Dynamic Trading Strategies in the Presence of Market Frictions

Mehmet Sağlam

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

under the Executive Committee

of the Graduate School of Arts and Sciencess

COLUMBIA UNIVERSITY

2012

ABSTRACT

Dynamic Trading Strategies in the Presence of Market Frictions

Mehmet Sağlam

This thesis studies the impact of various fundamental frictions in the microstructure of financial markets. Specific market frictions we consider are latency in high-frequency trading, transaction costs arising from price impact or commissions, unhedgeable inventory risks due to stochastic volatility and time-varying liquidity costs. We explore the implications of each of these frictions in rigorous theoretical models from an investor's point of view and derive analytical expressions or efficient computational procedures for dynamic strategies. Specific methodologies in computing these policies include stochastic control theory, dynamic programming and tools from applied probability and stochastic processes.

In the first chapter, we describe a theoretical model for the quantitative valuation of latency and its impact on the optimal dynamic trading strategy. Our model measures the trading frictions created by the presence of latency, by considering the optimal execution problem of a representative investor. Via a dynamic programming analysis, our model provides a closed-form expression for the cost of latency in terms of well-known parameters of the underlying asset. We implement our model by estimating the latency cost incurred by trading on a human time scale. Examining NYSE common stocks from 1995 to 2005 shows that median latency cost across our sample more than tripled during this time period.

In the second chapter, we provide a highly tractable dynamic trading policy for portfolio choice problems with return predictability and transaction costs. Our rebalancing rule is a linear function of the return predicting factors and can be utilized in a wide spectrum of portfolio choice models with minimal assumptions. Linear rebalancing rules enable to compute exact and efficient formulations of portfolio choice models with linear constraints, proportional and nonlinear transaction costs, and quadratic utility function on the terminal

wealth. We illustrate the implementation of the best linear rebalancing rule in the context of portfolio execution with positivity constraints in the presence of short-term predictability. We show that there exists a considerable performance gain in using linear rebalancing rules compared to static policies with shrinking horizon or a dynamic policy implied by the solution of the dynamic program without the constraints.

Finally, in the last chapter, we propose a factor-based model that incorporates common factor shocks for the security returns. Under these realistic factor dynamics, we solve for the dynamic trading policy in the class of linear policies analytically. Our model can accommodate stochastic volatility and liquidity costs as a function of factor exposures. Calibrating our model with empirical data, we show that our trading policy achieves superior performance in the presence of common factor shocks.

Table of Contents

List of Figures

List of Tables

Acknowledgments

The research in this thesis resulted from collaborations with my advisor Professor Ciamac C. Moallemi and my committee members Professor Collin-Dufresne and Professor Kent Daniel. While my long meetings with each of them have taught me a lot in theory and methodological tools, I am specifically indebted to their friendship and continuous support in establishing my academic training.

I would like to particularly thank Professor Ciamac Moallemi, who served as my academic advisor during my entire time at Columbia and provided constructive feedback for all my academic endeavors – from a standard homework question to an hour-long conference presentation. His professional attitude along with his sincere mentorship in all life matters is an exemplary figure that I will try to imitate in the rest of my career.

I am deeply grateful to my committee member Professor Pierre Collin-Dufresne, who taught me everything I know about dynamic asset pricing and continuous-time finance. From him, I learned the significance of building insights around mathematical formulas and the virtue of judgmental analysis without falling into the deception of number crunching.

I am further indebted to my committee members Professor Mark Broadie and Professor Paul Glasserman for reading my thesis carefully and their valuable suggestions and advice.

I would like to especially thank my classmates and friends from Columbia with whom my tenure as a student in New York was very enjoyable. I would like to specifically acknowledge Santiago Balseiro, Burak Başkurt, Soner Bilge, Berk Birand, Deniz Çiçek, Ezgi Demirdağ, Cem Dilmegani, Caner Göçmen, Neşet Güner, Damla Güneş, Çınar Kılçıoğlu, Serdar Kocaman, Paulita Pontiliano, Ahmet Serdar Şimsek, Erinç Tokluoğlu, Cengiz Üçbenli, and İzzet Yıldız.

I am very thankful to my wife, Merve Şehiraltı Sağlam, for her tremendous support, continuous patience and unlimited love that turned the difficult and stressful days of graduate

life into the most happiest and unforgettable.

Finally, I would like to thank my family, Nagihan Sağlam, Yusuf Sağlam, İlknur Sağlam Altun, Ümit Sağlam and İbrahim Altun for their invaluable support and unconditional love. I would like to specifically thank my parents for their utmost dedication and unbounded sacrifice that helped me reach this success. This thesis is dedicated to my family.

To my family

Chapter 1

Introduction

Classical finance models are based on an assumption of frictionless markets in one-period horizon. This simplicity usually provides ease in obtaining tractable models. However, it is not usually clear whether the one-period solution will have similar properties with the dynamic solution in the multi-period setting. Multi-period objective differs significantly from single-period objective by incorporating the ability to have decision with recourse which better reflects the actual objective of many investors in highly uncertain financial markets.

Incorporating financial frictions into the model is certainly a step forward to the "true" model of financial markets. Recent research that incorporates these frictions has shown us that these frictions may explain various anomalies observed in financial markets such as sudden liquidity dry-ups, the pricing of hard-to-borrow stocks, and valuation in over-the-counter markets.

Aiming to address these two perspectives, this thesis studies how various market frictions influence the investor's optimal decisions dynamically when underlying states of the economy are stochastic. Specific market frictions I have considered are latency in high-frequency trading, common and hidden factors in equity returns, transaction costs in portfolio rebalancing, unhedgeable inventory and residual risks due to stochastic volatility. I explored the implications of each of these frictions in rigorous theoretical models from an investor's point of view and derived analytical expressions or efficient computational procedures for dynamic strategies. Specific methodologies in computing these policies include stochastic

control theory, dynamic programming and tools from applied probability and stochastic processes.

This thesis theoretically concerns with optimal (or near-optimal) dynamic decision making in high-dimensional stochastic systems. My motivating research problems in this setting have originated from financial markets, yet, they are intrinsically operational questions: the impact of technological improvement in your trading system on your profit, the optimal control of transaction costs while trading with return predicting signals, and utilizing approximate trading rules when there are complex interactions between expected future returns and volatility and liquidity.

This thesis provides insightful contributions by enhancing our understanding of the implications of these frictions and suggests easy-to-implement strategies. In a nutshell, I believe that my research can help

- quantify the explicit cost of latency in high frequency trading and shed light on the very timely impact of speed in trading microstructure.

- characterize a near-optimal strategy to exploit return predictability while controlling transaction costs,

- propose a closed-form approximate policy for strategic asset allocation when returns exhibit factor driven covariance structure.

With these common distinguishing features, each chapter of my dissertation can be studied further in detail. In each chapter, the impact of the friction on the dynamic trading strategy is extensively studied, the dynamic problem is clearly posed and an optimal or near-optimal dynamic decision rule is derived.

1.1. The Cost of Latency

A very recent friction quoted extensively in the popular media has been latency, the delay between a trading decision and the resulting trade execution. As high frequency trading has flourished and subsequent regulatory questions about this trading activity have become a central focus of interest, thanks in part to the acclaimed "Flash Crash" on May 6th, 2010,

a growing interest has appeared in exploring the implications of latency to various market participants. Our first essay develops the first partial equilibrium model to concretely quantify the impact of latency on the optimal order submission policy and its resulting cost to the trader. In this essay, I first consider a stylized execution problem in the absence of latency as a benchmark, and I incorporate latency by not allowing the trader to continuously participate in the market. Trader's limit orders reach the market with a fixed latency, and the trader is forced to deviate from the benchmark policy in order to take into account the uncertainty introduced by this delay. I quantify the cost of latency as the normalized difference in expected payoffs between this model and the stylized model without latency. I obtain an explicit closed-form solution for the cost of latency in the most interesting regime of low-latency.

Our formulation of the latency model constitutes a powerful tool in computing the exact latency cost. Our model is the first theoretical approach in the literature to quantify the impact of latency on the optimal order submission policy and its resulting cost to the trader. I first characterize the optimal order submission policy in the model by providing an explicit recursion in a single variable. This recursion can efficiently be solved by numerical means and the exact latency cost can easily be computed. Due to the uncertainty introduced by latency, the optimal ordering policy becomes less aggressive compared to the benchmark solution. The extent to which the optimal quote is adjusted may be expressed in well-known market parameters, most evidently in the low latency regime. The highest order effect comes from the volatility of the stock movement and to a lesser degree from the average bid-ask spread. If the trader wishes to sell a share, the optimal premia that the trader sets decreases linearly with the volatility of the stock.

Since the latency values observed in modern electronic markets are on the order of milliseconds, I provide an asymptotic analysis for the low latency regime, in which I obtain an explicit closed-form solutions. In this case, the trader's optimal limit order policy becomes time-independent and the latency cost can be computed exactly without resorting to backward induction. If I interpret the cost of latency as a percentage of overall transaction costs in the absence of any latency (i.e., a normalized measure of latency), then the latency cost can be calculated in a simple closed-form expression. I find that latency cost is directly

proportional to the ratio of volatility and the average bid-ask spread. Thus, latency cost increases for more volatile or less liquid stocks. The dependence on the observed latency, is more complex with the first order contribution coming from the variance of the stock price during the latency interval and a second order adjustment that will enable to secure execution in the asymptotic limit. In order to derive this cost empirically, I only need to estimate the volatility, the average bid-ask spread of the stock and the intrinsic value of latency. This is an elegant and practical result as the estimation procedures for these quantities are readily abundant in the literature.

1.2. Linear Rebalancing Rules

One of the most well-studied market frictions is the impact of transaction costs on the optimal portfolio choice of the investor. Furthermore, when the investor has predictions for the expected future returns using return predicting factors such as market capitalization, book-to-market ratio, lagged returns, dividend yields, determining an optimal dynamic policy with realistic risk and trading constraints is almost certainly intractable.

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

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

techniques available from the sample average and stochastic approximation literature.

Finally, I implement the computation of the best linear policy in the context of portfolio execution, the execution of a large long position in a single security. For this purpose, I need positivity constraints on portfolio positions and the amount of shares sold in each period in order to achieve a feasible execution. In order to compare the performance of the best linear rebalancing rule, I use the identical discrete-time setup of Garleanu and Pedersen [2012] for which a closed-form solution is available in the lack of constraints. I calibrate the model parameters using two-days of transactions data on a liquid stock and construct two predictors in a high-frequency setting with different mean reversion speeds. The simulation implemented with these predictors and calibrated parameters reveal that the best linear policy performs better than the deterministic policy, model predictive control and a projected version of the optimal policy proposed by Garleanu and Pedersen [2012].

1.3. Common Factor Shocks in Strategic Asset Allocation

The foundations developed in the second chapter have been influential in analyzing the impact of common factor shocks when there are transaction costs and return predictability. In this essay, I take a deeper look at a particular dynamic portfolio choice problem with common factor shocks driving security returns. I propose a new factor model for security returns in which each security has its own return predicting factors based on short-term reversal, momentum, and long term reversal. In this model, I correctly account for the conditional variance of returns by allowing co-movements with factor exposures. I utilize linear decision rules in past returns and factor exposures for our dynamic trading strategy. I show that the optimal linear policy can be computed in closed-form in contrast to recent parametric approaches that rely on numerical optimization.

Garleanu and Pedersen [2012] has been a break-through by combining trading frictions with return predictability in a highly tractable model that actually allowed closed-form solution. However, this tractability has emerged with an obvious cost, a significant departure from standard dynamic portfolio choice literature. The simplifying assumption has been using number of shares in the portfolio decision vector in order to linearize the state dy-

namics. Using number of shares versus dollar holdings also required to model price changes in dollars instead of percentage terms. This is clearly problematic as it allows for negative prices. Furthermore, it is well-known that price changes are not stationary, cannot be estimated effectively using linear regression techniques. In this essay, I keep the nonlinear structure in the wealth evolution but instead of trying to solve the problem to optimality, I use linear policies in order to obtain a near-optimal policy. I obtain a closed-form solution for our policy parameters which allows us to expand the universe of parameters quite easily.

I evaluate the performance of our linear policy in a well-calibrated simulation. Our simulation study shows that best linear policy provides significant benefits compared to other approximate policies recently studied in the literature, especially when the transaction costs are high and returns evolve according to factor dependent covariance structure. Unlike other parametric approaches, our modeling provides a closed form solution instead of statistical fitting procedure. Analytical tractability allows us to expand our universe of parameters which allows for greater flexibility in obtaining different policy rules for different asset classes.

1.4. Organization of the Thesis

The balance of this thesis is organized as follows:

Chapter 2 provides a formal model to quantify the cost of latency. I present a stylized, continuous-time trade execution problem in the absence of latency. I develop a variation of the model with latency and provide a mathematical analysis of the optimal policy for our problem. By contrasting the results in the presence and absence of latency, I am able to quantitatively assess the cost of latency. In a later section, I consider some empirical applications of the model.

Chapter 3 presents the abstract form of a dynamic portfolio choice model and provide various specific problems that satisfy the assumptions of the abstract model. I formally describe the class of linear decision rules and discuss solution techniques in order to find the optimal parameters of the linear policy. I provide efficient and exact formulations of dynamic portfolio choice models using linear decision rules. In this generalized approach, I incorporate

linear equality and inequality constraints, proportional and nonlinear transaction costs and a measure of terminal wealth risk. Finally, I apply our methodology in an optimal execution problem and evaluate the performance of the best linear policy.

Chapter 4 provides a methodology that can address complex return predictability models in multi-period settings with transaction costs. Our return predicting factors does not need to follow any pre-specified model but instead can have arbitrary dynamics. I allow for factor dependent covariance structure in returns driven by common factor shocks and illustrate in a simulation study that linear policies perform very well in these intractable models.

Chapter 2

The Cost of Latency

2.1. Introduction

In the past decade, electronic markets have become pervasive. Technological advances in these markets have led to dramatic improvements in latency, or, the delay between a trading decision and the resulting trade execution. In the past 30 years, the time scale over which a trade is processed has gone from minutes[1]

One factor behind this trend has been competition between exchanges, as one mechanism for differentiation between exchanges is latency. This competition is driven by a significant demand amongst a class of investors, sometimes called "high frequency" traders, for low latency trade execution. High frequency traders are thought to account for more than half of all US equity trades.[3] They expend significant resources in order to develop algorithms and systems that are able to trade quickly. For example, on the time scale of milliseconds, the speed of light can become a binding constraint on the delay in communications. Hence, traders seeking low latency will "co-locate", or house their computers in the same facility as the exchange, in order eliminate delays due to a lack of physical proximity. This co-location

---

[1]NYSE, pre-1980 upgrade [Easley *et al.*, 2008]. to milliseconds[2] — "low latency" in a contemporary electronic market would be qualified as under 10 milliseconds, "ultra low latency" as under 1 millisecond. This change represents a dramatic reduction by *five orders of magnitude.* To put this in perspective, human reaction time is thought to be in the hundreds of milliseconds.

[3]"Stock traders find speed pays, in milliseconds," *New York Times*, July 23, 2009.

comes at a significant expense, however it has been stated that a 1 millisecond advantage can be worth $100 million to a major brokerage firm.[4]

There has been much discussion of the importance of latency among various market participants, regulators, and academics. Despite the significant amount of recent interest, however, latency remains poorly understood from a theoretical perspective. For example, how does latency relate to transaction costs? Is latency only relevant to investors with short time horizons, such as high frequency traders, or does latency also affect long term investors such as pension funds and mutual funds? Many of these important questions have been considered in anecdotal or ad hoc discussions. My goal here is to provide a framework for quantitative analysis of these issues.

In particular, I wish to understand the benefit to a single trader in the marketplace of lowering their latency, while holding everything else fixed. This is a different question than understanding the social costs of latency, i.e., whether in equilibrium the collective marketplace is better or worse off given lower latency. One might imagine, for example, that the benefit to a individual agents of lower latency may diminish in an equilibrium setting. Equilibrium or welfare analysis of low latency trading is a complex question with important policy and regulatory implications. I believe that understanding the single-agent effects of low latency trading, however, is an important first step which will inform my ultimate understanding of collective effects.

The cost that a trader bears due to latency can take many different forms, depending on the precise trading strategy. However, a number of broad themes can be identified,[5] sometimes overlapping, as to why the ability to trade with low latency might be valuable to an investor:

1. **Contemporaneous decision making.** A trader with significant latency will be making trading decisions based on information that is stale.

   For example, consider an automated trader implementing a market-making strategy in an electronic limit order book. The trader will maintain active limit orders to buy and sell. The prices at which the trader is willing to buy or sell will naturally depend

---

[4]"Wall Street's quest to process data at the speed of light," *Information Week*, April 21, 2007.

[5]See Cespa and Foucault [2008] for a related discussion.

on, say, the limit orders submitted by other investors, the price of the asset on other exchanges, the price of related assets, overall market factors, etc. If the trader cannot update his orders in a timely fashion in response to new information, he may end up trading at disadvantageous prices.

2. **Comparative advantage/disadvantage.** The ability to trade with low latency in absolute terms may not be as important as the ability to trade with low *relative* latency, that is, as compared to competitors.

   For example, consider a program trader implementing an index arbitrage strategy, seeking to profit on the difference between an index and its underlying components. There may be many market participants pursuing such strategies and identifying the same discrepancies. The challenge for the trader is to be able to act in the marketplace to exploit a discrepancy *before* a price correction takes place, i.e., before competitors are able to act. The means having a low relative latency.

3. **Time priority rules.** Many modern markets treat orders differentially based on the time of arrival, and favor earlier orders.

   For example, in an electronic limit order book, the limit orders on each side of the market are prioritized in a particular way. When a market order to buy arrives, it is matched against the limit orders to sell according to their priorities. Priority is first determined by price, i.e., limit orders with more lower prices receive higher priority. In many markets, however, prices are mandated to be discrete with a minimum tick size. In these markets, there may be multiple limit orders at the same price, which are then prioritized according to the time of their arrival. While a trader can always increase the priority of his orders by decreasing price, this comes at an obvious cost. If a trader can submit orders in a faster fashion, however, he can increase priority while maintaining the *same* price. Higher priority can be valuable for two reasons: first, higher priority orders have a higher likelihood of execution over any given time horizon. To the extent that investors submitting limit orders have a desire to trade, and to trade sooner rather than later, this is desirable. Second, higher priority orders at the same price level experience less adverse selection [see, e.g., Glosten, 1994; Sandås,

2001]. Hence, all things being equal, an investor who submits orders with lower latency will benefit from higher priority than if that investor had higher latency. This can be particularly important (in that a small improvement in latency can result in a significant difference in priority) when an existing quote is about to change. For example, consider the situation where a stock price is about to move up because of trades or cancellations at the best offered price. One might expect the bid price to rise as well, there will be a race among traders reacting to the same order book events to establish time priority at the new bid.

In this chapter, I will quantify the cost of latency due to the first effect, a lack of contemporaneous decision making. I do not consider effects of latency that arise from strategic considerations, or from time priority rules or price discreteness. It is an open question as to whether the other effects are more or less significant than the first, and their relative importance may depend on the particular investor and their trading strategy. My analysis does not speak to this point. However, in what follows I will demonstrate that, by itself, the lack of contemporaneous decision making can induce trading costs that are of the same order of magnitude as other execution costs faced by large investors, and hence cannot be neglected.

Further, the importance of contemporaneous decision making will certainly vary from investor to investor. I will focus on an aspect of this that is universal, however, which is the importance of timely information for the execution of *contingent orders*. A contingent order, such as a limit order in an electronic limit order book or a resting order in a dark pool, presents the possibility of uncertain execution over an interval of time in exchange for price improvement relative to a market order, which executes immediately and with certainty. Specifically, when an investor employs a contingent order, the investor may be exposed to the realization of new information (for example, in the form of price movements, news, etc.) over the lifespan of the order. Latency, which prevents the investor from continuously and instantaneously accessing the market so as to update the order, can thus adversely impact the investor.

As a broad proxy for understanding the importance of latency in contingent order execution, I consider the effects of latency in an extremely simple yet fundamental trade execution

problem: that of a risk-neutral investor who wishes to sell 1 share of stock (i.e., an atomic unit) over a fixed, short time horizon (i.e., seconds) in a limit order book, and must decide between market orders and limit orders. My problem formulation is reminiscent of barrier-diffusion models for limit order execution [e.g., Harris, 1998]. It captures the fundamental *cost of immediacy* of trading [e.g., Grossman and Miller, 1988; Chacko *et al.*, 2008], that is, the premium due to a patient liquidity supplier (who submits limit orders) relative to an impatient demander of liquidity (who submits market orders). While this problem is quite stylized, I will argue that it is broadly relevant since, at some level, *all* investors make such a choice of immediacy. For example, it may not seem at first glance that my execution problem is relevant for a pension fund that trades large blocks of stock over multiple days. However, the execution of a block trade via algorithmic trading involves the division of a large "parent" order into many atomic orders over the course of a day, each of these atomic "child" orders can be executed as limit orders or as market orders.

In my problem, in the absence of latency, the optimal strategy of the seller is a "pegging" strategy: the seller maintains a limit order at a constant spread above the bid price at any instant in time. I consider this case as a benchmark. In the presence of latency, the seller can no longer maintain continuous contact with the market so as to track the bid price in the market. The seller is forced to deviate from the benchmark policy in order to take into account the uncertainty introduced by the latency delay by incorporating a safety margin and lowering his limit order prices. The friction introduced by latency thus results in a loss of value to the seller. I will establish the difference in value to the seller between the case with latency and the benchmark case via dynamic programming arguments, and thus provide a quantification of the effects of latency.

The contributions of this essay are as follows:

- *This essay mathematically quantifies the cost of latency.*

  The trading problem I consider (deciding between limit and market orders) is faced by all large investors in modern equity markets, either directly (e.g., high frequency traders) or indirectly (e.g., pension funds who execute large trades via providers of automated execution services). My analysis suggests that latency impacts all of these market participants, and that, all else being equal, the ability to trade with low

latency results in quantifiably lower transaction costs. Further, when calibrated with market data, the latency cost we measure can be significant. It is of the same order of magnitude as other trading costs (e.g., commissions, exchange fees, etc.) faced by the most cost efficient large investors. Moreover, it is consistent with the rents that are extracted by agents who have made the requisite technological investments to trade with ultra low latency. For example, the latency cost of my model is comparable to the execution commissions charged by providers that offer algorithmic trade execution services on an agency basis. It is also comparable to the reported profits of high frequency traders.

To my knowledge, my model is the first to provide a quantification of the costs of latency in trade execution.

- *I provide a closed-form expression for the cost of latency as a function of well-known parameters of the asset price process.*

  The cost of latency in my model can be computed numerically via dynamic programming. However, in the regime of greatest interest, where the latency is close to zero, I provide a closed-form asymptotic expression. In particular, define the *latency cost* associated with an asset as the costs incurred due to latency as a fraction of the overall cost of immediacy (the premium paid to a patient liquidity supplier by an impatient demander of liquidity). Given a latency of Δt, a price volatility of σ, and a bid-offer spread of δ, the latency cost takes the form

  $$(2.1) \qquad \frac{\sigma\sqrt{\Delta t}}{\delta}\sqrt{\log\frac{\delta^2}{2\pi\sigma^2\Delta t}}$$

  as $\Delta t \to 0$.

- *My method can provide qualitative insight into the importance of latency.*

  From (2.1), it is clear that the latency cost is an increasing function of the ratio of the standard deviation of prices over the latency interval (i.e., $\sigma\sqrt{\Delta t}$) to the bid-offer spread. Latency has a more important role when trading assets that are either more volatile (σ large) or, alternatively, more liquid (δ small). Further, as the latency approaches 0, the marginal benefit of latency reduction is increasing.

- *This chapter empirically demonstrates that latency cost incurred by trading on a human time scale has dramatically increased for U.S. equities and the implied latency of a representative trader in this market decreased by approximately two orders of magnitude.*

I consider the cost due to the latency of trading on the time scale of human interaction.Using the data-set of Aït-Sahalia and Yu [2009], I estimate the latency cost of NYSE common stocks over the 1995–2005 period. I show that the median latency cost more than tripled in this time. This coincides with a period of decreasing tick sizes and increasing algorithmic and high frequency trading activity [Hendershott *et al.*, 2010].

An alternative perspective is to consider a hypothetical investor who fixes a target level of cost due to latency, relative to the overall cost-of-immediacy. The representative trader maintains this target over time through continual technological upgrades to lower levels of latency. I determine the requisite level of *implied latency* for such a trader, over time and across the aggregate market. Using the same data-set, I observe that the median implied latency decreased by approximately two orders of magnitude over this time frame.

The rest of this chapter is organized as follows: In Section 4.1.1, I review the related literature. In Section 2.2, as a starting point, I present a stylized, continuous-time trade execution problem in the absence of latency. I develop a variation of the model with latency in Section 4.2. In Section 2.4, I provide a mathematical analysis of the optimal policy for my problem. By contrasting the results in the presence and absence of latency, I am able to quantitatively assess the cost of latency. In Section 2.5, I consider some empirical applications of the model. Finally, in Section 3.6 I conclude and discuss some future directions.

2.1.1. Related Literature

There has been a significant empirical literature studying, broadly speaking, the effects of improvements in trading technology. Closest to the aspect I consider is the work of Easley

et al. [2008]. They empirically test the hypothesis that latency affects asset prices and liquidity by examining the time period around an upgrade to the New York Stock Exchange technological infrastructure that reduced latency. Hendershott *et al.* [2010] explore the more general, overall effects of algorithmic and high frequency trading. Hasbrouck and Saar [2009] provide different evidence of changes in investor trading strategies that may be a result of improved technology. In subsequent work, they further consider the impact of measurements of low latency on market quality [Hasbrouck and Saar, 2010]. Hendershott and Riordan [2009] analyze the impact of algorithmic trading on the price formation process using a data set from Deutsche Börse and conclude that algorithmic trading assists in the efficient price discovery without increasing the volatility. Kirilenko *et al.* [2010] consider the impact of high frequency trading on the 'flash crash' of 2010, while Brogaard [2010] more broadly examines the impact of high frequency traders on market quality.

On the theoretical front, Cespa and Foucault [2008] consider a rational expectations equilibrium between investors with different access to past transaction data. Some investors observe transactions in real-time, while others only observe transactions with a delay. This model of latency focuses on latency of the price ticker of past transactions, as opposed to latency in execution, which I consider here. Moreover, the goals of the two models differ significantly: Cespa and Foucault [2008] seek to build intuition regarding the equilibrium welfare implications of differential access to information via a structural model. I, on the other hand, seek a reduced form model that can be used to directly estimate the value of execution latency in a particular real world instance, given readily available data. Also related is the work of Ready [1999] and Stoll and Schenzler [2006], who consider the ability of intermediaries (e.g., specialists or dealers) to delay customer orders for their own benefit, thus creating a "free option" in the presence of execution latency. Cohen and Szpruch [2011] show that latency arbitrage exists between two traders with different speeds of trading in the presence of a limit order book. Finally, Cvitanić and Kirilenko [2010] and Jarrow and Protter [2011] consider the effect of high frequency traders on asset prices.

The trade execution problem I consider is that of an investor who wishes to sell a single share of and must decide between market and limit orders. This problem has been considered by many others [e.g., Angel, 1994; Harris, 1998; Lo *et al.*, 2002]. My formulation

is similar to the class of barrier-diffusion models considered by these authors; Hasbrouck [2007] provides a good account of this line of work. For a broad survey on limit order markets, see Parlour and Seppi [2008]. In my model, the inability to trade continuously gives a limit order an option-like quality that relates execution cost, order duration, and asset volatility. This idea goes as far back as the work of Copeland and Galai [1983]. Closely related is the concept of the cost of immediacy, or, the premium paid by a liquidity demander via a market order to a liquidity supplier who posts a limit order. Grossman and Miller [1988] and Chacko *et al.* [2008] develop theoretical explanations of the cost of immediacy. For empirical evidence of the demand for immediacy in capital markets, see Bacidore *et al.* [2003] and Werner [2003].

Finally, also related is work on the discrete-time hedging of contingent claims with or without transaction costs [e.g., Boyle and Emanuel, 1980; Leland, 1985; Bertsimas *et al.*, 2000]. This literature addresses a different problem and draws different conclusions than my chapter, however both relate to implications of a lack of continuous access to the market.

2.2. A Stylized Execution Model without Latency

My goal is to understand the impact on the trade execution of latency. To this end, I will first describe a trade execution problem in the absence of latency. In Section 4.2, I will revisit this model in the presence of latency, so as to understand the resulting trade friction that is introduced. The spirit of my model it to consider an investor who wants to trade, but at a price that depends on an informational process that evolves stochastically and must be monitored continuously. I could directly consider such an abstract model of investor behavior. Instead, however, I will motivate the informational dependence of the trader through a specific optimal execution problem.

Consider the following stylized execution problem of an uninformed trader who must sell exactly one share[6] of a stock over a time horizon $[0, T]$. At any time $t \in [0, T)$, the

---

[6]Note that the trade quantity of a single share is meant to represent an atomic unit of the asset, or the smallest commonly traded lot size. The underlying assumption is that the desired trade execution will ultimately be accomplished by a single transaction. In typical U.S. equity markets, for example, this atomic unit might be a block of 100 shares.

trader can take one of two actions:

1. The trader can submit a market order to sell. This order will execute at the best bid price at time t, denoted by S_t. I assume that the bid price evolves according to

$$(2.2) \qquad\qquad S_t = S_0 + \sigma B_t,$$

   where the process $(B_t)_{t \in [0,T]}$ is a standard Brownian motion and $\sigma > 0$ is an (additive) volatility parameter. Here, the choice of Brownian motion is made for simplicity; my model can be extended to the more general class of Markovian martingales, as discussed in Section 2.4.4.

2. The trader can choose to submit a limit order to sell. In this case, the trader must also decide the limit price associated with the order, which I denoted by L_t.

Once the trader sells one share, he exits the market. If the trader is not able to sell 1 share before time T, however, I assume that he is forced sell via a market order at time T, and therefore receives S_T. Here, I imagine the time horizon T to be small, on the order of the typical trade execution time (i.e., seconds).

2.2.1. Limit Order Execution

It remains to describe the execution of limit orders. In my setting, a limit order can execute in one of the following two ways:

1. I assume that there are impatient buyers who arrive to the market according to a Poisson process with rate μ. Denote by $(N_t)_{t \in [0,T)}$ the cumulative arrival process for impatient buyers. Each impatient buyer seeks to buy a single share. An arriving impatient buyer arriving at time t has a reservation price $S_t + z_t$, expressed as a premium $z_t \geq 0$ above the bid price S_t that the buyer is willing to forgo in order to achieve immediate execution. I assume that the premium z_t is independent and identically distributed with cumulative distribution function $F \colon \mathbb{R}_+ \to [0,1]$. In this setting, the instantaneous arrival rate of impatient buyers at time t willing to pay a limit order price of L_t is given by

$$(2.3) \qquad\qquad \lambda(u_t) \triangleq \mu(1 - F(u_t)),$$

where $u_t \triangleq L_t - S_t$ is the instantaneous price premium of the limit order. In what follows, I will be particularly interested in the special case where

$$(2.4) \qquad \lambda(u_t) \triangleq \begin{cases} \mu & \text{if } u_t \leq \delta, \\ 0 & \text{otherwise.} \end{cases}$$

Here, I assume that every impatient buyer is willing to pay a price premium of at most $\delta > 0$. I assume that δ will be specific to the security and fixed for the trading horizon. I will discuss the extension to the general case (2.3) in Section 2.4.4.

Given (2.4), an impatient buyer is willing to buy 1 share at a fixed premium $\delta > 0$ to the bid price at the time of their arrival. Hence, if a buyer arrives at time $\tau \in [0, T)$, and the trader has placed a limit order with price L_τ, the limit order will execute if $L_\tau \leq S_\tau + \delta$.

2. Alternatively, a limit order will also execute at time τ if the bid price crosses the limit order price, i.e., $S_\tau \geq L_\tau$.

The execution of limit orders in the model is illustrated in Figure 2.1.

The limit order execution dynamics above can also be economically interpreted in the spirit of the non-informational trade model of Roll [1984]. In particular, imagine that the asset has a fundamental value V_t at time t, and that V_t evolves exogenously according to the additive random walk

$$V_t = V_0 + \sigma B_t.$$

If all investors observe this underlying value process and are symmetrically informed, competitive market makers will always be willing to sell shares at a price of $\delta/2$ above the fundamental value or buy shares at a spread of $\delta/2$ below the fundamental value. Here, the quantity δ captures the per share operating costs of trade to the market markers. The liquidating trader can thus sell at the bid price $S_t = V_t - \delta/2$ at any time t. I assume that all other traders in the market are impatient, and that these traders arrive according to the Poisson dynamics described above. An arriving impatient buyer will choose to purchase from the liquidating trader only at a price lower than that provided by the market makers, i.e., only below the price of $V_t + \delta/2 = S_t + \delta$. In this way, I can interpret the parameter δ as
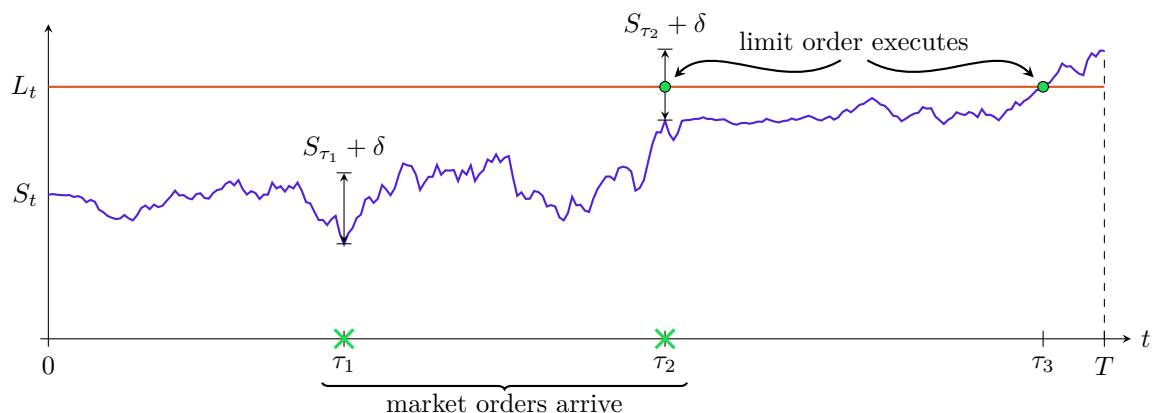
Figure 2.1: An illustration of the limit order execution in the stylized model over the time horizon $[0, T]$. Here, I assume the trader leaves a limit order with the (constant) price L_t and S_t is the bid price process. If market orders arrive at times τ_1 and τ_2, the limit order would execute at time τ_2 but not time τ_1, since the limit order price is in excess of δ to the best bid price. The limit order would also execute at time τ_3 in the absence of a market order arrival, since the bid price crosses the limit order price at this time.

the *prevailing bid-offer* spread, that is, the bid-offer spread in the absence of the liquidating trader.

2.2.2. Optimal Solution

Let P denote the random variable associated with the sale price. I assume the trader is risk-neutral and seeks to maximize the expected sale price. Equivalently, I assume the trader seeks to solve the optimization problem

$$(2.5) \qquad \bar{h}_0 \triangleq \text{maximize } \mathsf{E}\left[P\right] - S_0.$$

Here, the maximization is over policies of market orders and limit orders which are non-anticipating, i.e., which are adapted to the filtration generated by $(B_t, N_t)_{t \in [0,T]}$. This objective is equivalent to minimizing implementation shortfall [Perold, 1988].

Note that, while this stylized problem may seem quite simplified, it seeks to answer a fundamental question: at the level of an atomic unit of stock and over a short time horizon, how should a risk-neutral investor choose between limit orders and market orders? This problem is a central ingredient in more sophisticated optimal execution problems involving

risk averse investors selling large quantities over longer time horizons.[7] This is because, in a typical algorithmic trading setting, a large "parent" order will be scheduled across time into many very small "child" orders. Each of these "child" orders need to be executed optimally. Since each child order is small and since there are many such child orders, it is reasonable to view the investor as risk-neutral with respect to each child order.

The following lemma characterizes a simple strategy that is optimal for the execution problem I have described:

Lemma 1. *An optimal strategy is to employ only limit orders at times $t \in [0, T)$, with limit price $L_t = S_t + \delta$. In other words, the limit order price is "pegged" at a constant premium δ above the bid price. This pegging strategy achieves the optimal value*

$$(2.6) \qquad \bar{h}_0 = \delta \left(1 - e^{-\mu T} \right).$$

Proof. Consider a trader using an arbitrary strategy, and denote by $\tau \in [0, T]$ the (random) time at which the trader sells the share, and by $\tau_1 \in [0, \infty)$ the time at which the first impatient buyer arriving to the market. Let $\mathcal{E}$ be the event that the trader sells via a limit order to an impatient buyer at the price L_τ. Then, under the event $\mathcal{E}^c$, the trader sells at the bid price S_τ. Then, the sale price P can be written as[8]

$$(2.7) \qquad P = S_\tau \mathbb{I}_{\mathcal{E}^c} + L_\tau \mathbb{I}_{\mathcal{E}} \le S_\tau \mathbb{I}_{\mathcal{E}^c} + (S_\tau + \delta) \mathbb{I}_{\mathcal{E}} \le S_\tau + \delta \mathbb{I}_{\{\tau_1 < T\}}.$$

Here, for the first inequality, I used the fact that an impatient buyer will only buy at time τ is $L_\tau \le S_\tau + \delta$, and, for the second inequality, I used the fact that the event $\mathcal{E}$ can only occur if an impatient buyer arrives in the time interval $[0, \tau)$. Denote by $\bar{h}_0$ the value under an optimal strategy. Using the fact that τ is a bounded stopping time and the fact that S_t is a martingale, by the optional sampling theorem,

$$\bar{h}_0 \le \mathsf{E}[P] - S_0 \le \mathsf{E}[S_\tau + \delta \mathbb{I}_{\{\tau_1 < T\}}] - S_0 = \delta \mathsf{P}(\tau_1 < T) = \delta \left(1 - e^{-\mu T} \right).$$

On the other hand, the hypothesized strategy results in equality in (2.7). Thus, the result follows. ∎

---

[7]For example, see Bertsimas and Lo [1998] or Almgren and Chriss [2000]. These questions have also recently been addressed by Back and Baruch [2007] and Pagnotta [2010] in equilibrium settings.

[8]I denote by $\mathbb{I}_{\mathcal{E}}$ the indicator function of the event $\mathcal{E}$.
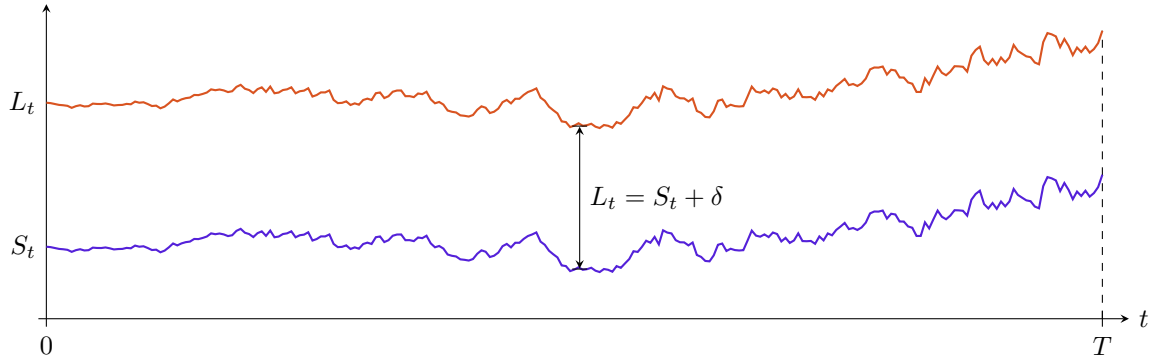
Figure 2.2: An illustration of an optimal strategy with no latency, over the time horizon $[0, T]$. The trader uses only limit orders prior to end of the time T. The limit order price L_t is pegged to the bid price S_t, with an additional premium corresponding to the bid-offer spread δ.

The optimal pegging strategy suggested by Lemma 1 is illustrated in Figure 2.2. This policy can be interpreted intuitively as follows: since the trader is risk-neutral and the bid price process is a martingale, the trader is indifferent between trading at time 0 at the bid price or trading at any other time at the bid price. Via a limit order, however, the trader can receive a price which is in excess of the bid price. The excess premium is limited to δ, since an impatient buyer will not pay more than this. Hence, the trader maintains a single limit order in the book, and continuously updates the price to track bid price, plus an additional premium of δ.

Note that my stylized execution model captures only the behavior of a single agent. My model does not capture the strategic response of other agents, either competing agents submitting limit orders to sell, or contra-side impatient buyers. Both of these types of agents might be expected to react to the activity of the limit order trader, and may diminish the gains of the limit order trader. Separately, my model also exaggerates the gains to be earned by placing limit orders rather than market orders, due to the fact I do not include adverse selection costs incurred by limit orders.

However, at a high level, a trader in my model with a mandate to trade over a fixed time horizon but with no private information as to the asset value prefers limit orders to market orders. I believe this is representative of the situation of algorithmic traders executing large "parent" orders in practice. When executing a "child" order over a short time horizon,

such traders typically first submit limit orders, and then "clean up" with market orders as time runs short. Hence, despite omissions of strategic considerations and other significant simplifications, the resulting policies do capture representative features of real world trading, if only at a stylized level. Moreover, my simplified single-agent mode enables us to address the dynamic nature of trade execution and obtain a closed-form expression highlighting the exact drivers of the latency cost.

2.3. A Model for Latency

The optimal policy for the stylized execution problem of Section 2.2 relied on the ability of a trader to continuously track an informational process, namely, the bid price in the market, and to update his order as the process evolves. Here, I will consider a variation of that problem where the trader is unable to continuously participate in the market, but faces a fixed latency $\Delta t > 0$. [9] I am interested in quantifying the cost of this latency by comparing the expected payoff in this model to that in the stylized model without latency. Note that the model at hand is quite basic with regards to some of primitives (e.g., the stochastic process describing the evolution of bid prices), I will discuss a number of tractable extensions in Section 2.4.4, including more complicated models of the bid price process and of limit order execution.

In general, latency that a trader experiences can take many forms. Minimally, for example, there is the delay of the data feeds that deliver market price information to the trader. There is the delay of the trader's own decision making. Finally, there is the delay of the trader's resulting order reaching the marketplace. I assume that the trader makes decisions instantaneously — we will see that this is reasonable since the optimal decision rule for the trader will take a very simple form. Further, from the trader's perspective, the roundtrip delay (the total delay for an order to be processed by an exchange and reflected in

---

[9]Note that many modern exchanges explicitly allow for pegged orders; these orders obviate the need for the trader to continually track the bid price in the manner I describe. However, more generally, when tracking an alternative informational process such as the price on a different exchange, the fundamental value (see Section 2.2), etc., a trader would still need to continuously monitor the market relative to the informational process, and latency would be important.
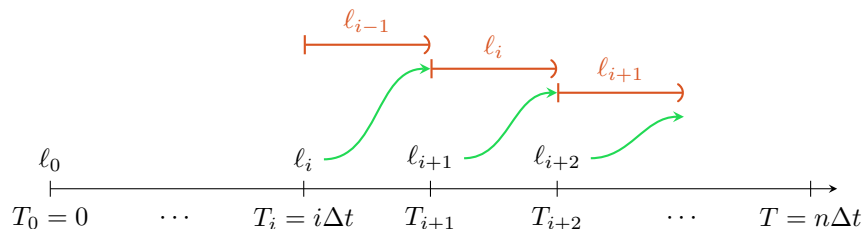
Figure 2.3: An illustration of the model of latency. Here, the time horizon $[0, T]$ is divided into n slots, each of duration equal to the latency Δt. The limit order price ℓ_i is decided at the start of the ith time slot, i.e., at time T_i. This price only takes effect Δt units of time later, and is active during the subsequent time interval $[T_{i+1}, T_{i+2})$.

the data feeds observed by the trader) cannot be decomposed into a delay to the exchange and a delay from the exchange. Hence, without loss of generality, I will assume that the trader is able to observe market price information with no delay or latency,[10] but that the trader's orders experience a latency Δt before they are processed by the exchange. This latency is meant to capture, for example, networking or routing delays that are specific to the trader, and that might be reduced through colocation or additional investment in networking technology.

In my latency model, I consider an investor who maintains a limit order to sell one share over the time horizon $[0, T]$ (the possibility of market orders will be discussed shortly), so that once the limit order is executed, the investor immediately exits the market. The time horizon $[0, T]$ is divided into n slots each of length Δt, i.e., $T = n\Delta t$. For each $i \in \{0, 1, \ldots, n\}$, define $T_i \triangleq i\Delta t$.

At each time T_i, based on all information observed thus far, I assume that the trader can instantaneously decide to update the limit order with a new price ℓ_i. Due to a latency of Δt, the updated price does not reach the market and take effect until the beginning of the next time slot, i.e., T_{i+1}. This limit order price remains active until time T_{i+2}, at which point it is superseded[11] by the next price ℓ_{i+1}. This sequence of events is illustrated in

---

[10]Equivalently, we can assume that my definition of time corresponds to the trader's clock.

[11]In practice, this ordering scheme might be achieved by a sequence of cancel-and-replace limit orders, each of which cancels the prior limit order, and inserts a new limit order with the updated price. If the prior limit order has already been filled when a subsequent cancel-and-replace order arrives, the new order will

Figure 2.3. Between the time T_i, when the price ℓ_i is decided, and the time T_{i+1}, when the updated order reaches the market, the following events can occur:

- $\mathcal{E}_i^{(1)}$: An impatient buyer arrives in the time interval (T_i, T_{i+1}) and $\ell_{i-1} \leq S_{T_i} + \delta$, i.e., the *prior* limit price ℓ_{i-1}, which is active at that time, is within a margin δ of the bid price at the start of the interval. In this case, the limit order executes at the price ℓ_{i-1}, and the investor leaves the market. Note that the updated limit price ℓ_i never takes effect.

  I assume that the probability that an impatient buyer arrives in any given time slot is $\mu \Delta t$, and that these arrivals occur independently of everything else.[12] I assume that $\Delta t < 1/\mu$ so that this probability is well-defined. The bid price process evolves according to the random walk (2.2).

- $\mathcal{E}_i^{(2)}$: Otherwise, if $S_{T_{i+1}} \geq \ell_i$, i.e., the bid price has crossed the order price ℓ_i at the instant the order reaches the market, then the order immediately executes at price $S_{T_{i+1}}$.

- $\mathcal{E}_i^{(3)}$: Otherwise, the limit order price ℓ_i is active over the time interval $[T_{i+1}, T_{i+2})$.

In order to consider the possibility of market orders, I allow the limit price $\ell_i = -\infty$. By picking this price, the trader can guarantee that the bid price at time T_{i+1} will cross the order price, i.e., $S_{T_{i+1}} \geq \ell_i$ with probability 1. Thus, the choice of $\ell_i = -\infty$ corresponds to a certain execution at the bid price $S_{T_{i+1}}$, i.e., a market order. Similarly, the trader can make the decision at time T_i not to trade by setting $\ell_i = \infty$. As in the model of Section 2.2, if the investor has been unable to sell the share by the end of the time horizon T, the investor is forced to sell via a 'clean-up' trade, i.e., a market order at time T. This is accomplished by enforcing the constraint that $\ell_{n-1} = -\infty$, which I will assume implicitly in what follows.

As before, if P is the random variable associated with the sale price, the trader is risk-neutral and seeks to solve the optimization problem

$$(2.8) \qquad\qquad h_0(\Delta t) \triangleq \underset{\ell_0,\dots,\ell_{n-1}}{\text{maximize}} \ \mathsf{E}\,[P] - S_0.$$

---

fail. Hence, the investor is guaranteed to sell at most one share.

[12]Note that this is simply a discrete-time Bernoulli arrival process that is analogous to the the Poisson arrival process of Section 2.2.

Here, the maximization is over the choice of limit order prices $(\ell_0, \ell_1, \ldots, \ell_{n-1})$. I assume that the price decisions are non-anticipating, i.e., each ℓ_i is adapted to the filtration generated by the bid price process and the arrival of impatient buyers up to and including time T_i. My goal is to analyze $h_0(\Delta t)$, which is the value under an optimal trading strategy when the latency is Δt.

Note that, as compared to the model of Section 2.2, my present model with latency differs in two ways: First, the trader makes decisions at the beginning of discrete-time intervals of length Δt, as opposed to continuously. Second, the orders of the trader incur a latency or delay of length Δt before they reach the marketplace. I am interested in studying the impact of the latter feature, latency, and I adopt the former feature, discrete-time decision making, so as to admit a tractable dynamic programming analysis. In Section 2.4.3, however, we will see that in the low latency regime in which we are most interested, the discrete-time nature of my model has a negligible impact.

2.4. Analysis

In this section, I solve for the optimal policy for the trader in the latency model of Section 4.2. This problem can be solved via a dynamic programming decomposition that is presented in Section 2.4.1. While the exact dynamic programming solution can be computed numerically, in Section 2.4.2 I will present an asymptotic analysis that provides a closed-form analytic expression for the cost of latency in the low latency regime, where $\Delta t \to 0$. In Section 2.4.3, I will consider the implications of the discrete-time nature of my latency model. Finally, in Section 2.4.4, I will discuss a number of extensions of my latency model.

2.4.1. Dynamic Programming Decomposition

The standard approach to solving the optimal control problem (2.8) is to employ dynamic programming arguments. In Appendix A.1, I formally derive the optimal control policy using these methods. In order to focus on the high level picture, however, for the moment I will be content with summarizing those results.

In particular, assume a fixed latency of Δt. For each decision time T_i with $0 \leq i < n$, define $\mathcal{U}_i$ to be the event that the trader's limit order remains unfulfilled prior to time T_{i+1}, i.e., none of the orders submitted at prices $\ell_0, \ldots, \ell_{i-1}$ are executed. Note that if the event $\mathcal{U}_i$ *does not* hold, then the limit order price ℓ_i to be decided at time T_i is irrelevant. This is because, by the time that order arrives to the market, the trader would have already sold a share. Define the quantity

$$(2.9) \qquad h_i \triangleq \underset{\ell_i, \ldots, \ell_{n-1}}{\text{maximize}} \; \mathsf{E}\left[P \mid S_{T_i}, \mathcal{U}_i \right] - S_{T_i}.$$

Note that $h_0 = h_0(\Delta t)$, where $h_0(\Delta t)$ is defined in (2.8), and thus my notation is consistent. More generally, for $i > 0$, I can interpret h_i to be the trader's expected payoff at time T_i relative to the current bid price S_{T_i} under the optimal policy, the order does not get filled prior to time T_{i+1}. Thus, h_i can be interpreted as a *continuation value* in the dynamic programming context.

The continuation values $\{h_i\}$ quantify the remaining value for a trader at each time period if his order remains unfulfilled. Given the continuation values, at each time T_i, the investor can make an optimal decision as to the limit order price ℓ_i by balancing the benefits of execution in the time slot $[T_{t+1}, T_{i+2})$ with the value h_{i+1} that will be obtained if the order is not executed. Moreover, the optimal decisions and continuation values can be jointly computed via backward induction of a Bellman equation. This result is captured in the following theorem. The proof, which is provided in Appendix A.1, follows from formal dynamic programming arguments.

Theorem 1. *Suppose $\{h_i\}$ satisfy, for $0 \leq i < n-1$,*

(2.10)
$$h_i = \max_{u_i} \; \left\{ \mu \Delta t \left[u_i \left(\Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) + \sigma \sqrt{\Delta t} \left(\phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) \right] \right.$$
$$\left. + h_{i+1} \left[(1 - \mu \Delta t) \Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) + \mu \Delta t \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right] \right\},$$

and

$$(2.11) \qquad\qquad\qquad\qquad h_{n-1} = 0.$$

Here, ϕ and Φ are, respectively, the p.d.f. and c.d.f. of the standard normal distribution. Then, $\{h_i\}$ correspond to the continuation values under the optimal policy.

Suppose further that, for $0 \leq i < n - 1$, u_i^* is a maximizer of (2.10). Then, a policy which chooses limit order prices which are pegged to the bid prices according to the premia defined by $\{u_i^*\}$, i.e.,

$$\ell_i^* = S_{T_i} + u_i^*, \quad \forall \, 0 \leq i < n - 1,$$

is optimal.

Theorem 1 suggests a computational strategy for determining continuation values and an optimal policy. Starting with the terminal condition $h_{n-1} = 0$, one proceeds via backward induction, solving the single variable optimization problem (2.10) over the decision variable u_i once per time slot. So long as optimal solutions exist, they will determine the continuation values and optimal policy. Moreover, the optimal policy is a pegging strategy. That is, the limit order price is pegged at a deterministic (but time varying) premium above the current bid price. These limit order premia are given by the maximizers $\{u_i^*\}$.

In the following theorem, whose proof is provided in Appendix A.2, I establish the existence and uniqueness of the optimal solutions to (2.10) and provide upper and lower bounds for the resulting limit price premia, for small values of latency Δt.

Theorem 2. *Fix $\alpha > 1$. If Δt is sufficiently small, then there exists a unique optimal solution $\{h_i\}$ to the dynamic programming equations (2.10)–(2.11). Moreover, the corresponding optimal policy $\{u_i^*\}$ is unique. For $0 \leq i < n - 1$, this strategy chooses limit prices in the range*

$$\ell_i^* \in \left(S_i + \delta - \sigma\sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, \ S_i + \delta - \sigma\sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}} \right),$$

where

$$L \triangleq \frac{\delta^2}{2\pi\sigma^2}, \qquad\qquad R(\Delta t) \triangleq \frac{\delta^2(1 - \mu\Delta t)^{2n}}{2\pi\sigma^2}.$$

Figure 2.4 illustrates the intuition behind Theorem 2, by considering the situation of a trader at time $t = 0$, when the bid price is S_0. In the absence of latency, the trader would peg the limit order price at a fixed premium of δ, i.e., $\ell_0 = S_0 + \delta$. This would result in a trade with the next impatient buyer with probability 1. If there is latency present, however, this limit price is not optimal. To see this, suppose that an impatient trader will arrive at
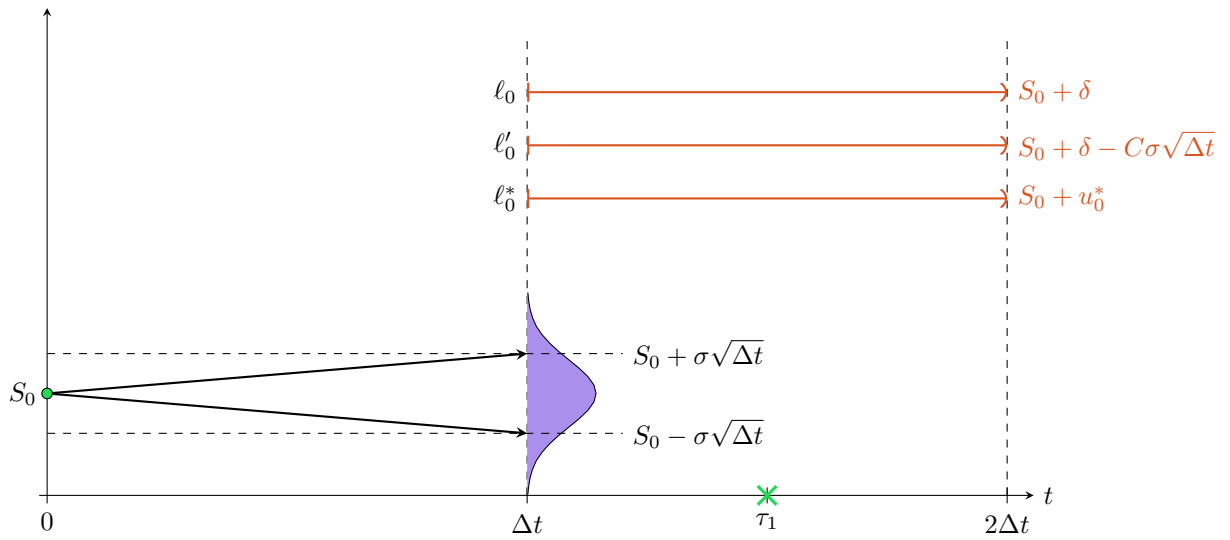
Figure 2.4: An illustration of the optimal policy of Theorem 2. In the absence of latency, at time $t = 0$, the trader would set the limit price at a premium of δ, i.e., $\ell_0 = S_0 + \delta$. In an environment with latency, the trader might set the limit price to be ℓ'_0, which lowers ℓ_0 by an additional safety margin of C standard deviations. This serves to increase the likelihood of trade execution in the interval $(\Delta t, 2\Delta t)$. The optimal limit price ℓ^*_0 utilizes a safety margin that is slightly larger.

time $\tau_1 \in (\Delta t, 2\Delta t)$. If the limit order price is set at ℓ_0, the probability that the trade *does not* get executed is

$$\mathsf{P}\left(\ell_0 \geq S_{\Delta t} + \delta\right) = \mathsf{P}\left(S_0 \geq S_{\Delta t}\right) = 1/2.$$

When Δt is small, the probability of missing an execution can be significantly lowered at a small cost by lowering ℓ_0 by an additional safety margin. If we set this safety margin to be C standard deviations of the one-period price change, i.e., $\ell_0' = S_0 + \delta - C\sigma\sqrt{\Delta t}$, then the probability of missing execution becomes

$$\mathsf{P}\left(\ell_0' \geq S_{\Delta t} + \delta\right) = \mathsf{P}\left(S_0 - C\sigma\sqrt{\Delta t} \geq S_{\Delta t}\right) = \Phi(-C).$$

This probability can be made close to 0 by the choice of C. However, given a fixed choice of C independent of Δt, the probability remains constant (i.e., independent of Δt) and non-zero. The additional safety margin corresponding to the log term in Theorem 2 is a second order adjustment. This is introduced so that, given the optimal limit price ℓ_0^*, the probability of execution tends to 1 as $\Delta t \to 0$.

2.4.2. Asymptotic Analysis

The dynamic programming decomposition developed in Section 2.4.1 allows the exact numerical computation of the value $h_0(\Delta t)$, the value under an optimal policy of the latency model introduced in Section 4.2, when the latency is Δt. As discussed earlier, the latency observed in modern electronic markets is extremely small, often on the time scale of milliseconds. Thus, we are most interested in the qualitative behavior of $h_0(\Delta t)$ in the asymptotic regime where $\Delta t \to 0$. The main result of this section is the following theorem, whose proof is provided in Appendix A.3. It provides a closed-form expression for $h_0(\Delta t)$, which holds asymptotically[13] as $\Delta t \to 0$.

---

[13]In what follows, given arbitrary functions f and g, and a positive function q, I will say that $f(\Delta t) = g(\Delta t) + O(q(\Delta t))$ if $\limsup_{\Delta t \to 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) < \infty$, i.e., if the difference between f and g, as $\Delta t \to 0$, is asymptotically bounded above by *some* positive multiple of q. Similarly, I will say that $f(\Delta t) = g(\Delta t) + o(q(\Delta t))$ if $\lim_{\Delta t \to 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) = 0$, i.e., if the difference between f and g, as $\Delta t \to 0$, is asymptotically dominated by *every* positive multiple of q. Finally, I will say that $f(\Delta t) = g(\Delta t) + \Theta(q(\Delta t))$ if $0 < \liminf_{\Delta t \to 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) \leq \limsup_{\Delta t \to 0} |f(\Delta t) - g(\Delta t)|/q(\Delta t) < \infty$, i.e., if the difference between f and g is asymptotically bounded above and below by positive multiples of q.

Theorem 3. *As $\Delta t \to 0$,*

$$h_0(\Delta t) = \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2 \Delta t}} \right) + o\left(\sqrt{\Delta t}\right),$$

where

$$\bar{h}_0 = \delta \left(1 - e^{-\mu T} \right)$$

is the optimal value for the stylized model without latency, i.e., the value defined by (2.5).

Theorem 3 is not surprising when considered in the context of Theorem 2. In the stylized model without latency, the optimal strategy is to peg the limit order price at a premium of δ, and this yields a value of $\bar{h}_0$. On the other hand, Theorem 2 suggests a trader facing latency Δt will lower this limit price premium by a factor of, approximately,

$$\frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2 \Delta t}} + o\left(\sqrt{\Delta t}\right).$$

If this lowers the ultimate value proportionally, then the value of the optimal policy in the presence of latency Δt should approximately be

$$\bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2 \Delta t}} \right) + o\left(\sqrt{\Delta t}\right).$$

The proof of Theorem 3, provided in Appendix A.3, makes this intuition precise.

One implication of Theorem 3 is that $h_0(\Delta t) \to \bar{h}_0$ as $\Delta t \to 0$, i.e., the value of the latency model converges to that of the stylized model without latency of Section 2.2. This suggests the following definition:

Definition 1. *Define the **latency cost** associated with latency Δt by*

(2.12) $$\mathsf{LC}(\Delta t) \triangleq \frac{\bar{h}_0 - h_0^*(\Delta t)}{\bar{h}_0}.$$

Latency cost has an easy interpretation. Using $\bar{h}_0$, the value obtained in the stylized model without latency as a benchmark, the numerator of (2.12) is the *lost revenue* incurred due the the presence of latency. On the other hand, we can regard the denominator as the *cost of immediacy* for an impatient investor in a time horizon of length T. This is because, in the stylized model without latency, it is the difference in revenue obtained by

a risk-neutral investor willing to patiently provide liquidity by employing limit orders over the length of the time horizon, and an impatient investor who demands immediate liquidity and sells at the bid price at time $t = 0$, cf. (2.5). Therefore, we can describe the latency cost as the amount a trader forgoes due to latency, as a percentage of the cost of immediacy.

The following corollary restates the asymptotic approximation of Theorem 3 in terms of latency cost.

Corollary 1. *As $\Delta t \to 0$,*

$$\mathsf{LC}(\Delta t) = \frac{\sigma\sqrt{\Delta t}}{\delta}\sqrt{\log \frac{\delta^2}{2\pi\sigma^2\Delta t}} + o\left(\sqrt{\Delta t}\right).$$

There are a number of interesting observations that can be made regarding the asymptotic approximation of Corollary 1. First of all, asymptotically, latency cost *does not* depend on the length of the time horizon T or the arrival rate of impatient traders μ. As a function of the remaining parameters, the asymptotic latency cost depends only on a composite parameter that is the ratio the one-period standard deviation of price changes $\sigma\sqrt{\Delta t}$ to the bid-offer spread δ. Both of these quantities are readily measurable empirically. Corollary 1 suggests that the latency cost increasing in this ratio. Thus, at the same level of latency, the latency cost is most significant for assets which are very volatile or very liquid. Further, Corollary 1 suggests that, when latency is low, there are increasing marginal benefits to further reductions in latency, i.e., $\mathsf{LC}''(\Delta t) < 0$. In Section 2.5.1, I illustrate some of facts numerically, as well as considering the accuracy of my approximation, as compared to the exact latency cost.

2.4.3. Discreteness of Time vs. Latency

The latency model introduced in Section 4.2 differs from the the stylized model without latency of Section 2.2 in two principal ways: (i) the trader faces a delay or latency between the time that trading decisions are made and when they reach the marketplace, and (ii) the latency model is formulated in discrete-time rather than continuous time. The latter point refers to the facts that, in the model with latency, a trader is only able to update his limit order at discrete intervals of time rather than continuously, impatient buyers arrive according to a Bernoulli process rather than a Poisson process, etc. In order to disentangle

these two effects, in this section I will briefly describe a trading model that is formulated in discrete time but *without* latency. By considering this model, I will demonstrate that the asymptotic latency cost derived in Section 2.4.2 is indeed due to latency effects and not due to the discreteness of time.

To this end, consider a model in the discrete-time setting of Section 4.2 but with no latency. Here, at each time $T_i \triangleq i\Delta t$, for $i = 0, 1, \ldots, n$, the investor sets a limit order price ℓ_i. This limit order price takes effect immediately. Between time T_i and time T_{i+1} the following events can occur:

- If $S_{T_i} \leq \ell_i$, i.e., the bid price is less than the limit order price, the limit order immediately executes at the price S_{T_i}.

- Otherwise, suppose that an impatient buyer arrives in the time interval (T_i, T_{i+1}) and $\ell_i \leq S_{T_i} + \delta$, i.e., the limit price ℓ_i is within a margin δ of the bid price at the start of the interval. In this case, the limit order executes at the price ℓ_i. I assume that an impatient buyer arrives with probability $\mu\Delta t$, independent of everything else.

As before, if the investor is unable to sell the share by the end of the time interval, he is forced to sell via a market order, i.e., $\ell_n = -\infty$. If P is the sale price, the optimal value for the trader in this discrete model is given by

$$h_0^{\mathsf{D}}(\Delta t) \triangleq \underset{\ell_0, \ldots, \ell_n}{\text{maximize}} \ \mathsf{E}\,[P] - S_0.$$

I have the following result, whose proof is identical to the martingale argument used to establish Lemma 1.

Lemma 2. *An optimal strategy for the discrete model is to place limit orders at the price* $\ell_i = S_{T_i} + \delta$, *for* $i = 0, 1, \ldots, n-1$. *This strategy achieves the value*

$$h_0^{\mathsf{D}}(\Delta t) \triangleq \delta\,(1 - (1 - \mu\Delta t)^n).$$

Now, note that, for all $0 < \Delta t < 1/\mu$,

$$e^{-\mu T - \frac{1}{2}\mu^2 T \Delta t} \leq (1 - \mu\Delta t)^{T/\Delta t} \leq e^{-\mu T}.$$

Therefore, the difference in value between the continuous model of Section 2.2 and the discrete model considered here is at most

$$|h_0^{\mathsf{D}}(\Delta t) - \bar{h}_0| \le \delta e^{-\mu T}\left(1 - e^{-\frac{1}{2}\mu^2 T \Delta t}\right) \le \tfrac{1}{2}\delta \mu^2 T e^{-\mu T}\Delta t.$$

In other words, this difference is asymptotically $O(\Delta t)$. By Theorem 3, however, the difference between the continuous model and the latency model is asymptotically

$$\Theta\left(\sqrt{\Delta t \log(1/\Delta t)}\right).$$

Hence, the asymptotic effect of latency dominates the asymptotic effect of the discreteness of time.

2.4.4.   Extensions

The analysis of the latency model that I have presented proceeded according to two high level steps:

(i) First, in Section 2.4.1, a simplified dynamic programming decomposition was developed. In this decomposition, at each time, the trader's value function is parameterized by a single scalar, rather than being an arbitrary function of state. This allows the Bellman equation to be solved through a system of n equations in n unknowns, given by (2.10)–(2.11).

(ii) Second, in Section 2.4.2, an asymptotic analysis of the simplified dynamic programming equations (2.10)–(2.11) was performed. This gave rise to the asymptotic latency cost expression of Corollary 1.

The dynamic programming decomposition step (i) that is at the heart of my analysis can be extended to a much broader set of stochastic primitives than the present setting. In each of these cases, a different set of simplified dynamic programming equations, analogous to (2.10)–(2.11) would arise, and would require a customized variation of asymptotic analysis step (ii). In particular, consider the following tractable generalizations:

- **Price process.** In my model, the price process S_t is a Brownian motion. My dynamic programming decomposition only requires that the S_t be a Markov process and a

martingale. It would be straightforward to extend the dynamic programming step (i) and consider other Markovian martingales, for example, allowing for non-Gaussian processes, time-inhomogeneous volatility, or for jump processes.

On the other hand, the asymptotic analysis step (ii) I have presented is quite sensitive to distributional assumptions of the price process, and would require specialized analysis for any such generalization. In Appendix A.4, I consider one generalization of particular interest, where the price dynamics also contain a jump component.

- **Limit order execution.** In my model, the execution of a limit order in the time slot (T_i, T_{i+1}) required that the limit order price ℓ_{i-1} be within a spread δ of the bid price S_{T_i}, and that an impatient trader arrive. More generally, my dynamic programming decomposition only requires that the execution of this limit order, conditional on the price difference $\ell_{i-1} - S_{T_i}$, be independent of everything else. This can accommodate a number of generalizations, for example, the arrival rate of impatient buyers can be time-varying. Further, the maximum premium above the bid price S_t that an impatient buyer is willing to pay can be randomly distributed, as in (2.3). This would allow models where a limit order that is priced aggressively low has a much higher probability of execution. Such models could alternatively be interpreted, as discussed in Section 2.2, as cases where the prevailing bid-offer spread is not constant, but is independent and identically distributed, varying from period to period.

2.5.   Empirical Estimation of Latency Cost

In this section, I will consider empirical applications of my model. First, I will illustrate the optimal trading policy and the corresponding value function when the model parameters are estimated from high frequency market data for a single stock. I will also compare the exact latency cost (numerically computed via dynamic programming) to the approximation provided by Corollary 1 in order to access the quality of my approximation. Subsequently, I show the historical evolution of latency cost and implied latency across a range of U.S. equities using cross-sectional data on volatilities and bid-offer spreads during the 1995–2005 period.

My empirical analysis should be regarded as a first-order study to obtain a rough calibration of my model. It will allow us to analyze the model in relevant parameter regimes, as well as gaining a broad understanding the implications of my model for the trading of U.S. equities. My empirical measurement of latency cost requires estimates of, in particular, the high frequency price volatility σ and the prevailing bid-offer spread δ. Here, I make a number of simplifications and rely on the recent empirical work of Aït-Sahalia and Yu [2009] to obtain these quantities:

- I estimate price volatility σ using the maximum likelihood estimates of the volatility of returns provided by Aït-Sahalia and Yu [2009]. Note that this estimation of high frequency volatility aims to filter out the impact of microstructure noise and obtain an unbiased estimate of daily volatility. However, for an investor with a trading horizon of 1 second, microstructure noise needs to be incorporated as well. Therefore, the high frequency volatility estimate that is used in my empirical analysis underestimates the actual volatility faced by a high frequency trader with a very short trading horizon.

- Recall that the prevailing bid-offer spread, δ, equals the bid-offer spread in the absence of the liquidating trader. In the empirical data, it is impossible to disentangle the presence of liquidating traders. Moreover, the bid-offer spread will not be constant, but will vary over the course of the trading day. As a proxy for δ, I use the average bid-offer spread over the trading day.

Despite these shortcomings, I believe that my empirical analysis can shed light on the importance of latency in the trading of U.S. equities.

2.5.1. The Optimal Policy and the Approximation Quality

In what follows, I will numerically evaluate the optimal policy in my model, the corresponding value function, and the latency cost approximation. These numerical experiments are meant to be illustrative of my model. I will use realistic model parameters estimated from recent market data for a single stock. My methodology here is not meant to be authoritative — there are many subtleties in the analysis of high frequency data; these are beyond the scope of the work at hand. However, I do seek to demonstrate that my model parameters

can be readily derived from commonly available data.

Specifically, the model parameters herein are estimated from trade-and-quote (TAQ) data for a stock that is a representative example of a liquid name, Goldman Sachs Group, Inc. (NYSE: GS), on the trading day of January 4, 2010. This data was obtained from the Wharton Research and Data Services (WRDS) consolidated TAQ database. Only trades and quotes originating from the primary exchange (NYSE) during regular trading hours were considered. The model parameters were estimated as follows:

- Initial bid price: $S_0^{\mathsf{GS}} = \$170.00$. This was chosen to be the first transaction price on the trading day.

- Bid-offer spread: $\delta^{\mathsf{GS}} = \$0.058$, i.e., equivalently, 3.4 basis points relative to the initial price S_0^{GS}. This was estimated by computing the average spread between bid and offer quotes over the course of the trading day and rounding to the nearest cent.

- Arrival rate of market orders: $\mu^{\mathsf{GS}} = 12.03$ (per minute). This was estimated by dividing the total number of NYSE trades by the length of the trading day.

- Price volatility: $\sigma^{\mathsf{GS}} = \$1.92$ (daily), i.e., approximately equivalent to an annualized volatility of returns of 17.9%. These were estimated from the time series of transaction prices over the course of the trading day, using maximum likelihood estimation as described in Aït-Sahalia and Yu [2009].

- Trading horizon: $T = 10$ (seconds).

Figure 2.5 illustrates the optimal limit order policy for GS under different values of latency. If there is no latency, the limit orders are submitted at a constant premium of δ. When there is latency, the optimal order policy is obtained using the exact dynamic programming solution of (2.10)–(2.11). As the latency increases, the limit order premium is reduced below δ so as to account for the increasing uncertainty of price movements over the latency interval. Theorem 2 suggests that this reduction is approximately equal to

$$(2.13) \qquad \sigma\sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2 \Delta t}}.$$
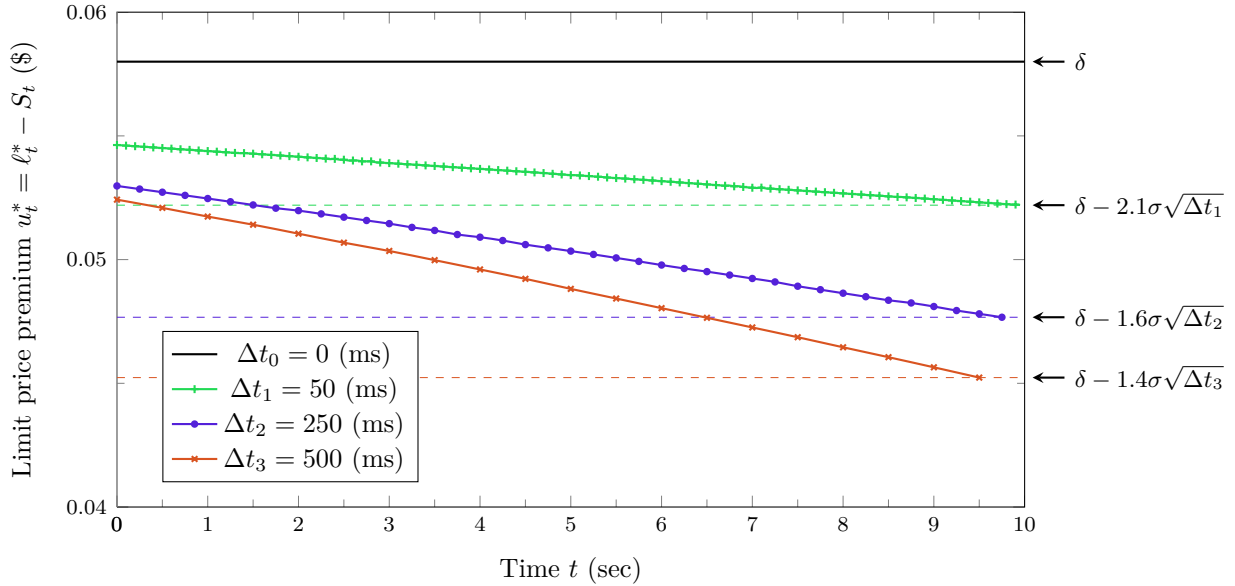
Figure 2.5: An illustration of the optimal strategy for GS, expressed in terms of limit price premium over the course of the time, for different choices of latency. In each case, the dashed line illustrates the relative distance below the bid-offer spread δ of the price premium of the final limit order, as a multiple of the standard deviation of prices over the latency interval.

In Figure 2.5, I see that, with a latency of 500 ms, this adjustment is up to approximately $1.4\sigma\sqrt{\Delta t}$, i.e., 1.4 times the standard deviation of prices over the latency interval. When the latency is reduced to 250 ms and to 50 ms, the adjustment increases to 1.6 and 2.1 standard deviations, respectively. The fact that this adjustment, when measured as a multiple of the uncertainty over the latency period, increases as the latency decreases is consistent with (2.13).

In Figure 2.5, I also observe that as t increases and the trading deadline approaches, the limit order premium u_t^* becomes lower. This makes intuitive sense: the trader faced with a terminal value of 0 since he is required to sell using market order at the end of the period. As the deadline approaches, the trader is more willing to sacrifice the potential profits of a limit order in order to increase the probability of execution.

Figure 2.6 illustrates the corresponding continuation value under the optimal policy for GS, for different values of latency. Clearly, the trader's expected payoff decreases as latency increases or the end of the trading horizon approaches.
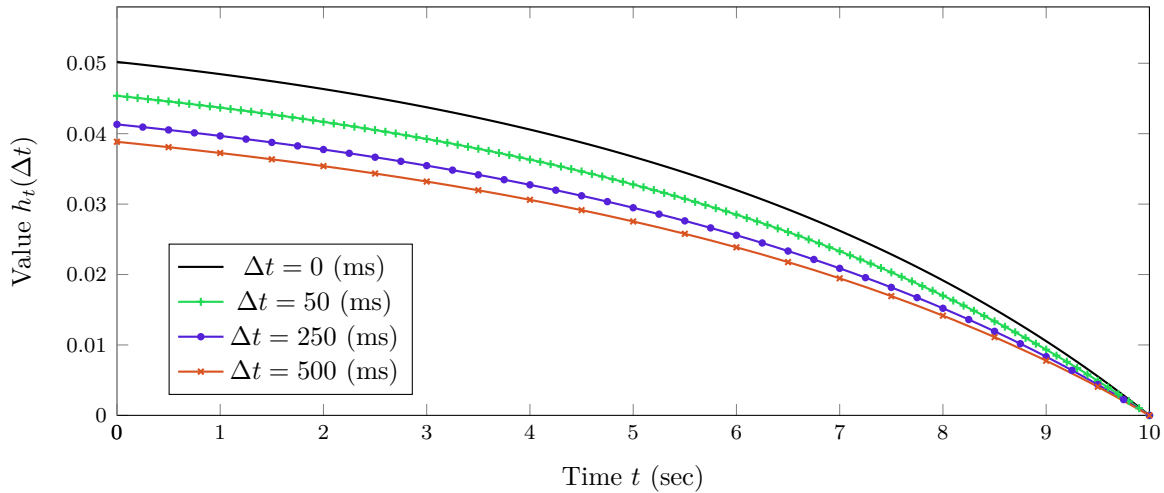
Figure 2.6: An illustration for the evolution of the continuation value of the optimal policy over time for GS, for different choices of latency. The expected value of the trader decreases as latency increases or as the end of the trading horizon approaches. As the latency increases from 0 ms to 500 ms, the trader loses more than 0.01 of the 0.05 cent spread, i.e., more than 20% of the spread.
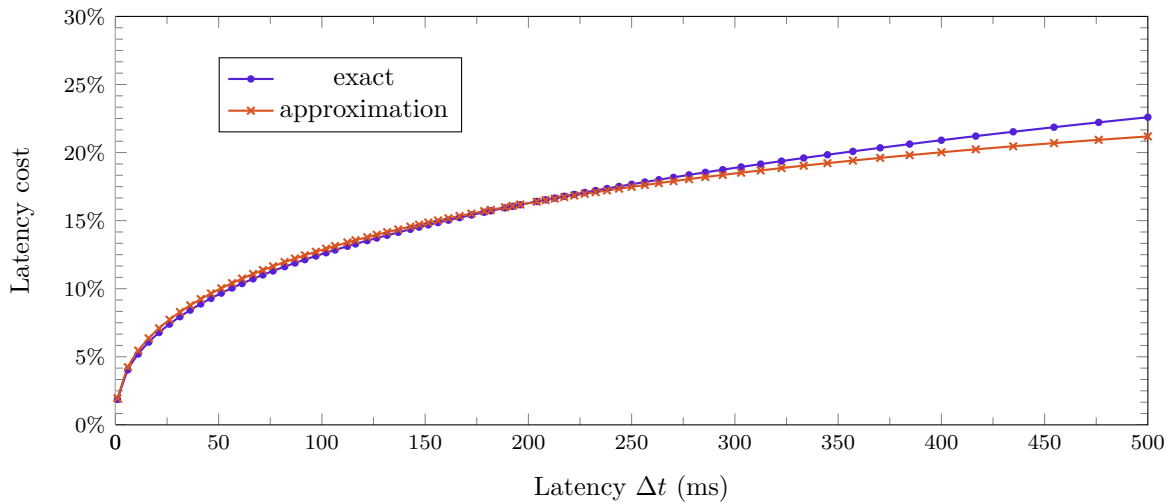
Figure 2.7: An illustration of the latency cost as a function of the latency. both the exact latency cost and the asymptotic approximation are shown. The approximate latency cost closely aligns with the exact latency cost across the entire range of latency values. This illustrates that my closed-form formula can accurately approximate the exact latency cost for low values of latency.

Finally, Figure 2.7 illustrates the latency cost as a function of latency. Both the exact value of the latency cost, computed numerically via the dynamic programming decomposition (2.10)–(2.11), and the asymptotic latency cost approximation provided by Corollary 1 are shown. The latency costs decrease from approximately 20% of the cost of immediacy to 5% of the cost of immediacy, as the latency decreases from 500 ms to 5 ms. Further, the marginal benefit of reducing latency increases as the latency approaches zero. Finally, I note that the approximate and exact latency costs are quite close across the entire range of latency values. This suggests that the approximation is of very high quality in this case.

2.5.2. Historical Evolution of Latency Cost

In this section, I will examine the historical evolution of latency cost in U.S. equities. Here, I consider the situation of a hypothetical investor with a fixed latency of 500 milliseconds. This choice of latency is made approximately to reflect the reaction time of a very fast human trader. I will use this as a proxy for the fastest possible trading on a "human time scale". By analyzing the evolution of the associated latency cost, I will get a sense of the importance of latency over time.

My empirical analysis relies on the data set of Aït-Sahalia and Yu [2009]. Their data set contains estimates for various liquidity measures for all NYSE common stocks on a daily basis during the sample period of June 1, 1995 to December 31, 2005. The estimates are derived from intraday transaction prices and quotes from the NYSE TAQ database. I utilized only the volatility and bid-offer spread data as we have seen both analytically (Corollary 1) and numerically (Figure 2.7) that latency cost can be approximated accurately for low values of latency using only these two measures.

The data set contain volatility and bid-offer spread estimates for given stock on a particular day if the number of transactions on that day exceeds 200. The minimum, average, and maximum number of stocks in the sample on any day are 61, 653, and 1,278, respectively. In particular, earlier periods in the data set contain fewer stocks due to a smaller number of firms and a lower volume of transactions. In this data set, the bid-offer spread is estimated using only NYSE quotes in the regular trading hours. The volatility estimate is obtained using maximum likelihood estimation in the presence of market microstructure

noise. Maximum likelihood estimation is preferred over other nonparametric estimation methods (e.g., "Two Scales Realized Volatility") as a simulation study shows that maximum likelihood estimation provides robust estimators under reasonable stochastic volatility and jump models in the underlying asset. The reader is urged to consult to Section 2.1 of Aït-Sahalia and Yu [2009] for full details of their estimation procedure.

For each stock in the data set, on a daily basis, I compute the latency cost facing an investor with a fixed latency of 500 ms using the asymptotic approximation of Corollary 1. These daily latency costs are then averaged over each month. Figure 2.8 displays percentiles of the monthly averages of latency cost over all of the stocks in the sample, as a function of time. As a representative example of a liquid name, I also report the monthly averages of latency cost of Goldman Sachs Group, Inc. (NYSE: GS). Note that the time series for GS begins from its initial public offering in 1999. For reference, I have added an additional point to this time series based on my estimation in Section 2.5.1 of the latency cost for GS on January 4, 2010.

Figure 2.8 illustrates that latency costs have had an increasing trend over the 1995–2005 period. In particular, we observe that the median latency cost incurred by trading on a human time scale roughly tripled, by increasing from approximately 3% to approximately 10%. One important factor in this increase has been the reduction of bid-offer spreads over this time period. Instances during the period when the NYSE reduced the tick size (from $1/8$ to $1/16$ in June 1997, and from $1/16$ to $0.01 in January 2001) coincide with spikes in latency cost. This is consistent with bid-offer spreads decreasing significantly and volatility maintaining the same level at these times. This suggests that any future reduction in tick sizes will result in increased latency costs.

Using a data set in a similar time-frame, from February 2001 to December 2005, Hendershott *et al.* [2010] conclude that in the post-decimalization era, the increase in algorithmic trading activity had a positive impact on the level of liquidity. This result suggests that the increase in algorithmic trading in and of itself elevated the importance of low latency trading and increased the cost of latency.
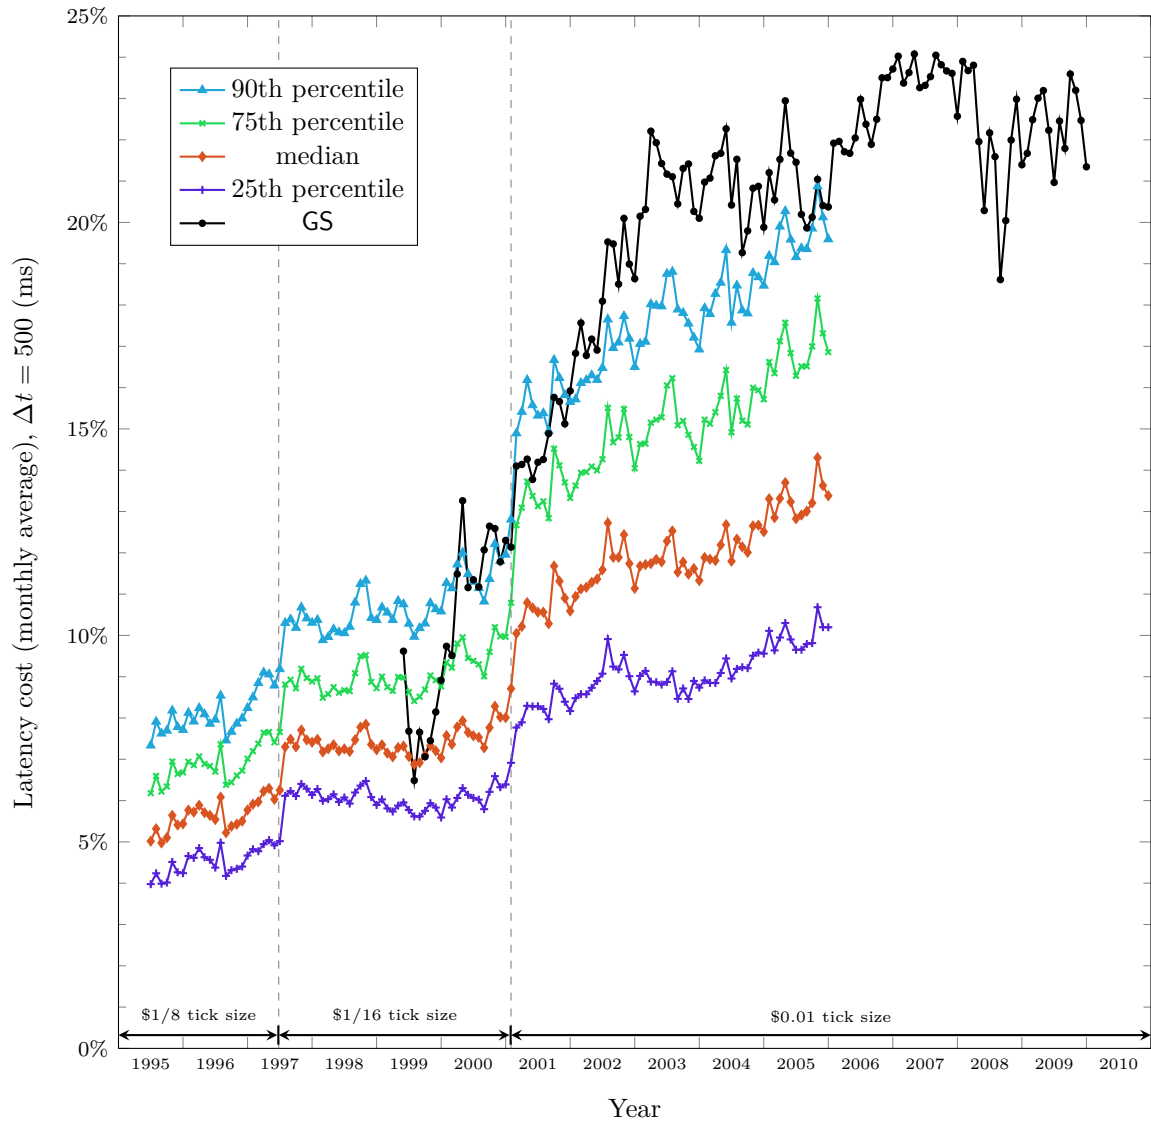
Figure 2.8: An illustration of the historical evolution of latency cost over the 1995–2005 time period. Here, I consider a hypothetical "human time scale" investor with a fixed latency of $\Delta t = 500$ (ms). Percentiles for the resulting latency cost are reported across NYSE common stocks. The latency costs are computed from data set of Aït-Sahalia and Yu [2009]. The latency cost for GS is also reported, beginning from its IPO. The dashed lines correspond to dates where the NYSE tick size was reduced. We observe that latency cost had a consistent increasing trend over the 1995–2005 period. Specifically, the median latency cost approximately increased three-fold by reaching roughly to 10% from 3%.

2.5.3. Historical Evolution of Implied Latency

An alternative perspective on the historical importance of latency comes from considering a hypothetical investor with a target level for the cost of latency, relative to the overall cost-of-immediacy. The representative trader maintains this target over time through continual technological upgrades to lower levels of latency. I determine the requisite level of latency for such a trader, over time and across the aggregate market. In other words, fixing the latency cost percentage LC to the target level, we can solve the asymptotic approximation (2.12) for the level of latency required at each time to achieve latency cost LC. I call this the *implied latency*.

Figure 2.9 illustrates the implied latency values over the 1995–2005 period assuming that the target level LC = 10% of overall transaction costs result from latency. We observe that the median implied latency decreased by approximately two orders of magnitude over this time frame. The 90th percentile of U.S. equities, for example, went from an implied latency on the scale of seconds to an implied latency on the scale of tens of milliseconds.

2.5.4. Empirical Importance of Latency

My model captures the cost of latency due to a lack of contemporaneous information. Figure 2.8 suggests that, when my model is calibrated to the topmost quartile of U.S. equities, a investor with latency on the human time scale faces a latency cost of at 15% to 25%. In order to assess the significance of this, we can compare it to other trading costs. Suppose we normalize the cost of immediacy to $0.01, which is the typical bid-offer spread for a liquid U.S. equity. Then, my model suggests that the benefit of reducing latency from a human time scale of 500 ms to an ultra low latency time scale of less than 1 ms is approximately $0.0015–$0.0025 per share traded.

While this might seem very small as an absolute number, note that is of the same order of magnitude as other trading costs faced by the most cost efficient institutional investors. For example, a hedge fund would pay an average commission of $0.0007 per share for market access.[14] Furthermore, investors may pay an SEC fee of $0.0005 per share

---

[14]"U.S. Equity Trading: Low Touch Trends," TABB Group, July 2010.
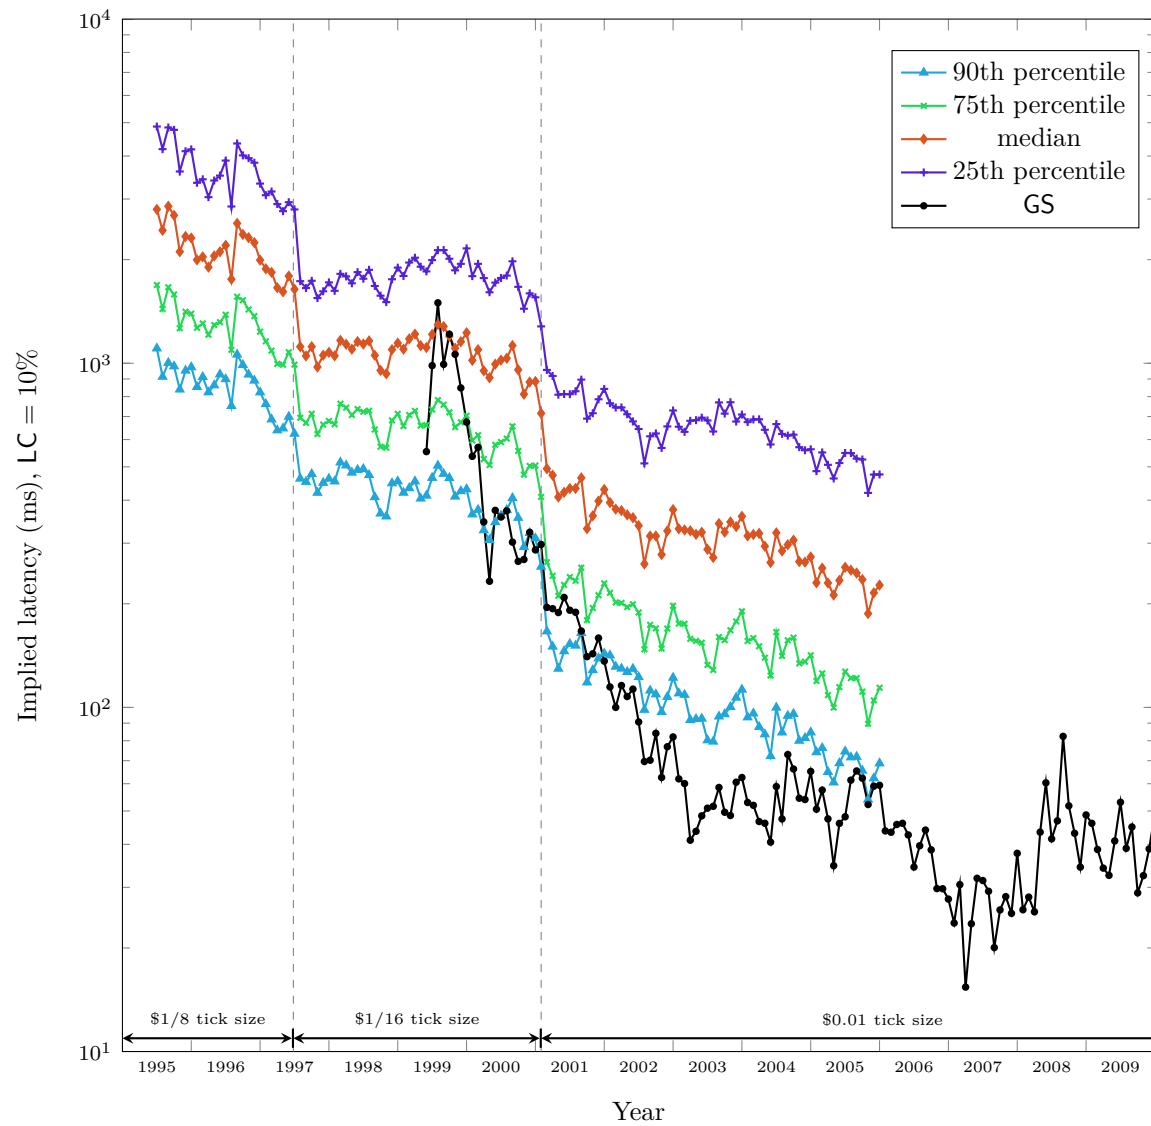
Figure 2.9: An illustration of the historical evolution of implied latency over the 1995–2005 time period. Here, I consider a hypothetical investor who makes sufficient technological investments to ensure a constant latency cost of 10%. The implied latency is the level of latency required to achieve this latency cost. Percentiles for the implied latency are reported across NYSE common stocks. The implied latencies are computed from data set of Aït-Sahalia and Yu [2009]. The implied latency for GS is also reported, beginning from its IPO. We observe that implied latency has had a decreasing trend over the 1995–2005 period. Specifically, the median implied latency decreased by approximately two orders of magnitude over this time frame.

traded,[15] and exchange fees or rebates of $0.0020–$0.0030 per share traded. To the extent that a sophisticated institutional investor is cost sensitive and wishes to optimize these other execution costs, they should also be concerned with latency. This isn't to suggest that latency cost is important to all investors. A typical retail investor, for example, may pay a brokerage fee that is up to $0.10 per share traded.[16] For this latter type of investor, the cost of latency as described here is not a significant component of overall trading costs.

Alternatively, we can compare the $0.0015–$0.0025 per share traded latency cost to the rents extracted by agents that have made the required technological investments to trade on an ultra low latency time scale. For example, providers of automated algorithmic trade execution services charge an average commission of $0.0033 per share traded for their execution services, which leverage sophisticated low latency technological infrastructure.[17] Note that this cost is comparable to the latency cost. Another class of agents with ultra low latency trading capabilities are high frequency traders. Reported net profit numbers for high frequency traders are in the range of $0.0010–$0.0020 per share traded.[18] This is of the same order of magnitude as the latency cost.

2.6. Conclusion and Future Directions

This chapter provides a model to quantify the cost of latency on transaction costs. I consider a stylized execution problem, where a trader must sell an atomic unit of stock over a fixed time horizon. I consider this model in the absence of latency as a benchmark, and

---

[15]As of January 21, 2011, the SEC fee is a fraction $0.0000192 of the proceeds of an equity sale. If we assume a typical stock price of $50, this is approximately $0.0010 per share sold. Amortizing this cost equally between buys and sells results in $0.0005 per share traded.

[16]For example, at the time of writing, the brokerage firm E-TRADE charges $10 per trade. Assuming a typical trade of 100 shares, this cost is $0.10 per share traded.

[17]"U.S. Equity Trading: Low Touch Trends," TABB Group, July 2010. Note that some institutional investors pay significantly larger commissions for trade execution in order to compensate their brokers for trading ideas or research services. The commission I quote here is for "non-idea driven" services that relate purely to trade execution using the algorithms and technological platform of the broker.

[18]"Tradeworx, Inc. Public Commentary on SEC Market Structure Concept Release," Tradeworx, Inc., April 2010.

I incorporate latency by not allowing the trader to continuously participate in the market. Orders submitted by the trader reach the market with a fixed latency, and the trader is forced to deviate from the benchmark policy in order to take into account the uncertainty introduced by this delay. I quantify the cost of latency as the normalized difference in expected payoffs between this model and the stylized model without latency.

Since the latency values observed in modern electronic markets are on the order of milliseconds, I provide an asymptotic analysis for the low latency regime, in which I obtain an explicit closed-form solution. In order to compute this asymptotic latency cost empirically, I only need to estimate the volatility and the average bid-offer spread of the stock. This is an elegant and practical result as data sets and estimation procedures for these quantities are readily abundant in the literature. Indeed, using an existing data set, I show that the cost of latency incurred by trading on a human time scale (500 ms) increased three-fold over the 1995–2005 time-frame. In addition, using the alternative approach of keeping a fixed level of latency cost through continuous technological improvements, I compute the various percentiles of the implied latency over this time frame. Using the same data set, I observe that the median implied latency decreased by approximately two orders of magnitude.

My empirical analysis can also be utilized to compare the magnitude of latency cost to other trading costs incurred by institutional investors. My results suggest that the difference in payoff between trading with a human time scale (500 ms) and an automated trading platform with ultra low latency (1 ms) is approximately of the same order of magnitude as other trading costs faced by institutional investors. This observation certainly underlines the significance of latency for such investors. In conclusion, my model is the first theoretical approach in the literature to concretely quantify the impact of latency on the optimal order submission policy and its resulting cost to the trader.

There are a number of interesting future directions for research. First, as discussed in Section 2.4.4, there are a number of tractable extensions to the present model that can be analyzed. One particularly interesting case would be where the bid price process is a jump process. Here, my suspicion is that the cost of latency would *decrease*. This is because, even in the absence of any latency, the trader cannot adjust his limit prices ahead of a jump.

More generally, in the introduction, I identified a number of broad themes to the costs

that arise from latency. The model I have presented captures mainly costs due to a lack of contemporaneous decision making. It does not capture the latency costs due to strategic effects (i.e., comparative advantage/disadvantage relative to other investors) or due to time priority rules. These remain important questions for future research.

Chapter 3

Dynamic Portfolio Choice with Linear Rebalancing Rules

3.1. 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 features:

- **Return predicability.** At the heart of active portfolio management is the fact that a manager will seek to predict future asset returns [Grinold and Kahn, 1999]. Such predictions are not limited to simple unconditional estimates of expected future returns. A typical asset manager will make predictions on short- and long-term expected returns using complex models, for example, including return predicting factors such as market capitalization, book-to-market ratio, lagged returns, dividend yields, gross industrial production, and other security specific or macroeconomic variables [see, e.g., Chen *et al.*, 1986; Fama and French, 1996; Goetzmann and Jorion, 1993].

- **Transaction costs.** Trading costs in dynamic portfolio management 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. The efficient management of such costs is an important issue broadly, but becomes especially

crucial in the setting of *optimal execution*. This particular class of portfolio optimization problems seeks to optimally liquidate a given portfolio over a fixed time horizon [Bertsimas and Lo, 1998; Almgren and Chriss, 2000].

- **Portfolio or trade constraints.** Often times managers cannot make arbitrary investment decisions, but rather face exogenous constraints on their trades or their resulting portfolio. Examples of this include short-sale constraints, leverage constraints, or restrictions requiring market neutrality (or specific industry neutrality).

- **Risk aversion.** 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 which are easy to interpret and are known to be important. For example, a manager may both be interested in the risk of the portfolio value changing over various intervals of time, including for example, both short intervals (e.g., daily or weekly risk), as well as risk associated with the terminal value of the portfolio. Such single period risk can be measured a number of ways (e.g., variance, value-at-risk). A manager might further be interested in multiperiod measures of portfolio risk, for example, the maximum drawdown of the portfolio.

Significantly complicating the analysis of portfolio choice is that the underlying problem is *multiperiod*. Here, in general, the decision made by a manager at a given instant of time might 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 focussed on the analytically determining the optimal dynamic policy. While this work has brought forth important structural insights, it is fundamentally quite restrictive: exact analytical solutions require very specific assumptions 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, 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 as well in cases of practical interest, where the universe of tradeable assets is large.

In search of tractable alternatives, many practitioners eschew multiperiod formulations. Instead, they consider portfolio choice problems in a myopic, single period setting, even when 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. Here, a fixed set of deterministic decisions for the entire time horizon are made at the initial time. Both single period and deterministic portfolio choice formulations are quite flexible, and can accommodate many of the features described above. They are typically applied in a quasi-dynamic fashion through the method of *model predictive control.* Here, at each time period, the simplified single period or deterministic portfolio choice problem is resolved based on the latest available information. In general, such methods are heuristics; in order to achieve tractability, they neglect the explicit consideration of the possibility of future recourse. Hence, these methods may be significantly suboptimal.

A second tractable alternative is the formulation of portfolio choice problems as a *linear quadratic control* [e.g., Hora, 2006; Garleanu and Pedersen, 2012]. Since at least the 1950's, linear quadratic control problems have been an important class of tractable multiperiod optimal control problems. In the setting of portfolio choice, if the return dynamics are linear and the transaction costs and risk aversion penalties decomposed into per period quadratic functions, and positions and trading decision are unconstrained, then these methods apply. However, there are many important problem cases that simply do not fall into the linear quadratic framework.

In this chapter, my central innovation is to propose a framework for multiperiod portfolio optimization, which admits a broad class of problems including many with features as described earlier. My 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 chapter

are as follows:

1. *I define a flexible, general setting for portfolio optimization.* My setting allows for very general dynamics of asset prices, which an arbitrary dependence on a history of 'return-predictive factors'. I 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. My framework admits, for example, many complex models for transaction costs or risk aversion.

2. *My portfolio optimization problem is computationally tractable.* In my setting, determining the optimal linear rebalancing policy is a convex program. Convexity guarantees that the globally optimal policy can be tractably found in general. This is in contrast to non-convex portfolio choice parametrizations [e.g., Brandt *et al.*, 2009a], where only local optimality can be guaranteed.

   In my case, numerical solutions can be obtained via, for example, sample average approximation or stochastic approximation methods [see, e.g., Shapiro, 2003; Nemirovski *et al.*, 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 where the factor and return dynamics are driven by Gaussian uncertainty, I illustrate that my portfolio optimization problem can be reduced to a standard form of convex optimization program, such as a quadratic program or a second order cone program. In such cases, the problem can be solved with off-the-shelf commercial optimization solvers.

3. *My class of linear rebalancing policies subsumes many common heuristic portfolio policies.* Both single period and deterministic policies are special cases of linear rebalancing polices, however linear rebalancing polices are a broader class. Hence, the optimal linear rebalancing policy will outperform policies from these more restricted classes. Further, my method can also be applied in the context of model predictive control. Also, portfolio optimization problems that can be formulated as linear quadratic control also fit in my setting, and their optimal policies are linear rebalancing rules.

4. *I demonstrate the practical benefits of my method in an optimal execution example.*
   I consider an optimal execution problem where an investor seeks to liquidate a posi-
   tion In order to highlight the performance gain using linear decision rules, I use the
   discrete-time linear quadratic control formulation of Garleanu and Pedersen [2012].
   However, I further introduce linear inequality constraints that allow the trading de-
   cisions to only be sales; such sale-only constraints are common in agency algorithmic
   trading. I demonstrate that the best linear policy performs better than the best de-
   terministic policy, model predictive control and a projected version of the optimal
   policy proposed by Garleanu and Pedersen [2012]. Further, the performance of the
   best linear policy is shown to be near optimal, by comparison to upper bounds on
   optimal policy performance computed for the same problem.

The balance of this chapter is organized as follows: In Section 4.1.1, I review the related
literature. In Section 4.2, I present the abstract form of a dynamic portfolio choice model
and provide various specific problems that satisfy the assumptions of the abstract model.
I formally describe the class of linear decision rules in Section 3.3 and discuss solution
techniques in order to find the optimal parameters of the linear policy. In Section 3.4, I
provide efficient and exact formulations of dynamic portfolio choice models with Gaussian
uncertainty using linear decision models while incorporating linear equality and inequality
constraints, proportional and nonlinear transaction costs and a measure of terminal wealth
risk. In Section 3.5, I apply my methodology in an optimal execution problem and evaluate
the performance of the best linear policy. Finally, in Section 3.6 I conclude and discuss
some future directions.

3.1.1.   Related Literature

My chapter is related to two different strands of literature: the literature of dynamic port-
folio choice with return predictability and transaction costs, and the literature on the use
of linear decision rules in the optimal control problems.

First, I consider the literature on dynamic portfolio choice. This vast body of work
begins choice starts with the seminal chapter of Merton [1971]. Following this chapter,
there has been a significant literature aiming to incorporate the impact of various frictions

such as transaction costs on the optimal portfolio choice. For a survey on this literature, see Cvitanic [2001]. The work of Constantinides [1986] is an early example that studies the impact of proportional transaction costs on the optimal investment decision and the liquidity premium in the context of CAPM. Davis and Norman [1990], Dumas and Luciano [1991a], and Shreve and Soner [1994] provide the exact solution for the optimal investment and consumption decision by formally characterizing the trade and no-trade regions. One drawback of these papers is that the optimal solution is only computed in the case of a single stock and bond. Liu [2004] extends these results to multiple assets with fixed and proportional transaction costs in the case of uncorrelated asset prices. Detemple *et al.* [2003] develop a simulation-based methodology for optimal portfolio choice in the presence of return predictability.

There is also a significant literature on portfolio optimization that incorporates return predictability [see, e.g., Campbell and Viceira, 2002]. My chapter is related to the literature that incorporates both return predictability with transaction costs. Balduzzi and Lynch [1999] and Lynch and Balduzzi [2000] illustrate the impact of return predictability and transaction costs on the utility costs and the optimal rebalancing rule by discretizing the state space of the dynamic program. With a similar state space discretization, Lynch and Tan [2010] model the dynamic portfolio decision with multiple risky assets under return predictability and transaction costs and provide numerical experiments with two risky assets. One significant issue with this line of work is that discretization suffers from the curse-of-dimensionality: the computational effort to determine an optimal policy scales exponentially with the dimension of the state space. When there are more than a few assets or return predicting factors, discretization cannot be applied.

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. On the other hand, my work is in the spirit of Brandt *et al.* [2009a], allow for broader modeling flexibility at the expense of considering a restricted class of rebalancing policies. The parameterize the rebalancing rule as a function of security characteristics and estimates the parameters of the rule from empirical data without modeling the distribution of the returns and the return predicting factors. Even though my

approach is also a linear parametrization of return predicting factors, there are fundamental differences between my approach and that of Brandt *et al.* [2009a]. First, the class of linear polices I consider is much larger than the specific linear functional form in Brandt *et al.* [2009a]. In my approach the parameters are time-varying and cross-sectionally different for each security. Second, the extensions provided in Brandt *et al.* [2009a] for imposing positivity constraints and transaction costs are ad-hoc and cannot be generalized to arbitrary convex constraints or transaction cost functions. Finally, the objective function of Brandt *et al.* [2009a] is a non-convex function of the policy parameters. Hence, it is not possible, in general to obtain the globally optimal set of parameters. My setting, on the other hand, is convex, and hence globally optimal policies can be determined efficiently.

Garleanu and Pedersen [2012] achieve a closed-form solution for a model with linear dynamics in return predictors and quadratic function for transaction costs and quadratic penalty term for risk. However, the analytic solution is highly sensitive to the quadratic cost structure with linear dynamics [see, e.g., Bertsekas, 2000]. This special case cannot handle any inequality constraints on portfolio positions, non-quadratic transactions costs, or more complicated risk considerations. On the other hand, my approach can be implemented efficiently in these realistic scenarios and provides more flexibility in the objective function of the investor and the constraints that the investor faces. Boyd *et al.* [2012] consider an alternative generalization of the linear-quadratic case, using ideas from approximate dynamic programming. Glasserman and Xu [2011] develop a linear-quadratic formulation for portfolio optimization that offers robustness to modeling errors or mis-specifications.

Second, there is also a literature on the use of linear decision rules in optimal control problems. This approximation technique has attracted considerable interest recently in robust and two-stage adaptive optimization context [see, e.g., Ben-Tal *et al.*, 2004, 2005; Chen *et al.*, 2007, 2008; Bertsimas *et al.*, 2010; Bertsimas and Goyal, 2011]. Shapiro and Nemirovski [2005] illustrate that linear decision rules can reduce the complexity of multi-stage stochastic programming problems. Kuhn *et al.* [2009] proposes an efficient method to estimate the loss of optimality incurred by linear decision rule approximation.

In this strand of literature, I believe the closest works to the methodology described in my chapter are Calafiore [2009] and Skaf and Boyd [2010]. Both of these papers use linear

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

3.2. Dynamic Portfolio Choice with Return Predictability and Transaction Costs

I consider a dynamic portfolio choice problem with allowing general models for the predictability of security returns and for trading frictions. The number of investable securities is N, time is discrete and indexed by $t = 1, \ldots, T$, where T is the investment horizon. Each security i has a price change[1] of $r_{i,t+1}$ from time t to $t + 1$.

I collect these price changes in the return vector $r_{t+1} \triangleq (r_{1,t+1}, \ldots, r_{N,t+1})$. I assume that the investor has a predictive model of future security returns, and that these predictions are made through a set of 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, the lagged twelve month return of the stock [see, e.g., Fama and French, 1996; Goetzmann and Jorion, 1993]. Alternatively, they could be macroeconomic signals that affect the return of each security, such as inflation, treasury bill rate, industrial

---

[1]I choose to describe the evolution of asset prices in my framework in terms of absolute price changes, and I will also refer to these as (absolute) returns. Note that this is without loss of generality: since the return dynamics specified by Assumption 1 allow for an arbitrary dependence on history, my framework also admits, for example, models which describe the rate of return of each security.

production [see, e.g., Chen *et al.*, 1986]. Denote by $f_t \in \mathbb{R}^K$ the vector of factor values at time t. I assume 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* $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathsf{P})$, *I 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. *Here,* $G_{t+1}(\cdot)$ *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 the exogenous, i.i.d. disturbances* ϵ_{t+1}. *I assume that the filtration* $\mathbb{F} \triangleq \{\mathcal{F}_t\}_{t \geq 0}$ *is the natural filtration generated by the exogenous noise terms* $\{\epsilon_t\}$.

Let $x_{i,t}$ denote the number of shares that the investor holds in ith security over the time period t. I collect the portfolio holdings across all securities at time t in the vector $x_t \triangleq (x_{1,t}, \ldots, x_{N,t})$, and I 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 he 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 my 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 $\mathbf{x} \triangleq (x_1, \ldots, x_T)$, $\triangleq (\mathbf{f_1}, \ldots, \mathbf{f_T})$, and $\mathbf{r} \triangleq (r_2, \ldots, r_{T+1})$, respectively. Similarly, the sample path of trades over time is denoted by $\mathbf{u} = (u_1, \ldots, u_T)$. I make the following assumption on feasible sample paths of trades:

Assumption 2 (Convex trading constraints). *The sample path of trades* $\mathbf{u}$ *are restricted to the non-empty, closed, and convex set* $U \subseteq \mathbb{R}^N \times \ldots \times \mathbb{R}^N$.

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 $\mathbf{f}$. I let $\mathcal{U}$ be the set of all policies. I assume that the investor's trading decisions are non-anticipating in that the trade vector u_t in period t depends only on what is known at the beginning of period t. Formally, I 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 non-anticipating policies.

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

$$(3.1) \qquad \sup_{\pi \in \mathcal{U}_\mathbb{F}} \; \mathsf{E}_\pi[p(\mathbf{x}, \mathbf{f}, \mathbf{r})],$$

where the real-valued reward function $p(\cdot)$ 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, $p(\cdot)$ may have the form

$$(3.2) \qquad p(\mathbf{x}, \mathbf{f}, \mathbf{r}) \triangleq W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) - \mathsf{RA}(\mathbf{x}, \mathbf{f}, \mathbf{r}).$$

Here, W denotes the terminal wealth (total trading gains ignoring of transaction costs), i.e.,

$$(3.3) \qquad W(\mathbf{x}, \mathbf{r}) \triangleq W_0 + \sum_{t=1}^{T} x_t^\top r_{t+1},$$

where W_0 is the 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.

I make the following assumption about my objective function:

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

If $p(\cdot)$ has the specified form in (3.2), 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.

3.2.1. Examples

In this chapter, I consider dynamic portfolio choice models that satisfy Assumptions 1–3. In order to illustrate the generality of this setting, I will now provide a number of specific examples that satisfy these assumptions.

Example 1 (Garleanu and Pedersen 2012)**.** *This model has the following dynamics, where returns are driven by mean-reverting factors, that fit into my general framework:*

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

for each time $t \geq 0$. Here, μ_t is the deterministic 'fair return', e.g., derived from the CAPM, while B is a matrix of constant factor loadings. The factor process f_t is a vector mean-reverting process, with Φ a matrix of mean reversion coefficients for the factors. It is assumed that the i.i.d. disturbances $\epsilon_{t+1} \triangleq (\epsilon_{t+1}^{(1)}, \epsilon_{t+1}^{(2)})$ are zero-mean with covariance given by $\mathrm{Var}(\epsilon_{t+1}^{(1)}) = \Psi$ and $\mathrm{Var}(\epsilon_{t+1}^{(2)}) = \Sigma$.

Trading is costly, and the transaction cost to execute $u_t = x_t - x_{t-1}$ shares is given by $\mathsf{TC}_t(u_t) \triangleq \frac{1}{2} u_t \Lambda u_t$, 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 maximize discounted future expected excess return, while accounting for transaction costs and adding a per-period penalty for risk, i.e.,

$$(3.4) \qquad \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \ \mathsf{E}_\pi \left[\sum_{t=1}^{T} \left(x_t^\top B f_t - \mathsf{TC}_t(u_t) - \mathsf{RA}_t(x_t) \right) \right].$$

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. Garleanu and Pedersen [2012] suggest this objective function for an investor who is compensated based on his performance relative to a benchmark. Each $x_t^\top B f_t$ term measures the excess return over the benchmark, while each $\mathsf{RA}_t(x_t)$ term measures the variance of the tracking error relative to the benchmark.[2]

The problem (3.4) clearly falls into my framework. The objective function is similar to that of (3.2) with the minor variation expected excess return rather than expected wealth is considered. Further, (3.4) has the further special property that total transaction costs and penalty for risk aversion decompose over time:

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

---

[2]See Garleanu and Pedersen [2012] for other interpretations.

Note that this problem can be handled easily using the classical theory from the linear-quadratic control (LQC) literature [see, e.g., Bertsekas, 2000]. This theory provides analytical characterization of optimal solution, for example, that the value function at any time t is quadratic function 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.

On the other hand, the tractability of this model rests critically on three key requirements:

- The state variables (x_t, f_t) at each time t must evolve as linear functions of the control u_t and the i.i.d. disturbances ϵ_t (i.e., linear dynamics).

- Each control decision u_t is 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 may involve portfolio position or trade constraints, different forms of transaction costs and risk measures, and more complicated return dynamics. In the following examples, I will provide concrete examples of many such cases that do not admit optimal solutions via the LQC methodology, but remain within my framework.

Example 2 (Portfolio or trade constraints). *In practice, a common constraint in constructing equity portfolios is the short-sale restriction. Most of the mutual funds are enforced not to have any short positions by law. This requires the portfolio optimization problem to include the linear constraint*

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

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

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

execution desk needs to impose the negativity constraint

$$u_t \leq \mathbf{0},$$

for each time t.

    Simple linear constraints such as these fit easily in my framework, but cannot be addressed via traditional LQC methods.

Example 3 (Non-quadratic transaction costs). *In practice, many trading costs such as the bid-ask spread, broker commissions, and exchange fees are intrinsically proportional to the trade size. Letting χ_i be the the proportional transaction cost rate (an aggregate sum of bid-ask cost and commission fees, for example) for trading security i, the investor will incur a total cost of*

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

The proportional transaction costs are a classical cost structure that is well studied in the literature [see, e.g., Constantinides, 1986].

    Furthermore, other trading costs occur due to disadvantageous transaction price caused by the price impact of the trade. The management of the trading costs due to price impact has recently attracted considerable interest [see, e.g., Obizhaeva and Wang, 2005; Almgren and Chriss, 2000]). Many models of transaction costs due to price impact imply a nonlinear relationship between trade size and the resulting transaction cost, for example

$$\mathsf{TC}(\mathbf{u}) \triangleq \sum_{t=1}^{T} \sum_{i=1}^{N} \chi_i |u_{i,t}|^{\beta}.$$

Here, $\beta \geq 1$[3] and χ_i is a security specific proportionality constant.

    In general, when the trade size is small relative to the total traded volume, proportional costs will dominate. On the other hand, when the trade size is large, costs due to price impact will dominate. Hence, both of these types of trading are important. However, the LQC framework of Example 1 only allows quadratic transaction costs (i.e., $\beta = 2$).

Example 4 (Terminal wealth risk). *The objective function of Example 1 includes a term to penalize excessive risk. In particular, the per-period quadratic penalty, $x_t^\top \Sigma x_t$, is used, in*

---

[3]Gatheral [2010] notes that $\beta = \frac{3}{2}$ is a typical assumption in practice.

order to satisfy the requirements of the LQC model. However, penalizing additively risk in a per-period fashion is nonstandard. Such a risk penalty does not correspond to traditional forms of investor risk preferences, e.g., maximizing the utility of terminal wealth, and the economic meaning of such a penalty is not clear. An investor is typically more interested in the risk associated with the terminal wealth, rather than a sum of per-period penalties.

In order to account for terminal wealth risk, let $\rho\colon \mathbb{R} \to \mathbb{R}$ be a real-valued convex function meant to penalize for excessive risk of terminal wealth (e.g., $\rho(w) = \frac{1}{2}w^2$ for a quadratic penalty) and consider the optimization problem

$$(3.5) \qquad \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \ \mathsf{E}_{\pi}\left[W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) - \gamma\rho(W(\mathbf{x}, \mathbf{r}))\right],$$

where $\gamma > 0$ is a risk-proportionality constant.

It is not difficult to see that the objective in (3.5) satisfies Assumption 3 and hence fits into my model. However, even when the risk penalty function $\rho(\cdot)$ is quadratic, (3.5) does *not* admit a tractable LQC solution, since the quadratic objective does not decompose across time.

Example 5 (Maximum drawdown risk). *In addition to the terminal measures of risk described in Example 4, an investor might also be interested controlling intertemporal measures of risk defined over the entire time trajectory. For example, a fund manager might be sensitive to a string of successive losses that may lead to the withdrawal of assets under management. One way to limit such losses is to control the maximum drawdown, defined as the worst loss of the portfolio between any two points of time during the investment horizon* [4]. *Formally,*

$$\mathsf{MD}(\mathbf{x}, \mathbf{r}) \triangleq \max_{1 \leq t_1 \leq t_2 \leq T} \left(-\sum_{t=t_1}^{t_2} x_t^\top r_{t+1}, 0\right).$$

It is easy to see that the maximum drawdown is a convex function of $\mathbf{x}$. Hence, the portfolio optimization problem

$$(3.6) \qquad \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \ \mathsf{E}_{\pi}[W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) - \gamma\mathsf{MD}(\mathbf{x}, \mathbf{r})],$$

where $\gamma \geq 0$ is a constant controlling tradeoff between wealth and the maximum drawdown penalty, satisfies Assumption 3. Moreover, standard convex optimization theory yields that

---

[4]For example, see Grossman and Zhou [1993] for an earlier example.

the problem (3.6) *is equivalent to solving the constrained problem*

(3.7)
$$\underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \quad \mathsf{E}_\pi \left[W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) \right]$$
$$\text{subject to} \quad \mathsf{E}_\pi \left[\mathsf{MD}(\mathbf{x}, \mathbf{r}) \right] \leq C,$$

where C (which depends on the choice of γ) is a limit on the allowed expected maximum drawdown.

Example 6 (Complex dynamics). *I can also generalize the dynamics of Example 1. Consider factor and return dynamics given by*

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

for each time $t \geq 0$. Here, each $\xi_{t+1} \in \mathbb{R}^{N \times K}$ is an extra noise term which captures model uncertainty regarding the factor loadings. I assume that

$$\mathsf{E}\left[(B + \xi_{t+1}) f_t \mid \mathcal{F}_t \right] = B f_t, \qquad \text{Var}\left[(B + \xi_{t+1}) f_t \mid \mathcal{F}_t \right] = f_t^\top \Upsilon f_t.$$

With this model, the conditional variance of returns becomes dependent on the factor structure and is time-varying, i.e., $\text{Var}[r_{t+1}|\mathcal{F}_t] = f_t^\top \Upsilon f_t + \Sigma$. This is consistent with the empirical work of Fama and French [1996], for example. In this setting, the per-period penalty of risk analogous to that in (3.4) becomes $\mathsf{RA}_t(\mathbf{x}, \mathbf{f}) = x_t^\top \left(f_t^\top \Upsilon f_t + \Sigma \right) x_t$. The resulting optimal control problem no longer falls into the LQC framework.

The dynamics and the reward functions considered in these examples satisfy my basic requirements of Assumptions 1–3. These examples illustrate that in many real-world problems with complex primitives for return predictability, transaction costs, risk measures and constraints, the dynamic portfolio choice becomes difficult to solve analytically via LQC methods.

3.3. Optimal Linear Model

The examples of Section 3.2.1 illustrated a broad range of practically important portfolio optimization problems. Without special restrictions, such as those imposed in the LQC framework, the optimal dynamic policy for such a broad set of problems cannot be computed

either analytically or computationally. In this section, in order to obtain policies in a computationally tractable way, I will consider a more modest goal. Instead of finding the optimal policy amongst *all* admissible dynamic policies, I will restrict my search to a subset of policies that are parsimoniously parameterized. That is, instead of solving for a globally optimal policy, I will instead find an approximately optimal policy by finding the best policy over the restricted subset of policies.

In order to simplify, I will assume that reward function of the investor's optimization (3.1) is a function only of the sample path of portfolio positions $\mathbf{x}$ and of factor realizations $\mathbf{f}$, and does not depend on the security returns $\mathbf{r}$. In other words, I assume that the reward function takes the form $p(\mathbf{x}, \mathbf{f})$. This is without loss of generality — given my general specification for factors under Assumption 1, we can simply include each security return as a factor. With this assumption, investor's trading decisions will, in general, be a non-anticipating 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:

Definition 2 (Linear rebalancing policy). *A linear rebalancing policy π is a non-anticipating policy parameterized by 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*

$$(3.8) \qquad\qquad u_t \triangleq c_t + \sum_{s=1}^{t} E_{s,t} f_s,$$

for each time $t = 1, 2, \ldots, T$.

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, i.e., $\mathbf{u}$ is feasible. Denote by $\mathcal{L} \subset \mathcal{U}_{\mathbb{F}}$ the corresponding set of feasible linear policies.

Observe that linear rebalancing allow recourse, albeit in a restricted functional form. The affine specification (3.8) includes several classes of polices of particular interest as special cases:

- **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 1 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 $1 \leq t \leq T$, i.e., the optimal portfolio are linear in the previous position and the current factor values. Equivalently, by induction on t,

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

  Since $u_t = x_t - x_{t-1}$, 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 (3.8).

- **Linear portfolio polices.** Brandt *et al.* [2009a] suggest a class of policies where portfolios are determined by adjusting a deterministic benchmark portfolio according to a linear function of a vector of stochastic, time-varying firm characteristics. In my setting, the firm characteristics would be interpreted as stochastic return predicting factors. An analogous rule would determine the positions at each time t via $x_t = \bar{x}_t + \Theta_t^\top (f_t - \bar{f}_t)$. Here, $\bar{f}_t$ 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 be mean zero) to adjustments to the benchmark portfolio. Such a portfolio rule is clearly of the form (3.8).

- **Policies based on basis functions.** Instead of having policies that are directly affine function of factor realizations, it is also possible to introduce basis functions. One might consider, for example, $\varphi \colon \mathbb{R}^K \to \mathbb{R}^D$, a collection of D (non-linear) functions that capture particular features of the factor space that are important for good decision making. Consider a class of policies of the form

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

  Such policies belong to the linear rebalancing class, if the factors are augmented also to include the value of the basis functions. This is easily accommodated in my framework, given the flexibility of Assumption 1. Similarly, policies which depend on

the past security returns (in addition to factor realizations) can be accomodated by augmenting the factors with past returns.

An alternative to solving the original optimal control problem (3.1) is to consider the problem

$$(3.9) \qquad \sup_{\pi \in \mathcal{L}} \ \mathsf{E}_\pi[p(\mathbf{x}, \mathbf{f})],$$

which restricts to linear rebalancing rules. In general, (3.9) will not yield an optimal control. The exception is if the optimal control for the problem is indeed a linear rebalancing rule (e.g., in a LQC problem). However, (3.9) will yield the best possible linear rebalancing rule. Further, in contrast to the original optimal control problem, (3.9) has the great advantage of being *tractable*, as suggested by the following result:

Proposition 1. *The optimization problem given by*

$$(3.10) \qquad \begin{aligned} \underset{\mathbf{E},\mathbf{c}}{\text{maximize}} \quad & \mathsf{E}\left[p(\mathbf{x}, \mathbf{f})\right] \\ \text{subject to} \quad & x_t = x_{t-1} + u_t, \qquad \forall \ 1 \le t \le T, \\ & u_t = c_t + \sum_{s=1}^{t} E_{s,t} f_s, \quad \forall \ 1 \le t \le T, \\ & (\mathbf{E}, \mathbf{c}) \in \mathcal{C}. \end{aligned}$$

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

Proof. Note that $p(\cdot, \mathbf{f})$ is concave for a fixed $\mathbf{f}$ by Assumption 3, and since $\mathbf{x}$ 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 $\mathbf{f}$ preserves this concavity. Finally, the convexity of the constraint set $\mathcal{C}$ follows from the convexity of U, under Assumption 2. ∎

The problem (3.10) is a finite-dimensional, convex optimization problem that will yield parameters for the optimal linear rebalancing policy. It is also a stochastic optimization problem, in the sense that the objective is the expectation of a random quantity. In general, there are a number of effective numerical methods that can been applied to solve such problems:

- **Efficient exact formulation.** In many cases, with further assumptions on the problem primitives (the reward function $p(\cdot)$, the dynamics of the factor realizations $\mathbf{f}$, and the trading constraint set U), the objective $\mathsf{E}[p(\mathbf{x}, \mathbf{f})]$ and the constraint set $\mathcal{C}$ of the program (3.10) can be explicitly analytically expressed in terms of the decision variables $(\mathbf{E}, \mathbf{c})$. In some of these cases, the program (3.10) 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 to these standard forms [e.g., Grant and Boyd, 2011] can be used. Alternatively, generic methods for constrained convex optimization such as interior point methods [see, e.g., Boyd and Vandenberghe, 2004] can be applied to efficiently solve large-scale instances of (3.10). I will explore this topic further, developing a number of efficient exact formulations in Section 3.4, and providing a numerical example in Section 3.5.

- **Sample average approximation (SAA).** In the absence of further structure on the problem primitives, the program (3.10) can also be solved via Monte Carlo sampling. Specifically, supposed that $\mathbf{f}^{(1)}, \ldots, \mathbf{f}^{(S)}$ are S independent sample paths of factor realization. The objective and constraints of (3.10) can be replaced with sampled versions, to obtain

$$
\begin{aligned}
\underset{\mathbf{E}, \mathbf{c}}{\text{maximize}} \quad & \frac{1}{S} \sum_{\ell=1}^{S} p\left(\mathbf{x}^{(\ell)}, \mathbf{f}^{(\ell)}\right) \\
\text{(3.11)} \qquad \text{subject to} \quad & x_t^{(\ell)} = x_{t-1}^{(\ell)} + u_t^{(\ell)}, \qquad \forall\, 1 \le t \le T,\ 1 \le \ell \le S, \\
& u_t^{(\ell)} = c_t + \sum_{s=1}^{t} E_{s,t} f_s^{(\ell)}, \quad \forall\, 1 \le t \le T,\ 1 \le \ell \le S, \\
& \mathbf{u}^{(\ell)} \in U, \qquad\qquad\qquad \forall\, 1 \le \ell \le S.
\end{aligned}
$$

  The sample average approximation (3.11) can be solved via standard convex optimization methods (e.g., interior point methods). Moreover, under appropriate regularity conditions, convergence of the SAA (3.11) to the original program (3.10) can be established as $S \to \infty$, along with guarantees on the rate of convergence [Shapiro, 2003].

- **Stochastic approximation.** Denote the collection of decision variables in (3.10) by $\mathbf{z} \triangleq (\mathbf{E}, \mathbf{c})$, and, allowing a minor abuse of notation, define $p(\mathbf{z}, \mathbf{f})$ to be 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 $h(\mathbf{z}) \triangleq p(\mathbf{z}, \mathbf{f})$, the problem (3.10) is simply to maximize $h(z)$ subject to the constraint that $z \in \mathcal{C}$. Under suitable technical conditions, superdifferentials of h and p are related according to $\partial h(\mathbf{z}) = \mathsf{E}[\partial_{\mathbf{z}} p(\mathbf{z}, \mathbf{f})]$. Stochastic approximation methods are incremental methods that seek to estimate ascent directions for $h(\cdot)$ from sampled ascent directions for $p(\cdot, \mathbf{f})$. For example, given a sequence of i.i.d. 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

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

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

One attractive feature of the sample average approximation and stochastic approximation approaches is that they can be applied in a *data-driven* fashion. These methods need access only to simulated trajectories of factors and returns — they do not need explicit knowledge of the dynamics in Assumption 1 that drive these processes. Hence, an optimal linear rebalancing policy can be determined using, for example, historical data to construct simulated trajectories, without specifying and estimating an explicit functional form for the factor and return dynamics.

Finally, observe that optimal linear policies can also be applied in concert with *model predictive control* (MPC). Here, at each time step t, the program (3.10) is resolved beginning from time t. This determines the optimal linear rebalancing rule from time t forward, conditioned on the realized history up to time t. The resulting policy is only used to determine the trading decision at the then current time t, and (3.10) is subsequently resolved at each future time period. At the cost of an additional computational burden, the use of optimal linear policies with MPC subsumes standard MPC approaches, such as resolving a myopic variation of the portfolio optimization problem (and ignoring the true multiperiod

nature) or solving a deterministic variation of the portfolio optimization problem (and ignoring the possibility of future recourse).

3.4. Efficient Exact Formulations

In this section, I will provide efficient exact formulations of dynamic portfolio choice problems using the class of linear policies for my feasible set of policies. In particular, I will consider a number of the examples of dynamic portfolio choice problems discussed in Section 3.2.1. These examples include features such as constraints on portfolio holdings, transaction costs, and risk measures. In each case, I will demonstrate how the optimization problem (3.10) can be transformed into a deterministic convex program by explicit analytical evaluation of the objective function $\mathsf{E}[p(\cdot, \mathbf{f})]$ and the constraint set $\mathcal{C}$.

Exact formulations require the evaluation of expectations taken over the sample path of factor realizations $\mathbf{f}$. In order to do this, I will make the following assumption for the rest of this section:

Assumption 4 (Gaussian factors). *Assume that the sample path $\mathbf{f}$ of factor realizations is jointly Gaussian. In particular, denote by $F_t \triangleq (f_1, \ldots, f_t)^\top \in \mathbb{R}^{Kt}$ the vector of all factors observed by time t. I assume that $F_t \sim N(\theta_t, \Omega_t)$, where $\theta_t \in \mathbb{R}^{Kt}$ is the mean vector and $\Omega_t \in \mathbb{R}^{Kt \times Kt}$ is the covariance matrix.*

With this assumption, the trades of any linear policy will also be jointly normally distributed, as each such policy is affine transformations of the factors. Formally, let

$$(3.12) \qquad M_t \triangleq \begin{bmatrix} E_{1,t} & E_{2,t} & \ldots & E_{t,t} \end{bmatrix} \in \mathbb{R}^{N \times Kt}$$

be the matrix of time t policy coefficients, so that the trade vector is given by $u_t = c_t + M_t F_t$. With this representation, it is easy see that $u_t \sim N(\bar{u}_t, V_t)$, where the mean vector and covariance matrix are given by

$$(3.13) \qquad \bar{u}_t \triangleq \mathsf{E}[u_t] = c_t + M_t \theta_t, \qquad\qquad V_t \triangleq \mathrm{Var}(u_t) = M_t \Omega_t M_t^\top.$$

Similarly, the portfolio x_t at time t is normally distributed. I have that

$$(3.14) \qquad x_t = x_0 + \sum_{i=1}^t u_i = x_0 + \sum_{i=1}^t \left(c_i + \sum_{s=1}^i E_{s,i} f_s \right) = d_t + \sum_{s=1}^t J_{s,t} f_s,$$

where $d_t \triangleq x_0 + \sum_{i=1}^{t} c_i$ and $J_{s,t} \triangleq \sum_{\ell=s}^{t} E_{s,\ell}$. With this representation, it is easy see that $x_t \sim N(\kappa_t, Y_t)$, where

$$(3.15) \qquad \kappa_t \triangleq \mathsf{E}[x_t] = d_t + P_t \theta_t, \qquad Y_t \triangleq \mathrm{Var}(x_t) = P_t \Omega_t P_t^\top,$$

$$(3.16) \qquad P_t \triangleq \begin{bmatrix} J_{1,t} & J_{2,t} & \dots & J_{t,t} \end{bmatrix}.$$

3.4.1. Linear Constraints

I will provide formulations for linear equality or inequality constraints on trades or positions, in the context of linear rebalancing policies. These type of constraint appear frequently in portfolio choice due to regulatory reasons such as short sale restriction, liquidation purposes or diversification needs such as keeping a specific industry exposure under a certain limit.

3.4.1.1. Equality Constraints

Equality constraints appear often in portfolio choice, particularly in portfolio execution problems when the investor needs to liquidate a certain portfolio (i.e., $x_T = \mathbf{0}$) or construct a certain target portfolio by the end of the time horizon (i.e., $x_T = \bar{x}$).

Suppose that for some time t, have a linear equality constraint on the trade vector u_t, of the form $Au_t = b$. Here, $A \in \mathbb{R}^{M \times N}$ and $b \in \mathbb{R}^N$. This constraint can be written as

$$(3.17) \qquad Ac_t + AM_t F_t = b.$$

Under Assumption 4, the left hand side of the (3.17) is normally distributed. Therefore, for (3.17) to hold almost surely, I must have that the left hand side have mean b and zero covariance. Thus, I require that

$$(3.18) \qquad Ac_t = b, \qquad\qquad\qquad AM_t = \mathbf{0}.$$

Thus, the linear equality constraint (3.17) on the trade vector u_t is equivalent to the linear equality constraint (3.18) on the policy coefficients (c_t, M_t). Linear equality constraints on the portfolio position x_t can be handled similarly.

3.4.1.2. Inequality Constraints

Inequality constraints on trades or positions are common as well. One example is a short-sale constraint, which would require that $x_t \geq \mathbf{0}$ for all times t. When the factor realizations do not have bounded support, inequality constraints cannot be enforced almost surely. This is true in the Gaussian case: there is a chance, however small, that factors may take extreme values, and if the policy if a linear function of the factors, this may cause an inequality constraint to be violated.

In order to account for such constraints in a linear rebalancing policy, instead of enforcing inequality constraints almost surely, we will enforce them at a given level of confidence. For example, given a vector $a \in \mathbb{R}^N$ and a scalar b, instead of enforcing the linear constraint $a^\top u_t \leq b$, almost surely, we can consider a relaxation where seek to guarantee that it is violated with small probability. In other words, we can impose the chance constraint $\mathsf{P}(a^\top u_t > b) \leq \eta$, for a small value of the parameter η. The following lemma, whose proof can be found in the Online Supplement, illustrates that this can be accomplished explicitly:

Lemma 3. *Given $\eta \in [0, 1/2]$, a non-zero vector $a \in \mathbb{R}^N$, and a scalar b, the chance constraint $\mathsf{P}(a^\top u_t > b) \leq \eta$ is equivalent to the constraint*

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

on the policy coefficients (c_t, M_t), where $\Phi^{-1}(\cdot)$ is the inverse cumulative normal distribution.

A similar approach be applied to incorporate linear inequality constraints on the portfolio position x_t with high confidence.

In many situations (e.g., short-sale constraints), it may not be sufficient to enforce an inequality constraint only probabilistically. In such cases, when a linear rebalancing policy is applied, the resulting trades can be projected onto the constraint set so as to ensure that the constraints are always satisfied. When the linear policy is designed, however, it is helpful to incorporate the desired constraints probabilistically so as to account for their presence. I will demonstrate this idea in the application in Section 3.5.

3.4.2. Transaction Costs

In this section, I will provide efficient exact formulations for the transaction cost functions discussed in Section 3.2.1, in the context of linear rebalancing policies. In general, once might consider a total transaction cost of

$$\mathsf{TC}(\mathbf{u}) \triangleq \sum_{t=1}^{T} \mathsf{TC}_t(u_t)$$

for executing the sample path of trades $\mathbf{u}$, where $\mathsf{TC}_t(u_t)$ is the cost of executing the trade vector u_t at time t. As seen in Section 3.2.1, we typically wish to subtract an expected transaction cost term from investor's objective. Hence efficient exact formulations for transaction costs involve explicit analytical computation of $\mathsf{E}[TC(\mathbf{u})] = \sum_{t=1}^{T} \mathsf{E}[\mathsf{TC}_t(u_t)]$, when each trade vector u_t is specified by a linear policy.

Under a linear policy, $u_t \sim N(\bar{u}_t, V_t)$ is distributed as a normal random variable, with mean and covariance $(\bar{u}_t, V_t)$ specified from the policy $(\mathbf{E}, \mathbf{c})$ coefficients through (3.13). Then, the evaluation of expected transaction costs reduces to the evaluation of the expected value of the per period transaction cost function $\mathsf{TC}_t(\cdot)$ for a Gaussian argument. This can be handled on a case-by-case basis as follows:

- **Quadratic transaction costs.** In the case of quadratic transaction costs, as seen in Example 1, the per period transaction cost function 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 definite matrix. In this case, $\mathsf{E}[\mathsf{TC}_t(u_t)] = \frac{1}{2}\left(\bar{u}_t \Lambda \bar{u}_t + \mathrm{tr}(\Lambda V_t)\right)$.

- **Proportional transaction costs.** In the case of proportional transaction costs, as discussed in Example 3, the per period transaction cost function is given by

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

where $\chi_i > 0$ is a proportionality constant specific to security i. Using the properties of the folded normal distribution, I obtain

$$\mathsf{E}[\mathsf{TC}_t(u_t)] = \sum_{i=1}^{N} \chi_i \left(\sqrt{\frac{2V_{t,i}}{\pi}} \exp\left\{ -\frac{\bar{u}_{t,i}^2}{2V_{t,i}} \right\} + \bar{u}_{t,i} \left\{ 1 - 2\Phi\left(-\frac{\bar{u}_{t,i}}{\sqrt{V_{t,i}}} \right) \right\} \right),$$

where $\Phi(\cdot)$ is the cumulative distribution function of a standard normal random variable.

- **Nonlinear transaction costs.** In the case of nonlinear transaction costs, as discussed in Example 3, the per period transaction cost function is given by

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

where $\chi_i > 0$ is a proportionality constant specific to security i, and $\beta \geq 1$ is an exponent capturing the degree of nonlinearity. As in the proportional case, evaluating the Gaussian expectation explicitly results in

$$\mathsf{E}[\mathsf{TC}_t(u_t)] = \sum_{i=1}^{N} \chi_i \Gamma\left(\frac{1+\beta}{2}\right) \frac{(2V_{t,i})^{\frac{\beta}{2}}}{\sqrt{\pi}} \, {}_1F_1\left(-\frac{\beta}{2}; \frac{1}{2}; -\frac{\bar{u}_{t,i}^2}{2V_{t,i}}\right),$$

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

3.4.3. Terminal Wealth and Risk Aversion

In many of the portfolio choice examples in Section 3.2.1, an investor wishes to maximize expected wealth net of transaction costs, subject to a penalty for risk, i.e.,

$$(3.19) \qquad \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \; \mathsf{E}_\pi\left[W(\mathbf{x}, \mathbf{r}) - \mathsf{TC}(\mathbf{u}) - \mathsf{RA}(\mathbf{x}, \mathbf{f}, \mathbf{r}))\right].$$

Here, $W(\cdot)$ is the terminal wealth associated with a sample path, $\mathsf{TC}(\cdot)$ are the transaction costs, and $\mathsf{RA}(\cdot)$ is a penalty for risk aversions. Exact calculation of expected transaction costs for linear policies were discussed in Section 3.4.2. Here, I will discuss exact calculation of the expected terminal wealth and the risk aversion penalty.

To begin, note that the terminal wealth depends on realized returns in addition to factor realizations. Hence, I will make the following assumption:

Assumption 5 (Gaussian returns). *As in Example 1, assume that for each time $t \geq 0$, returns evolve according to*

$$(3.20) \qquad\qquad r_{t+1} = \mu_t + Bf_t + \epsilon_{t+1}^{(2)},$$

where μ_t is a deterministic vector, B is a matrix of factor loadings, and $\varepsilon_t^{(2)}$ are zero-mean i.i.d. Gaussian disturbances with covariance Σ.

Note that the critical assumption I am making here is that the factor realizations $\mathbf{f}$ and the sample path of security returns $\mathbf{r}$ are jointly Gaussian. The particular form (3.20) is chosen out of convenience but is not necessary.

We can calculate the expected terminal wealth as

$$\mathsf{E}[W(\mathbf{x}, \mathbf{r})] = W_0 + \sum_{t=1}^{T} \mathsf{E}[x_t^\top r_{t+1}] = W_0 + \sum_{t=1}^{T} \left(\mu_t^\top \kappa_t + \mathsf{E}[x_t^\top B f_t] \right),$$

$$= W_0 + \sum_{t=1}^{T} \left(\mu_t^\top \kappa_t + d_t^\top B \delta_t + \sum_{s=1}^{t} \left(\delta_s^\top (B(I - \Phi)^{t-s} J_{s,t}^\top) \delta_s + \mathrm{tr} \left(B(I - \Phi)^{t-s} J_{s,t}^\top \omega_s \right) \right) \right),$$

where ω_s is the sth $K \times K$ diagonal block matrix of Ω_t.

For the risk aversion penalty, I consider two cases:

- **Per period risk penalty.** Consider risk aversion penalties that decompose over time as

$$\mathsf{RA}(\mathbf{x}, \mathbf{f}, \mathbf{r}) = \sum_{t=1}^{N} \mathsf{RA}_t(x_t),$$

  where $\mathsf{RA}_t(\cdot)$ is a function which penalizes for risk aversion based on the positions held at time t. One such case is the quadratic penalty $\mathsf{RA}_t(x_t) \triangleq \frac{\gamma}{2} x_t^\top \Sigma x_t$ of Example 1, where $\gamma > 0$ is a risk penalty proportionality constant. Here, the investor seeks to penalize in proportion to the conditional per period variance of the portfolio value. So long as the expectation of $\mathsf{RA}_t(\cdot)$ can be calculated for Gaussian arguments, then the overall expected risk aversion penalty can be calculated exactly. This can be accomplished for a variety of functions. For example, quadratic penalties can be handled in a manner analogous to the quadratic transaction costs discussed in Section 3.4.2.

- **Terminal wealth risk penalty.** Alternatively, as discussed in Example 4, a more natural risk aversion criteria might be to penalize risk as a function of the *terminal* wealth. Specifically, an investor with a quadratic utility function would consider a risk aversion penalty $\mathsf{RA}(\mathbf{x}, \mathbf{f}, \mathbf{r}) \triangleq -\frac{\gamma}{2} W(\mathbf{x}, \mathbf{r})^2$, where $\gamma > 0$ is a risk penalty proportionality constant. I show in the Online Supplement that $\mathsf{E}[W(\mathbf{x}, \mathbf{r})^2]$ can be analytically computed and the resulting expression is a quadratic convex function of policy coefficients. Note that handling this quadratic penalty enables to accommodate mean-variance type objectives on the terminal wealth.

3.5. Application: Equity Agency Trading

In this section, I provide an empirical application to illustrate the implementation and the benefits of the optimal linear policy. As my example, I consider an important problem in equity agency trading. Equity agency trading seeks to address the problem faced by large investors such as pension funds, mutual funds, or hedge funds that need to update the holdings of large portfolios. 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 impact from trading. Efficient execution of large trades is accomplished via 'algorithmic trading', and requires significant technical expertise and infrastructure. For this reason, large investors 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 risk objectives of the client.

The problem of finding an optimal execution schedule has received a lot of attention in the literature since the initial chapter of Bertsimas and Lo [1998]. In their model, when price impact is proportional to the number of shares traded, the optimal execution schedule is to trade equal number of shares at each trading time. There are number of papers that extend this model to incorporate the risk of the execution strategy. For example, Almgren and Chriss [2000] derive that risk averse agents need to liquidate their portfolio faster in order to reduce the uncertainty of the execution cost.

The models described above seek mainly to minimize execution costs by accounting for the price impact and supply/demand imbalances caused by the investor's trading. Complementary to this, an investor may also seek to exploit short-term predicability of stock returns to inform the design of a trade schedule. As such, there is a growing interest to model 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, e.g., the Capital Asset Pricing Model (CAPM), or the Fama-French Three 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 *et al.* [2010] document that systematic trading as described in the examples above and institutional fund flows lead to predictable patterns in intraday returns of common stocks.

I will consider an agency trading optimal execution problem in the presence of short-term predictability. One issue that arises here is that, due to the 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 may result in both buy and sell trades for the same security as it seeks to exploit short-term signals. Hence, it is necessary to impose constraints on the sign of trades, as in Example 2.

If an agency trading execution problem has price and factor dynamics which 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, i.e., the number of shares to trade at each time is a linear function of the previous return predicting factors. I will consider a particular formulation that involves linear price and factor dynamics and a quadratic objective function (as in Example 1). Note that this example does not highlight the full generality of my framework — more interesting cases would involve non-linear factor dynamics (e.g., microstructure-based order imbalance signals) or a non-quadratic objective (e.g., transaction costs as in Example 3). However, this example is intentionally chosen since, in the absence of the trade sign constraint, the problem can be solved exactly with LQC methods. Hence, 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. I present my optimal execution problem formulation in Section 3.5.1. An exact, analytical solution is not available to this problem, hence, in Section 3.5.2, I describe several approximate solution techniques, including finding the best linear policy. In order to evaluate the quality of the approximate methods, in Section 3.5.3, I describe several techniques for computing upper bounds on the performance

of any policy for my execution problem. In Section 3.5.4, I describe the empirical calibration of the parameters of my problem. Finally, in Section 3.5.5, I present and discuss the numerical results.

3.5.1. Formulation

I follow the general framework of Section 4.2. 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. I 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, I will assume that the initial position is positive, i.e., $x_0 > \mathbf{0}$. The execution schedule must liquidate the entire initial position by the end of the time horizon, thus

$$
(3.21) \qquad x_T = x_0 + \sum_{t=1}^{T} u_t = \mathbf{0}.
$$

  Further, agency trading regulations allow only sell trades, thus

$$
(3.22) \qquad u_t \leq \mathbf{0}, \qquad t = 1, \ldots, T.
$$

  Note that any schedule satisfying (3.21)–(3.22) will also satisfy

$$
(3.23) \qquad x_t = x_0 + \sum_{s=1}^{t} u_s \geq \mathbf{0}, \qquad t = 1, \ldots, T.
$$

  I denote by $\mathcal{U}_{\mathbb{F}}^0$ the set of non-anticipating policies satisfying (3.21) almost surely, and by $\mathcal{U}_{\mathbb{F}}$ the set of non-anticipating policies satisfying (3.21)–(3.23) almost surely.

- **Return and factor dynamics.** I follow the discrete time linear dynamics of Garleanu and Pedersen [2012],[5] as described in Example 1. I assume that the price change of

---

[5]Note that Garleanu and Pedersen [2012] consider an infinite horizon setting, while my setting is finite horizon. Further, Garleanu and Pedersen [2012] solve for dynamic policies in the absence of the constraints (3.21)–(3.23).

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 follow a mean reverting process. Formally, I have the following dynamics for price changes and factor realizations:

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

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 mean reversion coefficients for the factors, and $\mu \in \mathbb{R}^N$ is the mean return. I assume that the noise terms are i.i.d., and normally distributed with zero-mean and with covariance matrices given by. $\mathrm{Var}(\epsilon_{t+1}^{(1)}) = \Psi \in \mathbb{R}^{N \times N}$ and $\mathrm{Var}(\epsilon_{t+1}^{(2)}) = \Sigma \in \mathbb{R}^{K \times K}$. I discuss the precise choice of return predicting factors and the calibration of the dynamics shortly in Section 3.5.4.

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

$$(3.24) \qquad \mathcal{V}^* \triangleq \underset{\pi \in \mathcal{U}_{\mathbb{F}}}{\text{maximize}} \ \mathsf{E}_\pi \left[\sum_{t=1}^{T} \left(x_t^\top B f_t - \tfrac{1}{2} u^\top \Lambda u_t \right) \right],$$

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

Note that the problem (3.24) is a special case of the optimization program in Example 1, with the exception of the constraints (3.21)–(3.23).

3.5.2. Approximate Policies

Since an exact, analytical solution is not available, I compare four approximate solution techniques to solve the optimal execution problem in (3.24):

- **Deterministic.** Instead of allowing for a non-anticipating 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, i.e., a deterministic sequence of trades over the entire time horizon that is decided at the begininning of the time horizon. Here, observe that at the beginning of the time horizon, the expected future factor vector is given by $\mathsf{E}[f_t | f_0] = (I - \Phi)^t f_0$. Therefore, in order to find the optimal deterministic

policy, given f_0, I maximize the conditional expected value of the stochastic objective in (3.24) by solving the quadratic program

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

(3.25)

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

- **Model predictive control.** In this approximation, at each trading time, I 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 resolving at each trading time. Formally, at time t, I solve the quadratic program

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

(3.26)

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

- **Projected LQC.** If the inequality constraints (3.22)–(3.23) are eliminated, the program would reduce to the classical linear quadratic control problem

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

(3.27)

The optimal dynamic policy for the program in (3.27) yields the trade

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

at each time t as a function of the previous position x_{t-1} and the current factor values f_t. Here, the matrices $A_{xx,t+1}$ and $A_{xf,t+1}$ are derived in the Online Supplement. The dynamic rule for u_t in (3.28) of course will not be feasible for the constrained

program (3.24), in general. This, the projected LQC policy seeks a trade decision, $\hat{u}_t$, which is the projection of u_t onto the constraint set (3.22)–(3.23), i.e., $\hat{u}_{i,t} = \max\{-x_{i,t-1}, \min\{0, u_{i,t}\}\}$, for each time $t < T$ and for each security i.

- **Optimal linear.** As formulated in Definition 2, a linear rebalancing policy specifies trades according to

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

for each time $t = 1, 2, \ldots, T$, given parameters $(\mathbf{E}, \mathbf{c})$. Due to the linear relationship between position and trade vectors, I can represent the position vector in the similar form, i.e., $x_t = d_t + \sum_{s=1}^{t} J_{s,t} f_s$ where $d_t \triangleq x_0 + \sum_{i=1}^{t} c_i$ and $J_{s,t} \triangleq \sum_{i=s}^{t} E_{s,i}$. As in Section 3.4.1.1, I implement the almost sure equality constraint (3.21) via equality constraints on the policy parameters by setting $d_T = \mathbf{0}$, and $J_{t,T} = \mathbf{0}$ for all t. I replace the almost sure inequality constraints (3.22)–(3.23) with probabilistic relaxations, as in Section 3.4.1.2. With these assumptions, I compute the parameters of the optimal linear policy by solving the following stochastic program:

(3.29)

$$\underset{(\mathbf{E}, \mathbf{c})}{\text{maximize}} \quad \mathsf{E}\left[\sum_{t=1}^{T}\left(\left(d_t + \sum_{s=1}^{t} J_{s,t} f_s\right)^{\top} B f_t - \tfrac{1}{2}\left(c_t + \sum_{s=1}^{t} E_{s,t} f_s\right)^{\top} \Lambda \left(c_t + \sum_{s=1}^{t} E_{s,t} f_s\right)\right)\right]$$

$$\text{subject to} \quad d_t = x_0 + \sum_{i=1}^{t} c_i, \qquad\qquad 1 \le t \le T,$$

$$J_{s,t} = \sum_{i=s}^{t} E_{s,i}, \qquad\qquad 1 \le s \le t \le T,$$

$$\mathsf{P}\left(d_t + \sum_{s=1}^{t} J_{s,t} f_s < \mathbf{0}\right) \le \eta, \quad 1 \le t \le T,$$

$$\mathsf{P}\left(c_t + \sum_{s=1}^{t} E_{s,t} f_s > \mathbf{0}\right) \le \eta, \quad 1 \le t \le T,$$

$$d_T = \mathbf{0},$$

$$J_{t,T} = \mathbf{0}, \qquad\qquad 1 \le t \le T.$$

Here, the parameter $\eta \in (0, 1/2)$ controls the probability that the constraints (3.22)–(3.23) are violated.[6] Using the fact that the objective is an expectation of a quadratic expression in Gaussian random variables and the fact that the chance constraints can

---

[6]I used the value $\eta = 0.2$ in my simulation results.

be handled using Lemma 3, (3.29) can be explicitly written as a second-order cone program. This calculation is detailed in the Online Supplement. Then, (3.29) can be solved using an off-the-shelf convex optimization solver.

The solution of the (3.29) provides the desired linear policy, $u_t = c_t + \sum_{s=1}^{t} E_{s,t} f_s$, in the return predicting factors. However, due to the fact that some of the constraints of the original program in (3.24) are only probabilistically enforced, u_t may not be feasible for the original program. The projected optimal linear policy seeks a trade decision, $\hat{u}_t$, which is the projection of u_t onto the constraint set (3.22)–(3.23), i.e., $\hat{u}_{i,t} = \max \{-x_{i,t-1}, \min \{0, u_{i,t}\}\}$, for each time $t < T$ and security i.

3.5.3. Upper Bounds

In order to evaluate the quality of the policies described in Section 3.5.2, I compute a number of upper bounds on the performance of the *any* policy for the program (3.24), as follows:

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

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

(3.30)

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 (3.30). Since the non-anticipating policies of the original program (3.24) are not able to utilize future factor information in making trading decisions, I have 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 (3.27), where the inequality

constraints (3.22)–(3.23) are relaxed, also provides an upper bound to (3.24). The expected value of the relaxed program can be exactly computed and yields the upper bound

$$(3.31) \quad \mathcal{V}^* \leq -\tfrac{1}{2}x_0^\top A_{xx,0}x_0 + \tfrac{1}{2}\left(\text{tr}\left(\Omega_0(I - \Phi)^\top A_{ff,0}(I - \Phi)\right) + \sum_{t=0}^{T-2}\text{tr}(\Psi A_{ff,t})\right),$$

where the matrices $A_{xx,0}$ and $A_{ff,t}$ are derived in the Online Supplement.

- **Pathwise optimization.** Given a sample path $\mathbf{f}$ of factor realizations and a sequence $\boldsymbol{\zeta} \triangleq (\zeta_1, \ldots, \zeta_T)$ of vectors $\zeta_t \in \mathbb{R}^K$ for each t, consider the quadratic optimization program

$$\mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta}) \triangleq \underset{\mathbf{u}}{\text{maximize}} \quad \sum_{t=1}^{T}\left(x_t^\top B f_t - \zeta_t^\top \epsilon_t^{(1)} - \tfrac{1}{2}u_t^\top \Lambda u_t\right)$$

$$(3.32) \qquad \text{subject to} \quad \begin{aligned} \epsilon_t^{(1)} &= f_t - (I - \Phi)f_{t-1}, & t &= 1, \ldots, T, \\ u_t &= x_t - x_{t-1}, & t &= 1, \ldots, T, \\ u_t &\leq \mathbf{0}, \ x_t \geq \mathbf{0}, & t &= 1, \ldots, T, \\ x_T &= \mathbf{0}. \end{aligned}$$

It can be established [Desai *et al.*, 2011; Brown and Smith, 2010] that for any $\boldsymbol{\zeta}$, the upper bound $\mathcal{V}^* \leq \mathsf{E}[\mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta})]$ holds — observe that the perfect foresight upper bound is a special case of this when $\boldsymbol{\zeta}$ is zero. Roughly speaking, this upper bound corresponds to a relaxation of the non-anticipating policy requirement, and $\boldsymbol{\zeta}$ correspond to a choice of Lagrange multipliers for this relaxation. The pathwise optimization upper bound corresponds to making a choice for $\boldsymbol{\zeta}$ that results in an *optimal* upper bound, i.e., $\mathcal{V}^* \leq \min_{\boldsymbol{\zeta}} \ \mathsf{E}[\mathcal{V}^{\mathsf{PO}}(\mathbf{f}, \boldsymbol{\zeta})]$. This minimization involves a convex objective function and can be computed via stochastic gradient descent; I refer the reader to Desai *et al.* [2011] for details.

3.5.4. Model Calibration

In this section, I describe the calibration the parameters of the optimal execution problem formulated in Section 3.5.1. I chose one of the most liquid stocks, Apple, Inc. (NASDAQ: AAPL), for my empirical study. I set the execution horizon to be 1 hour and trade intervals to be 5 minutes. Thus, setting a trade interval to be a one unit of time, I have a time

horizon of $T = 12$, I assume that the the the initial position to be liquidated is $x_0 = 100,000$ shares.

In trade execution problems, the time horizon is typically a day, thus I will construct a factor model in the same time-frequency. I will use the intraday transaction prices of AAPL from the NYSE TAQ database on the trading days of January 4, 2010 (day 0) and January 5, 2010 (day 1) to construct $K = 2$ return predicting factors, each with a different mean reversion speed. I first divide each trading day into 78 time intervals, each 5 minutes in length. For each 5 minute interval, I 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, while $f_{2,t}$ is the average price change relative to the previous day. Here, I can interpret the factors as the representations of value and momentum signals. Intuitively, the first factor can be considered as a 'momentum'-type signal with fast mean reversion and the second factor as a 'value'-type signal with slow mean reversion.

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

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

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

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

Similarly, I obtain the mean reversion rates for the factors,

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

Thus,

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

The variance of the error terms is estimated to be

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

The distribution of the initial factor realization, f_0, is set to the stationary distribution under the given factor dynamics, i.e., f_0 is normally distributed with zero mean and covariance

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

A rough estimate of the transaction cost coefficient $\Lambda = 2.14 \times 10^{-5}$ is used — this implies a transaction cost of \$10 on a typical trade of 1,000 shares.

3.5.5.   Numerical Results

Using the calibrated parameters from Section 3.5.4, I run a simulation with 50,000 trials to estimate the performance of each of the approximate policies of Section 3.5.2. In each trial, I sample the initial factor f_0, solve for the resulting policy of each approximate method, and compute its corresponding payoff. In order to evaluate the performance of each policy effectively, I use the same set of simulation paths in each policy's computation of average payoff. I used CVX [Grant and Boyd, 2011], a package for solving convex optimization problems in Matlab, to solve the optimization problems that occur in the computation of the deterministic, model predictive control, and optimal linear policies.

Table 3.1 summarizes the performance of each policy. For each policy, I divide the total payoff into two components, the alpha gains (i.e., $\sum_{t=1}^{T} x_t^\top B f_t$) and the transaction costs (i.e., $\sum_{t=1}^{T} -u_t^\top \Lambda u_t$). For each component as well as the total, I report the mean value over all simulation trials and the associated standard error. Finally, I report the average computation time (in seconds) required to evaluate each policy for a single simulation trial.

I observe that the optimal linear policy achieves the best performance. The gain of the optimal linear policy is approximately 7% over the next closest policy, which is the projected LQC policy. The performance of the other two policies is significantly worse. Since the projected LQC policy has a closed form expression (given a one time solution of recursive equations), its computation time per sample path is much smaller than that of the

other policies, each of which involve solving at least one optimization problem per sample path. The remaining policies have roughly the same order of magnitude in computation time, with model predictive control (which solves a different optimization problem at every time step) having the longest running time.

Despite the higher total payoff for the optimal linear policy as compared to the projected LQC policy in Table 3.1, the relatively high standard errors preclude the immediate conclusion that the optimal linear policy achieves a statistically significant higher total payoff. Thus, in order to provide a more careful comparison, for each simulation trial, I consider the difference in alpha gains, transaction costs, and total payoff between these two policies. Table 3.2 show the statistics of these differences, and establishes that the performance benefit of the optimal linear policy is statistically significant. Moreover, Table 3.2 reveals that the optimal linear policy achieves a better result by more carefully managing transaction costs, at the expense of not achieving the alpha gains of the projected LQC policy.

Finally, observe that the bottom half of Table 3.1 reports upper bounds on the total payoff of any policy, as computed using the methods described in Section 3.5.3. The pathwise optimization method achieves the tightest upper bound. Comparing this with the performance of the optimal linear policy, I conclude that the optimality gap of employing the optimal linear policy is less than 5% of the optimal value of the original program in (3.24).

3.6. Conclusion

This chapter provides a highly tractable formulation for determining rebalancing rules in dynamic portfolio choice problems with involving complex models of return predictability. My rebalancing rule is a linear function of past return predicting factors and can be utilized in a wide spectrum of portfolio choice models with realistic considerations for risk measures, transaction costs, and trading constraints. I illustrate the broad utility of my method by showing its applicability across a broad range of modeling assumptions on these portfolio optimization primitives. As long as the underlying dynamic portfolio optimization problem is a convex programming problem (i.e., concave objective and convex decision constraints),

			Alpha	TC	Total	CPU time (sec)
Policies	Deterministic	Mean	19.34	-15.81	3.53	0.82
		S.E.	0.229	0.025	0.224	
	Model predictive control	Mean	21.25	-16.54	4.71	5.79
		S.E.	0.233	0.023	0.225	
	Projected LQC	Mean	25.13	-19.40	5.73	0.02
		S.E.	0.227	0.039	0.229	
	Optimal linear	Mean	23.24	-17.11	6.13	4.23
		S.E.	0.233	0.025	0.224	
Upper Bounds	Pathwise optimization	Mean			6.46	
		S.E.			0.04	
	Perfect foresight	Mean			8.57	
		S.E.			0.223	
	Unconstrained LQC	Mean			12.58	
		S.E.			n/a	

Table 3.1: Summary of the performance statistics of each policy, along with upper bounds. In the upper half of the table I consider the approximate policies. For each approximate policy, I divide the total payoff into two components, the alpha gains and the transaction costs. For each performance statistic, I report the mean value and the associated standard error. Finally, I report the average computation time (in seconds) for each policy per simulation trial. In the bottom half of the table, I report the computed upper bounds on the total payoff. For those methods which involve Monte Carlo simulation, standard errors are also reported.

| (**Optimal Linear**) − (**Projected LQC**) | | |
	Alpha	TC	Total
Mean	-1.89	2.29	0.40
S.E.	0.0137	0.0196	0.0095

Table 3.2: Detailed comparison between the alpha gains, transaction costs, and total performance of the optimal linear policy and projected dynamic policy. I observe that the standard error for the difference in total payoff is very small, thus, the performance gain by employing the optimal linear policy is statistically significant.

the modified optimization problem seeking the optimal parameters of the linear decision rule will be a convex programming problem that is tractable numerically. I demonstrate in an optimal execution problem that such modeling flexibility can offer significant practical benefits.

Chapter 4

Common Factor Shocks in Strategic Asset Allocation

4.1. Introduction

Strategic asset allocation has been a central objective for institutional investors in active asset management due to changes in the estimates of expected future returns. With the new estimates for the future returns, the asset manager needs to update the holdings of the portfolio while aligning with the risk objectives of the fund and keeping trading costs to a minimum. On top of these tradeoffs, expected future returns are often correlated with various market conditions such as volatility and liquidity. Characterizing an optimal rebalancing rule under these complex dynamics, interactions and restrictions is a daunting task if not impossible.

Many dynamic portfolio choice models need to impose restrictive assumptions, yet often unrealistic, about return generating model in order to achieve a tractable solution. A recent simplifying assumption has been using number of shares in the portfolio decision vector in order to linearize the state dynamics. Using number of shares versus dollar holdings also required to use price changes in dollars instead of percentage terms. However, it is well-known that price changes are not stationary, cannot be estimated effectively using linear regression techniques. In this essay, I keep the nonlinear structure in the wealth evolution but instead of trying to solve the problem to optimality, I use linear policies in order to

obtain a near-optimal policy. I obtain a closed-form solution for our policy parameters which allows us to expand the universe of parameters quite easily.

I have tremendous freedom in modeling the dynamics of the return predicting factors. In a realistic framework, I allow for factor dependent covariance structure in returns driven by common factor shocks i.e., stochastic volatility. Furthermore, I can also have time-varying liquidity costs which are correlated with the expected returns of the factors. Our model involves the standard wealth equation in dollars and nonlinear dynamics for the position holdings due to the shocks to the existing wealth with current returns.

I provide a well-calibrated simulation study to analyze the performance metrics of our approach. Our simulation study shows that best linear policy provides significant benefits compared to other frequently used policies in the literature, especially when the transaction costs are high and returns evolve according to factor dependent covariance structure. Unlike other parametric approaches studied so far, our approach provides a closed form solution and the driver of the policy dynamics can be analyzed in full detail.

4.1.1. Related literature

I addressed a similar review in the previous chapter but I will re-emphasize some of the references again in this chapter's context.

The vast literature on dynamic portfolio choice starts with the seminal paper by Merton [1971] which studies the optimal dynamic allocation of one risky asset and one bond in the portfolio in a continuous-time setting. Following this seminal paper, there has been a significant literature aiming to incorporate the impact of various frictions on the optimal portfolio choice. For a survey on this literature, see Cvitanic [2001]. Constantinides [1986] studies the impact of proportional transaction costs on the optimal investment decision and observes path dependence in the optimal policy. Similarly, Davis and Norman [1990], and Dumas and Luciano [1991a] study the impact of transaction costs on the the optimal investment and consumption decision by formally characterizing the trade and no-trade regions. One drawback of all these papers is that the optimal solution is only computed in the case of a single stock and bond. Liu [2004] extends this result to multiple assets but assumes that asset returns are not correlated.

There is a growing literature on portfolio selection that incorporates return predictability with transaction costs. Balduzzi and Lynch [1999] and Lynch and Balduzzi [2000] illustrate the impact of return predictability and transaction costs on the utility costs and the optimal rebalancing rule by discretizing the state space of the dynamic program. Recently, Brown and Smith [2010] provides heuristic trading strategies and dual bounds for a general dynamic portfolio optimization problem with transaction costs and return predictability. Brandt *et al.* [2009a] parameterizes the rebalancing rule as a function of security characteristics and estimates the parameters of the rule from empirical data without modeling the distribution of the returns and the return predicting factors. Our approach is also a linear parametrization of return predicting factors, but at the micro-level, I seek to obtain a policy that is coherent with the update of the position holdings in a nonlinear fashion. Thus, our linear policy uses the convolution of the factors with their corresponding returns in order to correctly satisfy the wealth equation at all times. On a separate note, I solve for the optimal policy in closed-form using a deterministic linear quadratic control and can achieve greater flexibility in parameterizing the trading rule.

Garleanu and Pedersen [2012] achieve a closed-form solution for a model with linear dynamics in return predictors and quadratic function for transaction costs and quadratic penalty term for risk. However, the model for the security returns is given in price changes which suffers highly from non-stationarity. This use of price changes is highly nonstandard and cannot be accommodated with the existing models for return predictability that almost always uses percentage returns.

4.2. Model

4.2.1. Security and factor dynamics

I consider a dynamic portfolio optimization problem with K factors and N securities. Let $S_{i,t}$ be the discrete time dynamics for the price of the security that pays a dividend $D_{i,t}$ at time t. I assume that the gross return to our security defined by $R_{i,t+1} = \frac{S_{i,t+1}+D_{i,t+1}}{S_{i,t}}$ have the following form:

$$R_{i,t+1} = g(t, B_{i,t}^{\top}(F_{t+1} + \lambda) + \epsilon_{i,t+1}) \qquad i = 1, \ldots, N$$

for some family of functions $g(t, \cdot) : \mathbb{R} \to \mathbb{R}$, increasing in their second argument, and where I further introduce the following notation:

- $B_{i,t}$ is the $(K, 1)$ vector of exposures to the factors.

- F_{t+1} is the $(K, 1)$ vector of random (as of time t) factor realizations, with mean 0 and conditional covariance matrix $\Omega_{t,t+1}$.

- $\epsilon_{i,t+1}$ is the idiosyncratic risk of stock i.

  I assume that $\epsilon_{\cdot,t+1}$ are mean zero, have a time-invariant covariance matrix Σ_ϵ, and are uncorrelated with the contemporaneous factor realizations.

- λ_t is the $(K, 1)$ vector of conditional expected factor returns

I assume that $B_{i,t}$ and λ_t are observable and follow some known dynamics, which for now I leave unspecified (when I solve a special example below, I assume that λ_t is constant and that the $B_{i,t}$ follow a Gaussian AR(1) process, but our approach could apply to more complex dynamics). As I show below, our approach can be extended to account for time varying factor expected returns (i.e., λ_t could be stochastic), and non-normal factor or idiosyncratic risk distributions (*e.g.*, GARCH features can easily be added).

Note that this setting captures two standard return generating processes:

1. The **"discrete exponential affine"** model for security returns in which log-returns are affine in factor realizations:[1]

$$\log R_{i,t+1} = \alpha_i + B_{i,t}^\top (F_{t+1} + \lambda) + \epsilon_{i,t+1} - \frac{1}{2} \left(\sigma_i^2 + B_{i,t}^\top \Omega B_{i,t} \right)$$

2. The **"linear affine factor model"** where returns (and therefore also excess returns) are affine in factor exposures:

$$r_{i,t+1} = \alpha_i + B_{i,t}^\top (F_{t+1} + \lambda) + \epsilon_{i,t+1}$$

As I show below, our portfolio optimization approach is equally tractable for both these return generating processes.

---

[1]The continuous time version of this model is due to Vasicek [1977], Cox *et al.* [1985], and generalized in Duffie and Kan [1996]. The discrete time version is due to Gourieroux *et al.* [1993] and Le *et al.* [2010].

4.2.2.   Cash and stock position dynamics

I will assume discrete time dynamics for our cash $(w(t))$ position and dollar holdings $(x_i(t))$ in stocks. I assume that

$$
\begin{aligned}
x_{i,t+1} &= x_{i,t}R_{i,t+1} + u_{i,t+1} \qquad i = 1,\ldots,N \\
w_{t+1} &= w_t R_{0,t+1} - \sum_{i=1}^{N} u_{i,t+1} - \frac{1}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} u_{i,t+1}\Lambda_{t+1}(i,j)u_{j,t+1}
\end{aligned}
$$

where $R_{i,t+1} = \frac{S_{i,t+1}+D_{i,t+1}}{S_{i,t}}$ is the total gross return (capital gains plus dividends) on the security i. I am here effectively assuming that each position in security $i = 1,\ldots,n$ is financed by a short position in a (*e.g.*, risk-free) benchmark security 0, which I assume can be traded with no transaction costs. I denote by $x_{i,t}$ the dollar investment in asset i, by w_t the total cash balances (invested in the risk-free security S_0), and $u_{i,t+1}$ is the dollar amount of security i I will trade at price $S_{i,t+1}$. In vector notation,

(4.1) $$x_{t+1} = x_t \circ R_{t+1} + u_{t+1}$$

(4.2) $$w_{t+1} = w_t R_{0,t+1} - \mathbf{1}^\top u_{t+1} - \frac{1}{2}u_{t+1}^\top \Lambda_{t+1} u_{t+1}$$

where the operator $\circ$ denotes element by element multiplication if the matrices are of same size or if the operation involves a scalar and a matrix, then that scalar multiplies every entry of the matrix.

The matrix Λ_t captures (possibly time-varying) quadratic transaction/price-impact costs, so that $\frac{1}{2}u_t^\top \Lambda_t u_t$ is the dollar cost paid when realizing a trade at time t of size u_t. For simplicity I assume this matrix is symmetric.[2] Garleanu and Pedersen [2012] present some micro-economic foundations for such quadratic costs. As they show, the quadratic form is analytically very convenient.

4.2.3.   Objective function

I assume that the investor's objective function is to maximize a linear quadratic function of his terminal cash and stock positions $F(w_T, x_T) = w_T + a^\top x_T - \frac{1}{2}x_T^\top b\, x_T$, net of a risk-penalty which I take to be proportional to the per-period variance of the portfolio. I assume

---

[2]The symmetry assumption could easily be relaxed.

a is a $(N, 1)$ vector and b a (N, N) symmetric matrix.[3] So I assume the objective function is simply:

$$(4.3) \qquad \max_{u_1,\ldots,u_T} \mathsf{E}\left[F(w_T, x_T) - \sum_{t=0}^{T-1} \frac{\gamma}{2} x_t^\top \Sigma_{t \to t+1} x_t \right]$$

I define $\Sigma_{t \to t+1} = E_t[(R_{t+1} - E_t[R_{t+1}])(R_{t+1} - \mathsf{E}_t[R_{t+1}])']$ to be the conditional one-period variance-covariance matrix of returns and γ can be interpreted as the coefficient of risk aversion.

The $F(\cdot, \cdot)$ function parameters can be chosen to capture different objectives, such as maximizing the terminal gross value of the position $(w_T + \mathbf{1}^\top x_T)$ or the terminal liquidation (i.e., net of transaction costs) value of the portfolio $(w_T + \mathbf{1}^\top x_T - \frac{1}{2} x_T^\top \Lambda_T x_T)$, or any intermediate situation.

Assuming the investor starts with some initial cash balances w_0 and an initial position in individual stocks x_0, note that x_T and w_T can be rewritten as:

$$(4.4) \qquad x_T = x_0 \circ R_{0 \to T} + \sum_{i=1}^{T} u_t \circ R_{t \to T}$$

$$(4.5) \qquad w_T = w_0 R_{0,0 \to T} - \sum_{i=1}^{T} \left(u_t^\top \mathbf{1} R_{0,t \to T} + \frac{1}{2} u_t^\top \Lambda_t u_t R_{0,t \to T} \right)$$

where I have defined the cumulative return between date t and T on security i as:

$$(4.6) \qquad R_{i, t \to T} = \prod_{s=t+1}^{T} R_{i,s}$$

(with the convention that $R_{i, t \to t} = 1$) and the corresponding N-dimensional vector $R_{t \to T} = [R_{1, t \to T}; \ldots; R_{N, t \to T}]$.

Now note that

$$(4.7) \qquad a^\top x_T = (a \circ R_{0 \to T})^\top x_0 + \sum_{i=1}^{T} (a \circ R_{t \to T})^\top u_t$$

$$(4.8) \qquad x_T^\top b x_T = x_0^\top \overline{RbR}_0 x_0 + \sum_{t=1}^{T} u_t^\top \overline{RbR}_t u_t + 2 \sum_{t=1}^{T} x_0 \circ R_{0, t \to T} \overline{bR}_t u_t$$

where I define the (N, N)-matrix $\overline{RbR}_t$ and $\overline{bR}_t$ with respective element:

$$(4.9) \qquad \{\overline{RbR}_t\}_{ij} = R_{i, t \to T} b_{ij} R_{j, t \to T}$$

$$(4.10) \qquad \{\overline{bR}_t\}_{ij} = b_{ij} R_{j, t \to T}$$

---

[3]The symmetry assumption on b could easily be relaxed.

Substituting I obtain the following:

$$(4.11) \qquad F(w_T, x_T) \;=\; F_0 + \sum_{i=1}^{T} \left\{ G_t^\top u_t - \frac{1}{2} u_t^\top P_t u_t \right\}$$

$$(4.12) \qquad F_0 \;=\; w_0 R_{0,0\to T} + (a \circ R_{0\to T})^\top x_0 - \frac{1}{2} x_0^\top \overline{RbR}_0 x_0$$

$$(4.13) \qquad G_t \;=\; a \circ R_{t\to T} + \mathbf{1} \circ R_{0,t\to T} - x_0 \circ R_{0,t\to T} \overline{bR}_t$$

$$(4.14) \qquad P_t \;=\; (\overline{RbR}_t + \Lambda_t \circ R_{0,t\to T})$$

Substituting into the objective function given in equation 4.3 it can be rewritten as:

$$(4.15) \qquad F_0 + \max_{u_1,\dots,u_T} \sum_{t=0}^{T-1} \mathsf{E} \left[G_{t+1}^\top u_{t+1} - \frac{1}{2} u_{t+1}^\top P_{t+1} u_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma_{t\to t+1} x_t \right]$$

subject to the non-linear dynamics given in equations 4.1 and 4.2.

I next describe our set of *linear policies*, which make this problem tractable. At this stage it is convenient to introduce the following notation (inspired from matlab): I write $[A; B]$ (respectively $[A\,B]$) to denote the vertical (respectively horizontal) concatenation of two matrices.

4.2.4. Linear policies

I consider a class of parametric linear policies that is richer than the one previously considered in the literature (see, *e.g.*, Brandt *et al.* [2009b]), but nevertheless has the advantage of leading to an explicit solution for the portfolio choice problem with transaction costs. Thus, in contrast to the approach proposed in Brandt *et al.* [2009b]), I do not need to perform a numerical optimization, and can handle transaction costs efficiently. Further, in contrast to Garleanu and Pedersen [2012] I can handle more complex asset return dynamics and explicitly formulate the problem in terms of dollar returns (as opposed to number of shares), and yet retain the analytical flexibility of the linear-quadratic framework.

These benefits come at a cost, namely that of restricting our optimization to a specific set of parametrized trading strategies. It is an empirical question whether the set I work with is sufficiently large to deliver useful results. I present some empirical tests of our approach in the next section. First, I describe the strategy set I consider. Then I explain how the portfolio optimization can be done in closed-form, within that restricted set.

I define our set of linear policies with a set of $(K+1)$-dimensional vectors of parameters, $\pi_{i,s,t}$ and $\theta_{i,s,t}$, defined for all $i = 1, \ldots, N$ and for all $s \leq t$. The (previously defined) time t trade of asset i ($u_{i,t}$) of and dollar investment in asset i ($x_{i,t}$) are given by:

$$(4.16) \qquad\qquad u_{i,t} = \sum_{u=1}^{t} \pi_{i,u,t}^{\top} \mathcal{B}_{i,u,t}$$

and

$$(4.17) \qquad\qquad x_{i,t} = \sum_{u=1}^{t} \theta_{i,u,t}^{\top} \mathcal{B}_{i,u,t}$$

are then vector products of $\pi_{i,s,t}$ and $\theta_{i,s,t}$ and a $(K+1)$ vector

$$(4.18) \qquad\qquad \mathcal{B}_{i,u,t} = [1; B_{i,t}] R_{i,u \to t}.$$

$\mathcal{B}_{i,u,t}$ is seen to be the (K) vector of time t factor exposures, augmented with a "1", and all weighted by the cumulative return earned by security i between time u and t. In other words, these policies allow trades at time t to depend on current factor exposures $B_{i,t}$, but also on all past exposures weighted by their past holding period returns.

Intuitively, the dependence on current exposures, unweighted by lagged returns, is clearly important. In fact, in a no-transaction cost affine portfolio optimization problem where the optimal solution is well-known, the optimal solution will involve only current exposures (see, *e.g.*,**?**). Note that this is also the choice made by Brandt *et al.* [2009b] for their 'parameteric portfolio policies.' However, while Brandt *et al.* [2009b] specify the loadings on exposure of individual stocks to be identical, I allow two stocks with identical exposures (and with perhaps different levels of idiosyncratic variance) to have different weights and trades.[4]

With transaction costs, allowing portfolio weights and trades to depend on past reutrns interacted with past exposures seems useful. The intuition for this comes from the path-dependence I observe in known closed-form solutions [see Constantinides, 1986; Davis and Norman, 1990; Dumas and Luciano, 1991b; Liu and Loewenstein, 2002, and others]

To proceed, I note that the assumed linear position and trading strategies in equations 4.42 and 4.41 have to satisfy the dynamics given in equations 4.1 and 4.2. It follows

---

[4](Note, for the Brandt *et al.* [2009b] econometric approach it is useful to have fewer parameters. This is not an issue with our approach as our solution is closed-form.

that the parameter vectors $\pi_{i,s,t}$ and $\theta_{i,s,t}$ have to satisfy the following restrictions, for all $i = 1, \ldots, N$:

(4.19) $$\pi_{i,s,t} = \theta_{i,s,t} - \theta_{i,s,t-1} \qquad \text{for } s < t$$

(4.20) $$\pi_{i,t,t} = \theta_{i,t,t}$$

I can rewrite these policies in a concise matrix form. First, define the $(N(K+1)t, 1)$ vectors π_t and θ_t as

(4.21) $$\pi_t = [\pi_{1,1,t}; \ldots; \pi_{n,1,t}; \pi_{1,2,t}; \ldots; \pi_{n,2,t}; \ldots; \pi_{1,t,t}; \ldots; \pi_{n,t,t}]$$

(4.22) $$\theta_t = [\theta_{1,1,t}; \ldots; \theta_{n,1,t}; \theta_{1,2,t}; \ldots; \theta_{n,2,t}; \ldots; \theta_{1,t,t}; \ldots; \theta_{n,t,t}]$$

Further, let's define the following $(N(K+1), N)$ matrices (defined for all $1 \leq s \leq t \leq T$) as the diagonal concatenations of the N vectors $\mathcal{B}_{i,s,t} \; \forall i = 1, \ldots, N$:

$$\mathcal{B}_{s,t} = \left\{ \begin{array}{ccccc} \mathcal{B}_{1,s,t} & 0 & 0 \ldots & & 0 \\ 0 & \mathcal{B}_{2,s,t} & 0 \ldots & & 0 \\ & & \ldots & & \\ 0 & & \ldots & 0 & \mathcal{B}_{n,s,t} \end{array} \right\}$$

Then I can define the $(N(K+1)t, N)$ matrix $\mathcal{B}_t$ by stacking the t matrices $\mathcal{B}_{s,t} \; \forall s = 1, \ldots, t$:

$$\mathcal{B}_t = [\mathcal{B}_{1,t}; \mathcal{B}_{2,t}, \ldots, \mathcal{B}_{t,t}]$$

It is then straightforward to check that:

(4.23) $$u_t = \mathcal{B}_t^\top \pi_t$$

(4.24) $$x_t = \mathcal{B}_t^\top \theta_t$$

Further, in terms of these definitions the constraints on the parameter vector in 4.19 can be rewritten concisely as:

(4.25) $$\theta_{t+1} - \theta_t^0 = \pi_{t+1}$$

where I define $\theta_t^0 = [\theta_t; \mathbf{0}_{K+1}]$ to be the vector θ_t stacked on top of a $(K+1, 1)$ vector of zeros $\mathbf{0}_{K+1}$.

The usefulness of restricting ourselves to this set of 'linear trading strategies' is that optimizing over this set amounts to optimizing over the parameter vectors π_t and θ_t, and that, as I show next, that problem reduces to a **deterministic** linear-quadratic control problem, which can be solved in closed form.

Indeed, substituting the definition of our linear trading strategies from equation 4.23 into our objective function I may rewrite the original problem given in equation 4.15 as follows.

$$(4.26) \qquad F_0 + \max_{\pi_1,\ldots,\pi_T} \sum_{t=0}^{T-1} \mathcal{G}_{t+1}^\top \pi_{t+1} - \frac{1}{2}\pi_{t+1}^\top \mathcal{P}_{t+1}\pi_{t+1} - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t\theta_t$$

$$(4.27) \qquad s.t. \ \theta_{t+1} - \theta_t^0 = \pi_{t+1}$$

and where I define the vectors $\mathcal{G}_t$ and the matrices $\mathcal{P}_t$ and $\mathcal{Q}_t$ defined for all $t = 0, \ldots, T$ by

$$(4.28) \qquad \mathcal{G}_t = \mathsf{E}[\mathcal{B}_t G_t]$$

$$(4.29) \qquad \mathcal{P}_t = \mathsf{E}[\mathcal{B}_t P_t \mathcal{B}_t^\top]$$

$$(4.30) \qquad \mathcal{Q}_t = \mathsf{E}[\mathcal{B}_t \Sigma_t \mathcal{B}_t^\top]$$

Note that I choose the time indices for the matrices $\mathcal{G}_t, \mathcal{P}_t, \mathcal{Q}_t$ to reflect their (identical) size (index t denotes a square-matrix or vector of row-length $N(K+1)t$. The matrices $\mathcal{G}_t, \mathcal{P}_t, \mathcal{Q}_t$ can be solved for explicitly or by simulation depending on the assumptions made about the return generating process R_t and the factor dynamics $B_{i,t}$. But once these expressions have been computed or simulated (and this only needs to be done once), then the explicit solution for the optimal strategy can be derived using standard deterministic linear-quadratic dynamic programming. I derive the solution next.

4.2.5. Closed form solution

Define the value function

$$V(n) = \max_{\pi_{n+1},\ldots,\pi_T} \sum_{t=n}^{T-1} \mathcal{G}_{t+1}^\top \pi_{t+1} - \frac{1}{2}\pi_{t+1}^\top \mathcal{P}_{t+1}\pi_{t+1} - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t\theta_t$$

Now at $n = T - 1$ I have

$$V(T-1) = \max_{\pi_T} \ \mathcal{G}_T^\top \pi_T - \frac{1}{2}\pi_T^\top \mathcal{P}_T\pi_T - \frac{\gamma}{2}\theta_{T-1}^\top \mathcal{Q}_{T-1}\theta_{T-1},$$

which yields the solution $\pi_T^* = \mathcal{P}_T^{-1}\mathcal{G}_T$ and the value function $V(T-1) = \frac{1}{2}\mathcal{G}_T^\top \mathcal{P}_{T-1}^{-1}\mathcal{G}_T - \frac{\gamma}{2}\theta_{T-1}^\top \mathcal{Q}_{T-1}\theta_{T-1}$. I therefore guess that the value function is of the form:

$$(4.31) \qquad V(n) = -\frac{1}{2}\theta_n^\top M_n \theta_n + L_n^\top \theta_n + H_n$$

The Hamilton-Jacobi-Bellman equation is

$$(4.32) \qquad V(t) = \max_{\pi_{t+1}} \left\{ \mathcal{G}_{t+1}^\top \pi_{t+1} - \frac{1}{2}\pi_{t+1}^\top \mathcal{P}_{t+1}\pi_{t+1} - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t\theta_t + V(t+1) \right\}$$

$$(4.33) \qquad s.t. \qquad \theta_{t+1} - \theta_t^0 = \pi_{t+1}$$

The first order condition is:

$$\mathcal{G}_{t+1} + L_{t+1} - (\mathcal{P}_{t+1} + M_{t+1})\pi_{t+1} = M_{t+1}\theta_t^0$$

which gives the optimal trade (and corresponding) state equation:

$$(4.34) \qquad \pi_{t+1} = [\mathcal{P}_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1} - M_{t+1}\theta_t^0)$$

$$(4.35) \qquad \theta_{t+1} = [\mathcal{P}_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1} + \mathcal{P}_{t+1}\theta_t^0)$$

The HJB equation can be rewritten with our guess as

$$V(t) = \pi_{t+1}^\top \left(\mathcal{G}_{t+1} + L_{t+1} - \frac{1}{2}(\mathcal{P}_{t+1} + M_{t+1})\pi_{t+1} \right) - \frac{1}{2}\theta_t^0 M_{t+1}(\theta_t^0)^\top - \pi_{t+1}^\top M_{t+1}\theta_t^0 - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t\theta_t + H_{t+1} + L_{t+1}^\top \theta_t^0$$

Now, for a $[N(K+1)t, N(K+1)t]$ dimensional square matrix X_t I define $\overline{X}_t$ to be the upper left-hand corner square submatrix with dimensions $[N(K+1)(t-1), N(K+1)(t-1)]$. Using this definition and substituting the FOC I get:

$$\begin{aligned} V(t) &= \frac{1}{2}(\mathcal{G}_{t+1} + L_{t+1} - M_{t+1}\theta_t^0)^\top [P_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1} - M_{t+1}\theta_t^0) - \frac{1}{2}\theta_t^\top (\overline{M}_{t+1} + \gamma \mathcal{Q}_t)\theta_t \\ &\quad + H_{t+1} + L_{t+1}^\top \theta_t^0 \end{aligned}$$

which I can simplify further:

$$\begin{aligned} V(t) &= \frac{1}{2}(\mathcal{G}_{t+1} + L_{t+1})^\top [\mathcal{P}_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1}) - \frac{1}{2}\theta_t^\top \left(\overline{M}_{t+1} + \gamma \mathcal{Q}_t \right. \\ &\quad \left. - \overline{M_{t+1}[\mathcal{P}_{t+1} + M_{t+1}]^{-1}M_{t+1}} \right)\theta_t + H_{t+1} + (L_{t+1} + M_{t+1}[\mathcal{P}_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1}))^\top \theta_t^0 \end{aligned}$$

Thus I confirm our guess for the value function and find the system of recursive equations:

$$(4.36) \qquad M_t = \overline{M_{t+1} + \gamma \mathcal{Q}_t - \overline{M_{t+1}[\mathcal{P}_{t+1} + M_{t+1}]^{-1} M_{t+1}}}$$

$$(4.37) \qquad L_t = \overline{L_{t+1} + M_{t+1}^\top [\mathcal{P}_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1})}$$

$$(4.38) \qquad H_t = H_{t+1} + \frac{1}{2}(\mathcal{G}_{t+1} + L_{t+1})^\top [\mathcal{P}_{t+1} + M_{t+1}]^{-1}(\mathcal{G}_{t+1} + L_{t+1})$$

4.3. Experiment

In this section I present several experiments to illustrate the usefulness of our portfolio selection approach. I compare portfolio selection in a characteristics-based versus factors-based return generating environment. As I show below the standard linear-quadratic portfolio approach is well-suited to the characteristics-based environment, but in a factor-based environment, since it cannot adequately capture the systematic variation in the covariance matrix due to variations in the exposures it is less successful. Instead, our approach can handle this feature.

4.3.1. Characteristics versus Factor-based return generating model

I wish to compare the following two environments:

- The factor-based return generating process

$$(4.39) \qquad R_{i,t+1} = \alpha_i + B_{i,t}^\top (F_{t+1} + \lambda) + \epsilon_{i,t+1}$$

- The characteristics based return generating process:

$$(4.40) \qquad R_{i,t+1} = \alpha_i + B_{i,t}^\top \lambda + \omega_{i,t+1}$$

where in both cases I assume that there are three return generating factors corresponding to (1) short term (5-day) reversal, (2) medium term (1 year) momentum, (3) long-term (5 year) reversal (and potentially a common market factor).

Note the difference between the two frameworks. In the characteristics based framework, the conditional covariance of returns is constant $\Sigma_{t \to t+1} = \Sigma_\omega$ and is therefore not affected

by the factor exposures. Instead, in the factor-based framework, the conditional covariance matrix of returns is time varying: $\Sigma_{t \to t+1} = B_t \Omega B_t^\top + \Sigma_\epsilon$ where $B_t = [B_{1,t}^\top; B_{2,t}^\top; \ldots; B_{n,t}^\top]$ is the (N, K) matrix of factor exposures.

I assume that the half-life of the 5-day factor is 3 days, that of the one-year factor is 150 days, that of the 5-year factor is 700 days. I define the exposure dynamics using the simple auto-regressive process:

$$B_{i,t+1}^k = (1 - \phi_k)B_{i,t}^k + \epsilon_{i,t+1}.$$

The value of ϕ_k is tied to its half-life (expressed in number of days) $\hat{h}_k$ by the simple relation $\phi_k = (\frac{1}{2})^{\hat{h}_k}$.

For the case, where I investigate the 'Characteristics based' model I set the constant covariance matrix Σ_ω so that it matches the unconditional covariance matrix of the factor based return generating process, i.e., I set

$$\Sigma_\omega = \mathsf{E}[B_t \Omega B_t^\top + \Sigma_\epsilon]$$

Note that

$$B_t \Omega B_t^\top = \sum_{l,m=1}^K \Omega_{l,m} B_{:,t}^l (B_{:,t}^m)^\top$$

where $B_{:,t}^k$ is the factor values of each asset corresponding to the kth factor at time t.

4.3.2. Calibration of main parameters

The number of assets in our experiment is 15. One can think of these as a collection of portfolios instead of individual stocks, e.g., stock or commodity indices. Our trading horizon is 26 weeks with weekly rebalancing. Our objective is to maximize net terminal wealth minus penalty terms for excessive risk. This requires us to set $a = 1$ and $b = 0$ in our objective function.

I calibrate the factor mean, λ, and covariance matrix, Ω, using Fama-French 10 portfolios sorted on short-term reversal, momentum, and long term reversal. Using monthly returns, I compute the performance of the long-short portfolio for the highest and lowest decile in each factor data. Obtaining 3 long-short portfolios, I set λ to be its mean and Ω to be its covariance matrix. Table 4.1 illustrates the estimated values for λ and Ω.

	Fama-French Moments
λ_1	-0.00726
λ_2	0.00182
λ_3	-0.00323
Ω_{11}	0.00103
Ω_{12}	0.00051
Ω_{13}	0.00154
Ω_{22}	0.00050
Ω_{23}	0.00081
Ω_{33}	0.00162

Table 4.1: Calibration results for λ and Ω.

For our simulations, I assume that both F and ϵ vectors are serially independent and normally distributed with zero mean and covariance matrix Ω and Σ_ϵ, respectively. I assume that Σ_ϵ is a diagonal matrix e.g., $\mathsf{diag}(\sigma_\epsilon)$. Each entry in σ_ϵ is set randomly at the beginning of the simulation according to a normal distribution with mean 0.20 and standard deviation 0.05.

Initial distribution for $B_{i,0}^k$ is given by the unconditional stationary distribution of $B_{i,t}^k$ which is given by a normal distribution with mean zero and variance $\frac{\sigma_{\epsilon,i}^2}{2\phi - \phi^2}$.

Transaction cost matrix, Λ is assumed to be a constant multiple of Σ_ω or Σ_ϵ with proportionality constant η in characteristics or factor-based return generating model respectively. I use a rough estimate of η according to widely used transaction cost estimates reported in the algorithmic trading community. I provide two regimes: low and high transaction cost environment. The slippage values for these two regimes are assumed to be around 4bps and 400bps respectively. Therefore, I expect that a trade with a notional value of $100,000$ results in $40 and $4000 of transaction costs in these regimes. In our model, $\eta \sigma_\epsilon^2 u^2$ measures the corresponding transaction cost of trading u dollars. Using $u = \$100,000$ and $\sigma_\epsilon = 0.20$, this yields that η is roughly around 5×10^{-6} and 5×10^{-4} for the low and high transaction cost regimes respectively.

Finally, I assume that the coefficient of risk aversion, γ equals 10^{-6}, which I can think of as corresponding to a relative risk aversion of 1 for an agent with 1 million dollars under management.

4.3.3. Approximate policies

Due to the nonlinear dynamics in our wealth function, solving for the optimal policy even in the case of concave objective function is intractable due to the curse of dimensionality. In this section, I will provide various policies that will help us compare the performance of the best linear policy to the existing approaches in the literature.

Garleanu & Pedersen Policy (GP): Using the methodology in Garleanu and Pedersen [2012] , I can construct an approximate trading policy that will work in our current set-up. A closed-form solution can be obtained if one works with linear dynamics in state and control variables:

$$\bar{r}_{t+1} = C_t f_t + \epsilon_{t+1}$$

$$f_{t+1} = (I - \Phi) f_t + \varepsilon_{t+1}$$

where $\bar{r}_{t+1}$ stores dollar price changes. Then, our problem can be cast in their notation with

$$\max \mathsf{E} \left[\sum_{t=1}^{T} \left(x_{t-1}^{\top} \bar{r}_t - \frac{\gamma}{2} x_t^{\top} \bar{\Sigma}_t x_t - \frac{1}{2} u_t^{\top} \bar{\Lambda} u_t \right) \right]$$

where $\bar{\Lambda}$ and $\bar{\Sigma}_t$ are deterministic and measured in dollars. The optimal solution to this is given by

$$x_t = \left(\bar{\Lambda} + \gamma \bar{\Sigma}_t + A_{xx}^t \right)^{-1} \left(\bar{\Lambda} x_{t-1} + \left(A_{xf}^t (I - \Phi) \right) f_t \right)$$

with the following recursions:

$$A_{xx}^{t-1} = -\bar{\Lambda} \left(\bar{\Lambda} + \gamma \bar{\Sigma}_t + A_{xx}^t \right)^{-1} \bar{\Lambda} + \bar{\Lambda}$$

$$A_{xf}^{t-1} = \bar{\Lambda} \left(\bar{\Lambda} + \gamma \bar{\Sigma}_t + A_{xx}^t \right)^{-1} \left(A_{xf}^t (I - \Phi) \right) + C_t$$

I use the following transformations in order to address our dynamics:

$$C_t = \mathsf{E}[\mathrm{diag}(S_t)]\left(\lambda^\top \otimes I_{N\times N}\right)$$

$$\bar{\Lambda} = \mathsf{E}[S_t S_t^\top]\Lambda$$

$$\bar{\Sigma}_t = \mathrm{Var}(\bar{r}_{t+1})$$

Myopic Policy (MP): I can solve for the myopic policy using only one-period data. I solve the myopic problem given by

$$\max \mathsf{E}\left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2}x_t^\top \Sigma_t x_t - \frac{1}{2}u_t^\top \Lambda u_t\right)\right].$$

Using the dynamics for r_{t+1}, the optimal myopic policy is given by

$$x_t = \left(\Lambda + \gamma\left(B_t \Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}\left(B_t \lambda + \Lambda\left(x_{t-1} \circ R_t\right)\right)$$

Myopic Policy with Transaction Cost Aversion (MP-TC): Since myopic policy only considers the current state of the return predicting factors, it realizes substantial transaction costs. This policy can be significantly improved by considering an another optimization problem on the transaction cost matrix which ultimately tries to control the amount of transaction costs incurred by the policy. Thus, this policy uses

$$x_t = \left(\tau^* \Lambda + \gamma\left(B_t \Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}\left(B_t \lambda + \tau^* \Lambda\left(x_{t-1} \circ R_t\right)\right)$$

where τ^* is given by

$$\underset{\tau}{\mathrm{argmax}}\, \mathsf{E}\left[\left(x_t^\top r_{t+1} - \frac{\gamma}{2}x_t^\top \Sigma_t x_t - \frac{1}{2}u_t^\top \Lambda u_t\right)\right]$$

with

$$x_t = \left(\tau \Lambda + \gamma\left(B_t \Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}\left(B_t \lambda + \tau \Lambda\left(x_{t-1} \circ R_t\right)\right)$$

Best Linear Policy (BL): Using the methodology in Section 4.2.4, I can find the optimal linear policy that satisfies our nonlinear state evolution:

$$u_t = \mathcal{B}_t^\top \pi_t^*$$
$$x_t = \mathcal{B}_t^\top \theta_t^*$$

where π_t^* and θ_t^* solve the following program

$$\max_{\pi_1,\ldots,\pi_T} \sum_{t=0}^{T-1} \mathcal{G}_{t+1}^\top \pi_{t+1} - \frac{1}{2}\pi_{t+1}^\top \mathcal{P}_{t+1}\pi_{t+1} - \frac{\gamma}{2}\theta_t^\top \mathcal{Q}_t\theta_t$$

$$s.t.\ \theta_{t+1} - \theta_t^0 = \pi_{t+1}$$

Restricted Best Linear Policy (RBL): Instead of using the whole history of stochastic factors in our policy, I can restrict the best linear policy to use only a fixed number of periods. In this experiment, I will use only the last observed exposures in our position vector, x_t, and the last two period's exposures and the last period's return in our trade vector, u_t. Formally, I will let

$$(4.41) \qquad\qquad\qquad x_{i,t} = \theta_{i,t}^\top \mathcal{B}_{i,t,t},$$

and

$$(4.42) \qquad\qquad\qquad u_{i,t} = \pi_{i,1,t}^\top \mathcal{B}_{i,t-1,t} + \pi_{i,2,t}^\top \mathcal{B}_{i,t,t},$$

where I need

$$\pi_{i,2,t} = \theta_{i,t}$$

$$\pi_{i,1,t} = -\theta_{i,t-1}$$

$$\pi_{i,1,1} = 0,$$

in order to satisfy the nonlinear state dynamics in (4.1) and (4.2).

Myopic Policy without Transaction Costs (NTC): Without transaction costs, our trading problem is easy to solve, namely, the myopic policy will be optimal. Thus, using the myopic policy in the absence of transaction costs, i.e.,

$$x_t = \left(\gamma\left(B_t\Omega B_t^\top + \Sigma_\epsilon\right)\right)^{-1}(B_t\lambda)$$

and applying it to the objective function without the transaction cost terms will provide us an upper bound for the optimal objective value of the original dynamic program. This policy will help us to evaluate how suboptimal the approximate policies are in the worst case.

4.3.4. Simulation Results

I run the performance statistics of our approximate policies in the presence and lack of factor noise and low and high transaction costs. I observe that in all of these cases, best linear policy performs very well compared to the other approximate policies and when compared to the upper bound it achieves near-optimal performance.

	GP	MP	MP-TC	RBL	BL	NTC
Avg Wealth	269	573	574.6	547.5	568.5	594.3
Avg Objective	108.1	281.9	282.4	281.1	291.0	297.0
Variance	1.21e+05	1.37e+05	1.37e+05	1.23e+05	1.30e+05	1.44e+05
TC	4.846	9.967	11.71	14.66	13.45	0
Sharpe with TC	1.094	2.188	2.196	2.207	2.231	2.215
Sharpe w/o TC	1.07	2.194	2.204	2.22	2.244	2.215
Weekly Sharpe with TC	2.205	3.39	3.393	3.383	3.443	3.453

Table 4.2: Summary of the performance statistics of each policy in the case of no common factor noise and low transaction cost environment. For each policy, I report average terminal wealth, average objective value, variance of the terminal wealth, average terminal sharpe ratio in the presence and lack of transaction costs and average weekly sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

Table 4.2 illustrates that when transaction costs are relatively small, myopic policies are also near-optimal but even in this case best linear policy dominates in terms of performance. Garleanu & Pedersen policy does not perform very well mainly due to the return dynamics expressed in percentage terms versus dollar units. Table 4.3 underlines the amount of improvement introduced with the best linear policy. In this case, myopic policies perform significantly worse than the best linear policy.

Table 4.4 and Table 4.5 depict the impact of common factor shocks in the terminal wealth statistics. It is important to note that in this regime, sharpe ratios are significantly lower. In both cases, best linear policy achieves the best objective value statistics.

	GP	MP	MP-TC	RBL	BL	NTC
Avg Wealth	127.2	180.9	52.19	74.76	232.2	594.3
Avg Objective	29.51	-98.41	25.31	59.58	138.1	297
Variance	6.37e+04	1.72e+05	2.28e+04	3.77e+03	3.14e+04	1.44e+05
TC	29.16	7.744	0.2971	44.43	43.97	0
Sharpe with TC	0.713	0.6168	0.4886	1.722	1.853	2.215
Sharpe w/o TC	0.7758	0.5421	0.4896	2.222	1.94	2.215
Weekly Sharpe with TC	1.8	2.132	2.098	2.003	2.517	3.453

Table 4.3: Summary of the performance statistics of each policy in the case of no common factor noise and high transaction cost environment. For each policy, I report average terminal wealth, average objective value, variance of the terminal wealth, average terminal sharpe ratio in the presence and lack of transaction costs and average weekly sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

4.4. Conclusion and Future Directions

In this essay, I provide a methodology that accommodates complex return predictability models studied in the literature in multi-period models with transaction costs. Our return predicting factors does not need to follow any pre-specified model but instead can have arbitrary dynamics. I allow for factor dependent covariance structure in returns driven by common factor shocks which is prevalent in the asset management literature. On an interesting further study, I can also have time-varying liquidity costs which are correlated with the expected returns of the factors.

Our simulation study shows that best linear policy provides significant benefits compared to other frequently used policies in the literature, especially when the transaction costs are high and returns evolve according to factor dependent covariance structure. Unlike other parametric approaches studied so far, our approach provides a closed form solution and the driver of the policy dynamics can be analyzed in full detail.

	GP	MP	MP-TC	RBL	BL	NTC
Avg Wealth	39.9	38	39.44	19.3	39.23	41.81
Avg Objective	-144.2	15.38	19.28	9.785	20.51	20.75
Variance	3.77e+04	2.25e+04	4.07e+03	2.04e+03	9.19e+03	4.21e+03
TC	0.9264	1.186	0.98	0.2683	1.785	0
Sharpe with TC	0.2907	0.3586	0.87	0.604	0.5786	0.911
Sharpe w/o TC	0.2932	0.8848	0.9	0.6121	0.586	0.911
Weekly Sharpe with TC	0.3693	0.9058	0.92	0.7347	0.8756	0.9436

Table 4.4: Summary of the performance statistics of each policy in the case of common factor noise and low transaction cost environment. For each policy, I report average terminal wealth, average objective value, variance of the terminal wealth, average terminal sharpe ratio in the presence and lack of transaction costs and average weekly sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

	GP	MP	MP-TC	RBL	BL	NTC
Avg Wealth	15.2	14.52	15.66	9.822	16.21	41.81
Avg Objective	-50.36	3.77	5.68	5.851	9.133	20.75
Variance	1.32e+04	4.94e+04	1.07e+04	8.12e+02	1.63e+03	4.21e+03
TC	5.939	1.916	2.91	1.881	2.059	0
Sharpe with TC	0.187	0.2919	0.21	0.4873	0.5674	0.911
Sharpe w/o TC	0.2693	0.477	0.57	0.5738	0.6065	0.911
Weekly Sharpe with TC	0.3619	0.5267	0.57	0.619	0.7386	0.9436

Table 4.5: Summary of the performance statistics of each policy in the case of common factor noise and low transaction cost environment. For each policy, I report average terminal wealth, average objective value, variance of the terminal wealth, average terminal sharpe ratio in the presence and lack of transaction costs and average weekly sharpe ratio in the presence of transaction costs. (Dollar values are in thousands of dollars.)

Bibliography

Y. Aït-Sahalia and J. Jacod. Is brownian motion necessary to model high-frequency data? *The Annals of Statistics*, 38(5):3093–3128, 2010.

Y. Aït-Sahalia and J. Yu. High frequency market microstructure noise estimates and liquidity measures. *The Annals of Applied Statistics*, 3(1):422–457, 2009.

R. F. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2):5–39, 2000.

J. Angel. Limit versus market orders. Working Paper No. FINC-1377-01-293, School of Business Administration, Georgetown University, Washington, D.C., 1994.

J. Bacidore, R. Battalio, and R. Jennings. Order submission strategies, liquidity supply, and trading in pennies on the New York Stock Exchange. *Journal of Financial Markets*, 6:337–362, 2003.

K. Back and S. Baruch. Working orders in limit order markets and floor exchanges. *The Journal of Finance*, 62(4):1589–1621, 2007.

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

O.E. Barndorff-Nielsen, D.G. Pollard, and N. Shephard. Discrete-valued levy processes and low latency financial econometrics. Technical report, Department of Economics (University of Oxford), 2010.

A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

A. Ben-Tal, B. Golany, A. Nemirovski, and J. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7:248–271, July 2005.

D. P. Bertsekas. *Dynamic Programming and Optimal Control.* Athena Scientific, 2nd edition, 2000.

D. Bertsimas and V. Goyal. On the power and limitations of affine policies in two-stage adaptive optimization. *Mathematical Programming*, pages 1–41, 2011.

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

D. Bertsimas, L. Kogan, and A. W. Lo. When is time continuous? *Journal of Financial Economics*, 55:173–204, 2000.

D. Bertsimas, D. A. Iancu, and P. A. Parrilo. Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.*, 35, May 2010.

S. Boyd and L. Vandenberghe. *Convex Optimization.* Cambridge University Press, March 2004.

S. Boyd, M. T. Mueller, B. O'Donoghue, and Wang Y. Performance bounds and suboptimal policies for multi-period investment. Working paper, 2012.

P. P. Boyle and D. Emanuel. Discretely adjusted option hedges. *Journal of Financial Economics*, 8:259–282, 1980.

M. W. Brandt, P. Santa-Clara, and R. Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22(9):3411–3447, September 2009.

Michael W. Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *Review of Financial Studies*, 22(9):3411–3447, 2009.

J. Brogaard. High frequency trading and its impact on market quality. Working paper, 2010.

D. B. Brown and J. E. Smith. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. Working paper, 2010.

G. C. Calafiore. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control Optim.*, 48:2254–2274, June 2009.

J.Y. Campbell and L.M. Viceira. *Strategic asset allocation.* Oxford University Press Oxford, 2002.

G. Cespa and T. Foucault. Insiders-outsiders, transparency, and the value of the ticker. Working Paper No. 628, Department of Economics, Queen Mary, University of London, April 2008.

G. C. Chacko, J. W. Jurek, and E. Stafford. The price of immediacy. *Journal of Finance*, 63(3):1253–1290, 2008.

N. Chen, R. Roll, and S. A. Ross. Economic forces and the stock market. *The Journal of Business*, 59(3):pp. 383–403, 1986.

X. Chen, M. Sim, and P. Sun. A robust optimization perspective on stochastic programming. *Operations Research*, 55:1058–1071, November 2007.

X. Chen, M. Sim, P. Sun, and J. Zhang. A linear decision-based approximation approach to stochastic programming. *Operations Research*, 56:344–357, March 2008.

S.N. Cohen and L. Szpruch. A limit order book model for latency arbitrage. *Arxiv preprint arXiv:1110.4811*, 2011.

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

T. Copeland and D. Galai. Information effects and the bid-ask spread. *Journal of Finance*, 38(5):1457–1469, 1983.

John C. Cox, Jonathan E. Ingersoll, Jr., and Stephen A. Ross. An intertemporal general equilibrium model of asset prices. *Econometrica*, 53:363–384, 1985.

J. Cvitanić and A. Kirilenko. High frequency traders and asset prices. Working paper, March 2010.

J. Cvitanic. Theory of portfolio optimization in markets with frictions. In E. Jouini, J. Cvitanic, and Marek Musiela, editors, *Handbooks in Mathematical Finance*. Cambridge University Press, 2001.

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

V.V. Desai, V.F. Farias, and C.C. Moallemi. Pathwise optimization for linear convex systems. *Working paper*, 2011.

J. B. Detemple, R. Garcia, and M. Rindisbacher. A monte carlo method for optimal portfolios. *The Journal of Finance*, 58(1):401–446, February 2003.

Darrell Duffie and Rui Kan. A yield-factor model of interest rates. *Mathematical finance*, 6(4):379–406, 1996.

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

Bernard Dumas and Elisa Luciano. An exact solution to a dynamic portfolio choice problem under transactions costs. *Journal of Finance*, pages 577–595, 1991.

R. Durrett. *Probability: Theory and Examples*. Duxbury Press, 3rd edition, 2004.

D. Easley, T. Hendershott, and T. Ramadorai. The price of latency. Working paper, May 2008.

E. F. Fama and K. R. French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1):pp. 55–84, 1996.

N. Garleanu and L.H. Pedersen. Dynamic trading with predictable returns and transaction costs. Working paper, April 2012.

J. Gatheral. No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7):749–759, 2010.

P. Glasserman and X. Xu. Robust portfolio control with stochastic factor dynamics. Working paper, 2011.

L. R. Glosten. Is the electronic open limit order book inevitable? *Journal of Finance*, 49(4):1127–1161, 1994.

W. N. Goetzmann and P. Jorion. Testing the predictive power of dividend yields. *Journal of Finance*, 48(2):663–679, 1993.

C. Gourieroux, A. Monfort, and E. Renault. Indirect inference. *Journal of applied econometrics*, 8(S1):S85–S118, 1993.

M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 1.21. `http://cvxr.com/cvx`, April 2011.

R. C. Grinold and R. N. Kahn. *Active Portfolio Management*. McGraw-Hill, 2nd edition, 1999.

S. Grossman and M. Miller. Liquidity and market structure. *Journal of Finance*, 43:617–633, 1988.

S.J. Grossman and Z. Zhou. Optimal investment strategies for controlling drawdowns. *Mathematical Finance*, 3(3):241–276, 1993.

L. Harris. Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions & Instruments*, 7(2):1–76, April 1998.

J. Hasbrouck and G. Saar. Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12:143–172, 2009.

J. Hasbrouck and G. Saar. Low-latency trading. Working paper, 2010.

J. Hasbrouck. *Empirical Market Microstructure*. Oxford University Press, 2007.

T. Hendershott and R. Riordan. Algorithmic trading and information. Working paper, September 2009.

T. Hendershott, C. M. Jones, and A. Menkveld. Does algorithmic trading improve liquidity? To appear in *Journal of Finance*, April 2010.

S.L. Heston, R.A. Korajczyk, and R. Sadka. Intraday patterns in the cross-section of stock returns. *Journal of Finance*, 65(4):1369–1407, 2010.

M. Hora. Tactical liquidity trading and intraday volume. *Working Paper*, 2006.

R. A. Jarrow and P. Protter. A dysfunctional role of high frequency trading in electronic markets. Working paper, 2011.

A. Kirilenko, A. Kyle, M. Samadi, and T. Tuzun. The flash crash: The impact of high frequency trading on an electronic market. Working paper, 2010.

D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, pages 1–33, 2009.

Ahn Le, Kenneth J. Singleton, and Qiang Dai. Discrete-time affine$^{\mathbb{Q}}$ term structure models with generalized market prices of risk. *Review of Financial Studies*, 23(5):2184–2227, 2010.

H. Leland. Option pricing and replication with transactions costs. *Journal of Finance*, 40:1283–1301, 1985.

Han Liu and Mark Loewenstein. Optimal portfolio selection with transaction costs and finite horizons. *Review of Financial Studies*, 15(3):805–835, 2002.

H. Liu. Optimal consumption and investment with transaction costs and multiple risky assets. *Journal of Finance*, 59(1):289–338, 02 2004.

A. W. Lo, A. C. MacKinlay, and J. Zhang. Econometric models of limit order execution. *Journal of Financial Economics*, 65:31–71, 2002.

A. W. Lynch and P. Balduzzi. Predictability and transaction costs: The impact on rebalancing rules and behavior. *Journal of Finance*, 55(5):2285–2309, October 2000.

A. W. Lynch 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(04):1015–1053, August 2010.

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

A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.

A. A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. Working paper, 2005.

E. Pagnotta. Information and liquidity trading at optimal frequencies. Technical report, Working paper, NYU Stern, 2010.

C.A. Parlour and D.J. Seppi. Limit order markets: A survey. *Handbook of Financial Intermediation and Banking*, 5, 2008.

A. F. Perold. The implementation shortfall: Paper versus reality. *Journal of Portfolio Management*, 14(3):4–9, 1988.

M. J. Ready. The specialist's discretion: Stopped orders and price improvement. *Review of Financial Studies*, 12(5):1075–112, Winter 1999.

R. Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 39:1127–1139, 1984.

P. Sandås. Adverse selection and competitive market making: Empirical evidence from a limit order market. *Review of Financial Studies*, 14(3):705–734, 2001.

A. Shapiro and A. Nemirovski. On complexity of stochastic programming problems. *Continuous optimization*, pages 111–146, 2005.

A. Shapiro. Monte carlo sampling methods. In A. Ruszczynski and A. Shapiro, editors, *Stochastic Programming*, volume 10 of *Handbooks in Operations Research and Management Science*, pages 353 – 425. Elsevier, 2003.

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

J. Skaf and S.P. Boyd. Design of affine controllers via convex optimization. *Automatic Control, IEEE Transactions on*, 55(11):2476 –2487, November 2010.

H. R. Stoll and C. Schenzler. Trades outside the quotes: Reporting delay, trading option or trade size? *Journal of Financial Economics*, 79:615–653, 2006.

Oldrich A. Vasicek. An equilibrium characterization of the term structure. *Journal of Financial Economics*, 5(2):177–188, 1977.

I. Werner. NYSE order flow, spreads, and information. *Journal of Financial Markets*, 6:309–335, 2003.

Appendix A

The Cost of Latency

A.1. Dynamic Programming Decomposition

In order to solve the optimal control problem (2.8) via dynamic programming, note that we can equivalently consider the objective of maximizing the sale price P. Consider a decision time T_i with $0 \leq i < n$, and assume that the trader's limit order remains unfilled at time T_i. The state of the system consists of the current price, S_{T_i} as well as the previously chosen limit price,[1] ℓ_{i-1}, since this price will become active at time T_i. We can define an optimal value function $J_i(S_{T_i}, \ell_{i-1})$, as a function of this state, by optimizing the eventual sale price over all future decisions. In other words,

$$(A.1) \qquad J_i(S_{T_i}, \ell_{i-1}) \triangleq \underset{\ell_i, \dots, \ell_{n-1}}{\text{maximize}} \; \mathsf{E}\left[P \mid S_{T_i}, \ell_{i-1}\right].$$

At time $T = T_n$, the trader must sell via a market order, hence

$$(A.2) \qquad J_n(S_{T_n}, \ell_{n-1}) = S_{T_n}.$$

Now, for $0 \leq i < n$, there are three mutually exclusive events one of which must occur between time T_i and time T_{i+1}. These are the events $\mathcal{E}_i^{(1)}$, $\mathcal{E}_i^{(2)}$, and $\mathcal{E}_i^{(3)}$ described in Section 4.2. By considering cases corresponding to these events, we have the Bellman equation

$$(A.3) \qquad J_i(S_{T_i}, \ell_{i-1}) \triangleq \max_{\ell_i} \; \mathsf{E}\left[\mathbb{I}_{\mathcal{E}_i^{(1)}} \ell_{i-1} + \mathbb{I}_{\mathcal{E}_i^{(2)}} S_{T_{i+1}} + \mathbb{I}_{\mathcal{E}_i^{(3)}} J_{i+1}(S_{T_{i+1}}, \ell_i) \;\middle|\; S_{T_i}, \ell_{i-1}\right].$$

---

[1] I will assume that $\ell_{-1} = \infty$, i.e., there is no limit order active at the beginning of the time horizon.

Here, the first term corresponds to an execution at the prior price ℓ_{i-1}, the second term corresponds to the price ℓ_i being crossed by the bid price upon arrival to the market, and the third term corresponds to all other cases.

Define the function Q_i, for $0 \leq i \leq n$, by

$$Q_i(S_{T_i}, v_{i-1}) \triangleq J_i(S_{T_i}, S_{T_i} + v_{i-1}) - S_{T_i}.$$

The function Q_i is the premium of the value at time T_i, *relative* to the current bid price S_{T_i}. Similarly, $v_{i-1} \triangleq \ell_{i-1} - S_{T_i}$ is the premium of limit price decided at time T_{i-1} *relative* to the current bid price at time T_i. Then, applying (A.3), we have for $0 \leq i < n$,

$$\begin{aligned}
Q_i&(S_{T_i}, v_{i-1}) \\
&= \max_{\ell_i} \; \mathsf{E}\left[\mathbb{I}_{\mathcal{E}_i^{(1)}}(S_{T_i} + v_{i-1}) + \mathbb{I}_{\mathcal{E}_i^{(2)}} S_{T_{i+1}} + \mathbb{I}_{\mathcal{E}_i^{(3)}} J_{i+1}(S_{T_{i+1}}, \ell_i) \;\Big|\; S_{T_i}, v_{i-1} \right] - S_{T_i} \\
&= \max_{u_i} \; \mathsf{E}\left[\mathbb{I}_{\mathcal{E}_i^{(1)}} v_{i-1} + \mathbb{I}_{\mathcal{E}_i^{(2)} \cup \mathcal{E}_i^{(3)}} X_{i+1} + \mathbb{I}_{\mathcal{E}_i^{(3)}} Q_{i+1}(S_{T_i} + X_{i+1}, u_i - X_{i+1}) \;\Big|\; S_{T_i}, v_{i-1} \right].
\end{aligned}$$

Here, $X_{i+1} \triangleq S_{T_{i+1}} - S_{T_i} \sim N(0, \sigma^2 \Delta t)$ is the change in bid price from time T_i to time T_{i+1}. I define $u_i \triangleq \ell_i - S_{T_i}$ as the premium of the limit price at time T_i (i.e., the decision variable) *relative* to the current bid price S_{T_i}. Note that the price change X_{i+1} is zero mean under the event

$$\mathcal{E}_i^{(2)} \cup \mathcal{E}_i^{(3)} = \left(\mathcal{E}_i^{(1)} \right)^c,$$

by the assumption that the arrival of impatient buyers is independent of the bid price process, hence

$$(A.4) \quad Q_i(S_{T_i}, v_{i-1}) = \max_{u_i} \; \mathsf{E}\left[\mathbb{I}_{\mathcal{E}_i^{(1)}} v_{i-1} + \mathbb{I}_{\mathcal{E}_i^{(3)}} Q_{i+1}(S_{T_i} + X_{i+1}, u_i - X_{i+1}) \;\Big|\; S_{T_i}, v_{i-1} \right].$$

Finally, by (A.2),

$$(A.5) \qquad\qquad\qquad\qquad Q_n(S_{T_n}, v_{n-1}) = 0.$$

As should be clear from the above discussion, the Bellman equation (A.3) with terminal condition (A.2) and the backward recursion (A.4) with terminal condition (A.5) are completely equivalent, up to a change in variables. Expressing these equations in the latter form, however, brings significant simplifications, as the following lemma shows.

Lemma 4. *Suppose a collection of functions $\{Q_i\}$ satisfies the dynamic programming equations* (A.4)–(A.5). *Then, for each $0 \leq i < n$, Q_i does not depend on the price S_{T_i}, and takes the form*

$$(A.6) \qquad Q_i(v_{i-1}) = \mathbb{I}_{\{v_{i-1} \leq \delta\}} [\mu \Delta t v_{i-1} + (1 - \mu \Delta t) h_i] + \mathbb{I}_{\{v_{i-1} > \delta\}} h_i,$$

where the scalar h_i satisfies

$$(A.7) \qquad h_i = \max_{u_i} \; \mathsf{P}(X_{i+1} < u_i) \mathsf{E} \left[Q_{i+1}(u_i - X_{i+1}) \mid X_{i+1} < u_i \right].$$

Proof. Observe that, for $0 \leq i < n$, (A.4) can be simplified according to

(A.8)
$$Q_i(S_{T_i}, v_{i-1})$$
$$= \max_{u_i} \; \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}}$$
$$+ \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} \right) \mathsf{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(S_{T_i} + X_{i+1}, u_i - X_{i+1}) \;\middle|\; S_{T_i} \right],$$

where I have used the definitions of the events $\mathcal{E}_i^{(1)}$ and and $\mathcal{E}_i^{(3)}$.

Now, we proceed by backward induction. For the terminal case $i = n - 1$, from (A.8) and the fact that $Q_n = 0$ and $u_{n-1} = -\infty$ (i.e., the trader must use a market order at the last time slot), we have that

$$Q_{n-1}(S_{T_{n-1}}, v_{n-2}) = \mu \Delta t v_{n-2} \mathbb{I}_{\{v_{n-2} \leq \delta\}}.$$

In other words, Q_{n-1} satisfies the hypotheses of the lemma, with $h_{n-1} = 0$.

Now, suppose that the result holds for some $0 \leq i + 1 < n$. By (A.8), and since Q_{i+1} does not depend on $S_{T_{i+1}}$,

$$Q_i(S_{T_i}, v_{i-1})$$
$$= \max_{u_i} \; \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} \right) \mathsf{E} \left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(u_i - X_{i+1}) \right]$$
$$= \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \leq \delta\}} \right) h_i$$
$$= \mathbb{I}_{\{v_{i-1} \leq \delta\}} [\mu \Delta t v_{i-1} + (1 - \mu \Delta t v_{i-1}) h_i] + \mathbb{I}_{\{v_{i-1} > \delta\}} h_i.$$

Here, in the second equality, I define h_i through (A.7). The result then follows. ∎

Notice that, at the beginning of the trading horizon, there is no active limit order, i.e., $u_{-1} = \infty$. From Lemma 4, I have that

$$h_0 = Q_0(\infty) = \underset{\ell_0, \ldots, \ell_{n-1}}{\text{maximize}} \; \mathsf{E}\left[P \mid S_0\right] - S_0.$$

In other words, $h_0 = h_0(\Delta t)$, as defined in (2.8), and the notation is consistent. More generally, for $i > 0$, from (A.7), I can interpret h_i to be the trader's expected payoff at time T_i relative to the current bid price under the optimal policy, assuming that the limit order does not get executed in that time slot. Thus, h_i can be interpreted as a *continuation value* in the dynamic programming context, as in (2.9).

The continuation values $\{h_i\}$ allow for a compact representation of the value function, since they consist of only a single real number for each time slot, rather than a function of the entire state space. Theorem 1 directly expresses the dynamic programming equations (A.4)–(A.5) in terms of this representation. The proof follows by explicitly computing the expectations in Lemma 4.

Theorem 1. *Suppose $\{h_i\}$ satisfy, for $0 \le i < n - 1$,*

(A.9)
$$h_i = \max_{u_i} \left\{ \mu \Delta t \left[u_i \left(\Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)\right) + \sigma\sqrt{\Delta t}\left(\phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)\right)\right] \right.$$
$$\left. + h_{i+1}\left[(1 - \mu\Delta t)\Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) + \mu\Delta t \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)\right]\right\},$$

and

(A.10)
$$h_{n-1} = 0.$$

Here, ϕ and Φ are, respectively, the p.d.f. and c.d.f. of the standard normal distribution. Then, $\{h_i\}$ correspond to the continuation values under the optimal policy. In other words, the value functions $\{Q_i\}$ defined by $\{h_i\}$ via (A.6) solve the dynamic programming equations (A.4)–(A.5).

Suppose further that, for $0 \le i < n - 1$, u_i^* is a maximizer of (A.9). Then, a policy which chooses limit prices according to the premia defined by $\{u_i^*\}$, i.e.,

$$\ell_i^* = S_{T_i} + u_i^*, \quad \forall \, 0 \le i < n - 1,$$

is optimal.

Proof. Suppose that I am given $\{h_i\}$ that satisfy the hypotheses of the theorem. Define $\{Q_i\}$ by setting, for $0 \le i \le n-1$,

$$(A.11) \qquad Q_i(v_{i-1}) \triangleq \mathbb{I}_{\{v_{i-1} \le \delta\}} [\mu \Delta t v_{i-1} + (1 - \mu \Delta t) h_i] + \mathbb{I}_{\{v_{i-1} > \delta\}} h_i,$$

and $Q_n \triangleq 0$. I wish to show that $\{Q_i\}$ solve the dynamic programming equations (A.4)–(A.5).

Note that (A.5) holds by definition. For $0 \le i < n$, we have that (A.4) is equivalent to (A.8). Define $\hat{Q}_i$ to be the right side of (A.8), i.e.,

$$\hat{Q}_i(v_{i-1}) \triangleq \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \le \delta\}} + \left(1 - \mu \Delta t v_{i-1} \mathbb{I}_{\{v_{i-1} \le \delta\}}\right) \max_{u_i} \mathsf{E}\left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(u_i - X_{i+1})\right].$$

Comparing with (A.11), in order that the dynamic programming equation (A.8) hold (i.e., that $\hat{Q}_i = Q_i$), we must have that

$$(A.12) \qquad h_i = \max_{u_i} \mathsf{E}\left[\mathbb{I}_{\{X_{i+1} < u_i\}} Q_{i+1}(u_i - X_{i+1})\right]$$

Using the definition of Q_{i+1} from (A.11), this is equivalent to

$$h_i$$
$$= \max_{u_i} \mathsf{E}\left[\mathbb{I}_{\{X_{i+1} < u_i\}} \left(\mathbb{I}_{\{u_i - X_{i+1} \le \delta\}} [\mu \Delta t(u_i - X_{i+1}) + (1 - \mu \Delta t) h_{i+1}] + \mathbb{I}_{\{u_i - X_{i+1} > \delta\}} h_{i+1}\right)\right]$$
$$= \max_{u_i} \mathsf{E}\left[\mathbb{I}_{\{0 < u_i - X_{i+1} \le \delta\}} \mu \Delta t(u_i - X_{i+1}) + \mathbb{I}_{\{X_{i+1} < u_i\}}(1 - \mu \Delta t) h_{i+1} + \mathbb{I}_{\{u_i - X_{i+1} > \delta\}} \mu \Delta t h_{i+1}\right].$$

For the first term in the expectation, we have

$$\mathsf{E}\left[\mathbb{I}_{\{0 < u_i - X_{i+1} \le \delta\}} \mu \Delta t(u_i - X_{i+1})\right]$$
$$= \mu \Delta t \int_{-\infty}^{u_i} (u_i - x) \mathbb{I}_{\{u_i - x \le \delta\}} \frac{1}{\sigma\sqrt{\Delta t}} \phi\left(\frac{x}{\sigma\sqrt{\Delta t}}\right) dx$$
$$= \mu \Delta t \int_{u_i-\delta}^{u_i} (u_i - x) \frac{1}{\sigma\sqrt{\Delta t}} \phi\left(\frac{x}{\sigma\sqrt{\Delta t}}\right) dx$$
$$= \mu \Delta t \left[u_i \left(\Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)\right) + \int_{u_i-\delta}^{u_i} \frac{-x}{\sigma\sqrt{\Delta t}} \phi\left(\frac{x}{\sigma\sqrt{\Delta t}}\right) dx\right]$$
$$= \mu \Delta t \left[u_i \left(\Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)\right) + \sigma\sqrt{\Delta t} \left(\phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right) - \phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)\right)\right].$$

For the second term in the expectation, we have

$$\mathsf{E}\left[\mathbb{I}_{\{X_{i+1} < u_i\}}(1 - \mu \Delta t) h_{i+1}\right] = (1 - \mu \Delta t) h_{i+1} \Phi\left(\frac{u_i}{\sigma\sqrt{\Delta t}}\right).$$

finally, for the last term in the expectation, we have

$$\mathsf{E}\left[\mathbb{I}_{\{u_i - X_{i+1} > \delta\}} \mu \Delta t h_{i+1}\right] = \mu \Delta t h_{i+1} \Phi\left(\frac{u_i - \delta}{\sigma\sqrt{\Delta t}}\right)$$

Combining all the terms, I obtain the desired recursion for h_i.

The balance of the theorem (i.e., the optimality of the $\{u_i^*\}$ policy) follows from standard dynamic programming arguments. ∎

A.2. Proof of Theorem 2

I begin with a preliminary lemma.

Lemma 5. *Suppose that $\{h_i : 0 \leq i < n\}$ solves the dynamic programming recursion* (2.10)–(2.11). *Then, for $0 \leq i < n$,*

(A.13) $$0 \leq h_i \leq \delta\left(1 - (1 - \mu\Delta t)^n\right) < \delta.$$

Proof. First, note that the result is trivially true for $i = n - 1$, since $h_{n-1} = 0$. Now, if $0 \leq i < n - 1$, we can always choose $u_i = -\infty$, i.e., a market order, and this results in a continuation value of 0. Thus, $h_i \geq 0$.

For the upper bound, consider the discrete model without latency described in Section 2.4.3. Any strategy for the latency model is also feasible for the discrete model, since the trader can simply delay the implementation of trading decisions by one period. Therefore, at time T_i (with $0 \leq i < n - 1$), a policy with latency cannot achieve more value than the optimal policy for the discrete model without latency. At time T_i, there are $n - i - 1$ trading decisions remaining. This corresponds to the initial time of a discrete model with a total time horizon of $(n - i - 1)\Delta t$. Then, with reference to Lemma 2, we have that

$$h_i \leq \delta\left(1 - (1 - \mu\Delta t)^{n-i-1}\right).$$

The result immediately follows. ∎

Theorem 2. *Fix $\alpha > 1$. If Δt is sufficiently small, then there exists a unique optimal solution $\{h_i\}$ to the dynamic programming equations* (2.10)–(2.11). *Moreover, the corresponding*

optimal policy $\{u_i^*\}$ is unique. For $0 \le i < n-1$, this strategy chooses limit prices in the range

$$\ell_i^* \in \left(S_i + \delta - \sigma\sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, \ S_i + \delta - \sigma\sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}}, \right),$$

where

$$L \triangleq \frac{\delta^2}{2\pi\sigma^2}, \qquad\qquad R(\Delta t) \triangleq \frac{\delta^2 (1 - \mu\Delta t)^{2n}}{2\pi\sigma^2}.$$

Proof. Assume that, for some $0 \le i < n-1$, a solution $\{h_j \ : \ i+1 \le j < n\}$ exists to (2.10)–(2.11). I will establish that, for Δt sufficiently small (and not dependent on i), a solution h_i also exists and satisfies the conditions of the theorem. The result will follow by backward induction. Note that the base case of the induction (i.e., the existence of h_{n-1}) is trivial.

To this end, define the auxiliary function f by

(A.14)
$$f(u, h) \triangleq \mu\Delta t \Big[u(\Phi(A_u) - \Phi(B_u)) + \sigma\sqrt{\Delta t}(\phi(A_u) - \phi(B_u)) \Big]$$
$$+ h \Big[(1 - \mu\Delta t)\Phi(A_u) + \mu\Delta t \Phi(B_u) \Big],$$

where

(A.15)
$$A_u \triangleq \frac{u}{\sigma\sqrt{\Delta t}}, \qquad\qquad B_u \triangleq \frac{u - \delta}{\sigma\sqrt{\Delta t}}.$$

Then, from Theorem 1, for $0 \le i < n-1$, the dynamic programming recursion is given by

(A.16)
$$h_i = \max_{u_i} f(u_i, h_{i+1}),$$

and I can establish the present theorem by proving that, for Δt sufficiently small, (A.16) has a unique maximizer $u_i^* \in (\hat{u}_L, \hat{u}_R)$, where

(A.17)
$$\hat{u}_L \triangleq \delta - \sigma\sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, \qquad\qquad \hat{u}_R \triangleq \delta - \sigma\sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}}.$$

Note that

(A.18)
$$R(0) \triangleq \lim_{\Delta t \to 0} R(\Delta t) = \lim_{\Delta t \to 0} L(1 - \mu\Delta t)^{2T/dt} = Le^{-2\mu T} < \alpha L.$$

Hence, there exists some $\overline{\Delta t} > 0$ so that if $0 < \Delta t < \overline{\Delta t}$, then

$$\delta/2 < \hat{u}_L < \hat{u}_R < \delta, \quad \text{and} \quad 0 < 1 - \mu\Delta t < 1.$$

For the balance of the theorem, I will assume that $0 < \Delta t < \overline{\Delta t}$, in addition to whatever other assumptions are made regarding the magnitude of Δt.

The first and second derivatives of $f(\cdot, h)$ are given by

$$f_u(u, h) = \mu \Delta t \left[\Phi(A_u) - \Phi(B_u) + A_u (\phi(A_u) - \phi(B_u)) - \frac{u\phi(A_u) + (\delta - u)\phi(B_u)}{\sigma\sqrt{\Delta t}} \right]$$
$$+ \frac{h}{\sigma\sqrt{\Delta t}} \left[(1 - \mu\Delta t)\phi(A_u) + \mu\Delta t\phi(B_u) \right]$$
$$= \frac{(1 - \mu\Delta t)h}{\sigma\sqrt{\Delta t}} \phi(A_u) + \mu\Delta t \left[\Phi(A_u) - \Phi(B_u) - \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(B_u) \right] + \frac{h\mu\sqrt{\Delta t}}{\sigma} \phi(B_u)$$
$$= \frac{(1 - \mu\Delta t)h}{\sigma\sqrt{\Delta t}} \phi(A_u) + \mu\Delta t [\Phi(A_u) - \Phi(B_u)] + \frac{\mu\sqrt{\Delta t}}{\sigma} \phi(B_u)(h - \delta),$$

(A.19)

$$f_{uu}(u, h) = \frac{-u(1 - \mu\Delta t)h}{\sigma^3 \Delta t\sqrt{\Delta t}} \phi(A_u) + \frac{\mu\sqrt{\Delta t}}{\sigma} [\phi(A_u) - \phi(B_u)] + \frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}} \phi(B_u)(h - \delta)$$
$$= \phi(A_u) \left[\frac{\mu\sqrt{\Delta t}}{\sigma} - \frac{u(1 - \mu\Delta t)h}{\sigma^3 \Delta t^{3/2}} \right] + \phi(B_u) \left[\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h - \delta) - \frac{\mu\sqrt{\Delta t}}{\sigma} \right].$$

First, I will show that, for Δt sufficiently small, $f(\cdot, h_{i+1})$ has a *local* maximum u_i^* in the interval $(\hat{u}_L, \hat{u}_R)$, and that this is the unique maximizer over the larger interval $(\delta/2, \delta)$. That is, $u \in (\delta/2, \delta)$ and $u \neq u_i^*$, then

(A.20) $$f(u, h_{i+1}) < f(u_i^*, h_{i+1}), \quad \text{for all } u \in (\delta/2, \delta), \ u \neq u_i^*.$$

This is implied by the following claims, which I will demonstrate hold for Δt sufficiently small:

(i) $f_u(\hat{u}_L, h_{i+1}) > 0$.

(ii) $f_u(\hat{u}_R, h_{i+1}) < 0$.

(iii) $f_{uu}(u, h_{i+1}) < 0$, for all $u \in (\delta/2, \delta)$.

Claim (i): Note that

$$f_u(\hat{u}_L, h_{i+1}) = \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}} \phi(A_{\hat{u}_L}) + \mu\Delta t [\Phi(A_{\hat{u}_L}) - \Phi(B_{\hat{u}_L})] + \frac{\mu\Delta t}{\delta\sqrt{\alpha}}(h_{i+1} - \delta)$$

(A.21)

$$\geq \mu\Delta t [\Phi(A_{\hat{u}_L}) - \Phi(B_{\hat{u}_L})] - \frac{\mu\Delta t}{\sqrt{\alpha}},$$

where I use the fact that $h_{i+1} \geq 0$ (cf. Lemma 5). In order to calculate a lower bound for $\Phi(A_{\hat{u}_L}) - \Phi(B_{\hat{u}_L})$, we need the following standard bound on the tail probabilities of the normal distribution [see, e.g., Durrett, 2004]. Define Q to be the tail probability of a standard normal distribution, i.e.,

$$Q(x) \triangleq 1 - \Phi(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-\frac{1}{2}u^2} \, du.$$

Then, for all $x > 0$,

(A.22)
$$\frac{x^2 - 1}{x^3 \sqrt{2\pi}} e^{-\frac{1}{2}x^2} \leq Q(x) \leq \frac{1}{x\sqrt{2\pi}} e^{-\frac{1}{2}x^2}.$$

Applying this to (A.21),

$$
\begin{aligned}
f_u(\hat{u}_L, h_{i+1}) &\geq \mu \Delta t \Big[1 - Q(A_{\hat{u}_L}) - Q(-B_{\hat{u}_L}) \Big] + \frac{\mu \Delta t}{\sqrt{\alpha}} \\
&= \mu \Delta t \left(1 - \frac{1}{\sqrt{\alpha}} \right) - \mu \Delta t [Q(A_{\hat{u}_L}) + Q(-B_{\hat{u}_L})] \\
&> \mu \Delta t \left(1 - \frac{1}{\sqrt{\alpha}} \right) - 2\mu \Delta t Q(-B_{\hat{u}_L}) \\
&\geq \mu \Delta t \left(1 - \frac{1}{\sqrt{\alpha}} \right) - \frac{2\mu \Delta t^{3/2}}{\sqrt{2\pi \alpha L \log \frac{\alpha L}{\Delta t}}} \\
&> 0,
\end{aligned}
$$

for sufficiently small Δt. Here, I have used the fact that $Q(-B_{\hat{u}_L}) > Q(A_{\hat{u}_L})$.

Claim (ii): Similarly, for the other endpoint of the interval, I have

$$
\begin{aligned}
f_u(\hat{u}_R, h_{i+1}) &= \frac{(1 - \mu \Delta t)h_{i+1}}{\sigma \sqrt{\Delta t}} \phi(A_{\hat{u}_R}) + \mu \Delta t [1 - Q(A_{\hat{u}_R}) - Q(-B_{\hat{u}_R})] \\
&\quad + \frac{\mu \Delta t}{\delta(1 - \mu \Delta t)^n}(h_{i+1} - \delta) \\
&\leq \frac{(1 - \mu \Delta t)h_{i+1}}{\sigma \sqrt{\Delta t}} \phi(A_{\hat{u}_R}) + \mu \Delta t [1 - Q(A_{\hat{u}_R}) - Q(-B_{\hat{u}_R})] \\
&\quad + \frac{\mu \Delta t}{\delta(1 - \mu \Delta t)^n}[\delta(1 - (1 - \mu \Delta t)^n) - \delta] \\
&= \frac{(1 - \mu \Delta t)h_{i+1}}{\sigma \sqrt{\Delta t}} \phi(A_{\hat{u}_R}) + \mu \Delta t [1 - Q(A_{\hat{u}_R}) - Q(-B_{\hat{u}_R})] - \mu \Delta t \\
&= \frac{(1 - \mu \Delta t)h_{i+1}}{\sigma \sqrt{\Delta t}} \phi(A_{\hat{u}_R}) - \mu \Delta t [Q(A_{\hat{u}_R}) + Q(-B_{\hat{u}_R})] \\
&\leq \frac{\delta}{\sigma \sqrt{\Delta t}} \phi(A_{\hat{u}_R}) - \mu \Delta t Q(-B_{\hat{u}_R}),
\end{aligned}
$$

where I have used the upper bound on h_{i+1} from Lemma 5. Using (A.18), for sufficiently small Δt, I have

$$\sqrt{\log \frac{R(\Delta t)}{\Delta t}} < \frac{\delta}{2\sigma\sqrt{\Delta t}},$$

and thus

$$A_{\hat{u}_R} \geq \frac{\delta}{2\sigma\sqrt{\Delta t}}.$$

On the other hand, using (A.22),

$$Q(-B_{\hat{u}_R}) \geq \left[\left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-1/2} - \left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-3/2} \right] \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}}.$$

Thus,

$$\begin{aligned}
f_u(\hat{u}_R, h_{i+1}) \leq &\frac{\delta}{\sigma\sqrt{2\pi\Delta t}} \exp\left(\frac{-\delta^2}{8\sigma^2 \Delta t} \right) \\
&- \mu \Delta t \left[\left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-1/2} - \left(\log \frac{R(\Delta t)}{\Delta t} \right)^{-3/2} \right] \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}} \\
&< 0,
\end{aligned}$$

for sufficiently small Δt.

Claim (iii): Note that, for $u \in (\delta/2, \delta)$,

$$\phi\left(\frac{\delta}{\sigma\sqrt{\Delta t}} \right) < \phi(A_u) < \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}} \right) < \phi(B_u) < \phi(0).$$

Then, from (A.19), and using the fact that $0 \leq h_{i+1} < \delta$ (cf. Lemma 5), I have for Δt sufficiently small,

$$\begin{aligned}
f_{uu}(u, h_{i+1}) &\leq \phi(A_u)\frac{\mu\sqrt{\Delta t}}{\sigma} + \phi(B_u)\left[\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h_{i+1} - \delta) - \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \\
&\leq \phi(B_u)\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h_{i+1} - \delta) < 0.
\end{aligned}$$

In order to complete the proof, it suffices to demonstrate that the *local* maximum $u_i^* \in (\hat{u}_L, \hat{u}_R)$ is the unique *global* maximum. Since u_i^* achieves a higher value than any other $u \in (\delta/2, \delta)$, I will analyze cases where $u \notin (\delta/2, \delta)$ as follows:

- $u \in [0, \delta/2]$.

Here,

$$\phi\left(\frac{\delta}{\sigma\sqrt{\Delta t}}\right) \le \phi(B_u) \le \phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) \le \phi(A_u) \le \phi(0).$$

Further, for Δt sufficiently small,

$$\Phi(A_u) - \Phi(B_u) \ge \Phi(0) - \Phi\left(\frac{-\delta}{2\sigma\sqrt{\Delta t}}\right) \ge \frac{1}{4}.$$

Then, for Δt sufficiently small,

$$
\begin{aligned}
f_u(u, h_{i+1}) &\ge \frac{(1 - \mu\Delta t)h_{i+1}}{\sigma\sqrt{\Delta t}}\phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) + \mu\Delta t\left[\Phi(0) - \Phi\left(\frac{-\delta}{2\sigma\sqrt{\Delta t}}\right)\right] \\
&\quad + \frac{\mu\sqrt{\Delta t}}{\sigma}\phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right)(h_{i+1} - \delta) \\
&\ge \frac{h_{i+1}}{\sigma\sqrt{\Delta t}}\phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) + \mu\Delta t\left[\Phi(0) - \Phi\left(\frac{-\delta}{2\sigma\sqrt{\Delta t}}\right)\right] \\
&\quad - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi\left(\frac{\delta}{2\sigma\sqrt{\Delta t}}\right) \\
&\ge \frac{\mu\Delta t}{4} - \frac{\mu\delta\sqrt{\Delta t}}{\sqrt{2\pi}\sigma}\exp\left(\frac{-\delta^2}{8\sigma^2\Delta t}\right) > 0.
\end{aligned}
$$

(A.23)

Here, I have used the fact that $h_{i+1} \ge 0$. Using (A.20) and the fact that $f(\cdot, h_{i+1})$ is continuous, this implies that

$$(A.24) \qquad \sup_{u \in [0, \delta/2]} f(u, h_{i+1}) \le f(\delta/2, h_{i+1}) < f(u_i^*, h_{i+1}).$$

- $u \in (-\infty, 0)$.

  In this case, since $h_{i+1} \ge 0$ and $B_u < A_u < 0$,

$$
\begin{aligned}
f_u(u, h_{i+1}) &\ge \mu\Delta t\left[\Phi(A_u) - \Phi(B_u)\right] - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi(B_u) \\
&= \mu\Delta t\int_{B_u}^{A_u}\phi(z)\,dz - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi(B_u) \\
&> \mu\Delta t(A_u - B_u)\phi(B_u) - \frac{\delta\mu\sqrt{\Delta t}}{\sigma}\phi(B_u) = 0.
\end{aligned}
$$

  In conjunction with (A.24), this implies that

$$(A.25) \qquad \sup_{u \in (-\infty, 0)} f(u, h_{i+1}) \le f(0, h_{i+1}) < f(u_i^*, h_{i+1}).$$

- $u \in [\delta, \infty)$.

  In this case, using the upper bound on h_{i+1} from Lemma 5,

  (A.26)
  $$f_u(u, h_{i+1}) \leq \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(A_u) + \mu\Delta t \left[\Phi(A_u) - \Phi(B_u)\right]$$
  $$- \frac{\mu\delta(1 - \mu\Delta t)^n \sqrt{\Delta t}}{\sigma} \phi(B_u).$$

  Consider two cases. First, assume that $u > \delta + \sqrt{\Delta t}$. Then, applying (A.22),

  $$f_u(u, h_{i+1}) \leq \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(A_u) + \mu\Delta t Q(B_u) - \frac{\mu\delta(1 - \mu\Delta t)^n \sqrt{\Delta t}}{\sigma} \phi(B_u)$$
  $$\leq \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(A_u) + \frac{\mu\sigma\Delta t^{3/2}}{u - \delta} \phi(B_u) - \frac{\mu\delta(1 - \mu\Delta t)^n \sqrt{\Delta t}}{\sigma} \phi(B_u)$$
  $$\leq \phi(B_u) \left[\frac{\delta}{\sigma\sqrt{\Delta t}} \exp\left(\frac{-\delta^2}{2\sigma^2\Delta t}\right) + \mu\sigma\Delta t - \frac{\mu\delta(1 - \mu\Delta t)^n \sqrt{\Delta t}}{\sigma} \right].$$

  Note that $(1 - \mu\Delta t)^n \to e^{-\mu T}$ as $\Delta t \to 0$. Then, for Δt sufficiently small,

  (A.27)
  $$\tfrac{1}{2} e^{-\mu T} < (1 - \mu\Delta t)^n.$$

  Hence, for Δt sufficiently small,

  $$f_u(u, h_{i+1}) \leq \phi(B_u) \left[\frac{\delta}{\sigma\sqrt{\Delta t}} \exp\left(\frac{-\delta^2}{2\sigma^2\Delta t}\right) + \mu\sigma\Delta t - \frac{\mu\delta e^{-\mu T}\sqrt{\Delta t}}{2\sigma} \right] < 0.$$

  On the other hand, suppose that $u \in [\delta, \delta + \sqrt{\Delta t}]$. Then, from (A.26), (A.27), and since $0 < B_u < A_u$, I have for Δt sufficiently small,

  $$f_u(u, h_{i+1}) \leq \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(A_u) + \frac{\mu\Delta t}{2} - \frac{\mu\delta(1 - \mu\Delta t)^n \sqrt{\Delta t}}{\sigma} \phi(B_u)$$
  $$\leq \frac{\delta}{\sigma\sqrt{\Delta t}} \phi(A_u) + \frac{\mu\Delta t}{2} - \frac{\mu\delta e^{-\mu T}\sqrt{\Delta t}}{2\sigma} \phi(B_u)$$
  $$\leq \frac{\delta}{\sigma\sqrt{2\pi\Delta t}} \exp\left(-\frac{\delta^2}{2\sigma^2\Delta t}\right) + \frac{\mu\Delta t}{2} - \frac{\mu\delta e^{-\mu T}\sqrt{\Delta t}}{2\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2}\right) < 0.$$

  The above discussion, combined with (A.20) and the fact that $f(\cdot, h_{i+1})$ is continuous, implies that

  (A.28)
  $$\sup_{u \in [\delta, \infty)} f(u, h_{i+1}) \leq f(\delta, h_{i+1}) < f(u_i^*, h_{i+1}).$$

  ∎

A.3.   Proof of Theorem 3

I will establish Theorem 3 via a sequence of lemmas. First, recall the function $f(u, h)$ defined in (A.14) and the quantities $\hat{u}_L$ and $\hat{u}_R$ defined in (A.17).

Lemma 6.   *(i) As $\Delta t \to 0$,*

$$\max_{\substack{u \in [\hat{u}_L, \hat{u}_R] \\ 0 \leq i < n-1}} |f_{uu}(u, h_{i+1})| = O\left(\sqrt{\Delta t \log \frac{1}{\Delta t}}\right).$$

(ii) For all $h \in \mathbb{R}$ and Δt sufficiently small,

$$0 \leq f_h(\hat{u}_R, h) \leq 1.$$

Proof. I begin with (i). Recall A_u and B_u from (A.15). Let u be in the interval $[\hat{u}_L, \hat{u}_R]$. Then, for $0 \leq i < n - 1$, from (A.19),

$$|f_{uu}(u, h_{i+1})| \leq \left| \phi(A_u) \left[\frac{\mu\sqrt{\Delta t}}{\sigma} - \frac{u(1 - \mu\Delta t)h_{i+1}}{\sigma^3 \Delta t^{3/2}} \right] + \phi(B_u) \left[\frac{\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}}(h_{i+1} - \delta) - \frac{\mu\sqrt{\Delta t}}{\sigma} \right] \right|$$

$$\leq \phi(A_u)\frac{\mu\sqrt{\Delta t}}{\sigma} + \phi(A_u)\frac{\delta^2}{\sigma^3\Delta t^{3/2}} + \phi\left(\frac{\delta - u}{\sigma\sqrt{\Delta t}}\right)\left[\frac{\delta\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma}\right].$$

Here, I have used the fact that $0 \leq u \leq \delta$ and $0 \leq h_{i+1} < \delta$ (cf. Lemma 5). Note that, for Δt sufficiently small, $\hat{u}_L \geq \delta/2$. Then,

$$\max_{u \in [\hat{u}_L, \hat{u}_R]} \phi(A_u) \leq \phi(A_{\delta/2}) = \frac{1}{\sqrt{2\pi}}\exp\left(-\frac{\delta^2}{8\sigma^2\Delta t}\right) \leq c_0\Delta t^2,$$

for an appropriately chosen constant c_0. Thus,

$$|f_{uu}(u, h_{i+1})| \leq \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \phi\left(\frac{\delta - u}{\sigma\sqrt{\Delta t}}\right)\left[\frac{\delta\mu(\delta - u)}{\sigma^3\sqrt{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma}\right]$$

$$\leq \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \phi\left(\frac{\delta - \hat{u}_R}{\sigma\sqrt{\Delta t}}\right)\left[\frac{\delta\mu(\delta - \hat{u}_L)}{\sigma^3\sqrt{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma}\right]$$

$$= \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}}\left[\frac{\delta\mu}{\sigma^2}\sqrt{\log\frac{\alpha L}{\Delta t}} + \frac{\mu\sqrt{\Delta t}}{\sigma}\right]$$

$$\leq \frac{c_0\mu}{\sigma}\Delta t^{5/4} + \frac{c_0\delta^2}{\sigma^3}\sqrt{\Delta t} + \frac{\mu\Delta t}{\sigma\sqrt{2\pi R(\Delta t)}} + \frac{\delta\mu}{\sigma^2\sqrt{2\pi R(\Delta t)}}\sqrt{\Delta t \log\frac{\alpha L}{\Delta t}}.$$

Since $R(\Delta t) \to Le^{-2\mu T}$ as $\Delta t \to 0$, the last term asymptotically dominates and (i) follows.

For (ii), note that $\Phi(A_{\hat{u}_R}), \Phi(B_{\hat{u}_R}) \in (0,1)$, so if $\Delta t < 1/\mu$, then for all h,

$$f_h(\hat{u}_R, h) = (1 - \mu\Delta t)\Phi(A_{\hat{u}_R}) + \mu\Delta t\Phi(B_{\hat{u}_R}) \in (0,1).$$

■

Lemma 7. *As $\Delta t \to 0$,*

$$\hat{u}_R - \hat{u}_L = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Proof. Note that

$$\hat{u}_R - \hat{u}_L = \sigma\sqrt{\Delta t}\left(\sqrt{\log\frac{\alpha L}{\Delta t}} - \sqrt{\log\frac{R(\Delta t)}{\Delta t}}\right) = \sigma\sqrt{\Delta t}\left[g(\alpha L) - g(R(\Delta t))\right],$$

where $g(x) \triangleq \sqrt{\log\frac{x}{\Delta t}}$. Then, by mean value theorem, for some $z \in [R(\Delta t), \alpha L]$,

$$\hat{u}_R - \hat{u}_L = \sigma\sqrt{\Delta t}g'(z)[\alpha L - R(\Delta t)] = \frac{\sigma}{2z}[\alpha L - R(\Delta t)]\sqrt{\frac{\Delta t}{\log\frac{z}{\Delta t}}} \leq \frac{\sigma\alpha L}{2R(\Delta t)}\sqrt{\frac{\Delta t}{\log\frac{R(\Delta t)}{\Delta t}}}.$$

The result follows since $R(\Delta t) \to R(0) \triangleq Le^{-2\mu T}$ as $\Delta t \to 0$. ■

Let $\{h_i : 0 \leq i < n-1\}$ be the optimal solution to the dynamic programming recursion (2.10)–(2.11), and let $\{u_i^* : 0 \leq i < n-1\}$ define the corresponding optimal policy. Define $\{\hat{h}_i : 0 \leq i \leq n-1\}$ by the recursion

$$\hat{h}_i \triangleq \begin{cases} f(\hat{u}_R, \hat{h}_{i+1}) & \text{if } 0 \leq i < n-1, \\ 0 & \text{if } i = n-1. \end{cases}$$

Note that $\hat{h}_i$ is the continuation value of the suboptimal policy that always chooses $u_i = \hat{u}_R$, for $0 \leq i < n-1$. I am interested in quantifying its difference to the optimal continuation value.

Lemma 8. *As $\Delta t \to 0$,*

$$0 \leq h_0 - \hat{h}_0 = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Proof. For $0 \leq i < n-1$, define $\Delta_i \triangleq h_i - \hat{h}_i$. Clearly, $\Delta_i \geq 0$.

Using the mean value theorem,

$$
\begin{aligned}
\Delta_i &= f(u_i^*, h_{i+1}) - f(\hat{u}_R, \hat{h}_{i+1}) \\
&= \left[f(u_i^*, h_{i+1}) - f(\hat{u}_R, h_{i+1}) \right] + \left[f(\hat{u}_R, h_{i+1}) - f(\hat{u}_R, \hat{h}_{i+1}) \right] \\
&= -\tfrac{1}{2} f_{uu}(\bar{u}, h_{i+1})(\hat{u}_R - u_i^*)^2 + f_h(\hat{u}_R, \bar{h})\Delta_{i+1}.
\end{aligned}
$$

where $\bar{u}$ is some point on the interval $(u_i^*, \hat{u}_R)$ and $\bar{h}$ is some point on the interval $(\hat{h}_{i+1}, h_{i+1})$. Here, I have used the fact that the optimal solution u_i^* satisfies the first order condition $f_u(u_i^*, h_{i+1}) = 0$.

Using Lemmas 6 and 7, for Δt sufficiently small, there exist constants c_1 and c_2 so that

$$
\max_{\substack{u \in [\hat{u}_L, \hat{u}_R] \\ 0 \leq i < n-1}} |f_{uu}(u, h_{i+1})| \leq c_1 \sqrt{\Delta t \log \frac{1}{\Delta t}}, \qquad \hat{u}_R - \hat{u}_L \leq c_2 \sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}.
$$

Also, from Lemma 6, note that $0 \leq f_h(\hat{u}_R, \bar{h}) \leq 1$. Then, I obtain that, for Δt sufficiently small,

$$
\Delta_i \leq \frac{c_1(\hat{u}_R - u_i^*)^2}{2} \sqrt{\Delta t \log \frac{1}{\Delta t}} + \Delta_{i+1} \leq \frac{c_1 c_2}{2} \frac{\Delta t^{3/2}}{\sqrt{\log \frac{1}{\Delta t}}} + \Delta_{i+1}.
$$

Then, since $\Delta_{n-1} = 0$, I have that

$$
\Delta_0 \leq \left(\frac{T}{\Delta t} \right) \frac{c_1 c_2}{2} \frac{\Delta t^{3/2}}{\sqrt{\log \frac{1}{\Delta t}}} = \frac{c_1 c_2 T}{2} \sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}.
$$

∎

Define the sequence $\{\hat{\beta}_i \ : \ 0 \leq i \leq n-1\}$ by the linear recursion

(A.29)
$$
\hat{\beta}_i \triangleq \begin{cases} \mu \Delta t(\hat{u}_R - \hat{\beta}_{i+1}) + \hat{\beta}_{i+1} & \text{if } 0 \leq i < n-1, \\ 0 & \text{if } i = n-1. \end{cases}
$$

Here, $\hat{\beta}_i$ is an approximation to the value $\hat{h}_i$. The next lemma bounds the approximation error.

Lemma 9. *As $\Delta t \to 0$,*
$$
|\hat{h}_0 - \hat{\beta}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}} \right).
$$

Proof. For $0 \leq i < n-1$, define $\epsilon_i \triangleq \hat{h}_i - \hat{\beta}_i$. Recall the following definition from the proof of Theorem 2,

$$A_{\hat{u}_R} \triangleq \frac{\hat{u}_R}{\sigma\sqrt{\Delta t}}, \qquad\qquad B_{\hat{u}_R} \triangleq \frac{\hat{u}_R - \delta}{\sigma\sqrt{\Delta t}} = -\sqrt{\log\frac{R(\Delta t)}{\Delta t}}.$$

Then, by the recursive definitions of $\hat{h}_i$ and $\hat{\beta}_i$, $0 \leq i < n-1$,

$$\epsilon_i = (1 - \mu\Delta t)\epsilon_{i+1} - \mu\Delta t\hat{u}_R[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\Delta t^{3/2}[\phi(A_{\hat{u}_R}) - \phi(B_{\hat{u}_R})]$$

$$- (1 - \mu\Delta t)\hat{h}_{i+1}\left[1 - \Phi(A_{\hat{u}_R}) - \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right].$$

Since $\hat{\mu}_R$ is not the optimal policy, I have $\hat{h}_{i+1} \leq h_{i+1} < \delta$ (cf. Lemma 5). Further, for Δt sufficiently small, $0 < \phi(A_{\hat{u}_R}) \leq \phi(B_{\hat{u}_R})$. This implies that

$$|\epsilon_i| \leq (1 - \mu\Delta t)|\epsilon_{i+1}| + \delta\mu\Delta t[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\Delta t^{3/2}\phi(B_{\hat{u}_R})$$

$$+ \delta\left[1 - \Phi(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right].$$

Note that, except for the first term, there is no dependence on i in the right side of this equality. Then, I can define

$$C(\Delta t) \triangleq \delta\mu\Delta t[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\Delta t^{3/2}\phi(B_{\hat{u}_R})$$

$$+ \delta\left[1 - \Phi(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right],$$

and I have that

$$|\epsilon_i| \leq (1 - \mu\Delta t)|\epsilon_{i+1}| + C(\Delta t).$$

Since $\epsilon_{n-1} = 0$, it is easy to verify by backward induction on i that

$$|\epsilon_i| \leq \frac{1 - (1 - \mu\Delta t)^{n-i-1}}{\mu\Delta t}C(\Delta t).$$

Therefore,

(A.30)
$$|\epsilon_0| \leq \frac{1 - (1 - \mu\Delta t)^{n-1}}{\mu\Delta t}C(\Delta t) \leq \frac{C(\Delta t)}{\mu\Delta t}$$

$$= \delta[1 - \Phi(A_{\hat{u}_R}) + \Phi(B_{\hat{u}_R})] + \mu\sigma\sqrt{\Delta t}\phi(B_{\hat{u}_R}) + \frac{\delta}{\mu\Delta t}\left[1 - \Phi(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}\Phi(B_{\hat{u}_R})\right]$$

$$= \delta[Q(A_{\hat{u}_R}) + Q(-B_{\hat{u}_R})] + \mu\sigma\sqrt{\Delta t}\phi(B_{\hat{u}_R}) + \frac{\delta}{\mu\Delta t}\left[Q(A_{\hat{u}_R}) + \frac{\mu\Delta t}{1 - \mu\Delta t}Q(-B_{\hat{u}_R})\right].$$

From (A.22), however,

$$Q(A_{\hat{u}_R}) \leq \frac{\sigma}{\hat{u}_R} \sqrt{\frac{\Delta t}{2\pi}} \exp\left(-\frac{\hat{u}_R^2}{2\sigma \Delta t}\right).$$

Since $\hat{u}_R \to \delta$ as $\Delta t \to 0$, for Δt sufficiently small, there exists constants a_1 and a_2, with $0 < a_2 < \delta^2/2\sigma$, so that

$$Q(A_{\hat{u}_R}) \leq a_1 \sqrt{\Delta t} \exp\left(-\frac{a_2}{\Delta t}\right).$$

Also by (A.22),

$$Q(-B_{\hat{u}_R}) \leq \sqrt{\frac{\Delta t}{2\pi R(\Delta t) \log \frac{R(\Delta t)}{\Delta t}}}.$$

Since and $R(\Delta t) \to R(0) \triangleq Le^{-2\mu T}$ as $\Delta t \to 0$, for Δt sufficiently small, there exists a constant a_3 so that

$$Q(-B_{\hat{u}_R}) \leq a_3 \sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}.$$

Finally,

$$\phi(B_{\hat{u}_R}) = \sqrt{\frac{\Delta t}{2\pi R(\Delta t)}},$$

so for Δt sufficiently small, there exists a constant a_4 with

$$\phi(B_{\hat{u}_R}) \leq a_4 \sqrt{\Delta t}.$$

Applying these bounds to (A.30), the result follows.   ∎

Lemma 10. *As $\Delta t \to 0$,*

$$\hat{\beta}_0 = \hat{u}_R \left(1 - e^{-\mu T}\right) + O(\Delta t).$$

Proof. Note that the recurrence (A.29) can be explicitly solved to obtain

$$\hat{\beta}_0 = \sum_{i=0}^{n-2} (1 - \mu \Delta t)^i \mu \Delta t\, \hat{u}_R = \hat{u}_R \left(1 - (1 - \mu \Delta t)^{n-1}\right) = \hat{u}_R \left(1 - (1 - \mu \Delta t)^{T/\Delta t - 1}\right).$$

The result follows since $(1 - \mu \Delta t)^{T/\Delta t} = e^{-\mu T} + O(\Delta t)$ as $\Delta t \to 0$.   ∎

I am now ready to prove Theorem 3.

Theorem 3. *As $\Delta t \to 0$,*

$$h_0(\Delta t) = \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi \sigma^2 \Delta t}}\right) + o\left(\sqrt{\Delta t}\right),$$

where

$$\bar{h}_0 = \delta \left(1 - e^{-\mu T}\right)$$

is the optimal value for the stylized model without latency, i.e., the value defined by (2.5).

Proof. First, define

$$\hat{\gamma}_0 \triangleq \left(1 - e^{-\mu T}\right) \left(\delta - \sigma \sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2 \Delta t}}\right).$$

Then,

(A.31)
$$\left| h_0 - \bar{h}_0 \left(1 - \frac{\sigma}{\delta}\sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}}\right) \right| = |h_0 - \hat{\gamma}_0|$$

$$\hat{\text{E}}\hat{\text{E}} \leq |h_0 - \hat{h}_0| + |\hat{h}_0 - \hat{\beta}_0| + |\hat{\beta}_0 - \hat{\gamma}_0|.$$

I will bound each of the terms in the right side of (A.31). First, by Lemma 8,

(A.32)
$$|h_0 - \hat{h}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Next, by Lemma 9,

(A.33)
$$|\hat{h}_0 - \hat{\beta}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

Finally, by Lemma 10, for Δt sufficiently small, there exists a constant c_1 so that

$$|\hat{\beta}_0 - \hat{\gamma}_0| \leq \sigma \left(1 - e^{-\mu T}\right) \left(\hat{u}_R - \delta + \sqrt{\Delta t \log \frac{L}{\Delta t}}\right) + c_1 \Delta t$$

$$\leq \sigma \left(1 - e^{-\mu T}\right) \left(\hat{u}_R - \delta + \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}\right) + c_1 \Delta t,$$

$$= \sigma \left(1 - e^{-\mu T}\right) (\hat{u}_R - \hat{u}_L) + c_1 \Delta t,$$

where $\alpha > 1$ and L are defined by Theorem 2. Applying Lemma 7, I have that

(A.34)
$$|\hat{\beta}_0 - \hat{\gamma}_0| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right).$$

By applying (A.32)–(A.34) to (A.31), I have that

$$\left| h_0 - \bar{h}_0 \left(1 - \frac{\sigma}{\delta}\sqrt{\Delta t \log \frac{\delta^2}{2\pi\sigma^2\Delta t}}\right) \right| = O\left(\sqrt{\frac{\Delta t}{\log \frac{1}{\Delta t}}}\right),$$

which implies the desired result. ∎

A.4. Price Dynamics with Jumps

At a high level, the goal has been to understand and build intuition as to the impact of a latency friction introduced by the lack of contemporaneous information. The spirit of the model it to consider an investor who wants to trade, but at a price that depends on an informational process that evolves stochastically and must be monitored continuously. While I have principally interpreted the informational process to be the bid price process, the model can alternatively be interpreted (as discussed in Section 2.2.1) in terms of a fundamental value process.

Thus far, I have employed a diffusive model to describe informational innovations over a short time horizon. There is significant empirical evidence that this is insufficient, particularly when modeling price processes, and that it is important to also allow for the instantaneous arrival of information, i.e., jumps. For example, Barndorff-Nielsen *et al.* [2010] propose the following compound Poisson process for high frequency price dynamics:

$$S_t = S_0 + \sum_{i=1}^{M_t} Y_i,$$

where N_t is a Poisson process counting the number of trades up to time t and Y_i is the potential jump movement at the ith trade time, having a distribution G.

On a short time horizon, innovations to fundamental value can be both instantaneous or diffusive.[2] In a recent empirical study, Aït-Sahalia and Jacod [2010] construct two formal statistical tests to deduce whether there is a need for a Brownian motion in modeling high-frequency data. Using individual high-frequency stock data, they conclude that both tests suggest the necessity of including a continuous component driven by Brownian motion.

Motivated by these studies, I will generalize the price dynamics of Section 2.2 by including both a continuous component (Brownian motion) and a jump component (governed by a compound Poisson process). In particular, consider a price process that evolves according

---

[2]As an example, note that an instantaneous innovation may result from a news event. On the other hand, the value of a stock will have a component that is driven by the market factor, i.e., an average of returns across all stocks. Innovations to the market factor can have a diffusive component even if all individual stock prices are discrete, by virtue of cross-sectional averaging.

to

(A.35)
$$S_t = S_0 + \sigma B_t + \sum_{i=1}^{M_t} Y_i,$$

where the process $(B_t)_{t \in [0,T]}$ is a standard Brownian motion, $\sigma > 0$ is an (additive) volatility parameter, and $(M_t)_{t \in [0,T]}$ is a Poisson process with intensity λ. For now, I will further assume that each jump Y_i has an i.i.d. Gaussian distribution with zero mean and variance ν^2 — I revisit the assumption of Gaussian jump sizes at the end of this section.

In the context of the latency model of Section 4.2, I define the price increment $X_{i+1} \triangleq S_{T_{i+1}} - S_{T_i}$ by the discrete time analog of (A.35),

(A.36)
$$X_{i+1} \sim \begin{cases} N(0, \sigma^2 \Delta t) & \text{with probability } (1 - \lambda \Delta t), \\ N(0, \sigma^2 \Delta t + \nu^2) & \text{with probability } \lambda \Delta t. \end{cases}$$

With this definition, the dynamic programming decomposition outlined in Lemma 4 holds exactly as before. Incorporating jumps, I then obtain the following analog of Theorem 1, that expresses dynamic programming equations (A.4)–(A.5) in terms of the continuation values $\{h_i\}$. The proof of this theorem follows steps identical to the proof of Theorem 1, and is omitted.

Theorem 4. *Define* $v^2(\Delta t) \triangleq \sigma^2 \Delta t + \nu^2$. *Suppose* $\{h_i\}$ *satisfy, for* $0 \le i < n - 1$,

(A.37)
$$\begin{aligned} h_i = \max_{u_i} \Bigg\{ &(1 - \lambda \Delta t) \left(\mu \Delta t \left[u_i \left(\Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) + \sigma \sqrt{\Delta t} \left(\phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) - \phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right) \right] \right. \\ &\left. + h_{i+1} \left[(1 - \mu \Delta t) \Phi \left(\frac{u_i}{\sigma \sqrt{\Delta t}} \right) + \mu \Delta t \Phi \left(\frac{u_i - \delta}{\sigma \sqrt{\Delta t}} \right) \right] \right) \\ &+ \lambda \Delta t \left(\mu \Delta t \left[u_i \left(\Phi \left(\frac{u_i}{v(\Delta t)} \right) - \Phi \left(\frac{u_i - \delta}{v(\Delta t)} \right) \right) + v^2(\Delta t) \left(\phi \left(\frac{u_i}{v(\Delta t)} \right) - \phi \left(\frac{u_i - \delta}{v(\Delta t)} \right) \right) \right] \right. \\ &\left. + h_{i+1} \left[(1 - \mu \Delta t) \Phi \left(\frac{u_i}{v(\Delta t)} \right) + \mu \Delta t \Phi \left(\frac{u_i - \delta}{v(\Delta t)} \right) \right] \right) \Bigg\} \end{aligned}$$

and

(A.38)
$$h_{n-1} = 0.$$

Suppose further that, for $0 \le i < n - 1$, u_i^* *is a maximizer of* (A.37). *Then, a policy which chooses limit prices according to the premia defined by* $\{u_i^*\}$, *i.e.,*

$$\ell_i^* = S_{T_i} + u_i^*, \quad \forall \, 0 \le i < n - 1,$$

is optimal.

The following theorem provides an analog of Theorem 2 that characterizes the optimal solution for the dynamic programming equation in the low latency regime, with the presence of jumps. The proof is similar to that of Theorem 2, and is again omitted.

Theorem 5. *Fix $\alpha > 1$ and define*

$$\kappa \triangleq 1 + \frac{\lambda \delta}{\nu \mu \sqrt{2\pi}}.$$

If Δt is sufficiently small, then there exists a unique optimal solution $\{h_i\}$ to the dynamic programming equations (A.37)–(A.38). Moreover, the corresponding optimal policy $\{u_i^*\}$ is unique. For $0 \leq i < n - 1$, this strategy chooses limit prices in the range

$$\ell_i^* \in \left(S_i + \delta - \sigma \sqrt{\Delta t \log \frac{\alpha L}{\Delta t}}, \ S_i + \delta - \sigma \sqrt{\Delta t \log \frac{R(\Delta t)}{\Delta t}} \right),$$

where

$$L \triangleq \frac{\delta^2}{2\pi \sigma^2}, \qquad\qquad R(\Delta t) \triangleq \frac{\delta^2 (1 - \mu \Delta t)^{2n}}{2\pi \sigma^2 \kappa^2}.$$

Note that, when compared to Theorem 2, the addition of jump component in Theorem 5 causes $R(\Delta t)$ to decrease by a constant multiple. Thus, the range containing the optimal solution is gets larger. However, the upper bound of the range is of the same order asymptotically (as $\Delta t \to 0$) as before. Hence, I can again provide a asymptotic closed-form expression for $h_0(\Delta t)$, as is done by the following theorem, which is an analog of Theorem 3 and Corollary 1. (As before, I omit the proof.)

Theorem 6. *As $\Delta t \to 0$,*

$$h_0(\Delta t) = \bar{h}_0 \left(1 - \frac{\sigma}{\delta} \sqrt{\Delta t \log \frac{\delta^2}{2\pi \sigma^2 \Delta t}} \right) + o\left(\sqrt{\Delta t} \right),$$

where

(A.39) $$\bar{h}_0 \triangleq \frac{\delta \mu}{\mu + \lambda p} \left(1 - e^{-(\mu + \lambda p)T} \right),$$

is the zero latency limit of $h_0(\Delta t)$, and

$$p \triangleq 1 - \Phi \left(\frac{\delta}{\nu} \right),$$

the probability of a jump size greater than δ.

Furthermore, latency cost is unchanged with the introduction of the jump components in the bid price dynamics, i.e., as $\Delta t \to 0$,

$$\mathsf{LC}(\Delta t) = \frac{\sigma\sqrt{\Delta t}}{\delta}\sqrt{\log\frac{\delta^2}{2\pi\sigma^2\Delta t}} + o\left(\sqrt{\Delta t}\right).$$

The analysis with the jump-diffusion model can be interpreted as follows. Theorem 6 illustrates that, when there is a jump component (i.e., $\lambda > 0$), the zero latency limit $\bar{h}_0$ has a lower value than in the absence of jumps, (i.e., $\lambda = 0$), all else being equal. In other words, the presence of jumps is detrimental even in the absence of latency. To see why, note that jumps are zero mean innovations in the price process. In the model, an investor only earns excess value by waiting for an impatient buyer. Jumps may cause the bid price to cross the investor's limit order price and execute his share without giving him the chance to revise his order. Thus, jumps reduce the probability of trading with an impatient buyer.

This intuition can be made precise by interpreting the zero latency limit in (A.39). Observe that $\mu + \lambda p$ is the combined arrival rate of impatient buyers asking for an immediate execution, or positive jumps in the price of the stock that are larger than the bid-offer spread and would result in trade execution. The quantity

$$\frac{\mu}{\mu + \lambda p}\left(1 - e^{-(\mu + \lambda p)T}\right)$$

is the probability that there at least one such arrival, and that the first such arrival is that of an impatient buyer. In this case, the trader earns a relative spread of δ. In all other cases (i.e., no arrivals, or the case where the first arrival is a large positive jump), the trade occurs at the bid price and the trader earns no spread. These two cases yield the expression for $\bar{h}_0$.

Now, comparing with the earlier results, jumps also negatively impact the investor in the presence of latency, for similar reasons as in the zero latency case. However, when measured *relative* to the zero latency case, i.e., in term of latency cost, jumps create no additional impact. That is, the latency cost expressions in Theorem 6 and Corollary 1 are identical. Intuitively, in the model, jumps are instantaneous, and the investor cannot react to them even in the absence of latency. Hence, latency cost, measured relatively, only depends on the diffusive innovations.

Note that I have thus far assumed Gaussian jump sizes. In Theorem 6, the only place that this distribution or its parameter ν arises explicitly is the quantity p. This is the probability that the jump will be larger than the prevailing bid-offer spread, δ, and hence will cross with the limit order places by the investor. This leads us to conjecture (without proof) the result in the non-Gaussian case: if the jump size Y_i in (A.35) is an i.i.d. zero mean random variable that has a cumulative distribution function G, then Theorem 6 holds with $p \triangleq 1 - G(\delta)$.

Appendix B

Dynamic Portfolio Choice with Linear Rebalancing Rules

B.1. Proof of Lemma 3

Lemma 3. *Given $\eta \in [0, 1/2]$, a non-zero vector $a \in \mathbb{R}^N$, and a scalar b, the chance constraint $\mathsf{P}(a^\top u_t > b) \leq \eta$ is equivalent to the second order cone constraint*

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

on the policy coefficients (c_t, M_t), where $\Phi^{-1}(\cdot)$ is the inverse cumulative normal distribution.

Proof. This proof follows standard arguments in convex optimization [see, e.g., Boyd and Vandenberghe, 2004]. Let $\bar{u}_t$ and V_t be the mean and the variance of u_t as given in (3.13). Then,

$$\mathsf{P}(a^\top u_t > b) = \mathsf{P}(\beta_t + \sigma_t Z > 0),$$

where

$$\beta_t \triangleq a^\top \bar{u}_t - b, \qquad \sigma_t \triangleq \|V_t^{1/2} a\|_2 \neq 0,$$

and Z is a standard normal random variable. Thus,

$$\mathsf{P}(a^\top u_t > b) = 1 - \Phi(-\beta_t / \sigma_t).$$

Note that the this probability is less than or equal to η if and only if

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

Substituting (3.13) into definitions for β_t and σ_t, we obtain the desired result. ∎

B.2. Exact Formulation of the Terminal Wealth Objective

Following the notation of Section 3.4, I will compute $\mathsf{E}[W(\mathbf{x},\mathbf{r})^2]$ analytically and demonstrate that the resulting expression is a quadratic convex function of the policy coefficients. Without loss of generality, assume that $W_0 = 0$ and $\mu_t = \mathbf{0}$. Then, observe that

(B.1)

$$
\begin{aligned}
\mathsf{E}\left[W(\mathbf{x},\mathbf{r})^2\right] &= \sum_{t=1}^{T}\sum_{k=1}^{T}\mathsf{E}\left[(x_t^\top r_{t+1})(x_k^\top r_{k+1})\right] \\
&= \sum_{t=1}^{T}\mathsf{E}\left[x_t^\top \epsilon_{t+1}^{(2)}\left(\epsilon_{t+1}^{(2)}\right)^\top x_t\right] + \sum_{t=1}^{T}\sum_{k=1}^{T}\mathsf{E}\left[x_t^\top Bf_t f_k^\top B^\top x_k\right] \\
&= \sum_{t=1}^{T}\mathsf{E}\left[x_t^\top \epsilon_{t+1}^{(2)}\left(\epsilon_{t+1}^{(2)}\right)^\top x_t\right] + \sum_{t=1}^{T}\mathsf{E}\left[x_t^\top Bf_t f_t^\top B^\top x_t\right] + 2\sum_{t=1}^{T}\sum_{k=1}^{t}\mathsf{E}\left[x_t^\top Bf_t f_k^\top B^\top x_k\right].
\end{aligned}
$$

I will consider each of the three terms in (B.1) separately.

The first term can be evaluated he first expectation, $\mathsf{E}\left[x_t^\top \epsilon_{t+1}^{(2)}\left(\epsilon_{t+1}^{(2)}\right)^\top x_t\right]$, can be evaluated in the following form:

$$
\begin{aligned}
\mathsf{E}\left[x_t^\top \epsilon_{t+1}^{(2)}\left(\epsilon_{t+1}^{(2)}\right)^\top x_t\right] &= \mathsf{E}\left[x_t^\top \mathsf{E}\left[\epsilon_{t+1}^{(2)}\left(\epsilon_{t+1}^{(2)}\right)^\top |\mathcal{F}_t\right] x_t\right] \\
&= \mathsf{E}\left[x_t^\top \Sigma x_t\right] \\
&= (d_t + P_t\theta_t)^\top \Sigma (d_t + P_t\theta_t) + \mathrm{tr}\left(\Sigma P_t \Omega_t P_t^\top\right).
\end{aligned}
$$

Using the representation for x_t, I obtain

$$
\mathsf{E}\left[x_t^\top Bf_t f_t^\top B^\top x_t\right] = \mathsf{E}\left[\left(d_t + \sum_{s=1}^{t} J_{s,t}f_s\right)^\top Bf_t f_t^\top B^\top \left(d_t + \sum_{s=1}^{t} J_{s,t}f_s\right)\right],
$$

which can be simplified further to

$$
\begin{aligned}
\mathsf{E}\Bigg\{ &\left(d_t + \sum_{s=1}^{t-1} J_{s,t}f_s + J_{t,t}\left((I-\Phi)f_{t-1} + \epsilon_t^{(1)}\right)^\top\right) B\left((I-\Phi)f_{t-1} + \epsilon_t^{(1)}\right)\left((I-\Phi)f_{t-1} + \epsilon_t^{(1)}\right)^\top B^\top \\
&\left(d_t + \sum_{s=1}^{t-1} J_{s,t}f_s + J_{t,t}\left((I-\Phi)f_{t-1} + \epsilon_t^{(1)}\right)\right)\Bigg\}.
\end{aligned}
$$

The resulting computation deals with taking expectations of the product of jointly normal random variables with integer exponents:

$$\mathsf{E}\left[\left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t}(I - \Phi)f_{t-1}\right)^\top B(I - \Phi)f_{t-1}\left(B(I-\Phi)f_{t-1}\right)^\top \left(d_t + \sum_{s=1}^{t-1} J_{s,t}f_s + J_{t,t}(I-\Phi)f_{t-1}\right)\right]$$

$$+ \mathsf{E}\left[\left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t}(I - \Phi)f_{t-1}\right)^\top \epsilon_t^{(1)}\left(\epsilon_t^{(1)}\right)^\top \left(d_t + \sum_{s=1}^{t-1} J_{s,t}f_s + J_{t,t}(I-\Phi)f_{t-1}\right)\right]$$

$$+ \mathsf{E}\left[\left(\epsilon_t^{(1)}\right)^\top B(I-\Phi)f_{t-1}\left(B(I-\Phi)f_{t-1}\right)^\top \epsilon_t^{(1)}\right] + \mathsf{E}\left[\left(\epsilon_t^{(1)}\right)^\top \epsilon_t^{(1)}\left(\epsilon_t^{(1)}\right)^\top \epsilon_t^{(1)}\right]$$

First define

$$\psi^2 \triangleq \mathsf{E}\left[\left(\epsilon_t^{(1)}\right)^\top \epsilon_t^{(1)}\left(\epsilon_t^{(1)}\right)^\top \epsilon_t^{(1)}\right] = \sum_{i=1}^K 3\Psi_{ii} + 2\sum_{i=1}^K \sum_{j=1}^i \Psi_{ij},$$

where Ψ_{ij} is the (i,j)th entry of the covariance matrix of the error terms for factor dynamics, Ψ. Using iterated expectations by conditioning on the information up to $t-1$, $\mathcal{F}_{t-1}$, $\mathsf{E}\left[x_t^\top B f_t f_t^\top B^\top x_t\right]$ equals

$$\mathsf{E}\left[\left(d_t + \sum_{s=1}^{t-1} J_{s,t} f_s + J_{t,t}(I - \Phi)f_{t-1}\right)^\top B(I - \Phi)f_{t-1}\left(B(I-\Phi)f_{t-1}\right)^\top \left(d_t + \sum_{s=1}^{t-1} J_{s,t}f_s + J_{t,t}(I-\Phi)f_{t-1}\right)\right]$$

$$+ \left(d_t + \bar{P}_{t-1}\theta_{t-1}\right)^\top \Psi \left(d_t + \bar{P}_{t-1}\theta_{t-1}\right) + \mathrm{tr}\left(\bar{P}_{t-1}\Omega_{t-1}\bar{P}_{t-1}^\top\right)\mathrm{tr}\left(B(I-\Phi)\Omega_t^{t-1}\left(B(I-\Phi)\right)^\top \Psi\right) + \psi^2.$$

where

$$(\text{B.2}) \qquad \bar{P}_{t-i} \triangleq \begin{bmatrix} J_{1,t} & J_{2,t} & \ldots & J_{t-i-1,t} & \bar{J}_{t-i,t} \end{bmatrix}, \qquad \bar{J}_{t-i,t} \triangleq \left(\sum_{k=0}^i J_{t-k,t}(I-\Phi)^{i-k}\right).$$

I continue conditioning in a recursive fashion, and finally obtain the deterministic form

$$\mathsf{E}\left[x_t^\top B f_t f_t^\top B^\top x_t\right] = \left(d_t + \sum_{s=1}^t J_{s,t}(I-\Phi)^s f_0\right)^\top B(I-\Phi)^t f_0\left(B(I-\Phi)^t f_0\right)^\top \left(d_t + \sum_{s=1}^t J_{s,t}(I-\Phi)^s f_0\right)$$

$$+ \sum_{i=1}^t \left(\left(d_t + \bar{P}_{t-i}\theta_{t-i}\right)^\top \Psi \left(d_t + \bar{P}_{t-i}\theta_{t-i}\right) + \mathrm{tr}\left(\bar{P}_{t-i}\Omega_{t-i}\bar{P}_{t-i}^\top\right)\mathrm{tr}\left(B(I-\Phi)^i\Omega_t^{t-i}\left(B(I-\Phi)^i\right)^\top \Psi\right) + \psi^2\right),$$

which is convex quadratic function of the linear policy parameters.

The third expectation, $\mathsf{E}\left[x_t^\top B f_t f_k^\top B^\top x_k\right]$, can be computed using the same procedure.

B.3. Derivation of the LQC Policies

I can derive a closed form solution for the trading policy when the problem satisfies the LQC framework. I guess a functional form for the value function and show that this functional form is preserved at each time step.

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

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

I guess the following quadratic form for the value function:

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

Then,

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

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

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

which satisfies the functional form with

$$A_{xx,T-1} = \Lambda \qquad A_{xf,T-1} = \mathbf{0} \qquad A_{ff,T-1} = \mathbf{0} \qquad m_{T-1} = 0.$$

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

$$Q_t = \Lambda + A_{xx,t}$$

$$q_t = \Lambda x_{t-1} + (B + A_{xf,t} (I - \Phi)) f_t$$

$$b_t = -\frac{1}{2}x_{t-1}^\top \Lambda x_{t-1} + \frac{1}{2}f_t^\top (I - \Phi)^\top A_{ff,t} (I - \Phi) f_t + \text{tr}(\Psi A_{ff,t}) + m_t$$

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

$$x_t = (\Lambda + A_{xx,t})^{-1} (\Lambda x_{t-1} + (B + A_{xf,t} (I - \Phi)) f_t)$$

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

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

$$A_{xx,t-1} = -\Lambda \left(\Lambda + A_{xx,t}\right)^{-1} \Lambda + \Lambda$$

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

$$A_{ff,t-1} = \left(B + A_{xf,t} \left(I - \Phi\right)\right)^\top \left(\Lambda + A_{xx,t}\right)^{-1} \left(B + A_{xf,t} \left(I - \Phi\right)\right) + \left(I - \Phi\right)^\top A_{ff,t} \left(I - \Phi\right)$$

$$m_{t-1} = \text{tr}(\Psi A_{ff,t}) + m_t$$

Using these recursions, I can compute the optimal expected payoff of the dynamic program. Using $f_0 = N(0, \Omega_0)$,

$$\mathsf{E}[V_0(x_0, f_1)] = \mathsf{E}\left[\mathsf{E}[V_0(x_0, f_1) | f_0]\right]$$

$$= \mathsf{E}\left[-\frac{1}{2} x_0^\top A_{xx,0} x_0 + x_0^\top A_{xf,0} \left(I - \Phi\right) f_0 + \frac{1}{2} f_0^\top \left(I - \Phi\right)^\top A_{ff,0} \left(I - \Phi\right) f_0 + \frac{1}{2} \left(\text{tr}(\Omega_0 A_{ff,0}) + m_0\right)\right]$$

$$= -\frac{1}{2} x_0^\top A_{xx,0} x_0 + \frac{1}{2} \left(\text{tr}\left(\Omega_0 (I - \Phi)^\top A_{ff,0} (I - \Phi)\right) + \sum_{t=0}^{T-2} \text{tr}(\Psi A_{ff,0})\right).$$

B.4. Exact Formulation of Best Linear Execution Policy

I will first compute the expectation in the objective of (3.29) and write the equivalent deterministic form. I will then replace probabilistic constraints with deterministic constraints using Lemma 3, and finally obtain the deterministic version of the stochastic program in (3.29)

I start working with the expectation in the objective function. For each t, I have to compute the expectation of the following two terms, $\mathsf{E}\left[x_t^\top (B f_t)\right]$, and $\mathsf{E}\left[u_t^\top \Lambda u_t\right]$. First, I derive the statistics for f_t, u_t and x_t. I first note that

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

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

$$F_t = \begin{bmatrix} (I - \Phi) f_0 \\ (I - \Phi)^2 f_0 \\ \vdots \\ (I - \Phi)^{t-1} f_0 \\ (I - \Phi)^t f_0 \end{bmatrix} + \underbrace{\begin{bmatrix} I & 0 & \cdots & 0 & 0 \\ (I - \Phi) & I & 0 & \cdots & 0 \\ \vdots & (I - \Phi) & \ddots & \cdots & 0 \\ (I - \Phi)^{t-1} & \cdots & \ddots & I & 0 \\ (I - \Phi)^t & \cdots & & (I - \Phi) & I \end{bmatrix}}_{\triangleq A_t} \begin{bmatrix} \epsilon_1^{(1)} \\ \epsilon_2^{(1)} \\ \vdots \\ \epsilon_{t-1}^{(1)} \\ \epsilon_t^{(1)} \end{bmatrix}.$$

Using this representation, I compute the mean

$$\text{(B.3)} \qquad \theta_t \triangleq \mathsf{E}[F_t] = \begin{bmatrix} \delta_1 \\ \delta_2 \\ \vdots \\ \delta_{t-1} \\ \delta_t \end{bmatrix} \triangleq \begin{bmatrix} (I - \Phi)f_0 \\ (I - \Phi)^2 f_0 \\ \vdots \\ (I - \Phi)^{t-1} f_0 \\ (I - \Phi)^t f_0, \end{bmatrix}$$

and the covariance matrix

$$\text{(B.4)} \qquad \Omega_t \triangleq \text{Var}[F_t] = A_t \begin{bmatrix} \Psi & 0 & \dots & 0 & 0 \\ 0 & \Psi & & \dots & 0 \\ \vdots & & \ddots & \dots & 0 \\ & & \dots & \Psi & 0 \\ 0 & \dots & & 0 & \Psi \end{bmatrix} A_t^\top.$$

Note that Ω_t is a block diagonal matrix with t blocks of size $K \times K$. Recall that in Section 3.4, I defined

$$\text{(B.5)} \qquad M_t \triangleq \begin{bmatrix} E_{1,t} & E_{2,t} & \dots & E_{t,t} \end{bmatrix}$$

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

$$\text{(B.6)} \qquad \mu_t \triangleq \mathsf{E}[u_t] = c_t + M_t \theta_t$$

$$V_t \triangleq \text{Var}(u_t) = M_t \Omega_t M_t^\top.$$

Therefore, u_t is normally distributed with mean μ_t and covariance matrix V_t. Similarly, I can obtain the statistics for x_t. Using (3.15),

$$\kappa_t \triangleq \mathsf{E}[x_t] = d_t + P_t \theta_t$$

$$Y_t \triangleq \text{Var}(x_t) = P_t \Omega_t P_t^\top.$$

I note the following fact from multivariate statistics.

Fact 1. *If z is a random vector with mean μ and variance Σ, and Q is positive definite matrix, then*

$$\mathsf{E}[z^\top Q z] = tr(Q\Sigma) + \mu^\top Q \mu.$$

Using Fact 1, I can compute each term in the expectation.

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

where ω_s is the sth diagonal block matrix of Ω_t having a size of $K \times K$. Finally, for the transaction cost terms,

$$
\begin{aligned}
\mathsf{E}\left[u_t^\top \Lambda u_t\right] &= \mathsf{E}\left[(c_t + M_t F_t)^\top \Lambda (c_t + M_t F_t)\right] \\
&= (c_t + M_t\theta_t)^\top \Lambda (c_t + M_t\theta_t) + \mathrm{tr}\left(\Lambda M_t\Omega_t M_t^\top\right)
\end{aligned}
$$

Summing up all the terms, the final objective function in deterministic form equals

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

which is a quadratic function of the policy parameters.

I now rewrite the equality constraint, $x_T = 0$ in terms of policy parameters. In order to enforce this equality, I need

$$
d_T = 0 \text{ and } J_{s,T} = \mathbf{0} \qquad s = 1, \ldots, T.
$$

Lastly, I replace probabilistic constraints with deterministic constraints using Lemma 3. Note that $\mathsf{P}\left(x_t \le 0\right) \le \eta$ can be written as $\mathsf{P}\left(-x_t \ge 0\right) \le \eta$. Then, using Lemma 3,

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

Similarly, I obtain that $\mathsf{P}\left(u_t \ge 0\right) \le \eta$ can be replaced by

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

Combining all the results, I obtain the deterministic version of the stochastic program in (3.29), a second-order cone program:

$$
\text{(B.7)} \qquad \underset{c_t, E_{s,t}}{\text{maximize}} \sum_{t=1}^{T} \left\{ d_t^\top B \delta_t + \sum_{s=1}^{t} \left(\delta_s^\top (B(I - \Phi)^{t-s} J_{s,t}^\top) \delta_s + \text{tr} \left(B(I - \Phi)^{t-s} E_{s,t}^\top \omega_s \right) \right) \right.
$$

$$
\left. + \frac{1}{2} \left((c_t + M_t \theta_t)^\top \Lambda (c_t + M_t \theta_t) + \text{tr} \left(M_t^\top \Lambda M_t \Omega_t \right) \right) \right\}
$$

$$
\text{subject to } d_t = x_0 + \sum_{i=1}^{t} c_i \qquad t = 1, \ldots, T,
$$

$$
J_{s,t} = \sum_{i=s}^{t} E_{s,i} \qquad 1 \le s \le t \le T,
$$

$$
(-d_t - P_t \theta_t) + \Phi^{-1}(1 - \eta) \left\| \left(P_t \Omega_t P_t^\top \right)^{1/2} \right\|_2 \le 0 \qquad t = 1, \ldots, T,
$$

$$
(c_t + M_t \theta_t) + \Phi^{-1}(1 - \eta) \left\| \left(M_t \Omega_t M_t^\top \right)^{1/2} \right\|_2 \le 0 \qquad t = 1, \ldots, T,
$$

$$
d_T = 0 \text{ and } J_{s,T} = \mathbf{0}.
$$

Note that the number of decision variables is considerably greater than that of the original execution problem in (3.24). Total number of decision variables in a problem with N securities, K factors and T periods equals $2NT + NKT(T + 1)$ which is on the order of $O(NKT^2)$.