0.0947 \end{bmatrix}.$$

The distribution of the initial factor realization, f_0, is set to the stationary distribution under the given factor dynamics; that is, f_0 is normally distributed with zero mean and covariance:

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

https://doi.org/10.1017/S0021090917000345 Published online by Cambridge University Press

A rough estimate of the transaction cost coefficient $\Lambda = 2.14 \times 10^{-5}$ is used. This implies a transaction cost of \$10 or 0.5 basis points (bps) on a typical trade of 1,000 shares.

E. Numerical Results

Using the calibrated parameters from Section III.D, we run a simulation with 50,000 trials to estimate the performance of each of the approximate policies of Section III.B. In each trial, we sample the initial factor f_0, solve for the resulting policy of each approximate method, and compute its corresponding payoff. 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 (Grant and Boyd (2011)), a package for solving convex optimization problems in the software program Matlab, to solve the optimization problems that occur in the computation of the deterministic, MPC, and optimal linear policies.

The upper half of Table 1 summarizes the performance of each of the policies described in Section III.B. For each policy, we divide the total payoff into two components: alpha gains (i.e., $\sum_{t=1}^{T} x_t^\top B f_t$) and transaction costs (i.e., $\sum_{t=1}^{T} -u_t^\top \Lambda u_t$). For each of these components, as well as the total, we report the mean value over all simulation trials and the associated standard error.[8] In the lower half of Table 1, we report upper bounds on the total payoff of any policy, as computed using the methods described in Section III.C. The pathwise optimization method achieves the tightest upper bound. For each policy, we report an optimality gap relative to this tightest upper bound.

Comparing the performance of the various policies in Table 1, we see that accounting for predictable price movements can make a significant difference. Indeed, the TWAP policy, which minimizes transaction costs but ignores predictable price movements, has the worst performance. Other policies incur higher transaction costs than the TWAP policy, but they more than make up for this by opportunistically timing the liquidation relative to predictable price movements. Of the remaining policies, the projected LQC and optimal linear policies perform the best. These are the only policies that are constructed in a manner that explicitly accounts for the dynamic, multiperiod nature of the problem and allows for recourse.

The overall best policy is the optimal linear policy, which achieves a value that is within 5% of the value that can be achieved by any policy. This optimality gap is an improvement by a factor of 2 over the optimality gap of the next best policy, the projected LQC policy, and is significantly better than any other policy.

Note that despite the higher total payoff for the optimal linear policy compared with that of the projected LQC policy in Table 1, the relatively high standard errors preclude the immediate conclusion that the optimal linear policy achieves a statistically significant higher total payoff. Thus, to provide a more careful comparison, for each simulation trial, we consider the differences in alpha gains, transaction costs, and total payoff between these two policies. Table 2 shows the statistics of these differences and establishes the statistical significance of the

[8] Note that the values for the TWAP policy and the unconstrained LQC upper bound are computed exactly, without Monte Carlo simulation.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

TABLE 1

Performance Statistics of Approximate Policies

In Panel A of Table 1, we consider the approximate policies. For each approximate policy, we divide the total payoff into 2 components: alpha gains and transaction costs (TC). For each performance statistic, we report the mean value and the associated standard error (SE). Finally, we report the average computation time in seconds (CPU Time) for each policy per simulation trial. In Panel B, we report the computed upper bounds on the total payoff. Standard errors are also reported for those methods involving Monte Carlo simulation.

Trading Rules	Statistics	Alpha ($K)	TC ($K)	Total ($K)	Optimality Gap	CPU Time (sec.)
Panel A. Policies						
TWAP	Mean	0.03	−8.91	−8.88	238.0%	<0.01
	SE	0.210	0.000	0.210		
Deterministic	Mean	19.34	−15.81	3.53	45.4%	0.82
	SE	0.229	0.025	0.224		
MPC	Mean	21.25	−16.54	4.71	27.1%	5.79
	SE	0.233	0.023	0.225		
Projected LQC	Mean	25.13	−19.40	5.73	11.3%	0.02
	SE	0.227	0.039	0.229		
Optimal linear	Mean	23.24	−17.11	6.13	5.11%	4.23
	SE	0.233	0.025	0.224		
Panel B. Upper Bounds						
Pathwise optimization	Mean			6.46		
	SE			0.04		
Perfect foresight	Mean			8.57		
	SE			0.223		
Unconstrained LQC	Mean			12.58		

TABLE 2

Detailed Comparison

Table 2 provides a detailed comparison of the difference in alpha gains, transaction costs (TC), and total performance between the optimal linear policy and the projected dynamic policy in the optimal execution example. We observe that the standard error (SE) for the difference in total payoff is very small; therefore, the performance gain from employing the optimal linear policy is statistically significant.

	(Optimal Linear) – (Projected LQC)		
Statistics	Alpha ($K)	TC ($K)	Total ($K)
Mean	−1.89	2.29	0.40
SE	0.0137	0.0196	0.0095

performance benefit of the optimal linear policy. Moreover, Table 2 reveals that the optimal linear policy achieves a better result by managing transaction costs more carefully, at the expense of not achieving the alpha gains of the projected LQC policy.

In Table 1, we also report the average computation time (in seconds) required to evaluate each policy over a single simulated sample path. This provides a sense of the relative computational complexity of the various policies. The TWAP and projected LQC policies are the fastest to evaluate. The former is essentially trivial, whereas the latter has a closed-form expression (via a solution of recursive equations). The remaining policies involve solving at least one optimization problem per sample path. These policies have roughly the same orders of magnitude in computation time, with the MPC policy (which solves a different optimization problem at every time step) having the longest running time.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

F. Sensitivity Results

In Table 3, we report the sensitivity of our simulation results with respect to the main parameters of the optimal execution problem, for example, the length of the time horizon, the level of transaction costs, the level of factor persistence, and the relaxation probability. We only vary the parameter at hand while keeping the other parameters fixed. We report the average objective value and its standard error for the optimal linear and projected LQC policies, the top-performing policies in our baseline simulation.

We observe that the optimal linear policy outperforms the projected LQC in every sensitivity analysis, and the percentage improvement can increase by up to 18%. We conclude that our initial model calibration does not provide the

TABLE 3
Sensitivity Analysis

Table 3 reports the sensitivity of the results in the optimal execution example with respect to the model's main parameters, that is, the length of the time horizon, the level of transaction costs, factor persistence, and the relaxation probability.

Parameter	Statistics	Optimal Linear ($K)	Projected LQC ($K)	Difference ($K)	Improvement
Panel A. Time Horizon					
$T=6$	Mean	−14.573	−14.577	0.0043	0.04%
	SE	0.105	0.105	0.0003	
$T=12$	Mean	6.13	5.73	0.40	6.98%
	SE	0.229	0.224	0.0095	
$T=18$	Mean	22.96	21.43	1.53	7.14%
	SE	0.339	0.349	0.028	
$T=24$	Mean	37.10	34.26	2.84	8.29%
	SE	0.443	0.459	0.059	
Panel B. Transaction Costs					
$\Lambda=4.28\times10^{-6}$	Mean	24.01	22.73	1.28	5.63%
	SE	0.235	0.240	0.025	
$\Lambda=1.07\times10^{-5}$	Mean	15.84	14.74	1.10	7.46%
	SE	0.229	0.236	0.017	
$\Lambda=2.14\times10^{-5}$	Mean	6.13	5.73	0.40	6.98%
	SE	0.229	0.224	0.0095	
$\Lambda=2.67\times10^{-5}$	Mean	2.25	1.90	0.35	18.42%
	SE	0.222	0.225	0.0067	
Panel C. Factor Persistence					
$\Phi=\begin{bmatrix}0.3573 & 0\\ 0 & 0.0176\end{bmatrix}$	Mean	10.51	9.94	0.57	8.21%
	SE	0.267	0.272	0.0127	
$\Phi=\begin{bmatrix}0.7146 & 0\\ 0 & 0.0353\end{bmatrix}$	Mean	6.13	5.73	0.40	6.98%
	SE	0.229	0.224	0.0095	
$\Phi=\begin{bmatrix}0.8000 & 0\\ 0 & 0.0400\end{bmatrix}$	Mean	5.52	5.06	0.46	9.09%
	SE	0.217	0.221	0.0082	
$\Phi=\begin{bmatrix}0.9000 & 0\\ 0 & 0.0500\end{bmatrix}$	Mean	4.32	3.93	0.39	9.92%
	SE	0.206	0.209	0.0072	
Panel D. Relaxation Probability					
$\nu=0.1$	Mean	6.18	5.73	0.45	7.85%
	SE	0.233	0.224	0.0092	
$\nu=0.2$	Mean	6.13	5.73	0.40	6.98%
	SE	0.229	0.224	0.0095	
$\nu=0.3$	Mean	6.14	5.73	0.41	7.16%
	SE	0.233	0.228	0.0089	
$\nu=0.4$	Mean	5.89	5.73	0.16	2.79%
	SE	0.233	0.228	0.0088	

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

greatest improvement. Our results suggest that increasing the time horizon or the level of transaction costs greatly increases the percentage of improvement. The level of factor persistence strictly varies with total performance, but the overall improvement seems to be similar across low and high mean-reversion speeds. Finally, choosing a smaller value for the relaxation probability leads to better performance, but we observe that the values in the interval $0.1 \leq \nu \leq 0.3$ provide roughly the same objective values.

IV. Dynamic Trading with a Mean-Variance Objective

In this section, we analyze a dynamic trading problem for an investor with mean-variance preferences. Because of their foundational role in modern portfolio theory, mean-variance preferences have been widely studied in the literature on multiperiod portfolio choice. Basak and Chabakauri (2010) survey this literature and characterize the optimal mean-variance portfolios under various stochastic investment opportunities. However, their economic setup does not include transaction costs or portfolio constraints. We provide this empirical experiment to illustrate that mean-variance preferences are accommodated by our general modeling framework and emphasize the potentially great benefits of using an optimal linear policy as opposed to an LQC-based optimal trading rule.

Our model specification is inspired by the previous section, and we follow the same model calibration described in Section III.D. The main novelty of this application is the exact implementation of a mean-variance objective function. Although convex transaction costs and constraints can also be added to this objective function without losing tractability in solving for the optimal linear policy, we do not do so here, for the mere sake of comparing our approach with an approximate LQC policy in a simpler framework. Using the same calibration, we are interested in how to trade a single stock optimally over a short time horizon when its price changes can be predicted by 2 factors with different mean-reversion speeds. We set our trading horizon at 1 hour and implement trading decisions every 5 minutes, which translates into a time horizon of $T = 12$ periods. In our simulation results, we also consider three other choices for the time horizon, $T = 6$, $T = 18$, and $T = 24$, as robustness checks. Finally, we assume that the trader has 0 shares as the initial position.

We have the same dynamics for price changes and factor realizations as in our previous application:

$$f_{t+1} = (I - \Phi) f_t + \epsilon_{t+1}^{(1)}, \qquad r_{t+1} = B f_t + \epsilon_{t+1}^{(2)},$$

where the noise terms are IID and normally distributed, with zero mean and covariance matrices given by $\mathrm{var}(\epsilon_{t+1}^{(1)}) = \Psi$ and $\mathrm{var}(\epsilon_{t+1}^{(2)}) = \Sigma$. We use the calibrated values in Section III.D for Φ, Σ, Ψ, and B. We set the coefficient of risk aversion γ to 5×10^{-5}. We use three other choices for Φ in our simulations to assess the robustness of our results with respect to mean-reversion speeds.

We assume that the investor has mean-variance preferences and seeks to maximize expected terminal wealth subject to the penalty term from the variance

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

of terminal wealth:

$$\underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}}\ \mathsf{E}_{\pi}\left[W(\mathbf{x},\mathbf{r})\right] - \frac{\gamma}{2}\,\text{var}_{\pi}\left(W(\mathbf{x},\mathbf{r})\right), \tag{22}$$

where $W(\mathbf{x},\mathbf{r}) \triangleq W_0 + \sum_{t=1}^{T} x_t^{\top} r_{t+1}$ is terminal wealth and γ is the coefficient of risk aversion. The main difference between this objective function and the previous application of Section III is the replacement of per-period additive quadratic penalty terms in the inventory with a single penalty term for the variance of terminal wealth.

We find the optimal linear policy by solving the following optimization problem:

$$\begin{aligned}\underset{(\mathbf{J},\mathbf{d})}{\text{maximize}}\ & \sum_{t=1}^{T} \mathsf{E}\left[\left(d_t + \sum_{s=1}^{t} J_{s,t} f_s\right)^{\top} B f_t\right] \\ & - \frac{\gamma}{2}\,\text{var}\left(\sum_{t=1}^{T}\left(d_t + \sum_{s=1}^{t} J_{s,t} f_s\right)^{\top} r_{t+1}\right),\end{aligned} \tag{23}$$

where $x_t = d_t + \sum_{s=1}^{t} J_{s,t} f_s$ specifies the linear rebalancing rule. In Internet Appendix II.C, we show that this program can be reduced to an exact deterministic convex optimization problem.

For comparison, we define an approximate LQC policy. Because LQC cannot directly penalize for the variance of terminal wealth, we consider an alternative LQC formulation that penalizes per period according to the following:

$$\underset{g}{\text{maximize}}\ \mathsf{E}\left[\sum_{t=1}^{T} x_t^{\top} r_{t+1} - \frac{g}{2} x_t^{\top} \Sigma x_t\right], \tag{24}$$

where $g > 0$ is a risk-aversion multiplier that we will specify momentarily. It is easy to see that in the absence of transaction costs, this problem is separable over time. Then, the optimal solution is myopic and is given by $x_t^* \triangleq (g\Sigma)^{-1} B f_t$. We then compute the best LQC policy by solving for the optimal value of the multiplier g that maximizes the true objective function with the variance penalty:

$$\begin{aligned}\underset{g}{\text{maximize}}\ & \mathsf{E}\left[\sum_{t=1}^{T} x_t^{\top} r_{t+1}\right] - \frac{\gamma}{2}\,\text{var}\left(\sum_{t=1}^{T} x_t^{\top} r_{t+1}\right), \\ \text{subject to}\ & x_t = (g\Sigma)^{-1} B f_t, 1 \le t \le T.\end{aligned} \tag{25}$$

Then, letting $x_t = (g^*\Sigma)^{-1} B f_t$ provides an approximate LQC policy.

In Tables 4 and 5, we report the simulation results from 50,000 trials, illustrating the performance differential of the optimal linear and LQC policies under four different choices of time horizon and mean-reversion speeds for the return-predicting factors.[9]

[9] Our simulation results across different values of Σ and γ simply scale the objective value by the ratio of these parameter values. Hence, we omit a sensitivity analysis of these parameters.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

TABLE 4

Performance Summary with Mean-Variance Preferences and Sensitivity to Investment Horizon

Table 4 reports performance statistics in the mean-variance preferences example of the optimal linear policy and the LQC policy under different choices of time horizon. For each policy, we report the resulting objective value, average wealth, variance of wealth, and Sharpe ratio estimates, along with their corresponding standard errors.

Statistics	Optimal Linear	LQC	Increase
Panel A. T = *6*			
Objective	703.8	632.6	11.3%
SE	8.51	6.20	
Average wealth	1,613.6	1,235.5	30.6%
SE	8.53	6.95	
Variance of wealth	3.64×10^6	2.41×10^6	50.9%
SE	0.39×10^6	0.33×10^6	
Sharpe ratio	0.85	0.80	6.3%
SE	0.005	0.005	
Panel B. T = *12*			
Objective	1,332.5	978.0	36.3%
SE	13.45	9.92	
Average wealth	2,913.7	1,919.5	60.6%
SE	11.25	8.68	
Variance of wealth	6.32×10^6	3.36×10^6	88.0%
SE	0.58×10^6	0.48×10^6	
Sharpe ratio	1.16	0.99	17.1%
SE	0.006	0.006	
Panel C. T = *18*			
Objective	1,916.7	1,234	55.3%
SE	17.60	10.75	
Average wealth	4,196.4	2,300.3	82.4%
SE	13.51	9.24	
Variance of wealth	9.12×10^6	4.27×10^6	113.8%
SE	0.72×10^6	0.59×10^6	
Sharpe ratio	1.39	1.11	24.8%
SE	0.006	0.005	
Panel D. T = *24*			
Objective	2,532.3	1,472.8	71.9%
SE	22.7	14.2	
Average wealth	5,532.6	2,814.8	96.6%
SE	15.49	10.36	
Variance of wealth	12.0×10^6	5.37×10^6	123.6%
SE	0.94×10^6	0.76×10^6	
Sharpe ratio	1.60	1.21	31.5%
SE	0.006	0.006	

We observe that the optimal linear policy substantially outperforms the LQC policy in every case, and the percentage improvement in the objective value can increase by up to 72%. Our results suggest that increases in the time horizon and the level of persistence in the signals are positively related to the corresponding percentage improvement.

V. Conclusion

This paper provides a highly tractable formulation for determining rebalancing rules in problems of dynamic portfolio choice involving complex models of return predictability. Our rebalancing rule is a linear function of past return-predicting factors and can be utilized in a wide spectrum of portfolio choice

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

TABLE 5

Performance Summary with Mean-Variance Preferences and Sensitivity to Factor Persistence

Table 5 reports performance statistics in the mean-variance preferences example of the optimal linear policy and the LQC policy under different choices of factor persistence. For each policy, we report the resulting objective value, average wealth, variance of wealth, and Sharpe ratio estimates, along with their corresponding standard errors.

Factor Persistence (Φ)	Statistics	Optimal Linear	LQC	Increase
$\Phi = \begin{bmatrix} 0.3573 & 0 \\ 0 & 0.0176 \end{bmatrix}$	Objective	1,316.3	931.6	41.3%
	SE	13.71	8.46	
	Average wealth	2,945.5	1,727.5	70.5%
	SE	11.42	7.98	
	Variance of wealth	6.52×10^6	3.18×10^6	104.7%
	SE	0.62×10^6	0.47×10^6	
	Sharpe ratio	1.15	0.97	19.2%
	SE	0.006	0.005	
$\Phi = \begin{bmatrix} 0.7146 & 0 \\ 0 & 0.0353 \end{bmatrix}$	Objective	1,332.5	978.0	36.3%
	SE	13.45	9.92	
	Average wealth	2,913.7	1,919.5	60.6%
	SE	11.25	8.68	
	Variance of wealth	6.32×10^6	3.36×10^6	88.0%
	SE	0.58×10^6	0.48×10^6	
	Sharpe ratio	1.16	0.99	17.1%
	SE	0.006	0.006	
$\Phi = \begin{bmatrix} 0.8000 & 0 \\ 0 & 0.0400 \end{bmatrix}$	Objective	1,336.8	987.8	35.3%
	SE	13.45	8.96	
	Average wealth	2,916.4	1,842.7	58.3%
	SE	11.24	8.27	
	Variance of wealth	6.32×10^6	3.42×10^6	84.8%
	SE	0.57×10^6	0.48×10^6	
	Sharpe ratio	1.16	1.00	16.4%
	SE	0.006	0.005	
$\Phi = \begin{bmatrix} 0.9000 & 0 \\ 0 & 0.0500 \end{bmatrix}$	Objective	1,345.7	1,024.7	31.3%
	SE	13.49	9.24	
	Average wealth	2,931.8	1,914.1	53.2%
	SE	11.26	8.43	
	Variance of wealth	6.34×10^6	3.56×10^6	78.3%
	SE	0.57×10^6	0.49×10^6	
	Sharpe ratio	1.16	1.01	14.7%
	SE	0.006	0.006	

models with realistic considerations for risk measures, transaction costs, and trading constraints. We illustrate the utility of our method by showing its applicability across a broad range of modeling assumptions on these portfolio optimization primitives. As long as the underlying problem of dynamic portfolio optimization is a convex programming problem (i.e., with a concave objective and convex decision constraints), the modified optimization problem seeking the optimal parameters of the linear decision rule will be a convex programming problem that is numerically tractable. We demonstrate in realistic numerical experiments that such modeling flexibility can offer significant practical benefits.

References

Almgren, R. F., and N. Chriss. "Optimal Execution of Portfolio Transactions." *Journal of Risk*, 3 (2000), 5–39.

Balduzzi, P., and A. W. Lynch. "Transaction Costs and Predictability: Some Utility Cost Calculations." *Journal of Financial Economics*, 52 (1999), 47–78.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

Basak, S., and G. Chabakauri. "Dynamic Mean-Variance Asset Allocation." *Review of Financial Studies*, 23 (2010), 2970–3016.
Ben-Tal, A.; B. Golany; A. Nemirovski; and J. Vial. "Retailer-Supplier Flexible Commitments Contracts: A Robust Optimization Approach." *Manufacturing and Service Operations Management*, 7 (2005), 248–271.
Ben-Tal, A.; A. Goryashko; E. Guslitzer; and A. Nemirovski. "Adjustable Robust Solutions of Uncertain Linear Programs." *Mathematical Programming*, 99 (2004), 351–376.
Bertsekas, D. P. *Dynamic Programming and Optimal Control*, 2nd ed. Belmont, MA: Athena Scientific (2000).
Bertsimas, D., and V. Goyal. "On the Power and Limitations of Affine Policies in Two-Stage Adaptive Optimization." *Mathematical Programming*, 134 (2012), 491–531.
Bertsimas, D.; D. A. Iancu; and P. A. Parrilo. "Optimality of Affine Policies in Multistage Robust Optimization." *Mathematics of Operations Research*, 35 (2010), 363–394.
Bertsimas, D., and A. Lo. "Optimal Control of Execution Costs." *Journal of Financial Markets*, 1 (1998), 1–50.
Boyd, S.; M. T. Mueller; B. O'Donoghue; and Y. Wang. "Performance Bounds and Suboptimal Policies for Multi-Period Investment." *Foundations and Trends in Optimization*, 1 (2014), 1–72.
Boyd, S., and L. Vandenberghe. *Convex Optimization*. Cambridge, UK: Cambridge University Press (2004).
Brandt, M. W.; A. Goyal; P. Santa-Clara; and J. R. Stroud. "A Simulation Approach to Dynamic Portfolio Choice with an Application to Learning about Return Predictability." *Review of Financial Studies*, 18 (2005), 831–873.
Brandt, M. W., and P. Santa-Clara. "Dynamic Portfolio Selection by Augmenting the Asset Space." *Journal of Finance*, 61 (2006), 2187–2217.
Brandt, M. W.; P. Santa-Clara; and R. Valkanov. "Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns." *Review of Financial Studies*, 22 (2009), 3411–3447.
Brown, D. B., and J. E. Smith. "Dynamic Portfolio Optimization with Transaction Costs: Heuristics and Dual Bounds." *Management Science*, 57 (2011), 1752–1770.
Calafiore, G. C. "An Affine Control Method for Optimal Dynamic Asset Allocation with Transaction Costs." *SIAM Journal on Control and Optimization*, 48 (2009), 2254–2274.
Campbell, J. Y., and L. M. Viceira. *Strategic Asset Allocation*. Oxford, UK: Oxford University Press (2002).
Chen, N.; R. Roll; and S. A. Ross. "Economic Forces and the Stock Market." *Journal of Business*, 59 (1986), 383–403.
Chen, X.; M. Sim; and P. Sun. "A Robust Optimization Perspective on Stochastic Programming." *Operations Research*, 55 (2007), 1058–1071.
Chen, X.; M. Sim; P. Sun; and J. Zhang. "A Linear Decision-Based Approximation Approach to Stochastic Programming." *Operations Research*, 56 (2008), 344–357.
Constantinides, G. M. "Capital Market Equilibrium with Transaction Costs." *Journal of Political Economy*, 94 (1986), 842–862.
Cvitanic, J. "Theory of Portfolio Optimization in Markets with Frictions." In *Handbooks in Mathematical Finance*, E. Jouini, J. Cvitanic, and M. Musiela, eds. Cambridge, UK: Cambridge University Press (2001).
Davis, M. H. A., and A. R. Norman. "Portfolio Selection with Transaction Costs." *Mathematics of Operations Research*, 15 (1990), 676–713.
Desai, V. V.; V. F. Farias; and C. C. Moallemi. "Pathwise Optimization for Linear Convex Systems." Working Paper, Columbia University (2011).
Detemple, J. B.; R. Garcia; and M. Rindisbacher. "A Monte Carlo Method for Optimal Portfolios." *Journal of Finance*, 58 (2003), 401–446.
Dumas, B., and E. Luciano. "An Exact Solution to a Dynamic Portfolio Choice Problem under Transactions Costs." *Journal of Finance*, 46 (1991), 577–595.
Fama, E. F., and K. R. French. "Multifactor Explanations of Asset Pricing Anomalies." *Journal of Finance*, 51 (1996), 55–84.
Gârleanu, N., and L. H. Pedersen. "Dynamic Trading with Predictable Returns and Transaction Costs." *Journal of Finance*, 68 (2013), 2309–2340.
Glasserman, P., and X. Xu. "Robust Portfolio Control with Stochastic Factor Dynamics." *Operations Research*, 61 (2013), 874–893.
Goetzmann, W. N., and P. Jorion. "Testing the Predictive Power of Dividend Yields." *Journal of Finance*, 48 (1993), 663–679.
Grant, M., and S. Boyd. "CVX: Matlab Software for Disciplined Convex Programming, version 1.21." Available at http://cvxr.com/cvx (2011).

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press

Grinold, R. C., and R. N. Kahn. *Active Portfolio Management*, 2nd ed. New York, NY: McGraw-Hill (1999).
Heston, S. L.; R. A. Korajczyk; and R. Sadka. "Intraday Patterns in the Cross-Section of Stock Returns." *Journal of Finance*, 65 (2010), 1369–1407.
Hora, M. "The Practice of Optimal Execution." *Trading*, 2006 (2006), 52–60.
Kuhn, D.; W. Wiesemann; and A. Georghiou. "Primal and Dual Linear Decision Rules in Stochastic and Robust Optimization." *Mathematical Programming*, 130 (2011), 177–209.
Liu, H. "Optimal Consumption and Investment with Transaction Costs and Multiple Risky Assets." *Journal of Finance*, 59 (2004), 289–338.
Liu, H., and M. Loewenstein. "Optimal Portfolio Selection with Transaction Costs and Finite Horizons." *Review of Financial Studies*, 15 (2002), 805–835.
Lynch, A. W., and P. Balduzzi. "Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior." *Journal of Finance*, 55 (2000), 2285–2309.
Lynch, A. W., and S. Tan. "Multiple Risky Assets, Transaction Costs, and Return Predictability: Allocation Rules and Implications for U.S. Investors." *Journal of Financial and Quantitative Analysis*, 45 (2010), 1015–1053.
Merton, R. C. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory*, 3 (1971), 373–413.
Nemirovski, A.; A. Juditsky; G. Lan; and A. Shapiro. "Stochastic Approximation Approach to Stochastic Programming." *SIAM Journal on Optimization*, 19 (2009), 1574–1609.
Shapiro, A. "Monte Carlo Sampling Methods." In *Handbooks in Operations Research and Management Science: Stochastic Programming*, Vol. 10, A. Ruszczynski and A. Shapiro, eds. Amsterdam, Netherlands: Elsevier (2003).
Shapiro, A., and A. Nemirovski. "On Complexity of Stochastic Programming Problems." In *Applied Optimization: Continuous Optimization*, Vol. 99, P. M. Pardalos, D. Hearn, V. Jeyakumar, and A. Rubinov, eds. New York, NY: Springer (2005).
Shreve, S. E., and H. M. Soner. "Optimal Investment and Consumption with Transaction Costs." *Annals of Applied Probability*, 4 (1994), 609–692.
Skaf, J., and S. P. Boyd. "Nonlinear Q-Design for Convex Stochastic Control." *IEEE Transactions on Automatic Control*, 54 (2009), 2426–2430.
Skaf, J., and S. P. Boyd. "Design of Affine Controllers via Convex Optimization." *IEEE Transactions on Automatic Control*, 55 (2010), 2476–2487.

https://doi.org/10.1017/S0022109017000345 Published online by Cambridge University Press