Diss. ETH No. 17746

Optimal Trading Algorithms: Portfolio Transactions, Multiperiod Portfolio Selection, and Competitive Online Search

A dissertation submitted to the
ETH Zürich

for the degree of
Doctor of Sciences

presented by
Julian Michael Lorenz
Diplom-Mathematiker, Technische Universität München
born 19.12.1978
citizen of Germany

accepted on the recommendation of
Prof. Dr. Angelika Steger, examiner
Prof. Dr. Hans-Jakob Lüthi, co-examiner
Dr. Robert Almgren, co-examiner

2008

To my parents

Abstract

This thesis deals with optimal algorithms for trading of financial securities. It is divided into four parts: risk-averse execution with market impact, Bayesian adaptive trading with price appreciation, multiperiod portfolio selection, and the generic online search problem k-search.

Risk-averse execution with market impact. We consider the *execution of portfolio transactions* in a trading model with *market impact*. For an institutional investor, especially in equity markets, the size of his buy or sell order is often larger than the market can immediately supply or absorb, and his trading will move the price (*market impact*). His order must be worked across some period of time, exposing him to price volatility. The investor needs to find a trade-off between the market impact costs of rapid execution and the market risk of slow execution.

In a mean-variance framework, an optimal execution strategy minimizes variance for a specified maximum level of expected cost, or conversely. In this setup, Almgren and Chriss (2000) give path-independent (also called *static*) execution algorithms: their trade-schedules are deterministic and do not modify the execution speed in response to price motions during trading.

We show that the static execution strategies of Almgren and Chriss (2000) can be significantly improved by *adaptive trading*. We first illustrate this by constructing strategies that update exactly once during trading: at some intermediary time they may readjust in response to the stock price movement up to that moment. We show that such single-update strategies yield lower expected cost for the same level of variance than the static trajectories of Almgren and Chriss (2000), or lower variance for the same expected cost.

Extending this first result, we then show how *optimal dynamic strategies* can be computed to any desired degree of precision with a suitable application of the dynamic programming principle. In this technique the control variables are not only the shares traded at each time step, but also the maximum expected cost for the remainder of the program; the value function is the variance of the remaining program. This technique reduces the determination of optimal dynamic strategies to a series of single-period convex constrained optimization problems.

The resulting adaptive trading strategies are "aggressive-in-the-money": they accelerate the execution when the price moves in the trader's favor, spending parts of the trading gains to reduce risk. The relative improvement over static trade schedules is larger for large initial positions, expressed in terms of a new nondimensional parameter, the *market power* μ. For small portfolios, $\mu \to 0$, optimal adaptive trade schedules coincide with the static trade schedules of Almgren and Chriss (2000).

Bayesian adaptive trading with price appreciation. This part deals with another major driving factor of transaction costs for institutional investors, namely *price appreciation* (price trend) during the time of a buy (or sell) program. An investor wants to buy (sell) a stock before other market participants trade the same direction and push up (respectively, down) the price.

Thus, price appreciation compels him to complete his trade rapidly. However, an aggressive trade schedule incurs high market impact costs. Hence, it is vital to balance these two effects.

The magnitude of the price appreciation is uncertain and comes from increased trading by other large institutional traders. Institutional trading features a strong *daily cycle*. Market participants make large investment decisions overnight or in the early morning, and then trade across the entire day. Thus, price appreciation in early trading hours gives indication of large net positions being executed and ensuing price momentum. We construct a model in which the trader uses the observation of the price evolution during the day to estimate price momentum and to determine an optimal trade schedule to minimize total expected cost of trading. Using techniques from dynamic programming as well as the calculus of variations we give explicit optimal trading strategies.

Multiperiod portfolio selection. We discuss the well-known mean-variance portfolio selection problem (Markowitz, 1952, 1959) in a multiperiod setting. Markowitz's original model only considers an investment in one period. In recent years, multiperiod and continous-time versions have been considered and solved. In fact, in a multiperiod setting the portfolio selection problem is related to the optimal execution of portfolio transactions. Using the same dynamic programming technique as in the first part of this thesis, we explicitly derive the optimal dynamic investment strategy for discrete-time multiperiod portfolio selection. Our solution coincides with previous results obtained with other techniques (Li and Ng, 2000).

Generic online search problem k-search. We discuss the generic *online search problem k-search*. In this problem, a player wants to sell (respectively, buy) $k \geq 1$ units of an asset with the goal of maximizing his profit (minimizing his cost). He is presented a series of prices, and after each quotation he must *immediately* decide whether or not to sell (buy) one unit of the asset. We impose only the rather minimal modeling assumption that all prices are drawn from some finite interval. We use the *competitive ratio* as a performance measure, which measures the worst case performance of an online (sequential) algorithm vs. an optimal offline (batch) algorithm.

We present (asymptotically) optimal deterministic and randomized algorithms for both the maximization and the minimization problem. We show that the maximization and minimization problem behave substantially different, with the minimization problem allowing for rather poor competitive algorithms only, both deterministic and randomized. Our results generalize previous work of El-Yaniv, Fiat, Karp, and Turpin (2001).

Finally, we shall show that there is a natural connection between k-search and *lookback options*. A lookback call allows the holder to buy the underlying stock at time T from the option writer at the historical minimum price observed over $[0, T]$. The writer of a lookback call option can use algorithms for k-search to buy shares as cheaply as possible before expiry, and analogously for a lookback put. Hence, under a no-arbitrage condition the competitive ratio of these algorithms give a bound for the price of a lookback option, which in fact shows similar qualitative and quantitative behavior as pricing in the standard Black-Scholes model.

Zusammenfassung

Diese Arbeit beschäftigt sich mit optimalen Algorithmen für den Handel von Wertpapieren, und gliedert sich in vier Teile: risikoaverse Ausführung von Portfolio-Transaktionen, Bayes-adaptiver Aktienhandel unter Preistrend, Mehrperioden-Portfoliooptimierung und das generische Online-Suchproblem k-search.

Risikoaverse Ausführung von Portfoliotransaktionen. Der ersten Teil behandelt das Problem der *optimalen Ausführung von Portfoliotransaktionen* in einem Marktmodel mit *Preiseinfluss*. Die Größe der Transaktionen eines institutionellen Anlegers am Aktienmarkt ist oftmals so groß, dass sie den Aktienpreis beeinflussen. Eine solche große Order kann nicht unmittelbar ausgeführt werden, sondern muss über einen längeren Zeitraum verteilt werden, zum Beispiel über den gesamten Tag. Dies führt dazu, dass der Anleger währenddessen der Preisvolatilität ausgesetzt ist. Das Ziel einer Minimierung der *Preiseinfluss-Kosten* durch langsame Ausführung über einen ausgedehnten Zeitraum und das Ziel einer Minimierung des Marktrisikos durch möglichst rasche Ausführung sind damit gegenläufig, und der Händler muss einen Kompromiss finden.

In dem bekannten Mittelwert-Varianz-Ansatz werden diejenigen Ausführungsstrategien als optimal bezeichnet, die die Varianz für einen bestimmten maximalen Erwartungswert der Kosten minimieren, oder aber für eine bestimmte maximale Varianz der Kosten deren Erwartungswert. Almgren und Chriss (2000) bestimmen in einem solchen Modell pfadunabhängige (auch statisch genannte) Ausführungsalgorithmen, d.h. ihre Strategien sind deterministisch und passen den Ausführungsplan nicht an als Antwort auf einen steigenden oder fallenden Aktienkurs.

Diese Arbeit zeigt, dass die Ausführungsstrategien von Almgren und Chriss (2000) entscheidend verbessert werden können, und zwar mit Hilfe von adaptiven, dynamischen Strategien. Wir demonstrieren dies zunächst am Beispiel von sehr einfachen dynamischen Strategien, die *genau einmal* während der Ausführung darauf reagieren dürfen, ob der Aktienpreis gestiegen oder gefallen ist. Wir bezeichnen diese Strategien als *single-update Strategien.* Trotz ihrer Einfachheit liefern sie bereits eine gewaltige Verbesserung gegenüber den Ausführungsstrategien von Almgren und Chriss (2000).

Als nächstes zeigen wir, wie man voll-dynamische optimale Ausführungsstrategien berechnen kann, und zwar mittels dynamischer Programmierung für Mittelwert-Varianz-Probleme. Der entscheidende Schritt für diese Technik ist es, die maximal erwarteten Kosten neben der Anzahl der Aktien als Zustandsvariable zu verwenden, und die Varianz der Strategie als Wertfunktion des dynamischen Programms. Auf diese Weise wird die Bestimmung von voll-dynamischen optimalen Strategien auf eine Serie von konvexen Optimierungsproblemen zurückgeführt.

Die optimalen dynamischen Strategien sind „aggressiv im Geld“, d.h. sie beschleunigen die Ausführung, sobald sich der Preis vorteilhaft für den Händler entwickelt. Die beschleunigte Ausführung führt zu höheren Preiseinflusskosten, was den Händler einen Teil der Gewinne aus der

positiven Kursentwicklung kostet. Dafür kann der Händler aber sein Kaufs- bzw. Verkaufsprogramm schneller beenden, und geht weniger Marktrisiko ein – was zu einer Verbesserung im Mittelwert-Varianz Kompromiss führt. Wir zeigen, dass die relative Verbesserung der dynamischen Strategien gegenüber den Strategien von Almgren und Chriss (2000) umso größer ausfällt, je größer die Portfoliotransaktion ist. Für sehr kleine Transaktionen ergibt sich keine Verbesserung, und die optimalen dynamischen Strategien stimmen mit den statischen Strategien von Almgren und Chriss überein.

Bayes-adaptiver Aktienhandel unter Preistrend. Der zweite Teil der Arbeit beschäftigt sich mit einem weiteren entscheidenden Bestandteil der Transaktionskosten eines grossen institutionellen Anlegers, nämlich mit *Preistrend* während der Ausführung eines Kaufs- bzw. Verkaufsprogramms. Ein solcher Preistrend rührt daher, dass andere Marktteilnehmer ähnliche Kaufs- bzw. Verkaufsprogramme laufen haben, was den Preis verschlechtert. Dies veranlasst den Händler dazu, sein Kaufs- bzw. Verkaufsprogramm möglichst schnell durchzuführen; wiederum muss er jedoch einen Kompromiss finden mit höheren Preiseinfluss-Kosten, die durch eine aggressive Ausführung anfallen.
Institutioneller Aktienhandel unterliegt einem starken *Tageszyklus.* Marktteilnehmer fällen grosse Anlageentscheidungen morgens vor Handelsbeginn und führen diese Programme dann über den Tag aus. Ein Preistrend am Anfang des Tages kann damit auf ein Übergewicht an Kauf- bzw. Verkaufsinteresse hindeuten und lässt vermuten, dass dieses Preismoment auch im weiteren Tagesverlauf anhält. Wir betrachten ein Modell für diese Situation, in dem der Händler die Entwicklung des Preises zur Schätzung des Preistrends verwendet und damit eine optimale Strategie ermittelt, die den Erwartungswert seiner Kosten minimiert. Die mathematischen Techniken hierzu sind dynamische Programmierung und Variationsrechnung.

Mehrperioden-Portfoliooptimierung. Der dritte Teil dieser Arbeit diskutiert das wohlbekannte Problem der Portfolio-Optimierung im Mittelwert-Varianz-Ansatz von Markowitz (1952, 1959) in einem Mehrperiodenmodell. Das ursprüngliche Modell von Markowitz beschränkte sich auf ein Investment in nur einer Periode. In den letzten Jahren wurde das Mittelwert-Varianz Portfolioproblem in solchen Mehrperiodenmodellen oder in Modellen mit stetiger Zeit betrachtet und gelöst. Das Problem der optimalen Ausführung von Portfolio-Transaktionen ist in der Tat mit diesem Problem verwandt. Mit der Technik der dynamischen Programmierung für Mittelwert-Varianz-Probleme, das wir für das Portfoliotransaktionsproblem entwickelt haben, können explizite optimale dynamische Anlagestrategien in diskreter Zeit bestimmt werden. Unsere Formeln stimmen mit denen von Li und Ng (2000) überein, die diese mit Hilfe einer anderen Technik ermittelt hatten.

Generisches Online-Suchproblem k-search. Der vierte Teil dieser Arbeit beschäftigt sich mit dem generischen Online-Suchproblem *k-search*: Ein Online-Spieler möchte $k \geq 1$ Einheiten eines Gutes verkaufen (bzw. kaufen) mit dem Ziel, seinen Gewinn zu maximieren (bzw. seine Kosten zu minimieren). Er bekommt der Reihe nach Preise gestellt, und muss nach jedem Preis unmittelbar entscheiden ob er für diesen Preis eine Einheit verkaufen (bzw. kaufen) möchte. Die einzige Modellannahme für die Preissequenz ist, dass alle Preise aus einem vorgegebenen Intervall stammen. Wir verwenden die *competitive ratio* als Qualitätsmaß. Dieses beurteilt einen Online-Algorithmus relativ zu einem optimalen „Offline-Algorithmus“, der die gesamte Preissequenz im Voraus kennt.

Diese Arbeit ermittelt (asymptotisch) optimale deterministische und randomisierte Algorithmen sowohl für das Maximierungs- als auch das Minimierungsproblem. Erstaunlicherweise verhalten sich das Maximierungs- und das Minimierungsproblem deutlich unterschiedlich: optimale Algorithmen für das Minimierungsproblem erzielen eine deutlich schlechtere competitive ratio als im Maximierungsproblem. Diese Ergebnisse verallgemeinern Resultate von El-Yaniv, Fiat, Karp und Turpin (2001).

Wir zeigen abschliessend, dass es eine natürliche Beziehung gibt zwischen k-search und sogenannten *Lookback-Optionen*. Ein Lookback-Call gibt das Recht, die zugrundeliegende Aktie zur Zeit T zum historischen Minimums-Preis während des Zeitraums $[0, T]$ zu erwerben. Der Stillhalter eines Lookback Calls kann Algorithmen für k-search verwenden, um bis Ablauf der Option möglichst billig die Aktien zu erwerben, die er zur Erfüllung seiner Pflicht benötigt; analog kann sich der Stillhalter eines Lookback-Puts Algorithmen für k-search in der Maximierungsversion zu Nutze machen. Unter der Annahme der Arbitragefreiheit gibt die competitive ratio eines k-search Algorithmus damit eine Schranke für den Preis einer Lookback-Option. Diese Schranke zeigt ähnliches qualitatives und quantitatives Verhalten wie die Bewertung der Option im Black-Scholes Standardmodell.

Acknowledgments

This dissertation would not have been possible without the help and support from many people to whom I am greatly indebted.

First and foremost, I am particularly grateful to my advisor, Professor Angelika Steger for her support, patience and encouragement throughout my doctoral studies. On the one hand, she gave me the freedom to pursue my own research ideas, but on the other hand she was always there to listen and to give valuable advice when I needed it.

I am immensely indebted to Robert Almgren. His enthusiasm and embracement of my proposal about the issue of adaptive vs. static trading strategies played a vital role in defining this thesis. His responsiveness, patience and expertise were a boon throughout our collaboration. I feel very honored that he is co-examiner of this thesis. I am also grateful for his genuine hospitality during my stay in New York.

I sincerely appreciate the collaboration with my other coauthors Martin Marciniszyn, Konstantinos Panagiotou and Angelika Steger. I am grateful to Professor Hans-Jakob Lüthi for accepting to act as co-examiner and reading drafts of this dissertation.

This research was supported by UBS AG. The material and financial support is gratefully acknowledged. I also would like to thank the people at UBS for fruitful discussions – Clive Hawkins, Stefan Ulrych, Klaus Feda and Maria Minkoff.

Thanks to the many international PhD students hosted by Switzerland, life as a graduate student at ETH is colorful and rich. I am happy that I could share the PhD struggle with some great colleagues in Angelika Steger's research group, and I want to thank all of them – Nicla Bernasconi, Stefanie Gerke, Dan Hefetz, Christoph Krautz, Fabian Kuhn, Martin Marciniszyn, Torsten Mütze, Konstantinos Panagiotou, Jan Remy, Alexander Souza, Justus Schwartz, Reto Spöhel and Andreas Weissl.

Most of all, I want to express my sincere gratitude to my family and especially my parents for their love and invaluable support at any time during my life. I am glad you are there.

Julian Lorenz
Zürich, January 2008

Contents

CHAPTER 1

Introduction

"*I can calculate the movement of the stars,*
but not the madness of men."

– Sir Isaac Newton after losing a fortune in the *South Sea Bubble*[1].

Until about the middle of the last century the notion of finance as a scientific pursuit was inconceivable. Trading in financial securities was mostly left to gut feeling and bravado. An important cornerstone for a quantitative theory of finance was laid by Harry Markowitz's work on portfolio theory (Markowitz, 1952). One lifetime later, the disciplines of finance, mathematics, statistics and computer science are now heavily and fruitfully linked. The emergence of mathematical models not only transformed finance into a quantitative science, but also changed financial markets fundamentally.

A mathematical model of a financial market typically leads to an *algorithm* for sound, automated decision-making, replacing the error-prone subjective judgment of a trader. Certainly, not only do algorithms improve the quality of financial decisions, they also pave the ability to handle the sheer amount of transactions of today's financial world.

For instance, the celebrated Black-Scholes option pricing model (Black and Scholes, 1973) not only constitutes a pricing formula, it is essentially an algorithm (*delta hedging*) to create a riskless hedge of an option. Similarly, Markowitz's portfolio theory entails an algorithmic procedure to determine efficient portfolios.

This thesis deals with optimal algorithms for trading of financial securities. In the remaining part of this introduction, we briefly present the main results of this work. The remaining part of the thesis discusses each of the results in detail. The problems arise from different applications of trading algorithms and use different methods. Chapters 2 – 5 use methods from mathematical finance, whereas Chapter 6 uses methods from computer science and the analysis of algorithms.

Section 1.1 gives a brief overview of algorithmic trading and the *optimal execution of portfolio transactions*. This thesis addresses two specific problems in this field: *optimal risk-averse execution under market impact* and *Bayesian adaptive trading with price appreciation*. We introduce these problems in Sections 1.2 and 1.3, and discuss them in detail in Chapters 2, 3 and 4. In Section 1.4 we formulate *optimal mean-variance portfolio selection* in a multiperiod setting; Chapter 5 is devoted to this problem. Finally, in Section 1.5 we introduce a rather generic online search problem and its relation to *pricing of lookback options*. We discuss this problem in detail in Chapter 6.

[1]The South Sea Company was an English company granted a monopoly to trade with South America. Overheated speculation caused the stock price to surge over the course of a single year from one hundred to over one thousand pounds per share before it collapsed in September 1720. Allegedly, Sir Isaac Newton lost £20,000 in the bubble.

1.1. Optimal Execution of Portfolio Transactions

Portfolio theory delivers insight into optimal asset allocation and optimal portfolio construction in a frictionless market. However, in reality the acquisition (or liquidation) of a portfolio position does not come for free. *Transaction costs* refer to any costs incurred while implementing an investment decision.

Some transaction costs are observable directly – brokerage commissions (typically on a per share basis), fees (e.g. clearing and settlement costs) and taxes. These cost components constitute the overwhelming part of total transaction costs for a retail investor.

However, for an institutional investor, the main transaction cost components are more subtle and stem primarily from two effects: *price appreciation* during the time of the execution of his orders and *market impact* caused by his trades. These hidden transaction cost components are called *slippage*.

Slippage is typically not an issue for the individual investor since the size of his trade is usually minuscule compared to the market liquidity. But for institutional investors slippage plays a big role. For instance, a fund manager forecasts that a stock will rise from \$100 to \$105 by the end of the month, but in acquiring a sizable position he pays an average price of \$101. As a result, instead of the 5% profit the investor only gains about 4% and falls considerably short of his forecasted profit.

Indeed, empirical research shows that for institutional sized transactions, total implementation cost is around 1% for a typical trade and can be as high as 2-3% for very large orders in illiquid stocks (Kissell and Glantz, 2003). Typically, institutional investors have a portfolio turnover a couple of times a year. Poor execution therefore can erode portfolio performance substantially.

In the last decade, trading strategies to implement a certain portfolio transaction (in order to achieve a desired portfolio) flourish in the financial industry. Such strategies are typically referred to as *algorithmic trading*, or "robo-trading" (The Economist, 2005). They can be described as the automated, computer-based execution of equity orders with the goal of meeting a particular benchmark. Their proliferation has been remarkable: by 2007, algorithmic trading accounts for a third of all share trades in the USA and is expected to make up more than half the share volumes and a fifth of options trades by 2010 (The Economist, 2007).

The contribution of this thesis to the field of algorithmic trading is twofold. In the first contribution we show that the classic arrival price algorithms in the market impact models of Almgren and Chriss (2000) can be significantly improved by *adaptive trading strategies*; we shall discuss this in Section 1.2, and present the results in detail in Chapters 2 and 3. The second contribution is an algorithm that deals with the other main driving factor of transaction costs, namely price appreciation; we give a brief outline in Section 1.3 and the detailed results in Chapter 4.

1.2. Adaptive Trading with Market Impact

For an individual trader, market impact arises from his substantial supply or demand in terms of sell or buy orders, respectively. It is the effect that buying or selling moves the price against the trader, i.e. upward when buying and downward when selling. It can be interpreted as a price premium ("sweetener") paid to attract additional liquidity to the market to absorb an order.

An obvious way to avoid market impact would be to trade more slowly to allow market liquidity recover between trades. Instead of placing a huge 10,000-share order in one big chunk, a trader may break it down in smaller orders and incrementally feed them into the market over time. However, slower trading increases the trader's susceptibility to market volatility and prices may potentially move disadvantageous to the trader. Kissell and Glantz (2003) coin this the traders' dilemma: *"Do trade and push the market. Don't trade and the market pushes you."*. Optimal trade schedules seek to balance the market impact cost of rapid execution against the volatility risk of slow execution.

The classic market impact model in algorithmic trading is due to Almgren and Chriss (2000). There the execution benchmark is the pre-trade price. The difference between the pre-trade and the post-trade book value of the portfolio (including cash positions) is the implementation shortfall (Perold, 1988). For instance, the implementation shortfall of a sell program is the initial value of the position minus the actual amount captured. Algorithms of this type are usually referred to as *arrival price algorithms*.

In the simplest model the expected value of the implementation shortfall is entirely due to market impact as a deterministic function of the trading rate, integrated over the entire trade. The variance of the implementation shortfall is entirely due to the volatility of the price process, which is modeled as a random walk. In the mean-variance framework used by Almgren and Chriss (2000), "efficient" strategies minimize this variance for a specified maximum level of expected cost or conversely. The set of such strategies is summarized in the "efficient frontier of optimal trading" introduced by Almgren and Chriss (2000) (see also Almgren and Chriss (1999)), akin to the well-known Markowitz efficient frontier in portfolio theory. Prior to Almgren and Chriss, Bertsimas and Lo (1998) considered optimal risk-neutral execution strategies.

The execution strategies of Almgren and Chriss (2000) are *path-independent* (also called *static*): they do not modify the execution speed in response to price movement during trading. Almgren and Chriss determine these strategies by optimizing the tradeoff criterion $\mathbb{E}[C] + \lambda \operatorname{Var}[C]$ for the total implementation shortfall C. Then they argue that *re-evaluating* this criterion at each intermediate time with an unchanged risk aversion λ, using the information available at that time, yields a trade schedule that coincides with the trade schedule specified at the initial time. Hence, they claim that optimal dynamic strategies will in fact be static trajectories, deterministically fixed before trading begins, and are not updated in the course of trading in response to the stock price motion. Whether the price goes up or down, the same trade schedule is rigidly adhered to.

However, in this thesis we show that substantial improvement with respect to the mean-variance tradeoff measured at the initial time is possible if we allow *path-dependent* policies. We determine optimal dynamic trading strategies that adapt in response to changes in the price of the asset being traded, even with no momentum or mean reversion in the price process. As Almgren and Chriss (2000), we assume a pure random walk with no serial correlation, using pure classic mean-variance as risk-reward tradeoff.

A dynamic trading strategy is a precomputed *rule* specifying the trade rate as a function of price, time and shares remaining. Complications in determining such optimal dynamic strategies stem from the fact that due to the square in the expectation of the variance term, the tradeoff criterion $\mathbb{E}[C] + \lambda \operatorname{Var}[C]$ is not directly amenable to dynamic programming (as presumed in the specification of Almgren and Chriss).

Besides the articles by Almgren and Chriss (Almgren and Chriss, 1999, 2000, 2003; Almgren, 2003), other work on optimal execution of portfolio transactions was done by Konishi and Makimoto (2001), Huberman and Stanzl (2005), Engle and Ferstenberg (2007), Krokhmal and Uryasev (2007), Butenko, Golodnikov, and Uryasev (2005), He and Mamaysky (2005). However, none of them determines optimal adaptive execution strategies in a pure mean-variance setting without price momentum or mean-reversion. Schied and Schöneborn (2007) consider optimal execution for a investor with exponential utility function (CARA utility); for this specific utility function, optimal dynamic strategies are indeed the path-independent static trajectories of Almgren and Chriss (2000).

A problem related to the optimal execution of portfolio transactions is the classic problem of mean-variance portfolio optimization in a multiperiod setting. We will introduce this problem in Section 1.4 below and discuss it in detail in Chapter 5. While mean-variance portfolio selection was originally considered as a single-period problem by Markowitz (1952), recently techniques have been proposed to solve it in a multiperiod setting (Li and Ng, 2000; Zhou and Li, 2000; Richardson, 1989; Bielecki, Jin, Pliska, and Zhou, 2005). However, these techniques do not easily carry over to the portfolio transaction problem due to the market impact terms which significantly complicate the problem. Thus, in this thesis we follow a different approach to determine optimal dynamic execution strategies with respect to the specification of measuring risk and reward at the initial time.

In Chapter 2 we construct simple dynamic trading strategies in a continuous-time trading model, which update exactly once: at some intermediary time they may readjust in response to the stock price movement. Before and after that "intervention" time they may not respond to whether the price goes up or down. We determine optimal single-update strategies using a direct optimization approach for the trade-off criterion $\mathbb{E}[C] + \lambda \operatorname{Var}[C]$. We will show that already these simple adaptive strategies improve over the path-independent strategies of Almgren and Chriss (2000) with respect to the mean-variance tradeoff measured at the initial time. Certainly, the single-update strategies are not the fully optimal dynamic execution strategy. Unfortunately, the approach in Chapter 2 does not generalize to multiple decision times during trading.

In Chapter 3 we propose a dynamic programming principle for mean-variance optimization in discrete time to determine optimal Markovian strategies. This technique reduces the determination of optimal dynamic strategies to a series of single-period convex constrained optimization problems. We give an efficient scheme to obtain numerical solutions. In fact, the same dynamic programming principle is applicable to multiperiod mean-variance portfolio selection. In Chapter 5, we show how it can be used to determine optimal investment strategies. Unlike in

Chapter 3, for the portfolio problem in Chapter 5 we obtain an explicit closed-form solution, matching the optimal policies given by Li and Ng (2000).

We show that the relative improvement of a dynamic strategy over a static trade-scheduel is larger for large portfolio transactions. This will be expressed in terms of a new preference-free nondimensional parameter μ that measures portfolio size in terms of its ability to move the market relative to market volatility. We show that for small portfolios, $\mu \to 0$, optimal adaptive trade schedules coincide with the optimal static trade schedules of Almgren and Chriss (2000).

Furthermore, we show that the improvement stems from introducing a correlation between the trading gains (or losses) in earlier parts of the execution and market impact costs incurred in later parts: if the price moves in our favor in the early part of the trading, then we reduce risk by accelerating the remainder of the program, spending parts of the windfall gains on higher market impact costs. If the price moves against us, then we reduce future costs by trading more slowly, despite the increased exposure to risk of future fluctuations.

Hence, our new optimal trade schedules are "aggressive-in-the-money" (AIM) in the sense of Kissell and Malamut (2005). Another interpretation of this behavior is that the trader becomes *more risk-averse* after *positive performance*. This effect can also be observed in multiperiod mean-variance portfolio selection in Chapter 5.

As it is well known, mean-variance based preference criteria have received theoretical criticism (see for instance Maccheroni, Marinacci, Rustichini, and Taboga, 2004). Artzner, Delbaen, Eber, Heath, and Ku (2007) discuss coherent multiperiod risk adjusted values for stochastic processes; risk-measurement is constructed in terms of sets of test probabilities, especially so called *stable* sets, which ensure consistency with single-period risk measurement. Densing (2007) discusses an application of this type of risk-measurement for the optimal multi-period operation of a hydro-electric pumped storage plant with a constraint on risk.

However, mean-variance optimization retains great practical importance due to its intuitive meaning. As discussed by Almgren and Lorenz (2007), optimizing mean and variance corresponds to how trading results are typically reported in practice. Clients of an agency trading desk are provided with a post-trade report daily, weekly, or monthly depending on their trading activity. This report shows sample average and standard deviation of the implementation shortfall benchmark, across all trades executed for that client during the reporting period. Therefore the broker-dealer's goal is to design algorithms that optimize sample mean and variance *at the per-order level*. That is, the broker/dealer offers the efficient frontier of execution strategies, and lets the client pick according to his needs.

Compared to utility based criteria (von Neumann and Morgenstern, 1953), this approach has the advantage that no assumptions on the clients' utility, his wealth or his other investment activities are needed; we shall briefly outline in Appendix A how the optimization problem can be studied in an expected utility framework. The relationship between mean-variance and expected utility is discussed, for instance, by Kroll, Levy, and Markowitz (1984) and Markowitz (1991). For Gaussian random variables, mean-variance is consistent with expected utility maximization as well as stochastic dominance (see for instance Levy (1992); Bertsimas, Lauprete, and Samarov

(2004)). Thus, as long as the probability distribution of the payoff is not too far from Gaussian, we can expect mean-variance to give reasonable results.

This is joint work with Robert Almgren. The results from Chapter 2 were published as Almgren and Lorenz (2007).

1.3. Bayesian Adaptive Trading with Price Appreciation

Another substantial source of slippage is price appreciation during the time of a buy (respectively, sell) program. Price appreciation is also referred to as *price trend* and is usually a cost because investors tend to buy stocks that are rising and sell stocks that are falling. In the following this is implicitly assumed.

Similar to price volatility, price appreciation also compels the trader to complete the trade rapidly. For a buy program, he wants to get in the stock before everybody else jumps on the bandwagon and pushes up the price. However, an aggressive strategy will cause high market impact costs. Conversely, a very passive trading schedule (for instance, selling at constant rate) will cause little market impact, but high costs due to price appreciation, since a large fraction of the order will be executed later in the trading period when prices are possibly less favorable. Therefore, an optimal trading strategy will seek a compromise. In contrast to algorithms for the market impact model discussed in the previous section, urgency stems from the anticipated drift, not from the trader's risk aversion.

The magnitude of the price appreciation is uncertain, and we only have a-priori estimates. Typically, price appreciation during a short time horizon (e.g. a day), comes from increased trading by other large institutional traders.

Institutional trading features a strong *daily cycle*. Investment decisions are often made overnight or in the early morning, and then left to be executed during the day. Consequently, within each trading day price appreciation in the early trading hours may give an indication of large net positions being executed on that day by other institutional investors. An astute investor will collect that information, and use it to trade in the rest of the day.

In Chapter 4, we model such price momentum and the daily cycle by assuming that the stock price follows a Brownian motion with an underlying drift. The drift is constant throughout the day, but we do not know its value. It is caused by the orders executed by other large traders across the entire day. At the start of the day, we have a prior belief for the drift, and use price observations during the day to update this belief.

With techniques from dynamic programming as well as from the calculus of variations, we determine optimal dynamic trading strategies that minimize the expectation of total cost. This corresponds to an optimal strategy for a risk-neutral investor. We compare the performance of optimal dynamic trading strategies to the performance of optimal static (path-independent) strategies, and give an explicit expression for the improvement.

This is joint work with Robert Almgren and published as Almgren and Lorenz (2006).

1.4. Multiperiod Mean-Variance Portfolio Selection

Modern Portfolio Theory (Markowitz, 1952) was the first formalization of the conflicting investment objectives of high profit and low risk. Markowitz made unequivocally clear that the risk of a stock should not just be measured independently by its individual variance, but also by its covariance with all other securities in the portfolio. Investors use diversification to optimize their portfolios. Mathematically speaking, given the expectation $\mu = \mathbb{E}[e]$ of the return vector e of all available securities and their covariance matrix $\Sigma = \mathbb{E}[(e-\mu)(e-\mu)']$, the expected return of a portfolio, whose investment in each available security is given by the vector y, is simply $E(y) = \mu' y$ and its variance is $V(y) = y'\Sigma y$. The collection of all combinations $(V(y), E(y))$ for every possible portfolio y defines a region in the risk-return space. The line along the upper edge of this region is known as the *efficient frontier*, which is obtained by solving a quadratic programming problem. Efficient portfolios offer the lowest risk for a given level of return, or conversely the best possible return for a given risk. A rational investor will always choose his portfolio somewhere on this line, according to his risk-appetite.

Markowitz's work initiated a tremendous amount of research (see Steinbach (2001) for an excellent survey), and has seen widespread use in the financial industry. Later, it led to the elegant *capital asset pricing model* (CAPM), which introduces the notions of systematic and specific risk (Sharpe, 1964; Mossin, 1966; Lintner, 1965). Nowadays MPT, for which Markowitz was awarded the Nobel Prize in Economics in 1990, seems like a rather obvious idea. In some sense, this is testimonial to its success. His discovery was made in a context in which it was not obvious at all and challenged established paradigms.

Markowitz's original model only considers an investment in one period: the investor decides on his portfolio at the beginning of the investment period (for instance, a year) and patiently waits without intermediate intervention. In subsequent years there has been considerable work on multiperiod and continous-time models for portfolio management, pioneered by Merton (1971). However, in this line of work a different approach is used, namely the optimization in the framework of *expected utility* (von Neumann and Morgenstern, 1953): instead of the mean and variance of a portfolio, a single quantity, the expected terminal wealth $\mathbb{E}[u(w)]$ for a utility function $u(\cdot)$, is optimized. The technique used to determine optimal dynamic investment strategies is dynamic programming (Bellman, 1957). As mentioned above, we use this technique in Chapter 4 to determine optimal risk-neutral strategies in the market model with price appreciation.

As mentioned in Section 1.2, the relationship between mean-variance optimization and expected utility is discussed by Kroll et al. (1984) and Markowitz (1991). Despite theoretical criticism, mean-variance portfolio optimization retains great practical importance due to its intuitive meaning. Furthermore, from the perspective of an investment company, who does not know its clients' utility, the only way is to offer the efficient frontier of mean-variance optimal portfolio strategies, and let the clients pick according to their needs.

In contrast to the advancements in utility based dynamic portfolio management, multiperiod and continous-time versions of the mean-variance problem have not been considered and solved until recently. For multiperiod versions of mean-variance portfolio optimization, difficulties arise from the definition of the mean-variance objective, which does not allow a direct application of the dynamic programming principle due to the square of the expectation in the variance term.

Li and Ng (2000) circumvent this issue by embedding the original mean-variance problem into a tractable auxiliary problem with a utility function of quadratic type (i.e. using first and second moments). They obtain explicit, optimal solutions for the discrete-time, multiperiod mean-variance problem. The most important contribution of this paper was a general framework of stochastic linear-quadratic (LQ) optimal control in discrete time. Zhou and Li (2000) generalize the LQ control theory to continuous-time settings, and obtain a closed-form solution for the efficient frontier. This technique set off a series of papers on dynamic mean-variance optimization (Lim and Zhou, 2002; Li, Zhou, and Lim, 2002; Zhou and Yin, 2003; Leippold, Trojani, and Vanini, 2004; Jin and Zhou, 2007).

Parallel to this line of work, a second technique was employed by Bielecki, Jin, Pliska, and Zhou (2005) to study continuous-time mean-variance portfolio selection. They use a decomposition approach (Pliska, 1986) to reduce the problem into two subproblems: one is to solve a constrained static optimization problem on the terminal wealth, and the other is to replicate the optimal terminal wealth. Using such a technique, Richardson (1989) also tackled the mean-variance problem in a continuous-time setting. The related mean-variance hedging problem was studied by Duffie and Richardson (1991) and Schweizer (1995), where optimal dynamic strategies are determined to hedge contingent claims in an imperfect market.

For the optimal execution problem introduced in Section 1.2, obtaining solutions with these approaches is rather intricate due to the market impact terms. We use a more direct technique by a suitable application of the dynamic programming principle for discrete-time multiperiod mean-variance problems, as introduced in Chapter 3 for the portfolio transaction problem.

Using this technique, we explicitly derive the optimal dynamic investment strategy for the classical mean-variance portfolio selection problem in Chapter 5. Our solution coincides with the solutions for the discrete-time model given by Li and Ng (2000). As in the optimal execution problem in Chapter 3, we observe a change in risk-aversion in response to past performance. After a fortunate period profit the investor will try to conserve his realized gains and put less capital at risk in the remainder. This behavior of the optimal dynamic multiperiod strategy was already shown by Richardson (1989).

This chapter is based on joint work with Robert Almgren.

1.5. Optimal k-Search and Pricing of Lookback Options

Finally, in Chapter 6 we discuss a generic trading problem, called *k-search*, from a computer science perspective, and demonstrate its usefulness for the pricing of *lookback options*.

The generic search problem *k-search* is defined as follows: a player wants to sell (respectively, buy) $k \geq 1$ units of an asset with the goal of maximizing his profit (minimizing his cost). He is presented a series of price quotations p_i, and after each quotation he must *immediately* decide whether or not to sell (buy) one unit of the asset.

We analyze this problem under the rather minimal modeling assumption that all prices are drawn from a finite interval $[m, M]$, and we use the *competitive ratio* as a performance measure.

The competitive ratio measures the performance of an online (sequential) algorithm vs. the performance of an optimal offline (batch) algorithm. We shall give a more detailed introduction to *competitive analysis* of online algorithms in Chapter 5.

We present (asymptotically) optimal deterministic and randomized algorithms for both the maximization and the minimization problem. We show that the maximization and minimization problem behave substantially different, with the minimization problem allowing for rather poor competitive algorithms only, both deterministic and randomized. Our results generalize previous work of El-Yaniv et al. (2001), who considered the case $k = 1$ and the closely related *one-way-trading*. In the latter, the player can trade arbitrary fractional amounts at each price quotation. In a way, for $k \to \infty$, the *k-search* problem can be understood as an approximation to the one-way-trading problem; interestingly, the competitive ratios of the optimal algorithms for both problems coincide for $k \to \infty$.

In the second part of Chapter 6, we show that there is a natural connection to the pricing of *lookback options*. An option is a contract whereby the option holder has the right (but not obligation) to exercise a feature of the option contract on or before an exercise date, delivered by the other party – the writer of the option (Hull, 2002, for instance). Since the option gives the buyer a right, it will have a price that the buyer has to pay to the option writer. The most basic type of options are European options on a stock (Black and Scholes, 1973), which give the holder the right to buy (respectively, sell) the stock on a prespecified date T (expiry date) for a prespecified price K.

Lookback options are one instance of a whole plethora of so called *exotic* options with more complex features. A lookback call allows the holder to buy the underlying stock at time T from the option writer at the historical minimum price observed over $[0, T]$, and a lookback put to sell at the historical maximum.

The writer of a lookback call option can use algorithms for *k-search* to buy shares as cheap as possible before expiry (and analogously, the writer of a lookback put to sell shares as dear as possible). Hence, the competitive ratio of those algorithms give a bound for the price of a lookback option under a no-arbitrage condition, which in fact shows similar qualitative properties and quantitative values as pricing in the standard Black-Scholes model. The underlying model of a *trading range* is very different from the Black-Scholes assumption of Geometric Brownian motion. This is in the spirit of a number of attempts in recent years to price financial instruments by *relaxing* the Black-Scholes assumptions, and to provide *robust* bounds that work with (almost) any evolution of the stock price. For instance, DeMarzo, Kremer, and Mansour (2006) derive bounds for European options in a model of bounded quadratic variation, and Epstein and Wilmott (1998) propose non-probabilistic models for pricing interest rate securities.

This is joint work with K. Panagiotou and A. Steger, and is published as Lorenz, Panagiotou, and Steger (2008). An extended abstract appeared as Lorenz, Panagiotou, and Steger (2007).

CHAPTER 2

Adaptive Trading with Market Impact: Single Update

In this chapter we show that the arrival price algorithms for the execution of portfolio transactions in the market impact model of Almgren and Chriss (2000) can be significantly improved by adaptive trading. We present simple adaptive strategies that respond to the stock price process exactly once. Already these simple proof-of-concept-type strategies yield a large improvement with respect to the mean-variance trade-off evaluated at the initial time.

2.1. Introduction

In this and the next chapter, we discuss the optimal execution of portfolio transactions in the model with market impact introduced in Section 1.2. The most common case is purchasing or unwinding a large block of shares. Such large transactions cannot be easily executed with a single or a few market orders without considerable price concession. Instead, they must be worked across the day.

Recall that the *implementation shortfall* (Perold, 1988) of a portfolio transaction is the difference between the pre-trade and the post-trade book value of the portfolio (including cash positions), e.g. the implementation shortfall of a sell program is the initial value of the position minus the dollar amount captured.

In the model of Almgren and Chriss (2000), the expected value of the implementation shortfall is entirely due to market impact incurred by trading at a nonzero rate. This expected cost is minimized by trading as slowly as possible; since we require the transaction to be completed in a fixed period of time, this results in a strategy that trades at a constant rate throughout. Conversely, since market impact is assumed deterministic, the variance of the implementation shortfall is entirely due to price volatility, and this variance is minimized by trading rapidly.

In a mean-variance framework, "efficient" strategies minimize variance for a specified maximum level of expected cost or conversely; the set of such strategies is summarized in the "efficient frontier of optimal trading" introduced by Almgren and Chriss (2000) (see also Almgren and Chriss (1999)). In this work we will follow Almgren and Chriss (2000), and use mean-variance optimization; as mentioned in Chapter 1, *expected utility* optimization constitutes a different approach (see Appendix A for a brief discussion).

A central question is whether the trade schedule should be path-independent or path-dependent: should the list of shares to be executed in each interval of time be computed and fixed before trading begins, or should the trade list be updated in "real time" using information revealed during execution? In the following, we shall refer to *path-independent* trade schedules as *static* strategies, and to *path-dependent* trade schedules as *dynamic*. In algorithmic trading, dynamic strategies of this form are often also called "scaling" strategies. In some situations, implementing a dynamic strategy may not be possible; Huberman and Stanzl (2005) suggest that a reasonable example of this is insider trading, where trades must be announced in advance. However, in most cases dynamic trade schedules are valid, and a possible improvement over static schedules welcomed.

The specification used in Almgren and Chriss (2000) to determine optimal strategies is to *re-evaluate* a static trade schedule at each intermediate time with an unchanged parameter of risk aversion, using the information available at that moment. In this specification, the re-evaluated trade schedule coincides with the trade schedule specified at the initial time, and Almgren and Chriss (2000) propose static trade schedules as optimal trading strategies.

However, in this and the following chapter we shall show that substantial improvement with respect to the mean-variance tradeoff measured at the initial time is possible by precomputing a *rule* determining the trade rate as a function of price. Once trading begins, the rule may not be modified, even if the trader's preferences re-evaluated at an intermediate time (treating incurred costs until that time as "sunk cost") would lead him to choose a different strategy.

In this chapter we present a single-update framework for adaptive trading strategies in a continuous trading model, where the algorithm adjusts the strategy at some intermediary time in response to the stock price movement in the first part of the trading. We do so by a direct one-step optimization approach. Already these simple strategies improve significantly over the trading trajectories of Almgren and Chriss (2000) with respect to the mean-variance trade-off evaluated at the initial time. As mentioned in Section 1.2, we find that the improvement depends on the size of the transaction and stems from introducing a correlation between the trading gains/losses in the first period and the market impact costs in the second period: If the price moves in the trader's favor in the first part of the trading, the algorithm accelerates the program for the remainder, spending parts of the windfall gains on higher market impact costs to reduce the remaining risk. We obtain a strategy that is "aggressive-in-the-money".

The results in this chapter are to be understood more as a proof-of-concept of adaptive execution strategies. The approach taken here does not directly generalize to multi-update frameworks. In Chapter 3, we will show how a suitable application of the dynamic programming principle can be used to derive a scheme to determine fully optimal dynamic trading strategies.

From a practical point of view, the single-update framework is attractive. As we will see the improvement by a single update is already substantial. In fact, the additional improvement by multiple decision times in Chapter 3 seems to decline (see Section 3.5.5 in the next chapter).

The remainder of this chapter is organized as follows. In Section 2.2, we review the trading and market impact model of Almgren and Chriss (2000) in its continuous-time version. In Section 2.3 we present the optimal static trading trajectories of Almgren and Chriss. In Section 2.4 we discuss single update strategies, i.e. trading strategies that may adjust at some intermediary time T_* in response to the stock price path until T_*. In Section 2.5 we present numerical results for the single update strategies and their improvement over static trajectories. Section 2.6 concludes.

2.2. Continuous-time Market Model

We consider a continuous trading model in a single asset whose price is $S(t)$, obeying a random walk. Instead of the more traditional model of geometric Brownian motion, we will use an arithmetic model. Since our interest is in short-term trading (typically less than one day), the difference between arithmetic and geometric Brownian motion is negligible and the arithmetic process is much more convenient. In particular, the arithmetic process has the property that the expected size of future price changes, as absolute dollar quantities, does not depend on past price changes or the starting price level[1].

Thus, we model the stock price as

$$S(t) \;=\; S_0 \;+\; \sigma\, B(t) \tag{2.1}$$

where σ is an absolute volatility and $B(t)$ is a standard one-dimensional Brownian motion on a filtered probability space $(\Omega, \mathcal{F}, P)$ with filtration $\{\mathcal{F}_t\}$, which satisfies the usual conditions, i.e. $(\Omega, \mathcal{F}, P)$ is complete, $\mathcal{F}_0$ contains all P-null sets in $\mathcal{F}$, and $\mathcal{F}_t$ is right-continuous (see for instance Yong and Zhou (1999, p. 63)).

This process has neither momentum nor mean reversion: future price changes are completely independent of past changes. The Brownian motion $B(t)$ is the only source of randomness in the problem. The reader is referred to standard textbooks for background information on stochastic processes and stochastic calculus, e.g. Oksendal (2003).

The trader has an order of X shares, which begins at time $t = 0$ and must be completed by time $t = T < \infty$. We shall suppose $X > 0$ and interpret this as a sell order. A *trading trajectory* is a function $x(t)$ with $x(0) = X$ and $x(T) = 0$, representing the number of shares remaining to sell at time t. For simplicity, we shall always allow non-integral stock holdings, i.e. $X \in \mathbb{R}^+$ and $x(t)$ a real-valued function. For a static trajectory, $x(t)$ is determined at $t = 0$, but in general $x(t)$ may be any non-anticipating random functional of B. The *trading rate* is $v(t) = -dx/dt$, which will generally be positive as $x(t)$ decreases to zero.

As in Almgren and Chriss (2000), the exogenous evolution of the stock price (2.1) is modified to incorporate *market impact* caused by our trading. Market impact can be thought of the incentive that must be provided to attract liquidity. Two types of market impact are distinguished, both causing adverse price movement that scales with the volume traded by our strategy. *Temporary* impact is a short-lived disturbance in price followed by a rapid reversion as market liquidity returns, whereas the price movement by *permanent* impact stays at least until we finish our sell program.

The aim of this chapter is to give a first proof-of-concept that significant improvements over path-independent strategies are possible by adaptive trading, even already by updating the strategy only once at some intermediary time. Thus, for simplicity we confine ourselves here to *linear temporary market impact* (and no permanent impact). In the next chapter, where we present a approximation scheme to obtain fully optimal adaptive strategies, we will discuss extensions to nonlinear impact functions as well.

[1]Certainly, even on an intraday-timescale the arithmetic process has the disadvantage that there is a small, yet positive probability of the stock price becoming negative, and we neglect this effect.

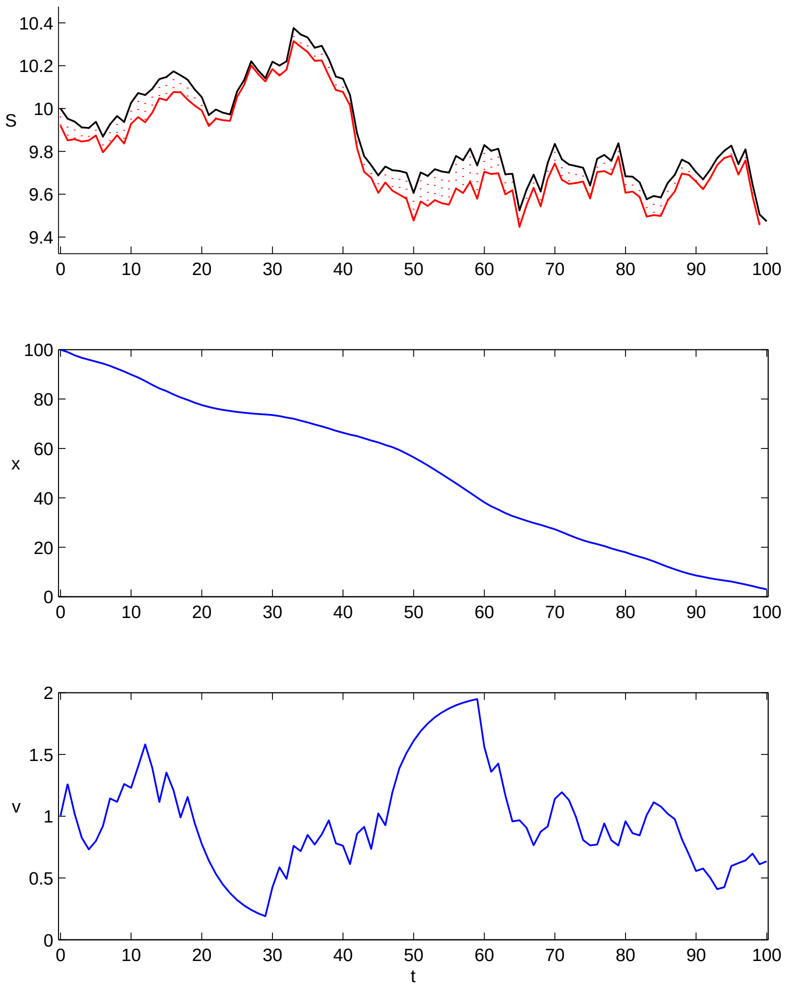

FIGURE 2.1. Illustration of temporary market impact for a sell program. In the upper panel, the black line is the exogenous price process $S(t)$. The red line is the resulting effective price $\tilde{S}(t)$ paid by the trading strategy $x(t)$ shown in the panel in the middle. The lower panel shows $v(t) = -dx/dt$ for this strategy.

Following the model of Almgren and Chriss (2000), with a linear market impact function the actual execution price at time t is

$$\tilde{S}(t) \;=\; S(t) \;-\; \eta\, v(t) \tag{2.2}$$

where $\eta > 0$ is the coefficient of temporary market impact, and $v(t)$ is the rate of trading at time t as defined above. Since we consider a sell program, $\tilde{S}(t) < S(t)$ is indeed less favorable than $S(t)$. Figure 2.1 illustrates how the exogenous price process is modified by temporary impact.

The *implementation shortfall* C is the total cost of executing the sell program relative to the initial value, i.e. the difference between the initial market value XS_0 and the total amount captured. We have the following Lemma.

LEMMA 2.1. *For a trading policy $x(t)$ the implementation shortfall is*

$$C \;= \eta \int_0^T v(t)^2\, dt - \sigma \int_0^T x(t)\, dB(t)\;\;, \tag{2.3}$$

with $v(t) = -dx/dt$.

PROOF. By the definition of the implementation shortfall C and because of (2.2),

$$\begin{aligned} C &= X\,S_0 - \int_0^T \tilde{S}(t)\,v(t)\,dt \\ &= X\,S_0 - \int_0^T S(t)\,v(t)\,dt \; + \eta \int_0^T v(t)^2\,dt \;\; . \end{aligned}$$

Integration by parts for the first integral yields

$$\begin{aligned} C &= X\,S_0 \; - \; \big[S(t)x(t)\big]_0^T - \sigma \int_0^T x(t)\,dB(t) + \eta \int_0^T v(t)^2\,dt \\ &= \eta \int_0^T v(t)^2\,dt - \sigma \int_0^T x(t)\,dB(t) \;\; . \end{aligned}$$

□

The implementation shortfall C consists of two parts. The first term represents the market impact cost. The second term represents the trading gains or losses: since we are selling, a positive price motion gives negative cost.

Note that C is in fact independent of the initial stock price S_0, and only depends on the dynamics of $S(t)$. Hence, optimal trading strategies will also be independent of S_0.

As C is a random variable, an "optimal" strategy will seek some risk-reward balance. Almgren and Chriss (2000) employ the well-known mean-variance framework: a strategy is called *efficient* if it minimizes the variance of a specified maximum level of expected cost or conversely.

The set of all efficient strategies is summarized in the *efficient frontier of optimal trading*, introduced by Almgren and Chriss (2000) in the style of the well-known Markowitz efficient frontier in portfolio theory.

2.3. Optimal Path-Independent Trajectories

If $x(t)$ is fixed independently of $B(t)$, then C is a Gaussian random variable with mean and variance

$$E \;=\; \eta \int_0^T v(t)^2\,dt \qquad \text{and} \qquad V \;=\; \sigma^2 \int_0^T x(t)^2\,dt \;\; . \tag{2.4}$$

Using (2.4), Almgren and Chriss (2000) explicitly give the family of efficient static trading trajectories:

THEOREM 2.2 (Almgren and Chriss (2000)). *The family of efficient path-independent ("static") trade trajectories is given by*

$$x(t) = X\,h(t,T,\kappa) \qquad \textit{with} \quad h(t,T,\kappa) = \frac{\sinh\big(\kappa(T-t)\big)}{\sinh\big(\kappa T\big)} \;\; , \tag{2.5}$$

parametrized by the static "urgency" parameter $\kappa \geq 0$.

The units of κ are inverse time, and $1/\kappa$ is a desired time scale for liquidation, the "half-life" of Almgren and Chriss (2000). The static trajectory is effectively an exponential $\exp(-\kappa t)$ with adjustments to reach $x = 0$ at $t = T$. For fixed κ, the optimal time scale is independent of portfolio size X since both expected costs and variance scale as X^2.

Almgren and Chriss determine these efficient trade schedules by optimizing

$$\min_{x(t)} \; E + \lambda V \tag{2.6}$$

for each $\lambda \geq 0$, where $E = \mathbb{E}[C]$ and $V = \mathrm{Var}[C]$ are the expected value and variance of C. If restricting to path-independent strategies, the solution of (2.6) is obtained by solving an optimization problem for the deterministic trade schedule $x(t)$. For given $\lambda \geq 0$, the solution of (2.6) is (2.5) with $\kappa = \sqrt{\lambda\sigma^2/\eta}$. By taking $\kappa \to 0$, we recover the linear profile $x(t) = X(T-t)/T$. This profile has expected cost and variance

$$E_{\text{lin}} = \eta X^2/T \quad \text{and} \quad V_{\text{lin}} = \sigma^2 X^2 T/3 \ . \tag{2.7}$$

2.4. Single Update Strategies

We now consider strategies that respond to the stock price movement exactly once, at some intermediary time T_* with $0 < T_* < T$. On the first trading period $0 \leq t \leq T_*$, we follow a static trading trajectory with initial urgency κ_0; that is, the trajectory is $x(t) = X\,h(t, T, \kappa_0)$ with h from (2.5). Let $X_*(\kappa_0, T_*) = X\,h(T_*, T, \kappa_0)$ be the remaining shares at time T_*. At this time, we switch to a static trading trajectory with one of n new urgencies $\kappa_1, \ldots, \kappa_n$: with urgency κ_i, we set $x(t) = X_*(\kappa_0, T_*)\,h(t - T_*, T - T_*, \kappa_i)$ for $T_* \leq t \leq T$. We choose the new urgency based on the realized cost up until T_*,

$$C_0 \;=\; \eta \int_0^{T_*} v(t)^2\,dt - \;\sigma \int_0^{T_*} x(t)\,dB(t) \ . \tag{2.8}$$

To that end, we partition the real line into n intervals $I_1, \ldots, I_n$, defined by $I_j = \{b_{j-1} < C_0 < b_j\}$ with $b_j = \mathbb{E}[C_0] + a_j\sqrt{\mathrm{Var}[C_0]}$ and $a_1, \ldots, a_{n-1}$ fixed constants ($a_0 = -\infty$, $a_n = \infty$). Then, on the second period we use κ_j if $C_0 \in I_j$.

Thus, a single-update strategy is defined by the parameters $T_*, a_1, \ldots, a_{n-1}$ and $\kappa_0, \ldots, \kappa_n$. The overall trajectory is summarized by

$$x(t) = \begin{cases} X h(t, T, \kappa_0) & \text{for } 0 \leq t \leq T_* \\ x(T_*)\,h(t - T_*, T - T_*, \kappa_i) & \text{for } T_* < t \leq T, \text{ if } C_0 \in I_i \ . \end{cases} \tag{2.9}$$

Note that we do not know which trajectory we shall actually execute in the second period until we observe C_0 at time T_*.

Straightforward calculation shows that h from (2.5) satisfies $h(t, T, \kappa) = h(s, T, \kappa)h(t-s, T-s, \kappa)$ for $0 \leq s \leq t \leq T$. Thus, if we choose $\kappa_0 = \kappa_1 = \cdots = \kappa_n =: \kappa$, the single-update strategy will degenerate to the static trajectory $x(t) = h(t, T, \kappa)$. That is, the set of static trajectories from Theorem 2.2 is indeed a subset of the set of single-update strategies.

Let $\phi(\cdot)$ be the standard normal density and $\Phi(\cdot)$ its cumulative. We need to prove the following Lemma about the conditional expectation of a normally distributed random variable.

Lemma 2.3. *Let $X \sim \mathcal{N}(\mu, \sigma^2)$ and $-\infty < a < b < \infty$. Then*

$$\mathbb{E}\left[X \,\middle|\, a \leq \frac{X - \mu}{\sigma} \leq b \right] = \mu + \sigma \frac{\phi(a) - \phi(b)}{\Phi(b) - \Phi(a)} \ . \tag{2.10}$$

PROOF. Let $Z = \frac{X-\mu}{\sigma}$, i.e. $Z \sim \mathcal{N}(0,1)$. Recall $\phi(x) = \exp(-x^2/2)/\sqrt{2\pi}$, and hence

$$\int_a^b z\phi(z)\,dz = \phi(a) - \phi(b) \ .$$

We have

$$\begin{aligned}
\mathbb{E}\left[X\,\mathbb{1}_{a\leq\frac{X-\mu}{\sigma}\leq b}\right] &= \mu\,\mathbb{E}\left[\mathbb{1}_{a\leq Z\leq b}\right] + \sigma\,\mathbb{E}\left[Z\,\mathbb{1}_{a\leq Z\leq b}\right] \\
&= \mu\,\mathbb{P}\left[a\leq Z\leq b\right] + \sigma\int_a^b z\phi(z)\,dz \\
&= \mu\big(\Phi(b)-\Phi(a)\big) + \sigma\big(\phi(a)-\phi(b)\big) \ .
\end{aligned}$$

Since

$$\mathbb{E}\left[X\,\middle|\,a\leq\frac{X-\mu}{\sigma}\leq b\right]\cdot\mathbb{P}\left[a\leq\frac{X-\mu}{\sigma}\leq b\right] = \mathbb{E}\left[X\,\mathbb{1}_{a\leq\frac{X-\mu}{\sigma}\leq b}\right] \ ,$$

equation (2.10) follows. □

We will now show that single-update strategies can actually significantly improve over path-independent strategies. The magnitude of the improvement will depend on a single market parameter, the *"market power"*

$$\mu = \frac{\eta X/T}{\sigma\sqrt{T}} \ . \tag{2.11}$$

Here the numerator is the price concession for trading at a constant rate, and the denominator is the typical size of price motion due to volatility over the same period. The ratio μ is a nondimensional preference-free measure of portfolio size, in terms of its ability to move the market.

The following Theorem gives mean and variance of a single-update strategy as a function of the strategy parameters. Below we will use these expressions to determine optimal single-update strategies, and demonstrate that they indeed improve over the static trajectories of Almgren and Chriss (2000).

THEOREM 2.4. *Let $\mu > 0$, and π be a single-update strategy given by the set of parameters $(T_*, a_1, \ldots, a_{n-1}, \kappa_0, \kappa_1, \ldots, \kappa_n)$. Then the mean $E = \mathbb{E}[C]$ and variance $V = \mathrm{Var}[C]$ of π, scaled by the respective values of the linear strategy (2.7), are given by*

$$\hat{E} = E/E_{lin} = E_0 + \overline{E} \tag{2.12}$$

and

$$\hat{V} = V/V_{lin} = V_0 + \overline{V} + 2\sqrt{3}\mu\sqrt{V_0}\sum_{i=1}^{n} q_i E_i + 3\mu^2\sum_{i=1}^{n} p_i\big(E_i - \overline{E}\big)^2 \ , \tag{2.13}$$

with $p_j = \Phi(a_j) - \Phi(a_{j-1})$, $q_j = \phi(a_{j-1}) - \phi(a_j)$, $\overline{E} = \sum_{i=1}^n p_i E_i$, $\overline{V} = \sum_{i=1}^n p_i V_i$ *and*

$$E_0 = \frac{\kappa_0 T\big(\sinh(2\kappa_0 T) - \sinh(2\kappa_0(T-T_*)) + 2\kappa_0 T_*\big)}{4\sinh^2(\kappa_0 T)} , \tag{2.14}$$

$$V_0 = \frac{3\big(\sinh(2\kappa_0 T) - \sinh(2\kappa_0(T-T_*)) - 2\kappa_0 T_*\big)}{4\kappa_0 T \sinh^2(\kappa_0 T)} , \tag{2.15}$$

$$E_i = \frac{\sinh^2(\kappa_0(T-T_*))}{\sinh^2(\kappa_i(T-T_*))} \frac{\kappa_i T\big(\sinh(2\kappa_i(T-T_*)) + 2\kappa_i(T-T_*)\big)}{4\sinh^2(\kappa_0 T)} , \tag{2.16}$$

$$V_i = \frac{3\sinh^2(\kappa_0(T-T_*))}{\sinh^2(\kappa_i(T-T_*))} \frac{\sinh(2\kappa_i(T-T_*)) - 2\kappa_i(T-T_*)}{4\kappa_i T \sinh^2(\kappa_0 T)} . \tag{2.17}$$

PROOF. Recall that C_0, defined in (2.8), is the realized cost on the first period. Furthermore, we denote by C_j ($j = 1, \dots, n$) the cost incurred on the second part of the trajectory, if urgency κ_j is used,

$$C_j = \eta \int_{T_*}^{T} v(t)^2\, dt - \sigma \int_{T_*}^{T} x(T_*)\, h(t - T_*, T - T_*, \kappa_j)\, dB(t) . \tag{2.18}$$

Each variable C_i ($i = 0, \dots, n$) is Gaussian. By (2.8), (2.18) and the definition of the single-update strategy (2.9), straightforward integration yields

$$\mathbb{E}[C_0] = \eta X^2 \int_0^{T_*} -h'(t, T, \kappa_0)^2\, dt = E_0 \cdot E_{lin} \tag{2.19}$$

$$\mathrm{Var}[C_0] = \sigma^2 X^2 \int_0^{T_*} h(t, T, \kappa_0)^2\, dt = V_0 \cdot V_{lin} \tag{2.20}$$

and for $i = 1, \dots, n$

$$\mathbb{E}[C_i] = \eta X_*(\kappa_0, T_*)^2 \int_{T_*}^{T} -h'(t - T_*, T - T_*, \kappa_i)^2\, dt = E_i \cdot E_{lin} \tag{2.21}$$

$$\mathrm{Var}[C_i] = \sigma^2 X_*(\kappa_0, T_*)^2 \int_{T_*}^{T} h(t - T_*, T - T_*, \kappa_i)^2\, dt = V_i \cdot V_{lin} \tag{2.22}$$

(2.23)

where E_0, V_0, E_i, V_i are defined by (2.14)–(2.17). E_{lin}, V_{lin} denote the expected cost and variance of the linear strategy given by (2.7), and $X_*(\kappa_0, T_*) = x(T_*) = X\, h(T_*, T, \kappa_0)$.

The total cost is

$$C = C_0 + C_{\jmath(C_0)} \tag{2.24}$$

where $\jmath(C_0) = i$ if and only if $C_0 \in I_i$.

With the fixed nondimensional quantities

$$p_j = \Phi(a_j) - \Phi(a_{j-1}) \qquad \text{and} \qquad q_j = \phi(a_{j-1}) - \phi(a_j) ,$$

for $j = 1, \dots, n$, we have $\mathbb{P}[C_0 \in I_j] = p_j$ and by Lemma 2.3,

$$\mathbb{E}\big[C_0 \big| C_0 \in I_j\big] = E_0 E_{lin} + (q_j/p_j)\sqrt{V_0 V_{lin}} .$$

By linearity of expectation, we readily get $E = \mathbb{E}[C] = E_{lin}\big(E_0 + \overline{E}\big)$ with $\overline{E} = \sum_{i=1}^n p_i E_i$. The variance is more complicated because of the dependence of the two terms in (2.24). We use the conditional variance formula $\mathrm{Var}[X] = \mathbb{E}[\mathrm{Var}[X|Y]] + \mathrm{Var}[\mathbb{E}[X|Y]]$ to write, with $\overline{V} = \sum_{i=1}^n p_i V_i$,

$$\begin{aligned}\mathrm{Var}[C] &= \mathbb{E}\left[\mathrm{Var}\left[\, C_0 + C_{\jmath(C_0)} \,\middle|\, C_0 \,\right]\right] \;+\; \mathrm{Var}\left[\mathbb{E}\left[\, C_0 + C_{\jmath(C_0)} \,\middle|\, C_0 \,\right]\right] \\ &= \mathbb{E}\left[V_{\jmath(C_0)} V_{lin}\right] + \;\mathrm{Var}\left[C_0 + E_{\jmath(C_0)} E_{lin}\right] \\ &= \overline{V} V_{lin} \;+\; \mathrm{Var}[C_0] \;+\; 2\, E_{lin}\, \mathrm{Cov}\left[C_0, E_{\jmath(C_0)}\right] \;+\; E_{lin}^2\, \mathrm{Var}\left[E_{\jmath(C_0)}\right] \;\;.\end{aligned}$$

By definition, $\mathrm{Var}[C_0] = V_0 V_{lin}$, and $\mathrm{Var}\left[E_{\jmath(C_0)}\right] = \sum_{i=1}^n p_i \big(E_i - \overline{E}\big)^2$. Also,

$$\begin{aligned}\mathrm{Cov}\left[C_0, E_{\jmath(C_0)}\right] &= \mathbb{E}\left[\, C_0\, E_{\jmath(C_0)} \,\right] \;-\; \mathbb{E}[C_0]\, \mathbb{E}\left[E_{\jmath(C_0)}\right] \\ &= \sum_{i=1}^n \mathbb{P}[C_0 \in I_i]\, \mathbb{E}\left[C_0 E_{\jmath(C_0)} \middle| C_0 \in I_i\right] - \mathbb{E}[C_0]\, \mathbb{E}\left[E_{\jmath(C_0)}\right] \\ &= \sum_{i=1}^n p_i E_i E_{lin}\, \mathbb{E}\left[C_0 \middle| C_0 \in I_i\right] \;-\; \mathbb{E}[C_0]\; \mathbb{E}\left[E_{\jmath(C_0)}\right] \\ &= \sum_{i=1}^n p_i E_i E_{lin} E_0 E_{lin} + \sum_{i=1}^n q_i E_i E_{lin} \sqrt{V_0 V_{lin}} \;-\; \mathbb{E}[C_0]\; \mathbb{E}\left[E_{\jmath(C_0)}\right] \\ &= \mathbb{E}[C_0] \sum_{i=1}^n p_i E_i E_{lin} + \sqrt{V_0 V_{lin}} E_{lin} \sum_{i=1}^n q_i E_i \;-\; \mathbb{E}[C_0]\; \mathbb{E}\left[E_{\jmath(C_0)}\right] \\ &= \sqrt{V_0 V_{lin}} E_{lin} \sum_{i=1}^n q_i E_i \;\;.\end{aligned}$$

Putting all this together, we have

$$V \;=\; \mathrm{Var}[C] \;=\; V_0 V_{lin} + \overline{V} V_{lin} \;+\; 2\sqrt{V_0 V_{lin}} E_{lin} \sum_{i=1}^n q_i E_i \;+\; E_{lin}^2 \sum_{i=1}^n p_i \big(E_i - \overline{E}\big)^2 \;\;.$$

Since $E_{lin}^2 / V_{lin} = 3\mu^2$ and $\sqrt{V_{lin}} E_{lin} / V_{lin} = \sqrt{3}\mu$, (2.13) follows. □

As can be seen, scaling expectation and variance of the total cost C by their respective values of the linear strategy yields expressions E/E_{lin} and V/V_{lin} as functions of the strategy parameters a_i, κ_i, and the single market parameter μ. That is, we reduce the four dimensional parameters X, T, σ and η to one nondimensional (i.e. scalar) parameter, the market power μ.

2.4.1. Density of Total Cost C for Single-update Strategies. In the following, to shorten notation we denote $\tilde{E}_i = \mathbb{E}[C_i] = E_i \cdot E_{lin}$ and $\tilde{V}_i = \mathrm{Var}[C_i] = V_i \cdot V_{lin}$ $(i = 0 \dots n)$ defined by (2.19, 2.20, 2.21, 2.22). Each C_i is Gaussian with mean $\tilde{E}_i$ and variance $\tilde{V}_i$, so its density is $f_i(c_i) = 1/\sqrt{2\pi \tilde{V}_i}\, \exp\Big(-(C_i - \tilde{E}_i)^2 / 2\tilde{V}_i\Big)$. Let $f(c)$ be the density of the total cost $C = C_0 + C_{\jmath(C_0)}$, see (2.24). That is, the total cost C is a certain mixture of Gaussians. We have

$$\begin{aligned}f(c) &= \sum_{i=1}^n \int_{I_i} f_0(c_0)\, f_i(c - c_0)\, dc_0 \\ &= \sum_{i=1}^n \frac{1}{2\pi \sqrt{\tilde{V}_0 \tilde{V}_i}} \int_{b_{i-1}}^{b_i} \exp\left(-\frac{(c_0 - \tilde{E}_0)^2}{2\tilde{V}_0} - \frac{(c - c_0 - \tilde{E}_i)^2}{2\tilde{V}_i} \right) dc_0 \;\;,\end{aligned}$$

i.e.

$$
\begin{aligned}
f(c) = \sum_{i=1}^{n} \frac{1}{2\pi\sqrt{\tilde{V}_0\tilde{V}_i}} \exp\Bigg(-\frac{1}{2}\Bigg[\frac{\tilde{E}_0^2}{\tilde{V}_0} + \frac{(c-\tilde{E}_i)^2}{\tilde{V}_i} - \frac{\big(\tilde{E}_0\tilde{V}_i + (c-\tilde{E}_i)\tilde{V}_0\big)^2}{\tilde{V}_0\tilde{V}_i(\tilde{V}_0+\tilde{V}_i)} \Bigg]\Bigg) \\
\times \int_{b_{i-1}}^{b_i} \exp\Bigg(-\frac{1}{2}\Bigg[\frac{\tilde{V}_0+\tilde{V}_i}{\tilde{V}_0\tilde{V}_i}\bigg(c_0 - \frac{\tilde{E}_0\tilde{V}_i + (c-\tilde{E}_i)\tilde{V}_0}{\tilde{V}_0+\tilde{V}_i}\bigg)^2\Bigg]\Bigg)\, dc_0 \\
= \sum_{i=1}^{n} \frac{1}{\sqrt{2\pi(\tilde{V}_0+\tilde{V}_i)}} \exp\bigg(-\frac{\big(c-\tilde{E}_0-\tilde{E}_i\big)^2}{2\,(\tilde{V}_0+\tilde{V}_i)}\bigg) \\
\times \Bigg[\Phi\Bigg(\frac{(c-\tilde{E}_i-b_{i-1})\tilde{V}_0 + (\tilde{E}_0-b_{i-1})\tilde{V}_i}{\sqrt{\tilde{V}_0\tilde{V}_i(\tilde{V}_0+\tilde{V}_i)}}\Bigg) \\
- \Phi\Bigg(\frac{(c-\tilde{E}_i-b_i)\tilde{V}_0 + (\tilde{E}_0-b_i)\tilde{V}_i}{\sqrt{\tilde{V}_0\tilde{V}_i(\tilde{V}_0+\tilde{V}_i)}}\Bigg)\Bigg] .
\end{aligned}
$$

2.4.2. Optimal Single-update Strategies. Let us fix μ, T_*, n and the interval breakpoints, given by $a_1, \ldots, a_{n-1}$. We want to compute optimal single-update strategies for the mean-variance trade-off criterion

$$\min_{\kappa_0,\ldots,\kappa_n} \hat{E} + \lambda\hat{V} \tag{2.25}$$

with risk-aversion $\lambda \geq 0$, where $\hat{E}$ and $\hat{V}$ are from (2.12–2.13). We assume that similar to (2.6), solving (2.25) for all $\lambda \geq 0$ traces out the efficient frontier of single-update strategies, but we leave this as an open question.

Let us fix $T_*, a_1, \ldots, a_{n-1}$. The solutions depend on the market power parameter μ. Because of (2.11), $\mu \to 0$ corresponds to portfolio transactions for a small initial position X. Thus, we refer to $\mu \to 0$ as the *small portfolio limit.*

For the small portfolio limit, $\mu \to 0$, (2.12, 2.13) reduce to $\hat{E} = E_0 + \overline{E}$ and $\hat{V} = V_0 + \overline{V}$, and the optimal solution to (2.25) becomes $\kappa_0 = \kappa_1 = \cdots = \kappa_n$. Thus, for $\mu \to 0$, the static trajectories from Theorem 2.2 are still optimal in the single-update framework, and the adaptive efficient frontier coincides with the static efficient frontier (see Figure 2.2). We will also observe this property in the multi-update framework in Chapter 3, where we will further elaborate and give a more formal proof.

Unfortunately, there is no closed-form solution to (2.25). Thus, for fixed $T_*, a_1, \ldots, a_{n-1}$, given market power $\mu > 0$ and $\lambda \geq 0$, we minimize (2.25) numerically over the urgencies $\kappa_0, \kappa_1, \ldots, \kappa_n$.

The reason single-update strategies can improve over the static trajectories is the $\mathcal{O}(\mu)$ term in (2.13), $\mathrm{Cov}\left[C_0, E_{\jmath(C_0)}\right]$, which can be made negative by making E_i negatively related to C_0. This corresponds to anticorrelation between second-period impact costs and first-period trading gains/losses. Increasing second-period impact costs after a first-period windfall trading gain (i.e. lower than expected cost) means trading faster in the second period. Thus, the strategy is "aggressive-in-the-money" (AIM) as mentioned in the introduction.

2.5. Numerical Results

In Figure 2.2 we plot E and V relative to E_{lin} and V_{lin}; each curve is computed by varying λ in (2.25) from 0 to ∞, for a fixed value of μ. The solution for each pair (λ, μ) is computed using $n = 2$, and $a_1 = 0$ ($a_0 = -\infty, a_2 = \infty$). That is, we consider single-update strategies which switch to one of two trajectories, depending on whether the realized cost in the first part was higher or lower than its expected value. We fix T_* to the half-life of the corresponding static trajectory. The plot also shows the frontier corresponding to the static trajectories from Theorem 2.2 (solid curve). As can be seen, this frontier indeed corresponds to the limit $\mu \to 0$.

We use these frontiers to obtain cost distributions for adaptive strategies that are better than the cost distributions for any static strategy. In Figure 2.2, the point marked on the static frontier (solid curve) corresponds to a static trajectory computed with parameter $\kappa = 14$, giving a Gaussian cost distribution. For a portfolio with $\mu = 0.15$, this distribution has expectation $E \approx 7.0 \times E_{lin}$ and variance $V \approx 0.11 \times V_{lin}$. The inset shows this distribution as a black dashed line. The inset also shows the cost distributions associated with the adaptive strategies corresponding to the two marks on the frontier for $\mu = 0.15$ (outermost dashed curve): one of them has the same mean, but lower variance ($V \approx 0.06 \times V_{lin}$); the other one has same variance, but lower mean ($E \approx 4.47 \times E_{lin}$). These distributions are the extreme points of a one-parameter family of distributions along the adaptive frontier, each of which is strictly preferable to the given static strategy for a mean-variance investor. The distribution plots of the adaptive strategies are generated from the density formula derived in Section 2.4.1. These cost distributions are slightly skewed toward positive costs. But as they are not too far from Gaussian, we can still expect mean-variance to gives reasonable results (Recall that for Gaussian random variables, mean-variance optimization is consistent with expected utility maximization and stochastic dominance; see for instance Levy (1992); Bertsimas et al. (2004)).

Figure 2.3 shows an example of a single-update trading strategy for $\mu = 0.15$. The dashed line is the static optimal trajectory with urgency $\kappa = 14$. We choose the efficient adaptive strategy, which achieves the same variance as this static trajectory, but has lower expected cost. The adaptive strategy initially trades more slowly than the optimal static trajectory. At T_*, if prices have moved in the trader's favor and the realized cost is lower than the expected value (conditional on $t = 0$), then the strategy accelerates, spending the investment gains on impact costs. If prices have moved against the trader, corresponding to higher than expected values of C_0, then the strategy decelerates to save impact costs in the remaining period. Thus, the single-update strategy is indeed "aggressive-in-the money" (AIM).

2.6. Conclusion

The single-update strategies presented in this chapter demonstrate that price adaptive scaling strategies can lead to significant improvements over static trade schedules (Almgren and Chriss, 2000) for the execution of portfolio transactions, and they illustrate the importance of the new "market power" parameter μ. Unfortunately, the single-update framework presented in this chapter does not directly generalize to multi-decision frameworks.

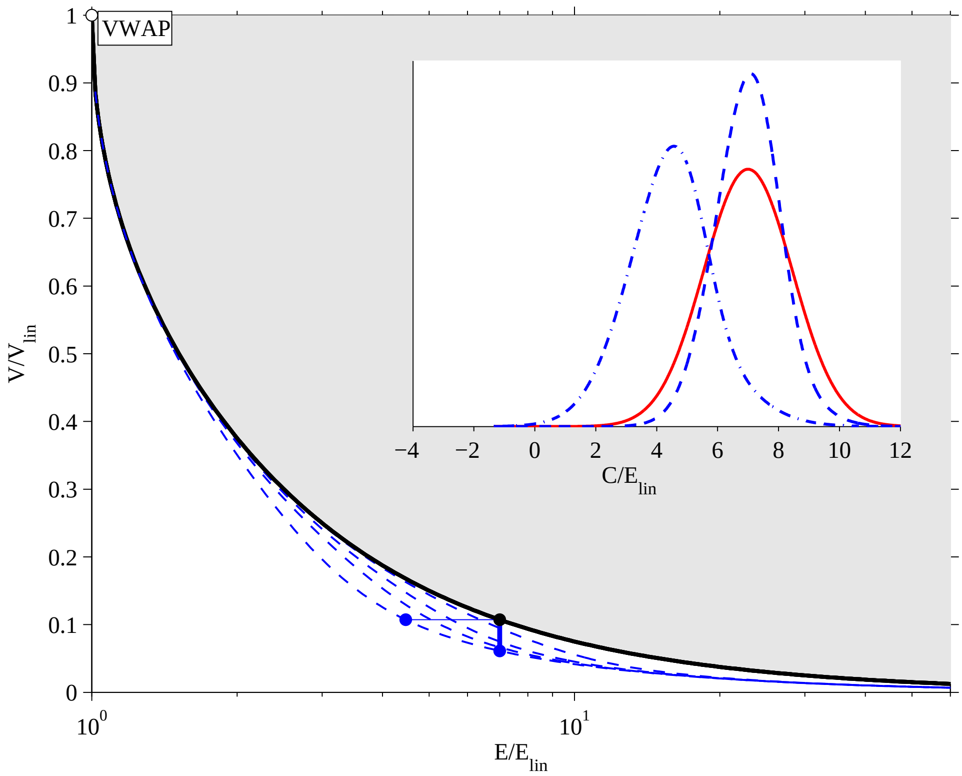

FIGURE 2.2. Adaptive efficient frontiers for different values of market power μ. The expectation of trading cost $E = \mathbb{E}[C]$ and its variance $V = \text{Var}[C]$ are normalized by their values for a linear trajectory (VWAP). The grey shaded region is the set of values accessible to a static trajectory and the black solid curve is the static frontier, which is also the limit $\mu \to 0$. The blue dashed curves are the improved values accessible to adaptive strategies, $\mu \in \{0.025, 0.05, 0.075, 0.15\}$; the improvement is greater for larger portfolios. The inset shows the actual distributions corresponding to the indicated points.

In the next chapter, we will show how a suitable application of the dynamic programming principle can be used to derive a scheme to determine fully dynamic, optimal trading strategies for a discrete version of the market impact model introduced in Section 2.2. Instead of optimizing the tradeoff function $\mathbb{E}[C] + \lambda \text{Var}[C]$, we will work with the constrained version of mean-variance optimization, $\min \mathbb{E}[C]$ s.t. $\text{Var}[C] \leq V_*$.

As mentioned in the introduction of this chapter, from a practical point of view the single-update framework is attractive as it requires only rather straightforward numerical calculation. As we will see in Section 3.5.5 of the next chapter, the single update strategies achieve improvements that are already comparable to the fully dynamic trading strategies which we derive there.

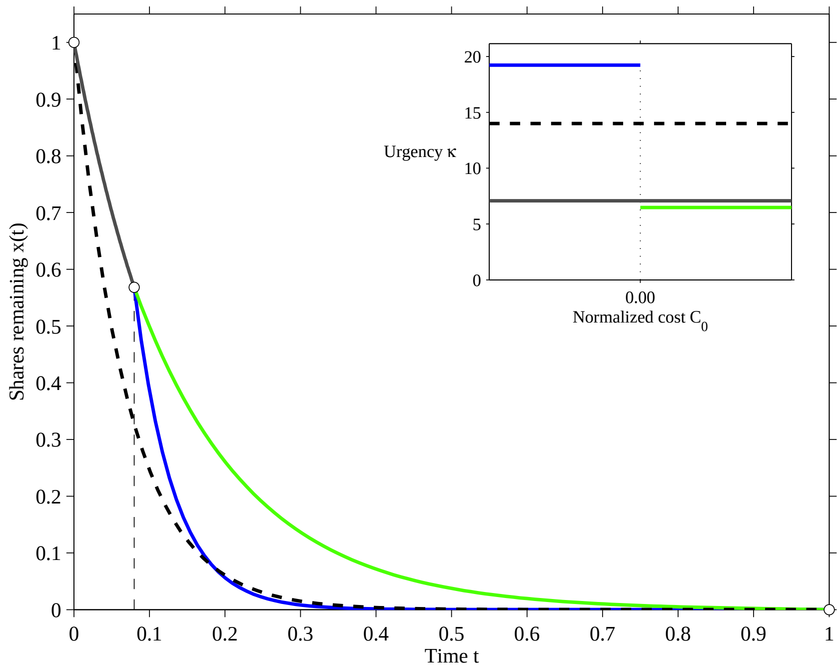

FIGURE 2.3. Example of an adaptive trading trajectories for market power $\mu = 0.15$. The dashed line is a static optimal trajectory with urgency $\kappa = 14$; the adaptive strategy with $n = 2$ and $a_1 = 0$ is chosen such that it achieves the same variance as this static trajectory, but lower expected cost. The inset shows the dependence of the new urgency on the initial trading cost C_0, normalized by the *ex ante* expectation and standard deviation of C_0.

CHAPTER 3

Optimal Adaptive Trading with Market Impact

In this chapter we further discuss the optimal dynamic execution of portfolio transactions in the market impact model of Almgren and Chriss (2000). Improving over the single update strategies presented in the previous chapter, in this chapter we show how a suitable application of the dynamic programming principle can be used to derive a scheme to determine fully optimal dynamic trading strategies for the execution problem in discrete time. Our adaptive strategies significantly improve over the static arrival price algorithms of Almgren and Chriss (2000) with respect to the mean-variance trade-off evaluated at the initial time.

3.1. Introduction

In this chapter we further study adaptive trading strategies for the market impact model introduced in Section 1.2 with the implementation shortfall as benchmark, i.e. the difference between the average execution price and the pre-trade price. In the previous chapter we showed that adaptive trading strategies can indeed significantly improve over the arrival price algorithms of Almgren and Chriss (2000) with respect to the mean-variance trade-off evaluated at the initial time. However, the single-update strategies considered there do not generalize to multiple updates or fully optimal adaptive trading.

As mentioned in Section 1.2, a related problem is mean-variance portfolio optimization in a multiperiod setting: executing a buy order can be seen as an optimal portfolio strategy in one risky asset (the stock to be purchased) and one riskless asset (cash) under transaction costs and with the additional terminal constraint that the entire wealth is in the stock at the end of the buy progam at $t = T$. Only recently Li and Ng (2000) gave a closed-form solution for the mean-variance problem in a discrete-time multiperiod setting. They use an embedding technique of the original problem into a family of tractable auxiliary problems with a utility function of quadratic type. While in principle this approach may be applied to the portfolio transaction problem as well, the market impact terms together with the need to introduce a second parameter in the auxiliary problem (next to the risk aversion) significantly complicate the problem.

We therefore follow a different approach to determine optimal trading strategies with respect to the specification of measuring risk and reward at the initial time. Using a dynamic programming principle for mean-variance optimization, we determine fully optimal Markovian trading strategies for the arrival price problem in a discrete time setting. The key idea is to use the variance of total cost as the value function, and consider the targeted expected return as a state variable. We give an efficient scheme to obtain numerical solutions by solving a series of convex constrained optimization problems.

We observe the same qualitative behavior as in the previous chapter. First, the improvement through adaptivity is larger for large transactions, expressed in terms of the market power μ, (2.11); for small portfolios, $\mu \to 0$, optimal adaptive trade schedules coincide with optimal static trade schedules. Second, the improvement through a dynamic strategy comes from introducing a correlation between the trading gains or losses and market impact costs incurred in the remainder. If the price moves in the trader's favor in the early part of the trading, then the algorithm spends parts of those gains on market impact costs by accelerating the remainder of the program. If the price moves against the trader, then the algorithm reduces future market impact costs by trading more slowly. Thus, our new optimal trade schedules are "aggressive-in-the-money" (AIM). A "passive-in-the-money" (PIM) strategy would react oppositely (Kissell and Malamut, 2005).

One important reason for using a AIM or PIM strategy would be the expectation of serial correlation in the price process. Our strategies arise in a pure random walk model with no serial correlation, using pure classic mean and variance. This provides an important caveat for our formulation. Our strategy suggests to "cut your gains and let your losses run." If the price process does have any significant momentum, then this strategy can cause much more serious losses than the gains it provides. Thus we do not advocate implementing them in practice before doing extensive empirical tests.

The remainder of this chapter is organized as follows: In Section 3.2 we present the market and trading model. In contrast to Chapter 2, we consider a discrete time setting. In Section 3.3 we review the concept of mean variance efficient strategies in discrete time, the efficient frontier of trading as well as the optimal static trajectories of Almgren and Chriss (2000). In Section 3.4, we show how to construct fully optimal adaptive policies by means of a dynamic programming principle for mean-variance optimization. In Section 3.5 we give numerical results.

3.2. Trading Model

Let us start by reviewing the discrete trading model considered by Almgren and Chriss (2000). We confine ourselves to sell programs in a single security. The definitions and results for a buy program are completely analogous. Furthermore, an extension to multiple securities ("basket trading") is possible.

Suppose we hold a block of X shares of a stock that we want to completely sell by time T. We divide the trading horizon T into N intervals of length $\tau = T/N$, and define discrete times $t_k = k\tau$, $k = 0, \ldots, N$. A trade schedule π is a list of stock holdings $(x_0, x_1, \ldots, x_N)$ where x_k is the number of shares we plan to hold at time t_k; we require $x_0 = X$ and $x_N = 0$. Thus we shall sell $x_0 - x_1$ shares between t_0 and t_1, $x_1 - x_2$ shares between times t_1 and t_2 and so on. We require a pure sell program which may never buy shares, that is $x_0 \geq x_1 \geq \cdots \geq x_N$. Our strategy comprises N sales, and effectively $N - 1$ decision variables $x_1, \ldots, x_{N-1}$. The average rate of trading during the time interval t_{k-1} to t_k is $v_k = (x_{k-1} - x_k)/\tau$. Crucially, this trade list need not be fixed in advance but may depend on the observed price evolution: π must be adapted to the filtration of the price process as described below.

The stock price follows a random walk, modified to incorporate permanent and temporary impact due to our trading. *Temporary* impact is a short-lived disturbance in price followed by a rapid reversion as market liquidity returns, whereas the price movement by *permanent* impact stays at least until we finish our sell program. As in Chapter 2, we employ an arithmetic model for the stock price. For the intraday time scales of interest to us, the difference between an arithmetic and a geometric model is negligible, and an arithmetic model is much simpler mathematically. Thus, we take the price to follow the process

$$S_k = S_{k-1} + \sigma\tau^{1/2}\xi_k - \tau g\Big(\frac{x_{k-1} - x_k}{\tau}\Big) \ , \tag{3.1}$$

for $k = 1, \ldots, N$. Here ξ_k are independent random variables on Ω with an arbitrary distribution, having $\mathbb{E}[\xi_k] = 0$ and $\mathrm{Var}[\xi_k] = 1$. Unless stated otherwise, we shall always assume that Ω is finite, $|\Omega| < \infty$. For example, for a binomial tree model we would set $\Omega = \{\pm 1\}$. $g(v)$ is a permanent impact function and will be described below.

Let $\mathcal{F}_k$ be the σ-algebra generated by $\{\xi_1, \ldots, \xi_k\}$, for $k = 0, \ldots, N$. For $k = 0, \ldots, N-1$, $\mathcal{F}_{k-1}$ is the information available to the investor before he makes his decision x_k, $1 \le k \le N-1$, i.e. x_k must be $\mathcal{F}_{k-1}$ measurable. That is, we allow each trade decision x_k to depend on prices $S_0, \ldots, S_{k-1}$ but not on S_k or ξ_k; the strategy π is adapted to the filtration of the price motion. Thus the first trade decision $x_0 - x_1$ must be nonrandom; the second decision $x_1 - x_2$ may depend on the realization of the first price increment ξ_1, *etc.*

σ is the absolute volatility of the stock per unit time, so $\sigma^2\tau$ is the variance of price change over a single time step, and the variance of price change over the entire trading period is $\sigma^2 T$. The permanent impact $g(v)$ is a measurable, convex and continuous function of the average rate of trading $v = (x_{k-1} - x_k)/\tau \ge 0$ during the interval t_{k-1} to t_k, with $g(v) \ge 0$ for all $v \ge 0$.

Temporary market impact is modeled by considering our trade price to be slightly less favorable than the "market" price S_k. Hence, the effective price per share when selling $x_{k-1} - x_k$ during the interval t_{k-1} to t_k is

$$\tilde{S}_k = S_{k-1} - h\Big(\frac{x_{k-1} - x_k}{\tau}\Big) \ , \tag{3.2}$$

for some measurable, continuous function $h(v) \ge 0$ for all $v \ge 0$. We shall require that $v\,h(v)$ is convex. Unlike permanent impact $g(v)$, the temporary impact effect $h(v)$ does not affect the next market price S_k. Figure 3.1 illustrates how the exogenous price process is modified by temporary market impact.

For a sell program $\pi = (x_0, x_1, \ldots, x_N)$, the capture upon completion is

$$\begin{aligned}\sum_{k=1}^{N}(x_{k-1} - x_k)\,\tilde{S}_k \;&=\; XS_0 \;+\; \sigma\tau^{1/2}\sum_{k=1}^{N}\xi_k x_k \\ &\quad -\; \tau\sum_{k=1}^{N} x_k\, g\Big(\frac{x_{k-1} - x_k}{\tau}\Big) \;-\; \sum_{k=1}^{N}(x_{k-1} - x_k)\,h\Big(\frac{x_{k-1} - x_k}{\tau}\Big) \ .\end{aligned}$$

The first term on the right side is the initial market value of the position. The second term represents trading gains or losses due to volatility; note that this term has strictly zero expected value for any adapted strategy. The third and fourth terms are the losses due to permanent and temporary impact.

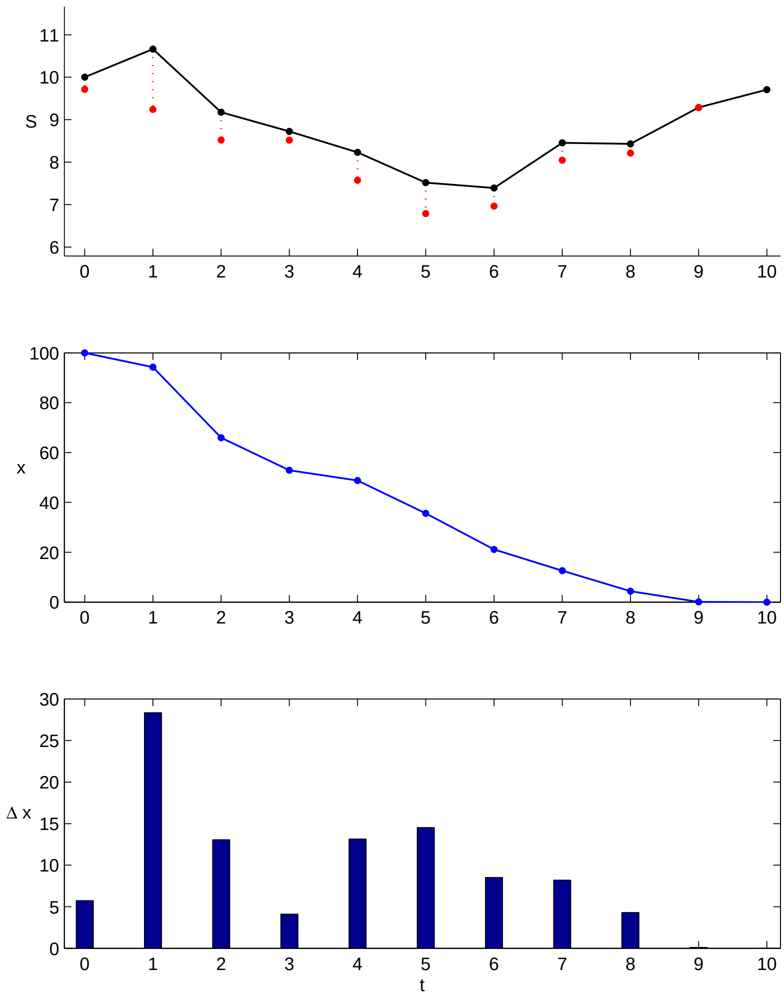

FIGURE 3.1. Illustration of temporary market impact for a sell program in the discrete trading model. In the upper panel, the black line is the exogenous price process S_t. The red line is the resulting effective price $\tilde{S}_t$ paid by the trading strategy x_t shown in the lower panels, assuming linear temporary market impact (and no permanent impact).

The total cost of trading, or *implementation shortfall*, for selling X shares over N periods of length τ by a trading policy $\pi = (x_0, \ldots, x_N)$ is the difference between the initial market value XS_0 and the final capture of the trade,

$$\sum_{k=1}^{N} \left\{ \tau\, g\left(\frac{x_{k-1} - x_k}{\tau}\right) x_k \;+\; (x_{k-1} \;-\; x_k)\, h\left(\frac{x_{k-1} - x_k}{\tau}\right) \;-\; \sigma\tau^{1/2}\xi_k x_k \right\} . \tag{3.3}$$

Because of our assumption of arithmetic random walk, this quantity is independent of the initial stock price S_0.

In the following, we focus on *linear* price impact

$$g(v) = \gamma\, v, \qquad h(v) = \eta v. \tag{3.4}$$

Then, as noted by Almgren and Chriss (2000) it suffices to consider temporary impact only, because with the linear impact functions (3.4) the expression for the cost (3.3) is

$$\sum_{k=1}^{N}\left\{\gamma(x_{k-1}-x_k)x_k+\frac{\eta}{\tau}(x_{k-1}-x_k)^2-\sigma\tau^{1/2}\xi_k x_k\right\} \ . \tag{3.5}$$

Since

$$\frac{1}{2}X^2=\frac{1}{2}\Big(\sum_{k=1}^{N}(x_{k-1}-x_k)\Big)^2=\sum_{k=1}^{N}(x_{k-1}-x_k)x_k+\frac{1}{2}\sum_{k=1}^{N}(x_{k-1}-x_k)^2 \ ,$$

(3.5) becomes

$$\begin{aligned}&\frac{1}{2}\gamma X^2+\sum_{k=1}^{N}\{-\frac{1}{2}\gamma(x_{k-1}-x_k)^2+\frac{\eta}{\tau}(x_{k-1}-x_k)^2-\sigma\tau^{1/2}\xi_k x_k\}\\ =\ &\frac{1}{2}\gamma X^2+\sum_{k=1}^{N}\Big\{\frac{\tilde{\eta}}{\tau}(x_{k-1}-x_k)^2-\sigma\tau^{1/2}\xi_k x_k\Big\}\end{aligned}$$

with $\tilde{\eta}=\eta-\frac{1}{2}\gamma\tau$. Neglecting the constant $\gamma X^2/2$ and replacing $\tilde{\eta}\leftarrow\eta$ (implicitly assuming that τ is small enough that $\tilde{\eta}>0$), we obtain

$$C(X,N,\pi)=\frac{\eta}{\tau}\sum_{k=1}^{N}(x_{k-1}-x_k)^2\ -\ \sigma\tau^{1/2}\sum_{k=1}^{N}\xi_k x_k \ , \tag{3.6}$$

In the remainder of this chapter, we shall always work with (3.6) as the expression for the cost of a N period strategy π to sell X shares.

We assume that volatility, as well as the dependence of permanent impact and on our trade decisions, are not only non-random and known in advance but have a uniform profile. Predictable intraday seasonality can largely be handled by interpreting time t as a "volume time" corresponding to the market's average rate of trading. Random variations in volatility and liquidity are more difficult to model properly (Walia, 2006).

As $\tau\to 0$, that is as $N\to\infty$ with T fixed, this discrete trading model converges to a continuous process: the exogenous price process is an arithmetic Brownian motion, the shares process is an adapted function $x(t)$, and the instantaneous trade rate is $v(t)=-dx/dt$. Hence, the scheme that we will give in the rest of this chapter to determine optimal strategies immediately yields a scheme to determine optimal strategies also for the arrival price problem in a continuous setting as introduced in Chapter 2. The same techniques would also work with nonlinear cost functions, with a drift term added to the price dynamics, or for multi-asset portfolios.

3.3. Efficient Frontier of Optimal Execution

For any trading strategy, the final implementation shortfall $C(X,N,\pi)$ is a random variable: not only do the price motions ξ_k directly affect our trading gains or losses, but for an adapted strategy the trade list itself may be different on each realization, according to the rule π. An "optimal" strategy will determine some balance between minimizing the expected cost and its variance.

Let

$$\mathscr{D}(X,N) = \left\{ (\pi,\overline{C}) \;\middle|\; \begin{array}{l} \pi = (x_0, x_1, \ldots, x_N) \text{ with } x_0 = X,\; x_N = 0 \\ x_0 \geq x_1 \geq \cdots \geq x_N \\ C(X,N,\pi) \leq \overline{C} \;\; a.e. \end{array} \right\} \tag{3.7}$$

be the set of all adapted trading policies that sell X shares in N periods. We require a pure sell program, $x_0 \geq x_1 \geq \cdots \geq x_N$, that is the strategy may never buy shares. π is the trade schedule and must be adapted to the filtration $\mathcal{F}_k$ of the stock price process S_k as described above. $\overline{C}$ is the cost strategy, which is a $\mathcal{F}_N$-measurable random variable that gives the trading cost associated with this policy for each path of the (random) stock price process. For each path of the stock price process, $\overline{C}$ must be *at least* the cost (3.6) of the execution strategy π, but the specification "$C(X,N,\pi) \leq \overline{C}$" allows the trader to deliberately incur extra costs if that improves his trade-off. Since $\overline{C}$ and $C(X,N,\pi)$ are random variables that give a cost for each realization of the stock price process, $\overline{C} > C(X,N,\pi)$ means that the trader may selectively account more than the actual cost (3.6) in *some* realizations of the stock price. We will discuss this issue in Section 3.4.3.

For given $E \in \mathbb{R}$, let

$$\mathscr{A}(X,N,E) = \left\{ \; (\pi,\overline{C}) \in \mathscr{D}(X,N) \;\middle|\; \mathbb{E}\left[\,\overline{C}\,\right] \leq E \; \right\} \tag{3.8}$$

be the subset of policies for which the expected cost is at most E (as described below, this set can be empty if E is too small).

Mean-variance optimization solves the constrained problem

$$\text{For } E \in \mathbb{R}, \text{ minimize } \operatorname{Var}\left[\,\overline{C}\,\right] \text{ over } (\pi,\overline{C}) \in \mathscr{A}(X,N,E) \;. \tag{3.9}$$

The solutions are *efficient* trading strategies such that no other strategy has lower variance for the same level (or lower) of expected costs. By varying $E \in \mathbb{R}$ we determine the family of all efficient strategies

$$\begin{aligned} \mathscr{E}(X,N) \;=\; \Big\{ \; (\pi,\overline{C}) \in \mathscr{D}(X,N) \;\Big|\; & \nexists\, (\tilde{\pi},\widetilde{C}) \in \mathscr{D}(X,N) \text{ s.t.} \\ & \Big(\mathbb{E}[\,\widetilde{C}\,] \leq \mathbb{E}\left[\,\overline{C}\,\right] \;\wedge\; \operatorname{Var}\left[\,\widetilde{C}\,\right] < \operatorname{Var}\left[\,\overline{C}\,\right] \;\Big) \Big\} \;. \end{aligned} \tag{3.10}$$

Plotting $\operatorname{Var}\left[\overline{C}\right]$ *vs* $\mathbb{E}\left[\overline{C}\right]$ for all $(\pi,\overline{C}) \in \mathscr{E}(X,N)$, we obtain a downward-sloping curve in the V-E plane, the *efficient frontier of optimal trading strategies* (Almgren and Chriss, 2000). The set of efficient strategies $\mathscr{E}(X,N)$ is independent of the initial stock price S_0. This is by virtue of the arithmetic random walk (3.1).

The feasible domain $\mathscr{A}(X,N,E)$ of (3.9) is empty if the target expected cost E is too small. To determine the minimum possible expected cost, note that the last term in (3.6) has strictly zero expected value. Minimizing the remaining term $\mathbb{E}\left[\sum(x_{k-1}-x_k)^2\right]$ gives the nonrandom linear strategy π_{lin} with $x_{k-1} - x_k = X/N$ for $k = 1 \ldots N$, and $\overline{C} = C(X,N,\pi_{\text{lin}})$. The expectation and variance of total cost of this strategy are

$$E_{\text{lin}}(X,N) \;=\; \frac{\eta X^2}{N\tau} \;=\; \frac{\eta X^2}{T} \;, \tag{3.11}$$

$$V_{\text{lin}}(X,N) \;=\; \frac{1}{3}\sigma^2 X^2 \, N\tau \left(1 - \frac{1}{N}\right)\left(1 - \frac{1}{2N}\right) \;. \tag{3.12}$$

Since $N\tau = T$ is the fixed trading duration, E_{lin} is independent of the discretization parameter N, and V_{lin} depends only weakly on N when N is large. In volume time, this strategy is equivalent to the popular VWAP profile (see for instance Madhavan (2002)).

The other extreme is to sell the entire position in the first time period: π_{inst} with $x_1 = \cdots = x_N = 0$. This yields $V_{\text{inst}}(X, N) = 0$ and an expectation of

$$E_{\text{inst}}(X) \;=\; \frac{\eta X^2}{\tau} \;. \tag{3.13}$$

We have $E_{\text{lin}} \geq E_{\text{inst}}$, and in fact $\overline{C} = C(X, N, \pi_{\text{inst}}) = E_{\text{inst}}(X)$ is non-random. For $\tau \to 0$, $E_{\text{inst}}(X)$ becomes arbitrarily large.

From this discussion, we conclude that $\mathscr{A}(X, N, E)$ is non-empty exactly if $E_{\text{lin}}(X, N) \leq E$.

For E_{lin}, the mean and variance are both proportional to the square of the initial position X. This is a general feature of linear impact functions, since the impact per share is linear in the trade size, and the total dollar cost is the per-share cost multiplied by the number of shares; the variance is always quadratic in the number of shares.

3.3.1. Static Trajectories. We distinguish two types of trading strategies: path-independent and path-dependent. Path-independent strategies are determined in advance of trading at time t_0, and are deterministic schedules that depend only on information available at time t_0. Path-dependent strategies, conversely, are arbitrary non-anticipating trading policies for which each x_k depends on all information up to and including time t_k. In the following, we shall refer to path-independent strategies as *static* strategies, and to path-dependent strategies as *adaptive* or *dynamic*.

More precisely, for a static trading strategy we require $\pi = (x_0, \ldots, x_N)$ to be fixed at the start of trading, independently of the ξ_k, and $\overline{C} = C(X, N, \pi)$. Then $C(X, N, \pi)$ has mean and variance

$$\mathbb{E}\left[C(X, N, \pi)\right] = \frac{\eta}{\tau} \sum_{k=1}^{N} (x_{k-1} - x_k)^2, \quad \text{Var}\left[C(X, N, \pi)\right] = \sigma^2 \tau \sum_{k=1}^{N} x_k^2 \;,$$

and (3.9) becomes

$$\min_{x_1 \geq \cdots \geq x_{k-1}} \left\{ \sigma^2 \tau \sum_{k=1}^{N} x_k^2 \;\middle|\; \frac{\eta}{\tau} \sum_{k=1}^{N} (x_{k-1} - x_k)^2 \leq E \right\} \tag{3.14}$$

with $x_0 = X$ and $x_N = 0$. As in the continuous time case in Theorem 2.2 in Chapter 2, the solutions are given by the family

$$x_j = X \, \frac{\sinh(\kappa(T - t_j))}{\sinh(\kappa T)}, \qquad j = 0, \ldots, N \;, \tag{3.15}$$

with the *urgency parameter* κ. For given $E \geq 0$, the explicit solution of (3.14) is then obtained by substituting (3.15) into (3.14) and solving for κ.

The execution strategy is independent of the portfolio size X except for an overall factor, and the expected value and variance of total cost are quadratic in portfolio size. This is an artifact of the linear impact model; for nonlinear models the trajectories do depend on portfolio size and the cost does not have a simple expression (Almgren, 2003).

3.4. Optimal Adaptive Strategies

As seen in the previous section, we can construct optimal static trading strategies by solving a straightforward optimization problem. On the other hand, determining optimal adaptive strategies is a difficult problem. In the previous chapter, we demonstrated that adaptive strategies can improve over static trajectories – even adaptive trade schedules which update only once at an intermediate time T^* and follow path-independent trade schedules before and after that "intervention" time. As we discussed there, this approach does not generalize to fully optimal adaptive strategies. We will now present a dynamic programming technique to determine optimal dynamic strategies to any degree of precision.

3.4.1. Dynamic programming. It is alluring to use dynamic programming (Bellman, 1957) to determine optimal trading strategies for the mean-variance criterion $\mathbb{E}[Y] + \lambda \operatorname{Var}[Y]$, since this technique works so well for objective functions of the form $\mathbb{E}[u(Y)]$. But dynamic programming for expected values relies on the "smoothing property" $\mathbb{E}[\mathbb{E}[u(Y)|X]] = \mathbb{E}[u(Y)]$. For the square of the expectation in the variance term $\operatorname{Var}[Y] = \mathbb{E}[Y^2] - \mathbb{E}[Y]^2$, there is no analog of this expression, and it is difficult to see how to design an iterative solution procedure.

However, with a suitable choice of the value function, mean-variance optimization is indeed amenable to dynamic programming. The dynamic programming principle for mean-variance relies on the following counterpart to the smoothing property of expectation, known as the law of total variance.

Lemma 3.1. *Let X and Y be random variables on the same probability space, and $\operatorname{Var}[Y] < \infty$. Then*

$$\operatorname{Var}[Y] = \operatorname{Var}[\,\mathbb{E}[Y \mid X]] + \mathbb{E}[\,\operatorname{Var}[Y \mid X]] \quad .$$

Proof. Follows from

$$\operatorname{Var}[\mathbb{E}[Y|X]] = \mathbb{E}\left[\mathbb{E}[Y|X]^2\right] - \mathbb{E}[\mathbb{E}[Y|X]]^2 = \mathbb{E}\left[\mathbb{E}[Y|X]^2\right] - \mathbb{E}[Y]^2$$

and

$$\mathbb{E}[\operatorname{Var}[Y|X]] = \mathbb{E}\left[\mathbb{E}\left[Y^2|X\right]\right] - \mathbb{E}\left[\mathbb{E}[Y|X]^2\right] = \mathbb{E}\left[Y^2\right] - \mathbb{E}\left[\mathbb{E}[Y|X]^2\right] \quad .$$

□

For $(\pi, \overline{C}) \in \mathscr{D}(X, N)$ (with $N \geq 2$), $\pi = (X, x_1, \dots, x_{N-1}, 0)$, denote

$$(\pi, \overline{C})_{\xi_1 = a_1} \qquad \text{the "tail" of the trading strategy } (\pi, \overline{C})$$

for the remaining $N-1$ trading periods *conditional* on the outcome $\xi_1 = a_1$ of the first period. This tail strategy has the trade schedule

$$\pi_{\xi_1 = a_1} = (x_1, \dots, x_{N-1}, 0) \in \mathscr{D}(x_1, N-1) \ ,$$

and the cost random variable $\overline{C}_{\xi_1 = a_1}$. That is, the cost $\overline{C}$ of the total strategy in terms of the cost strategies of its tails is

$$\overline{C} = \overline{C}_{\xi_1} + \frac{\eta}{\tau}(x - x_1)^2 - \sigma \tau^{1/2} \xi_1 x_1 \quad . \tag{3.16}$$

Note that an adaptive policy $(\pi, \overline{C})$ may use the information $\xi_1 = a_1$ from time t_1 onwards, hence in general $(\pi, \overline{C})_{\xi_1 = a_1}$ indeed depends on the realization a_1.

The key ingredient in dynamic programming is to write the time-dependent optimization problem on N periods as the combination of a single-step optimization with an optimization on the remaining $N-1$ periods. We must carefully define the parameters of the $(N-1)$-step problem so that it gives the same solution as the "tail" of the N-step problem.

In Almgren and Chriss (2000), the risk-aversion parameter λ is constant in time and is constant across realizations of the price process $\xi_1 \in \Omega$. For the expected utility function, Schied and Schöneborn (2007) hold constant the analogous parameter α. For our mean-variance formulation, the following Lemma asserts that the "tail" of the initial strategy is indeed an optimal strategy across $N-1$ steps, if it is defined to be the minimum-variance solution for an appropriate cost limit $\mathbb{E}[C]$. This cost limit is taken to be the expected value of the remainder of the initial strategy, and will be different in each realization.

LEMMA 3.2. *For $N \geq 2$, let $(\pi, \overline{C}) \in \mathscr{E}(X, N)$ be an efficient policy, $\pi = (X, x_1, \ldots, x_{N-1}, 0)$, for* (3.9)*. Then*

$$(\pi, \overline{C})_{\xi_1 = a} \in \mathscr{E}(x_1, N-1) \quad \textit{for almost all outcomes } a \in \Omega \textit{ of } \xi_1 \ ,$$

that is, $B = \{\, a \in \Omega \mid (\pi, \overline{C})_{\xi_1 = a} \notin \mathscr{E}(x_1, N-1) \,\}$ has probability zero.

PROOF. For each $a \in B$ (if B is empty the result is immediate), the tail-strategy $(\pi, \overline{C})_{\xi_1 = a}$ is not efficient, and thus there exists $(\pi_a^*, \overline{C}_a^*) \in \mathscr{D}\big(x_1, N-1\big)$ such that

$$\mathbb{E}\left[\, \overline{C}_a^* \,\right] = \mathbb{E}\left[\, \overline{C}_{\xi_1 = a} \,\right] \quad \wedge \quad \mathrm{Var}\left[\, \overline{C}_a^* \,\right] < \mathrm{Var}\left[\, \overline{C}_{\xi_1 = a} \,\right] \ .$$

Define $(\tilde{\pi}, \widetilde{C})$ by replacing the policy for t_1 to t_N in $(\pi, \overline{C})$ by $(\pi_a^*, \overline{C}_a^*)$ for all $a \in B$ if $\xi_1 = a$ (and identical to π for all other outcomes $\Omega \setminus B$ of ξ_1). Then by definition

$$\widetilde{C} \geq C(X, N, \widetilde{\pi}) \ \ a.e. \tag{3.17}$$

and hence $(\widetilde{\pi}, \widetilde{C}) \in \mathscr{D}(X, N)$. Also by construction, we have

$$\mathbb{E}[\, \widetilde{C} \,] = \mathbb{E}\left[\, \overline{C} \,\right] \ ,$$

and conditional on $\xi_1 \in B$

$$\mathrm{Var}[\, \widetilde{C} \,|\, \xi_1] < \mathrm{Var}\left[\, \overline{C} \,|\, \xi_1 \right] \ .$$

If B has positive probability then[1]

$$\mathbb{E}\left[\mathrm{Var}[\widetilde{C} \,|\, \xi_1 \in B] \right] < \mathbb{E}\left[\mathrm{Var}\left[\overline{C} \,|\, \xi_1 \in B \right] \right]$$

and hence, by the law of total variance,

$$\mathrm{Var}[\widetilde{C}] = \mathrm{Var}[\mathbb{E}[\widetilde{C} \,|\, \xi_1]] + \mathbb{E}[\mathrm{Var}[\widetilde{C} \,|\, \xi_1]] < \mathrm{Var}\left[\mathbb{E}\left[\overline{C} \,|\, \xi_1 \right] \right] + \mathbb{E}\left[\mathrm{Var}[\overline{C} \,|\, \xi_1] \right] = \mathrm{Var}\left[\overline{C} \right] \ ,$$

contradicting $(\pi, \overline{C}) \in \mathscr{E}(X, N)$. □

[1] In general, let $(\Omega, \mathcal{A}, \mu)$ be a measure space and f a measurable function with $f > 0$ a.e. Then

$$\int_B f d\mu > 0 \quad \text{for all measurable sets } B \text{ with } \mu(B) > 0 \ . \tag{3.18}$$

This follows from the Theorem of Beppo-Levi, which asserts that

$$\sup_{n \in \mathbb{N}} \int_\Omega f_n d\mu = \int_\Omega \sup_{n \in \mathbb{N}} f_n d\mu$$

for every monotone increasing sequence of non-negative functions. For (3.18), we set $B_n = B \cap \{x \,|\, f(x) \geq 1/n\}$, and define $f_n(x) = x$ for $x \in B_n$ and $f_n(x) = 0$ otherwise.

Lemma 3.2 shows a Markov property of mean-variance efficient strategies, and constitutes the foundation of why the dynamic programming principle is applicable. Dynamic programming can be used for all types of risk-measurement which satisfy such a Markov property. For instance, it naturally holds for the optimization of the expected utilty of a final payoff (see also Appendix A). Unfortunately, most other risk measures don't satisfy this property (see Artzner et al. (2007)).

For $k \geq 1$ and for fixed τ, we define the value function

$$J_k(x,c) = \min_{(\pi,\overline{C})\in\mathscr{A}(x,k,c)} \mathrm{Var}\left[\,\overline{C}\,\right] \tag{3.19}$$

for $x \geq 0$ (we will always assume a non-negative number of shares left to sell since we consider pure sell programs which are not allowed to buy shares). Note that for $|\Omega| < \infty$, the space $\mathscr{A}(x,k,c)$ is finite-dimensional and minimizing solutions will exist; we leave this issue for $|\Omega| = \infty$ as an open question.

If the cost limit c is below the cost $E_{\mathrm{lin}}(x,k) = \eta x^2/(k\tau)$ (3.11) of the linear strategy, then no admissible solution exists and we set $J_k = \infty$. If $c = E_{\mathrm{lin}}(x,k)$, then the linear strategy is the only solution, with variance given by (3.12). If the cost limit is above the cost $E_{\mathrm{inst}}(x) = \eta x^2/\tau$ (3.13), then instantaneous liquidation is admissible with variance zero and we have $J_k = 0$. Thus, we have (for $x \geq 0$)

$$J_k(x,c) \;=\; \begin{cases} \infty, & c < \eta x^2/(k\tau) \\ V_{\mathrm{lin}}(x,k), & c = \eta x^2/(k\tau) \\ \text{non-increasing in } c, & \eta x^2/k\tau \leq c \leq \eta x^2/\tau \\ 0, & c \geq \eta x^2/\tau \;\; . \end{cases} \tag{3.20}$$

For fixed x (and k), as we vary c, $J_k(x,c)$ traces the efficient frontier in the E-V-plane for trading x shares in k periods. That is, $J_k(x,c)$ gives a family of efficient frontiers, one for each level of portfolio size x. The solution to (3.9) is $J_N(X,E)$, together with the corresponding optimal policy.

For $k = 1$, $E_{\mathrm{inst}} = E_{\mathrm{lin}}$ and so (for $x \geq 0$)

$$J_1(x,c) \;=\; \begin{cases} \infty, & c < \eta x^2/\tau \\ 0, & c \geq \eta x^2/\tau \;\; , \end{cases} \tag{3.21}$$

and for given (x,c) an optimal strategy $(\pi_1^*, \overline{C}_1^*)$ is given by

$$\pi_1^*(x,c) = x \quad \text{and} \quad \overline{C}_1^*(x,c) = \max\{\eta x^2/\tau, c\} \;\; . \tag{3.22}$$

Note that (3.22) indeed means that the trader can account more than the actual trading cost $\eta x^2/\tau$ in the last period (corresponding to $\overline{C} \geq C(X,N,\pi)$ in the definition (3.7) of $\mathscr{D}$), if that improves his tradeoff. We will discuss that issue in Section 3.4.3.

By definitions (3.7,3.8,3.10,3.19), the value function $J_k(x,c)$ and the set $\mathscr{E}(x,k)$ are related by

$$(\pi^*, \overline{C}^*) = \operatorname*{argmin}_{(\pi,\overline{C})\in\mathscr{A}(x,k,c)} \mathrm{Var}\left[\,\overline{C}\,\right] \implies (\pi^*, \overline{C}^*) \in \mathscr{E}(x,k) \tag{3.23}$$

and

$$(\pi, \overline{C}) \in \mathscr{E}(x,k) \implies \mathrm{Var}\left[\,\overline{C}\,\right] = J_k(x, \mathbb{E}\left[\overline{C}\right]) \;\; . \tag{3.24}$$

In view of the known static solutions, and by inspection of the expressions (3.20) and (3.21), it is natural to conjecture that the value function and the cost limit should be proportional to the square of the number of shares: $J_k(x,c) = x^2 f_k(c/x^2)$. In fact for dynamic strategies this is *not* true, even for linear impact functions, except in the limit of small portfolio size (in a suitable nondimensional sense made clear below).

In the spirit of dynamic programming, we use the efficient frontier for trading over $k-1$ periods, plus an optimal one-period strategy, to determine the efficient frontier for trading over k periods. The key is to introduce an additional control parameter in addition to the number of shares we trade in the next period. This extra parameter is the expected cost limit for the remaining periods, which we denote by z; it is a real-valued integrable function $z \in L^1(\Omega;\mathbb{R})$ of the price change $\xi \in \Omega$ on that step; for $|\Omega| < \infty$, z is effectively a real-valued vector.

THEOREM 3.3. *Let the stock price change in the next trading period be* $\sigma\tau^{1/2}\xi$ *with* $\xi \in \Omega$ *the random return. Define*

$$\mathcal{G}_k(x,c) = \Big\{ \big(y,z\big) \in \mathbb{R} \times L^1(\Omega;\mathbb{R}) \;\Big|\; \mathbb{E}\left[z(\xi)\right] + \frac{\eta}{\tau}(x-y)^2 \le c,\; 0 \le y \le x \Big\}. \tag{3.25}$$

Then for $k \ge 2$,

$$J_k(x,c) = \min_{(y,z)\in\mathcal{G}_k(x,c)} \Big(\operatorname{Var}\Big[z(\xi) - \sigma\tau^{1/2}\xi y\Big] + \mathbb{E}\Big[J_{k-1}\big(y, z(\xi)\big)\Big] \Big). \tag{3.26}$$

and for given (x,c) *an optimal* k *period strategy* $(\pi_k^*(x,c), \overline{C}_k^*(x,c)) \in \mathscr{D}(x,k)$ *for* x *shares with maximal expected cost* c *is given by*

$$\pi_k^*(x,c) = \Big(y^*,\; \pi_{k-1}^*(y^*,z^*) \Big) \tag{3.27}$$

$$\overline{C}_k^*(x,c) = \overline{C}_{k-1}^*(y^*,z^*) + \frac{\eta}{\tau}(x-y^*)^2 - \sigma\tau^{1/2}\xi y^* \;, \tag{3.28}$$

where (y^*,z^*) *is an optimal control in* (3.26) *for* (x,c), *and for each* $a \in \Omega$ *the* $(k-1)$*-period strategy* $(\pi_{k-1}^*(y^*,z^*(a)), \overline{C}_{k-1}^*(y^*,z^*))$ *an optimal* $(k-1)$*-period strategy for selling* y^* *shares with maximal cost of* $z^*(a)$.

PROOF. Let ξ be the random price innovation in the first of the remaining k trading periods. For given $x \ge 0$ and $E_{\mathrm{lin}}(x,k) \le c$, let

$$(\pi^*, \overline{C}^*) = \operatorname*{argmin}_{(\pi,\overline{C})\in\mathscr{A}(x,k,c)} \operatorname{Var}\left[\,\overline{C}\,\right] \;,$$

That is, π^* is an optimal strategy to sell x shares in k time periods of length τ with expected cost at most c. By (3.23), we have $(\pi^*, \overline{C}^*) \in \mathscr{E}(X,N)$, and by the definition (3.19) of J_k we have $J_k(x,c) = \operatorname{Var}\left[\,\overline{C}^*\,\right]$. Let y be the number of shares held by π^* after the first trading period, so $\pi^* = (x, y, \ldots, 0)$.

The strategy π^* may be understood as consisting of two parts: First, the number of shares to be sold in the first period, $x-y$. This is a deterministic variable, and may not depend on the next period price change ξ. Second, the strategy for the remaining $k-1$ periods. When the trader proceeds with this $(k-1)$-period strategy, the outcome of ξ is known, and the strategy may depend on it. Conditional on $\xi = a$, let $(\pi^*, \overline{C}^*)_{\xi=a}$ be the $(k-1)$-period tail-strategy.

By Lemma 3.2, $(\pi^*, \overline{C}^*)_{\xi=a} \in \mathscr{E}(y, k-1)$ for almost all realizations $a \in \Omega$ of ξ. Thus, there exists $z \in L^1(\Omega;\mathbb{R})$ such that using (3.24) we have for each a

$$\begin{aligned} \mathbb{E}\left[\overline{C}^*_{\xi=a}\right] &= z(a) \\ \mathrm{Var}\left[\overline{C}^*_{\xi=a}\right] &= J_{k-1}\big(y, z(a)\big) \ . \end{aligned}$$

Since $(\pi^*, \overline{C}^*)_{\xi=a} \in \mathscr{E}(y, k-1)$, we must have

$$z(a) \geq E_{\mathrm{lin}}(y, k-1) \tag{3.29}$$

(the minimal expected cost is achieved by the linear profile π_{lin}). With (3.16), we conclude

$$\begin{aligned} \mathbb{E}\left[\,\overline{C}^* \,|\, \xi = a\right] &= z(a) + \frac{\eta}{\tau}(x-y)^2 - \sigma\tau^{1/2} a y \ , \\ \mathrm{Var}\left[\,\overline{C}^* \,|\, \xi = a\right] &= J_{k-1}(y, z(a)) \ , \end{aligned}$$

and by the law of total expectation and total variance (Lemma 3.1)

$$\mathbb{E}\left[\,\overline{C}^*\,\right] = \mathbb{E}\left[z(\xi)\right] + \frac{\eta}{\tau}(x-y)^2 \ , \tag{3.30}$$

$$\mathrm{Var}\left[\,\overline{C}^*\,\right] = \mathrm{Var}\left[z(\xi) - \sigma\tau^{1/2}\xi y\right] + \mathbb{E}\left[J_{k-1}\big(y, z(\xi)\big)\right] \ . \tag{3.31}$$

That is, an optimal strategy $(\pi^*, \overline{C}^*)$ for k periods is defined by y and the $(k-1)$-period tail specified by $\mathbb{E}\left[\overline{C}^*_{\xi=a}\right] = z(\xi)$; the expectation and variance of $(\pi^*, \overline{C}^*)$ are then given by (3.30,3.31). As noted in (3.29), not all $z(\xi)$ are possible, since the minimal possible expected cost of the tail strategy $\mathbb{E}\left[\overline{C}^*_{\xi=a}\right]$ is the expected cost of a linear profile.

Conversely, for given $0 \leq y \leq x$ and $z \in L^1(\Omega;\mathbb{R})$ with $z \geq E_{\mathrm{lin}}(y, k-1)$ *a.e.*, there exists a k-period strategy $(\pi, \overline{C})$ as follows: inductively by (3.27, 3.28) for $k \geq 3$ (and (3.22) for $k = 2$, respectively), for almost all $a \in \Omega$ there exists a $(k-1)$-period strategy $(\pi_{k-1}(y, z(a)), \overline{C}_{k-1}(y, z(a)))$ to sell y shares with at most expected costs of $z(a)$, i.e. $\mathbb{E}\left[\overline{C}_{k-1}(y, z)\right] = \mathbb{E}\left[z(\xi)\right]$. We can combine this set of tail strategies with a sale of $x - y$ shares in the current period to obtain a k-period strategy $(\pi, \overline{C})$,

$$\begin{aligned} \pi &= \Big(\, y,\ \pi_{k-1}(y, z) \,\Big) \\ \overline{C} &= \overline{C}_{k-1}(y, z) + \frac{\eta}{\tau}(x-y)^2 - \sigma\tau^{1/2}\xi y \ . \end{aligned}$$

Since $x \geq y$ and $\pi_{k-1}(y, z) \in \mathscr{D}(y, k-1)$ (and hence $\overline{C} \geq C(y, k-1, \pi_{k-1}(y, z)) + \frac{\eta}{\tau}(x-y)^2 - \sigma\tau^{1/2}\xi y = C(x, k, \pi)$ *a.e.*), we have $(\pi, \overline{C}) \in \mathscr{D}(x, k)$. Since $\mathbb{E}\left[\overline{C}\right] = \mathbb{E}\left[z(\xi)\right] + \frac{\eta}{\tau}(x-y)^2$ (compare to (3.30)), if $(y, z) \in \mathcal{G}_k(x, c)$, then $(\pi, \overline{C}) \in \mathscr{A}(x, k, c)$.

We conclude that an optimal solution $(\pi^*, \overline{C}^*)$ to

$$J_k(x, c) = \min_{(\pi, \overline{C}) \in \mathscr{A}(x,k,c)} \mathrm{Var}\left[\,\overline{C}\,\right] \ ,$$

can indeed be obtained by finding an optimal solution to

$$\min_{(z,y)} \operatorname{Var}\left[z(\xi) - \sigma\tau^{1/2}\xi y\right] + \mathbb{E}\left[J_{k-1}\big(y, z(\xi)\big)\right]$$

$$\text{s.t. } \mathbb{E}\left[z(\xi)\right] + \frac{\eta}{\tau}(x-y)^2 \le c$$

$$0 \le y \le x \tag{3.32}$$

$$E_{\text{lin}}(y, k-1) \le z(\xi) \;, \tag{3.33}$$

and constructing a k-period strategy from there by means of (3.27, 3.28).

The constraint (3.32) comes from our requirement that $(\pi^*, \overline{C}^*)$ must be a pure sell-program. Since $\mathscr{A}(x,k,c) = \emptyset$ for $c < E_{\text{lin}}(x,k)$ in (3.19), $J_{k-1}(y, z(\xi)) = \infty$ for $z(\xi) < E_{\text{lin}}(y, k-1)$ and thus the constraint (3.33) never becomes binding. Thus, the result (3.25,3.26) follows. □

Thus we can construct an optimal Markovian k-period policy by choosing an optimal control $\big(y, z(\xi)\big)$ and combining it with an optimal strategy for $k-1$ periods: we sell $x-y$ in the first of the k periods, and commit ourselves that if $\xi = a$ during this first period, then we sell the remaining y shares with the mean-variance optimal strategy with expected cost $z(a)$ and variance $J_{k-1}\big(y, z(a)\big)$. The *rule* $z(\xi)$ may be any $z \in L^1(\Omega; \mathbb{R})$ on the set Ω of possible values for ξ.

General Temporary and Permanent Impact. In fact, the dynamic programming principle described above can be extended to general temporary and permanent impact functions in the setting (3.1,3.2,3.3). For temporary impact $h(\cdot)$ and permanent impact $g(\cdot)$, instead of (3.25) we have

$$\mathcal{G}_k(x,c) = \left\{ (y,z) \in \mathbb{R} \times L^1(\Omega;\mathbb{R}) \;\middle|\; \begin{array}{l} \mathbb{E}\left[z\right] + (x-y)h\Big(\dfrac{x-y}{\tau}\Big) + y\tau g\Big(\dfrac{x-y}{\tau}\Big) \le c \\ 0 \le y \le x \end{array} \right\} \tag{3.34}$$

with the backwards step equation (3.26) unchanged. The expression for the terminal value function $J_1(x,c)$ changes accordingly.

3.4.2. Nondimensionalization. The optimization problem (3.9) (respectively the one-step optimization problem (3.25, 3.26)) depends on five dimensional constants: the initial shares X, the total time T (or the time step τ in conjunction with the number of steps N), the absolute volatility σ and the impact coefficient η. To simplify the structure of the problem, it is convenient to define scaled variables.

We measure shares x relative to the initial position X. We measure impact cost c and its limit z relative to the total dollar cost that would be incurred by liquidating X shares in time T using the linear strategy; the per-share cost of this strategy is $\eta v = \eta X/T$ so the total cost is $\eta X^2/T$. We measure variance J_k relative to the variance (in squared dollars) of holding X shares across time T with absolute volatility σ. The standard deviation of the price per share is $\sigma\sqrt{T}$, so the standard deviation of dollar value is $\sigma\sqrt{T}X$ and the variance scale is $\sigma^2 T X^2$.

We denote nondimensional values by a caret ˆ, so we write

$$x = X\,\hat{x}, \qquad c = \frac{\eta X^2}{T}\,\hat{c}, \qquad z = \frac{\eta X^2}{T}\,\hat{z} \tag{3.35}$$

and

$$J_k(x,c) \;=\; \sigma^2 T X^2\, \hat{J}_k\!\left(\frac{x}{X}, \frac{c}{\eta X^2/T} \right) \;. \tag{3.36}$$

Then $\hat{X} = \hat{x}_0 = 1$, so the trading strategy is $\hat{\pi} = (1, \hat{x}_1, \ldots, \hat{x}_{N-1}, 0)$.

The one-period value function is

$$\hat{J}_1(\hat{x},\hat{c}) \;=\; \begin{cases} \infty, & \hat{c} < N\hat{x}^2 \\ 0, & \hat{c} \ge N\hat{x}^2 \end{cases} \tag{3.37}$$

and in Theorem 3.3 we have the scaled set of admissible controls

$$\hat{\mathcal{G}}_k(\hat{x},\hat{c}) = \left\{ (\hat{y},\hat{z}) \in \mathbb{R} \times L^1(\Omega;\mathbb{R}) \;\middle|\; \mathbb{E}\left[\hat{z}(\xi)\right] + N(\hat{x}-\hat{y})^2 \le \hat{c},\; 0 \le \hat{y} \le \hat{x} \right\} \tag{3.38}$$

and the dynamic programming step

$$\hat{J}_k(\hat{x},\hat{c}) = \min_{(\hat{y},\hat{z}) \in \hat{\mathcal{G}}_k(\hat{x},\hat{c})} \Big(\operatorname{Var}\left[\mu \hat{z}(\xi) - N^{-1/2}\xi \hat{y}\right] + \mathbb{E}\left[\hat{J}_{k-1}\big(\hat{y}, \hat{z}(\xi)\big)\right] \Big). \tag{3.39}$$

Again, we have the nondimensional "market power" parameter

$$\mu \;=\; \frac{\eta X}{\sigma T^{3/2}} \;=\; \frac{\eta X/T}{\sigma\sqrt{T}} \tag{3.40}$$

that we already discovered in the previous chapter as a preference-free measure of portfolio size. The numerator is the per-share price impact that would be caused by liquidating the portfolio linearly across the available time; the denominator is the amount that the price would move on its own due to volatility in the same time.

To estimate realistic sizes for this parameter, we refer to Almgren, Thum, Hauptmann, and Li (2005); there, the nonlinear model $K/\sigma = \eta (X/VT)^\alpha$ is introduced, where K is temporary impact, σ is daily volatility, X is trade size, V is an average daily volume (ADV), and T is the fraction of a day over which the trade is executed. The coefficient was estimated empirically as $\eta = 0.142$, as was the exponent $\alpha = 3/5$. Therefore, a trade of 100% ADV executed across one full day gives $\mu = 0.142$. Although this is only an approximate parallel to the linear model used here, it does suggest that for realistic trade sizes, μ will be substantially smaller than one.

The nondimensional version (3.37,3.38,3.39) of the optimization problem now depends only on two nondimensional parameters: the time discretization parameter N and the new market power parameter μ. Especially for numerical treatment, this reduction is very useful. From now on, we shall drop the nondimensionalization mark ˆ, assuming that all variables have been nondimensionalized.

3.4.3. Convexity. We now show that the optimization problem at each step is convex, and that the value function J_k is a convex function of its two arguments.

We need the following lemma which is proved by an easy modification of the argument in Boyd and Vandenberghe (2004, sect. 3.2.5).

Lemma 3.4. *Let $f(v)$ and $h(u,v)$ be real-valued convex functions on vector spaces V and $U \times V$ respectively, possibly taking the value $+\infty$. Then $g : U \mapsto \mathbb{R}$ defined by*

$$g(u) = \inf_{v \in V} \left\{ \; f(v) \;\middle|\; h(u,v) \le 0 \; \right\}$$

is convex.

Proof. Let $u_1, u_2 \in U$ be such that $g(u_i) < \infty$. Then for every $\epsilon > 0$, there exist $v_1, v_2 \in V$ such that

$$f(v_i) \le g(u_i) + \epsilon \quad \text{and} \quad h(u_i, v_i) \le 0 \ .$$

Since $h(u,v)$ is convex, this implies

$$h\big(\,\theta u_1 + (1-\theta)u_2,\ \theta v_1 + (1-\theta)v_2\,\big) \ \le\ 0$$

for $0 \le \theta \le 1$. Then

$$\begin{aligned} g\big(\,\theta u_1 + (1-\theta)u_2\,\big) &= \inf_v \big\{ f\big(\,v\,\big) \;\big|\; h\big(\,\theta u_1 + (1-\theta)u_2,\ v\,\big) \le 0 \big\} \\ &\le\ f\big(\,\theta v_1 + (1-\theta)\,v_2\,\big) \\ &\le\ \theta\, f(v_1) + (1-\theta)\, f(v_2) \\ &\le\ \theta\, g(u_1) + (1-\theta)\, g(u_2) + \epsilon \ . \end{aligned}$$

Since this holds for any $\epsilon > 0$, we have

$$g(\,\theta u_1 + (1-\theta)u_2\,) \ \le\ \theta\, g(u_1) \ +\ (1-\theta)\, g(u_2) \ ,$$

and hence $g(u)$ is convex. □

Now we are ready to prove the convexity of $J_k(x,c)$ and the dynamic programming step.

Theorem 3.5. *The optimization problem (3.38,3.39) is convex. The value function $J_k(x,c)$ is convex for $k \ge 1$.*

Proof. We proceed by induction. Clearly, $J_1(x,c)$ in (3.37) is convex, since it is an indicator function on the convex domain $\{(x,c)\,|\,c \ge Nx^2\} \subseteq \mathbb{R}^2$.

The optimization problem (3.38,3.39) is of the form described in Lemma 3.4, with the identifications $u = (x,c)$ and $v = (y,z)$, and $h(u,v)$, $f(v)$ given by the functions appearing on the right side of (3.38) and (3.39) respectively. Thus we need only show that these functions are convex.

For each k, the constraint function in (3.38) is convex in (x,c,y,z), since the expectation operator is linear and the quadratic term is convex. In (3.39), the second term in the objective function is convex in (y,z), since J_{k-1} is assumed convex and expectation is linear.

The first term in (3.39) may be written $\mathrm{Var}\,[w(\xi)]$ where $w(\xi)$ depends linearly on y and $z(\xi)$. And it is easy to see that $\mathrm{Var}\,[w]$ is convex in w. (Indeed, this is certainly true for random variables w having $\mathbb{E}\,[w] = 0$ since then $\mathrm{Var}\,[w] = \int w^2$. For general w, one can write $w(\xi) = \overline{w} + u(\xi)$ where $\overline{w}$ is constant and $u(\xi)$ has mean zero; then $\mathrm{Var}\,[w] = \int u^2$ and convexity follows.) The result follows. □

In fact, the optimization problem (3.52) will be of convex type also in more general settings of temporary and permanent impact. As already briefly discussed at the end of Section 3.4.1, the objective function in (3.26) will be unchanged for different choices of impact functions $g(\cdot)$ and $h(\cdot)$, and only the constraint in (3.25) will change to

$$\mathbb{E}\,[z] + f(y,x) \le c \ , \tag{3.41}$$

where $f(y,x) = (x-y)h((x-y)/\tau) + y\tau g((x-y)/\tau)$ is the combined temporary and permanent expected cost induced by selling $x-y$ shares of x units held. For the linear temporary impact model (and no permament impact), we have $f(y,x) = \eta(x-y)^2/\tau$, which is convex in y. Obviously, (3.34) will be a convex set for any choice of $h(\cdot)$ and $g(\cdot)$ where f is convex in (y,x), resulting in a convex optimization problem.

Now let us return to the definitions (3.7,3.8,3.10), especially the constraint "$C(X,N,\pi) \le \overline{C}$ *a.e.*" in (3.7). Recall that $\overline{C}$ and $C(X,N,\pi)$ are random variables that give a cost for each realization of the stock price process. Thus, $\overline{C} \ge C(X,N,\pi)$ in (3.7) means that the trader can incur extra costs (by giving away money) in some realizations of the stock price process, if that improves his mean-variance tradeoff.

The dynamic program (3.21, 3.25, 3.26) indeed allows for $\overline{C} > C(X,N,\pi)$, ultimately by the trader's choice of the last period tail strategy: the trader has a certain number x of shares left to sell with actual cost of $\eta x^2/\tau$; there is no decision variable in reality. However, in the specification of the dynamic program (3.21, 3.25, 3.26), which is in line with the definition (3.10) of $\mathscr{E}$, the trader additionally specifies $\overline{C} \ge \eta x^2/\tau$; the difference $\overline{C} - \eta x^2/\tau$ is the money that the trader is giving away in that particular realization of the stock price process.

This is counterintuitive, but the trader may want to make use of it due to a rather undesirable property of the mean-variance criterion: A mean-variance optimizer can reduce his variance by making positive outcomes less so. Of course this also reduces his mean benefit, but depending on the parameters the tradeoff may be advantageous. Indeed, mean-variance comparison is not necessarily consistent with stochastic dominance and may fail to be monotone (Levy, 2006; Gandhi and Saunders, 1981; Maccheroni et al., 2004).

If we want to bar the trader from making use of this peculiarity, we replace (3.7) by

$$\mathscr{D}'(X,N) = \left\{ (\pi,\overline{C}) \;\middle|\; \begin{array}{l} \pi = (x_0, x_1, \ldots, x_N) \text{ with } x_0 = X,\ x_N = 0 \\ x_0 \ge x_1 \ge \cdots \ge x_N \\ C(X,N,\pi) = \overline{C} \quad a.e. \end{array} \right\} \tag{3.42}$$

Now the random variable $\overline{C}$ is required to be the *exact* actual cost (3.6) of the trade schedule π at all times. We change the definitions (3.8,3.10) of $\mathscr{A}$ and $\mathscr{E}$ accordingly, replacing $\mathscr{D}$ by $\mathscr{D}'$, and denote these new sets $\mathscr{A}'$ and $\mathscr{E}'$. We define

$$J'_k(x,c) = \min_{(\pi,\overline{C}) \in \mathscr{A}'(x,k,c)} \operatorname{Var}\left[\,\overline{C}\,\right] \ . \tag{3.43}$$

With these definitions, $\mathscr{D}'(X,1)$ is the single-element set $\mathscr{D}'(X,1) = \{(\pi_{\text{inst}}, E_{\text{inst}}(X))\}$ where $\pi_{\text{inst}} = (X,0)$ is the immediate liquidation of X shares and $E_{\text{inst}}(X)$ its cost. Contrary, we have $\mathscr{D}(X,1) = \{(\pi_{\text{inst}}, c) \,|\, c \ge E_{\text{inst}}(X)\}$.

It can be shown that Lemma 3.2 also holds for $\mathscr{E}'$.

LEMMA 3.6. *For $N \ge 2$, let $(\pi,\overline{C}) \in \mathscr{E}'(X,N)$ with $\pi = (X, x_1, \ldots, x_{N-1}, 0)$. Then $B = \{\, a \in \Omega \,|\, (\pi,\overline{C})_{\xi_1 = a} \notin \mathscr{E}'(x_1, N-1) \,\}$ has probability zero.*

The proof follows the proof of Lemma 3.6 word by word, with $\mathscr{D}$ and $\mathscr{E}$ replaced by $\mathscr{D}'$ and $\mathscr{E}'$, respectively; only (3.17) changes to "$\widetilde{C} = C(X,N,\widetilde{\pi})$ *a.e.*", which indeed implies $(\widetilde{\pi},\widetilde{C}) \in \mathscr{D}'(X,N)$ accordingly.

Using Lemma 3.6, we can then argue along the lines of the proof of Theorem 3.3 to obtain the (nondimensionalized) dynamic program

$$\mathcal{G}'_k(x,c) = \left\{ (y,z) \in \mathbb{R} \times L^1(\Omega;\mathbb{R}) \;\middle|\; \begin{array}{l} \mathbb{E}\left[z\right] + N(x-y)^2 \le c \\ y^2 \le z \le Ny^2 \text{ a.e.} \\ 0 \le y \le x \end{array} \right\} \tag{3.44}$$

and

$$J'_k(x,c) = \min_{(y,z)\in\mathcal{G}'_k(x,c)} \left(\operatorname{Var}\left[\mu z - N^{-1/2}\xi y\right] + \mathbb{E}\left[J_{k-1}(y,z)\right] \right) , \tag{3.45}$$

with $J'_1(x,c) = J_1(x,c)$ unchanged. The additional constraint "$z(\xi) \le Ny^2$ a.e." in (3.44) comes from the fact that $z(\xi)$ specifies the $(k-1)$-period tail strategy $(\pi^*, \overline{C}^*)_\xi$ by means of $\mathbb{E}\left[\overline{C}^*_\xi\right] = z(\xi)$ and not all $z(\xi)$ correspond to a $(\pi^*, \overline{C}^*)_\xi \in \mathscr{E}'(y, k-1) \subseteq \mathscr{D}'(y, k-1)$: the maximal expected cost $\mathbb{E}\left[\overline{C}^*_\xi\right]$ of a $(k-1)$-period tail strategy is the cost of immediate liquidation $E_{\text{inst}}(y) = \eta y^2/\tau$ (or $E_{\text{inst}}(y) = Ny^2$ after rescaling according to (3.35)). Since (3.42) requires $\overline{C} = C(y, k-1, \pi^*_\xi)$, we must have $z(\xi) = \mathbb{E}\left[\overline{C}^*_\xi\right] = \mathbb{E}\left[C(y, k-1, \pi^*_\xi)\right] \le E_{\text{inst}}(y) = Ny^2$. In Theorem 3.3, this upper bound does not apply: because $(\pi_{\text{inst}}, c) \in \mathscr{E}'(y, k-1) \subseteq \mathscr{D}'(y, k-1)$ for *all* $c \ge E_{\text{inst}} = Ny^2$, the trader may indeed sell all remaining shares in the next period, and additionally give away $c - \eta y^2/\tau$ (or $c - Ny^2$ in nondimensionalized variables) in cash to incur any cost c he wishes.

Obviously, we have $J'_k(x,c) \ge J_k(x,c)$ for all $k \ge 1$. Contrary to (3.38,3.39), the optimization problem (3.44,3.45) is *not* a convex optimization problem because the additional constraint breaks the convexity of the set $\mathcal{G}'_k(x,c)$ (since $h(z,y) = z - Ny^2$ is not convex, the sublevel set $h(z,y) \le 0$ is not convex), and Lemma 3.4 is not applicable.

In the following, we shall continue to work with the value function $J_k(x,c)$ as defined in Section 3.4.1. In Section 3.5 we will give numerical examples for both, $J_k(x,c)$ and $J'_k(x,c)$, and show that the difference between these two specifications diminishes very rapidly as we increase N.

3.4.4. Small-portfolio limit. In the previous chapter we observed that in the small-portfolio limit $\mu \to 0$ the optimal adaptive and optimal static efficient frontier coincide. We shall now formally prove this property in the context of general strategies. We denote by J^0_k the value function when $\mu = 0$. Note that this limiting case is perfectly natural in the nondimensional form (3.38,3.39). But in the original dimensional form, from the definition (3.40), this requires $\eta \to 0$, $X \to 0$, $\sigma \to \infty$, or $T \to \infty$, all of which pose conceptual problems for the model. We leave it as an open question to show that the solution of the problem for $\mu = 0$ is the same as the limit of the solutions for positive $\mu \searrow 0$.

THEOREM 3.7. *For $\mu = 0$, the optimal policy of (3.38,3.39) is path-independent (static) and the efficient frontier coincides with the static efficient frontier* (3.14).

PROOF. For $\mu = 0$, (3.39) becomes

$$J^0_k(x,c) = \min_{(y,z)\in\mathcal{G}_k(x,c)} \left(N^{-1}y^2 + \mathbb{E}\left[J^0_{k-1}(y, z(\xi))\right] \right) . \tag{3.46}$$

Inductively, we now show that for $k \geq 1$ (defining $x_k = 0$ to shorten notation)

$$J_k^0(x,c) = \min_{x_1 \geq \cdots \geq x_{k-1}} \left\{ \frac{1}{N} \sum_{j=1}^{k-1} x_j^2 \;\middle|\; (x - x_1)^2 + \sum_{j=2}^{k} (x_{j-1} - x_j)^2 \leq \frac{c}{N} \right\} \tag{3.47}$$

for $c \geq Nx^2/k$, and $J_k^0(x,c) = \infty$ otherwise. For $k = 1$, (3.47) reduces to

$$J_1^0(x,c) = 0 \quad \text{for } c \geq Nx^2, \quad \text{and } J_1^0(x,c) = \infty \quad \text{for } c < Nx^2 \ ,$$

proving the inductive hypothesis since by definition (3.37) indeed $J_1^0(x,c) = J_1(x,c)$. For the inductive step, let $k \geq 2$ and suppose that (3.47) holds for $k-1$. $J_{k-1}^0(x,c)$ is convex since (3.47) is the minimization of a convex function with convex constraints of the type in Lemma 3.4. Thus, for any nonconstant $z(\xi)$, Jensen's inequality implies $J_{k-1}^0\big(y, \mathbb{E}\left[z\right]\big) \leq \mathbb{E}\left[J_{k-1}^0(y,z)\right]$. Hence, there exists a constant nonadaptive optimal control $z(\xi) \equiv z$. Thus, (3.46) becomes

$$J_k^0(x,c) = \min_{0 \leq y \leq x, z \in \mathbb{R}} \left\{ N^{-1} y^2 + J_{k-1}^0\big(y,z\big)\Big) \;\middle|\; z + N(x-y)^2 \leq c \right\} \ ,$$

which concludes the proof of the inductive step. After undoing the nondimensionalization, for $k = N$ the optimization problem (3.47) is exactly the optimization problem (3.14) for the static trajectory. Hence, for $\mu = 0$ the adaptive efficient frontier does coincide with the static efficient frontier. □

The theorem also holds for the variant (3.44, 3.45), where we restrict the trader from giving away money. The reason is that for $c \geq Nx^2$, $J_{k-1}^0(x,c) = 0$ since then $x_1 = \cdots = x_{k-1} = 0$ in (3.47) is admissible. Hence, the constraint $z \leq Ny^2$ in (3.44, 3.46) will in fact never become binding for $\mu = 0$.

For $\mu > 0$, improvements over static strategies come from introducing anticorrelation between the two terms inside the variance in (3.39). This reduces the overall variance, which we can trade for a reduction in expected cost. Thus, following a positive investment return, we decrease our cost limit for the remaining part of the program.

3.5. Examples

3.5.1. Scenario Trees. If $|\Omega| < \infty$, the control $z \in L^1(\Omega, \mathbb{R})$ is effectively a real-valued vector $(z_1, \ldots, z_n) \in \mathbb{R}^n$, with

$$z(\xi) = z_j \quad \text{for} \quad \xi = a_j \ , \tag{3.48}$$

where $\Omega = \{a_1, \ldots, a_n\}$ are the possible values of ξ. We assume that $a_1 < a_2 < \cdots < a_n$.

Let $p_i = \mathbb{P}\left[\xi = a_i\right]$ and let $\jmath = \jmath(\xi)$ be the indicator random variable $\jmath \in \{1, \ldots, n\}$ such that $\jmath = j$ if and only if $\xi = a_j$.

Defining

$$\bar{z} = \sum_{i=1}^{n} p_i z_i \ , \tag{3.49}$$

$$E_x(y,z) = N\big(x-y\big)^2 + \bar{z} \ , \tag{3.50}$$

$$V(y,z) = \sum_{i=1}^{n} p_i \Big\{ \Big(\mu(z_i - \bar{z}) - \frac{a_i y}{\sqrt{N}} \Big)^2 + J_{k-1}(y, z_i) \Big\} \tag{3.51}$$

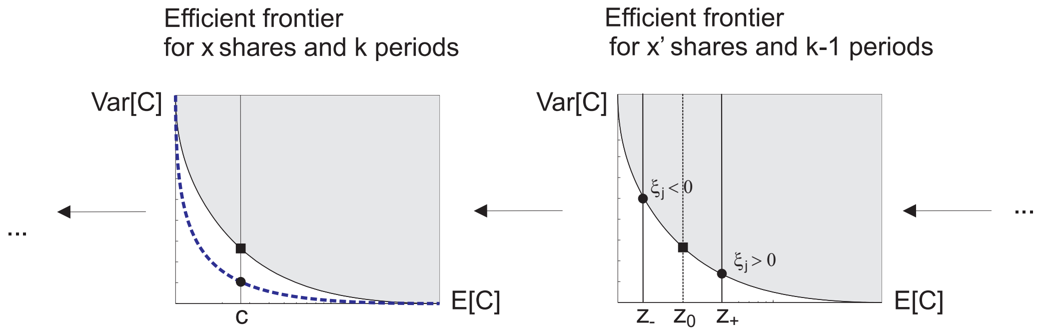

FIGURE 3.2. Illustration of the backwards optimization in the binomial model in Section 3.5. If we already have the efficient frontiers to sell in $(k-1)$ periods and want to compute the efficient frontier to sell in k periods, we need to determine an optimal control (y, z_+, z_-): we sell $x-y$ shares in the first period, and commit ourselves to trading strategies for the remaining $(k-1)$ periods – chosen from the set of $(k-1)$-period efficient strategies *depending on whether the stock goes up or down.* If the stock goes up, we follow the efficient strategy with expected cost z_+. If it goes down, we follow the efficient strategy with expected cost z_-. The coice $z_+ = z_- = z_0$ would lead to a path-independent strategy. By choosing $z_+ > z_-$ (and y) optimally, we can reduce the variance of the whole strategy – measured at the beginning of the k-period trading time window. Instead of the square shaped point on the static k-period frontier, we obtain a point on the improved k-period frontier (blue dashed line).

the optimization problem (3.38,3.39) reads

$$J_k(x,c) = \min_{(y,z_1,\ldots,z_n)\in\mathbb{R}^{n+1}} \left\{ V(y,z) \;\middle|\; \begin{array}{ll} E_x(y,z) \le c & (C1) \\ \dfrac{Ny^2}{k-1} \le z_i,\; i=1\ldots n & (C2) \\ 0 \le y \le x & (C3) \end{array} \right\} \tag{3.52}$$

in $\mathbf{dom}\, J_k = \{(x,c) \mid x \ge 0,\; c \ge Nx^2/k\}$, with the terminal value function $J_1(x,c)$ in (3.37).

Thus, we have to solve an optimization problem with $n+1$ variables in each step.

In the simplest example, we set $\Omega = \{\pm 1\}$, i.e. $n = 2$. That is, we consider a binomial tree where the stock price goes either up or down one tick during each of the N trading periods. Figure 3.2 illustrates the resulting binomial adaptivity model.

The corresponding dynamic program (3.44, 3.45) for $J'(x,c)$, defined in (3.43), is given by

$$J'_k(x,c) = \min_{(y,z_1,\ldots,z_n)\in\mathbb{R}^{n+1}} \left\{ V(y,z) \;\middle|\; \begin{array}{ll} E_x(y,z) \le c & (C1) \\ \dfrac{Ny^2}{k-1} \le z_i,\; i=1\ldots n & (C2) \\ 0 \le y \le x & (C3) \\ Ny^2 \ge z_i,\; i=1\ldots n & (C4) \end{array} \right\}, \tag{3.53}$$

where we now have the additional constraint (C4). The terminal value function $J'_1(x,c) = J_1(x,c)$ remains unchanged. As (3.44,3.45), (3.53) does not consitute a convex problem due to the non-convex constraint (C4).

3.5.2. Aggressiveness In-the-money (AIM) of Optimal Adaptive Policies. In this section we shall prove that for an optimal adaptive policy the control $z(\xi)$ is positively correlated to the stock price return.

In terms of the controls $z_1, \ldots, z_n$ this means that $z_1 \leq \cdots \leq z_n$ since by definition (3.48) for $i < j$, the control z_i corresponds to a price change ξ that is smaller than for z_j. The interpretation is that if the stock price goes up, we sell faster (higher expected cost z_i for the remainder). As in Chapter 2, we obtain a AIM strategy (aggressive in the money), which burns part of the windfall trading gains to sell faster and reduce the risk for the time left.

THEOREM 3.8. *Let $\mu > 0$ and $3 \leq k \leq N$. For $Nx^2/k < c < Nx^2$ the optimal control for $\widetilde{J}_k(x,c)$ in (3.49,3.50,3.51,3.52) satisfies*

$$y > 0 \quad \textit{and} \quad z_1 \leq z_2 \leq \cdots \leq z_k \ . \tag{3.54}$$

For $c \geq Nx^2$, the optimal control is $y = z_1 = \cdots = z_k = 0$ (immediate liquidation) and for $c = Nx^2/k$ it is $y = (k-1)x/k$ and $z_1 = \cdots = z_k = x^2(k-1)N/k^2$ (linear profile).

PROOF. It is easy to see that for $c \geq E_{\text{inst}} = Nx^2$ we have $\widetilde{J}_k(x,c) = 0$ with optimal control $y = z_1 = \cdots = z_k = 0$. For $c = Nx^2/k$, the only point that satisfies (C1)–(C3) in (3.52) is $y = (k-1)x/k$ and $z_1 = \cdots = z_k = x^2(k-1)N/k^2$, which indeed corresponds to the linear strategy.

For $c < Nx^2$, suppose $y = 0$. Then, by (3.50) we have $E_x(y,z) \geq Nx^2$, a contradiction.

We prove $z_1 \leq z_2 \leq \cdots \leq z_k$ by contradiction as well. Suppose $z_s > z_r$ for $r > s$. Let

$$\bar{z} = \frac{p_r z_r + p_s z_s}{p_r + p_s} \quad \text{and} \quad \delta = z_r - z_s \ .$$

Then $\delta < 0$, and

$$z_r = \bar{z} + \frac{p_s}{p_r + p_s}\delta \quad \text{and} \quad z_s = \bar{z} - \frac{p_r}{p_r + p_s}\delta \ . \tag{3.55}$$

Let

$$A = \mu \sum_{j=1}^{n} p_j z_j = (p_r + p_s)\mu\bar{z} + \mu \sum_{j \neq r,s} p_j z_j \ ,$$

and

$$\Delta = V(y, \tilde{z}_1, \ldots, \tilde{z}_n) - V(y, z_1, \ldots, z_n) \tag{3.56}$$

with $\tilde{z}_i = z_i$ for $i \notin \{r,s\}$ and $\tilde{z}_r = \tilde{z}_s = \bar{z}$. Since $E_x(y, \tilde{z}_1, \ldots, \tilde{z}_n) = E_x(y, z_1, \ldots, z_n)$, the control $(y, \tilde{z}_1, \ldots, \tilde{z}_n)$ satisfies (C1) in (3.52). Since z satisfies (C2) in (3.52), i.e. $z_r \geq y^2 N/(k-1)$, we have $\bar{z} > z_r \geq y^2 N/(k-1)$, and hence $\tilde{z}$ also satisfies (C2).

We shall prove that $\Delta < 0$, contradicting the optimality of $(y, z_1, \ldots, z_n)$. To shorten notation, let $J(x,c) = \widetilde{J}_{k-1}(x,c)$.

Since $J(x,c)$ is convex,

$$p_r J(y, z_r) + p_s J(y, z_s) \geq (p_r + p_s)\, J\Big(y, \frac{p_r z_r + p_s z_s}{p_r + p_s}\Big) = (p_r + p_s)\, J(y, \bar{z}) \ .$$

Hence,

$$\begin{aligned}\Delta &= \sum_{i=r,s} p_i\big[\big(\mu\bar{z} - \frac{ya_i}{\sqrt{N}} - A\big)^2 - \big(\mu z_i - \frac{ya_i}{\sqrt{N}} - A\big)^2 + J(y,\bar{z}) - J(y,z_i)\big] \ , \\ &\le \sum_{i=r,s} p_i\big[\big(\mu\bar{z} - \frac{ya_i}{\sqrt{N}} - A\big)^2 - \big(\mu z_i - \frac{ya_i}{\sqrt{N}} - A\big)^2\big] \ .\end{aligned}$$

Because of (3.55),

$$\big(\mu\bar{z} - \frac{ya_r}{\sqrt{N}} - A\big)^2 - \big(\mu z_r - \frac{ya_r}{\sqrt{N}} - A\big)^2 = -2\big(\mu\bar{z} - \frac{ya_r}{\sqrt{N}} - A\big)\frac{\mu p_s \delta}{p_r + p_s} - \frac{p_s^2\delta^2}{(p_r+p_s)^2}$$

and

$$\big(\mu\bar{z} - \frac{ya_s}{\sqrt{N}} - A\big)^2 - \big(\mu z_s - \frac{ya_s}{\sqrt{N}} - A\big)^2 = 2\big(\mu\bar{z} - \frac{ya_s}{\sqrt{N}} - A\big)\frac{\mu p_r \delta}{p_r + p_s} - \frac{p_r^2\delta^2}{(p_r+p_s)^2} \ .$$

Thus,

$$\begin{aligned}\Delta &\le -\frac{(p_r p_s^2 + p_s p_r^2)\delta^2\mu^2}{(p_r+p_s)^2} - \frac{2p_r p_s \delta\mu}{p_r+p_s}\frac{y(a_s - a_r)}{\sqrt{N}} \\ &\le \frac{2p_r p_s \mu}{(p_r+p_s)\sqrt{N}}\delta y(a_r - a_s) \ .\end{aligned}$$

Since $a_r > a_s$, $\delta < 0$ and $y > 0$ (as shown above), we conclude $\Delta < 0$, which completes the proof. □

In the Markowitz portfolio problem in a multiperiod setting that we shall consider in Chapter 5, where a closed-form analytical solution for the dynamic program can be obtained, the optimal control function $z : \Omega \to \mathbb{R}$ is perfectly anti-correlated to the portfolio return in the current period. Unfortunately, this argument cannot be used for (3.37, 3.38, 3.39) because we effectively have the constraint "$z \ge E_{\text{lin}}(x,k)$" (this is enforced by $J_k(x,c) = \infty$ for all $z < E_{\text{lin}}(x,k)$ by the definition the the value function (3.19, 3.20), which is also easily verified inductively for the dynamic program). Since in general the stock price change ξ may take arbitrarily large positive and negative values, a lower bounded z can never be *perfectly* correlated to ξ.

3.5.3. Approximation of Optimal Control. If $|\Omega|$ is very large, or we even have $|\Omega| = \infty$ (for instance, if $\xi \sim \mathcal{N}(0,1)$), approximate solutions can be obtained by restricting the space of admissible controls instead of considering all measurable functions on Ω. In the following, we assume $\Omega \subseteq \mathbb{R}$.

We partition the real line into n intervals $I_1, \dots, I_n$ with

$$I_1 = (-\infty, a_1), \quad I_2 = [a_1, a_2), \dots, I_{n-1} = [a_{n-1}, a_n), \quad I_n = [a_n, \infty)$$

for $a_1 < \cdots < a_n$. For given $(z_1, \dots, z_n) \in \mathbb{R}^n$ we define the step function $z : \mathbb{R} \to \mathbb{R}$ by

$$z(\xi) = z_j \quad \text{for} \quad \xi \in I_j \ .$$

We want to approximate the space of all measurable functions on Ω by considering only step functions of this type. Not least due to the recursive structure of the dynamic program, formal convergence is not obvious, and we leave this as an open question. In practice, this scheme will certainly yield good results.

We let $p_i = \mathbb{P}[\xi \in I_i]$, $E_i = \mathbb{E}[\xi \,|\, \xi \in I_i]$ and $V_i = \mathrm{Var}[\xi \,|\, \xi \in I_i]$, and now define the indicator random variable $\jmath = \jmath(\xi)$ such that $\jmath = j$ if and only if $\xi \in I_j$.

By the law of total variance in Lemma 3.1, we obtain

$$\begin{aligned}
\mathrm{Var}\left[\mu z_\jmath - \frac{\xi y}{\sqrt{N}}\right] &= \mathrm{Var}\left[\mathbb{E}\left[\mu z_\jmath - \frac{\xi y}{\sqrt{N}}\,\middle|\, \jmath\right]\right] + \mathbb{E}\left[\mathrm{Var}\left[\mu z_\jmath - \frac{\xi y}{\sqrt{N}}\,\middle|\, \jmath\right]\right] \\
&= \mathrm{Var}\left[\mu z_\jmath - \frac{y E_\jmath}{\sqrt{N}}\right] + \mathbb{E}\left[\frac{y^2 V_\jmath}{N}\right] \\
&= \sum_{i=1}^n p_i \left(\mu z_i - \frac{E_i y}{\sqrt{N}} - \mu \sum_{j=1}^n p_j z_j\right)^2 + \frac{y^2}{N}\sum_{i=1}^n p_i V_i \ .
\end{aligned}$$

Similar to (3.49, 3.50, 3.51, 3.52), we define

$$\begin{aligned}
\bar{z} &= \sum_{i=1}^n p_i z_i \ , \\
\widetilde{E}_x(y,z) &= N\big(x-y\big)^2 + \bar{z} \ , \\
\widetilde{V}(y,z) &= \sum_{i=1}^n p_i \Big\{ \Big(\mu(z_i - \bar{z}) - \frac{E_i y}{\sqrt{N}}\Big)^2 + \frac{y^2 V_i}{N} + \widetilde{J}_{k-1}(y, z_i) \Big\}
\end{aligned}$$

and the optimization problem now reads

$$\widetilde{J}_k(x,c) = \min_{(y,z_1,\dots,z_n)\in\mathbb{R}^{n+1}} \left\{ \widetilde{V}(y,z) \;\middle|\; \begin{array}{ll} \widetilde{E}_x(y,z) \le c & (C1) \\ \dfrac{Ny^2}{k-1} \le z_i,\ i = 1\dots n & (C2) \\ 0 \le y \le x & (C3) \end{array} \right\} \ .$$

An approximation of the variant $J'(x,c)$ is obtained analogously.

The simplest case of an approximation of this type is to consider a binomial framework: set $n = 2$ with $a_1 = 0$ ($I_1 = (-\infty, 0)$ and $I_2 = [0,\infty)$), i.e. only respond to whether $\xi \ge 0$ or $\xi < 0$.

3.5.4. Numerical Results. For numerical computations, we discretize the state space of the value functions $J_k(x,c)$. The figures presented in this section were generated for $T = 1$, $N = 50$ time steps (i.e. $\tau = 1/50$) with $N_x = 150$ grid points for the relative portfolio size $x \in [0,1]$ and $N_c = 150$ in the cost dimension (i.e. N_c points on the frontier for each value of x). We use $\xi \sim \mathcal{N}(0,1)$ and the approximation of the optimal controls described in Section 3.5.3 with $n = 2$ and $a_1 = 0$. Starting with $J_1(x,c)$, we successively determine $J_2(x,c), \dots, J_N(x,c)$ by means of (3.52), using interpolated values from the grid data of the previous value function. We use standard direct search methods to find the optimal control (y, z_+, z_-). Since the optimization problem is convex, local minima are global optimal solutions. For each level of x we have to trace an efficient frontier. The function value (3.11) for the linear strategy at the upper-left end of the frontier is readily available; from there, we work towards the right (increasing c) and compute optimal controls for each c by taking the optimal controls for the point $c - h$ (where h is the discretization along the cost dimension) as the starting point for the iteration. Note that the optimal control for $c - h$ is indeed a feasible starting point for the optimization problem with maximal cost c.

Figure 3.4 shows the set of efficient frontiers at the initial time $t = 0$ for the entire initial portfolio (i.e. relative portfolio size $x = 1$) for different values of the market power $0 \le \mu \le 0.15$. (Recall the discussion about the order of magnitude for μ in Section 3.4.2.) The x-axis is the expectation of total cost, and the y-axis its variance. We scale both, expectation and variance, by their values for the linear trajectories (see (3.11)). The two blue marks on the frontier for $\mu = 0.15$ correspond to optimal adaptive strategies with the same mean, but lower variance (below the black mark) and same variance, but lower mean (to the left of the black mark) as the static strategy corresponding to the black mark. The inset shows the cost distributions associated with these three strategies (trading costs increase to the right of the x-axis). The static cost distribution is readily available as a Gaussian with mean and variance according to its location along the frontier. For the adaptive strategies, this is not the case. Since the solution of the dynamic program is only the optimal strategy (given as a series of optimal one-step controls as a function of the state variables), we determine the associated cost distributions by Monte Carlo simulation. Suppose we want to simulate the strategy that has expected cost at most E (measured at time $t = 0$). Then, using (3.27, 3.28, 3.22), for each of $m = 10^5$ randomly sampled paths of the stock price $(\xi_1, \ldots, \xi_N)$ the final cost C_N is obtained by sequentially applying the optimal one-step controls $(y_k^*(x,c), z_{k\pm}^*(x,c))$ associated with $J_k(x_k, c)$ (respectively, their interpolated values over the discretized state space),

$$
\begin{aligned}
x_{i+1} &= y_i^*(x_i, c_i) \\
c_{i+1} &= \begin{cases} z_{i+}^*(x_i, c_i) & \xi_{i+1} \ge 0 \\ z_{i-}^*(x_i, c_i) & \xi_{i+1} < 0 \end{cases} \\
C_{i+1} &= C_i + \mu N (x_i - x_{i+1})^2 - x_{i+1}\xi_{i+1}/\sqrt{N}
\end{aligned}
$$

with $x_0 = 1$, $C_0 = 0$ and the initial limit $c_0 = E$ for the expected cost. C_i is measured in units of $\sigma\sqrt{T}X$, the standard deviation of the initial portfolio value due to the stock price volatility across the trading horizon.

The adaptive cost distributions are slightly skewed, suggesting that mean-variance optimization may not give the best possible solutions. Figure 3.5 shows four static and adaptive cost distributions along the frontier. In the upper left corner (near the linear strategy), the adaptive cost distributions are almost Gaussian (Point #1); indeed, for high values of V adaptive and static strategies coincide. As we move down the frontiers (towards less risk-averse strategies), the skewness first increases (Point #2). Interestingly, as we move further down – where the improvement of the adaptive strategy becomes larger – the adaptive distributions look more and more like Gaussian again (Point #3 and #4). All adaptive distributions are strictly preferable to their reference static strategy, since they have lower probability of high costs and higher probability of low costs. Table 1 compares the semi-variance, value-at-risk (VaR) and conditional value-at-risk (CVaR) (see Artzner, Delbaen, Eber, and Heath, 1999, for instance) for the four distribution pairs shown in Figure 3.5. For Gaussian random variables, mean-variance is consistent with expected utility maximization as well as stochastic dominance (see for instance Levy (1992); Bertsimas et al. (2004)). As the adaptive distributions are indeed not too far from Gaussian, we can expect mean-variance to give reasonable results.

One way to improve on this further would be to add constraints on the skewness to the optimization problem (3.9); for instance, we could require $\operatorname{Skew}[C(X,N,\pi)] \geq 0$. In fact, using a generalization of Lemma 3.1 for higher central moments ("law of total cumulance"), such an extension of the dynamic programming principle in Section 3.4 is indeed possible.

To illustrate the behavior of adaptive policies in the binomial framework, Figure 3.3 shows trajectories for two sample paths of the stock price in a small instance with $N = 4$. The inset shows the underlying binomial decision tree. As can be seen, the optimal adaptive policies are indeed "aggressive in the money". If the stock price goes up, we incur unexpectedly smaller total trading cost and react with selling faster (burning some of the gains), whereas for falling stock price, we slow down trading. Figure 3.6 shows optimal adaptive and static trading for $N = 50$. The static strategy is chosen such that it has the same expected cost (yet higher variance) as the adaptive policy.

The numerical results presented so far were obtained using the value function J_N. Let us briefly discuss the results for the value function definition J'_N. As mentioned there, $J'_N(x,c) \geq J_N(x,c)$. In the specification of J_N (respectively, $\mathscr{E}$) the trader can reduce his variance by destroying money in order to make a positive outcome less so (see discussion in Section 3.4.3). In the specification of J'_N (respectively, $\mathscr{E}'$) this is not possible. The numerical results show that while this effect is important for small values of N and large values of μ (see Figure 3.8), it diminishes very rapidly as N increases (see Figure 3.7). In fact, the value of $\mu = 2$ in Figure 3.7 is very large (recall our discussion for the order of μ in Section 3.4.2), and for realistic values of μ the difference is even smaller. That is, while the specification of J_N and $\mathscr{E}(X,N)$ allow for the undesirable peculiarity that the trader gives away money in positive outcomes, our numerical results show that this effect does not play a big role in practice. Compared to J'_N, the specification of J_N has the big advantage that the associated dynamic program is convex, which makes the numerical optimization significantly easier.

3.5.5. Comparison with Results of Single-update Strategies. It is interesting to compare the improvement of the optimal adaptive strategy with the improvement obtained by the single update strategies from Chapter 2. Before doing so, let us first stress that the single update framework presented in the previous chapter is formulated in the continuous-time version of the Almgren/Chriss model, whereas the solution obtained with the dynamic programming technique in this chapter works in the discrete-time version. As mentioned in Section 3.2, for $\tau \to 0$ the discrete time model converges to a continuous-time model, but for $\tau \gg 0$ the two models differ.

In the following let us refer to the optimal adaptive strategy from this chapter (in the binomial framework) with $N = 50$ time steps as strategy M ("Multi-Update"), and to the single update strategy (with $n = 2$ there) from the previous chapter as strategy S ("Single-Update").

Figures 2.2 and 3.4 show the improved efficient frontiers for the same values of μ for both strategies. Expectation and variance are measured relative to the cost of the respective linear strategy, i.e. the linear strategy in continuous-time for strategy S (Figure 2.2) and the linear strategy in discrete-time strategy M (Figure 3.4). The sample distributions shown in the insets in both figures use a static strategy with mean $E = 7.00 \cdot E_{lin}$, and variance $V = 0.107 \cdot V_{lin}$ (in case of strategy S, Figure 2.2) and $V = 0.0884 \cdot V_{lin}$ (in case of strategy M, Figure 3.4); note

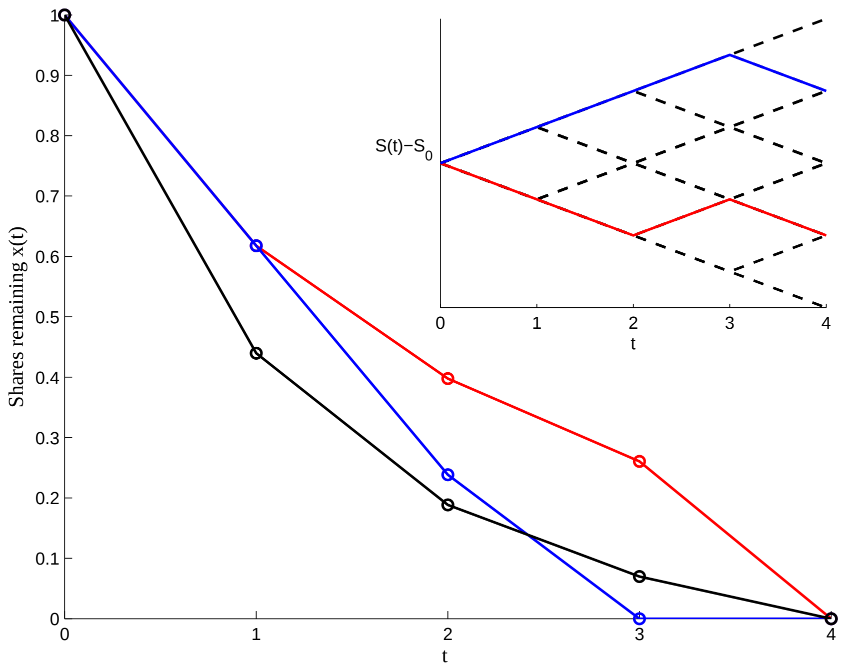

Figure 3.3. Optimal adaptive trading for $N = 4$ time steps, illustrating the binomial adaptivity model. The blue trajectory corresponds to the rising stock price path, and sells faster than red trajectory (falling stock price path). The inset shows the schematics of the stock price on the binomial tree. x_1 at $t = 1$ is the same for all adaptive trajectories because x_1 is determined at $t = 0$ with no information available. Only from $t = 1$ onwards the adaptive trajectories split.

that the difference in variance (for the same level of expected cost) is due to the slightly different models (discrete-time vs. continuous-time). For strategy M (Figure 3.4), the two improved adaptive distributions have mean and variance of $(E = 4.00 \cdot E_{lin}, V = 0.0884 \cdot V_{lin})$ and $(E = 7.00 \cdot E_{lin}, V = 0.0334 \cdot V_{lin})$. For strategy S (Figure 3.4), the improved adaptive distributions have mean and variance of $(E = 4.47 \cdot E_{lin}, V = 0.107 \cdot V_{lin})$ and $(E = 7.00 \cdot E_{lin}, V = 0.0612 \cdot V_{lin})$. Figure 3.9 shows the entire efficient frontiers (with $\mu = 0.15$) for the two models, single-update in continuous-time (Chapter 2) and multi-update discrete-time (Chapter 3) compared to their corresponding static frontiers.

Due to the difference between the continuous-time and the discrete-time model those values have to be approached with caution. However, it is evident that single update strategies achieve a relative improvement that is already comparable to the improvement of the adaptive strategies developed in this chapter. Thus, from a practical point of view the single-update framework is indeed attractive, as the computational cost for solving the single-step optimization problem is naturally much less than for solving the dynamic program.

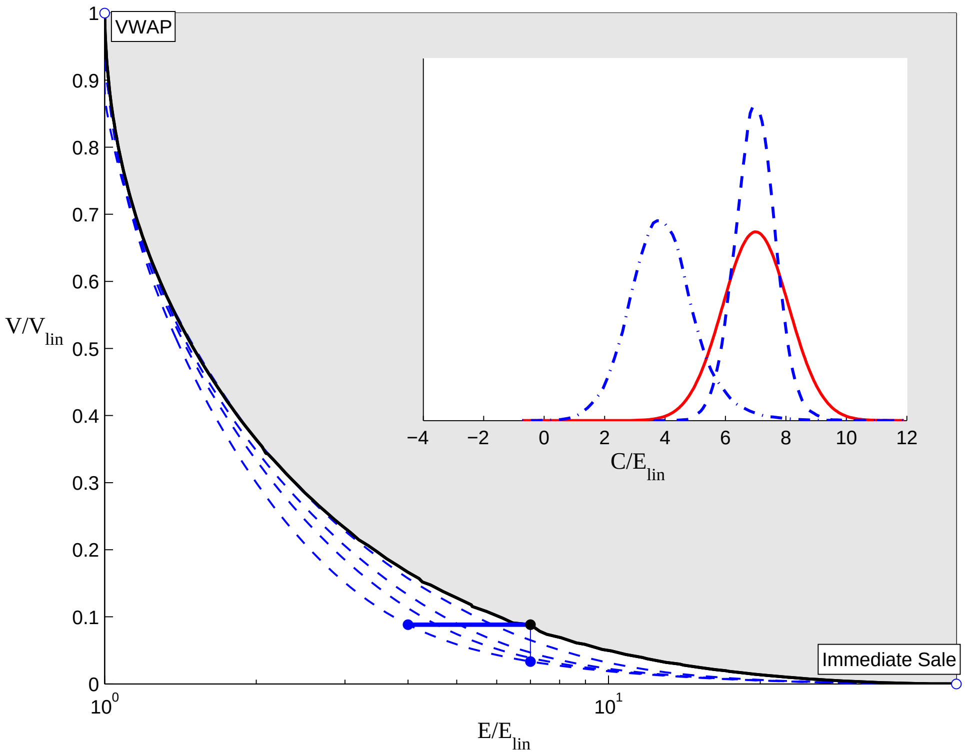

FIGURE 3.4. Adaptive efficient frontiers (in the binomial model described in Section 3.5 for different values of the market power $\mu \in \{0.025, 0.05, 0.075, 0.15\}$, and $N = 50$. The expectation and variance of the total trading cost are normalized by their values of a linear trajectory (VWAP) and plotted in a semilogarithmic scale. The grey shaded region is the set of values accessible to static trading trajectories and the black line is the static efficient frontier, which is also the limit $\mu \to 0$. The blue dashed curves are the improved efficient frontiers, with the improvement increasing with μ. The inset shows the distributions of total cost corresponding to the three points marked on the frontiers.

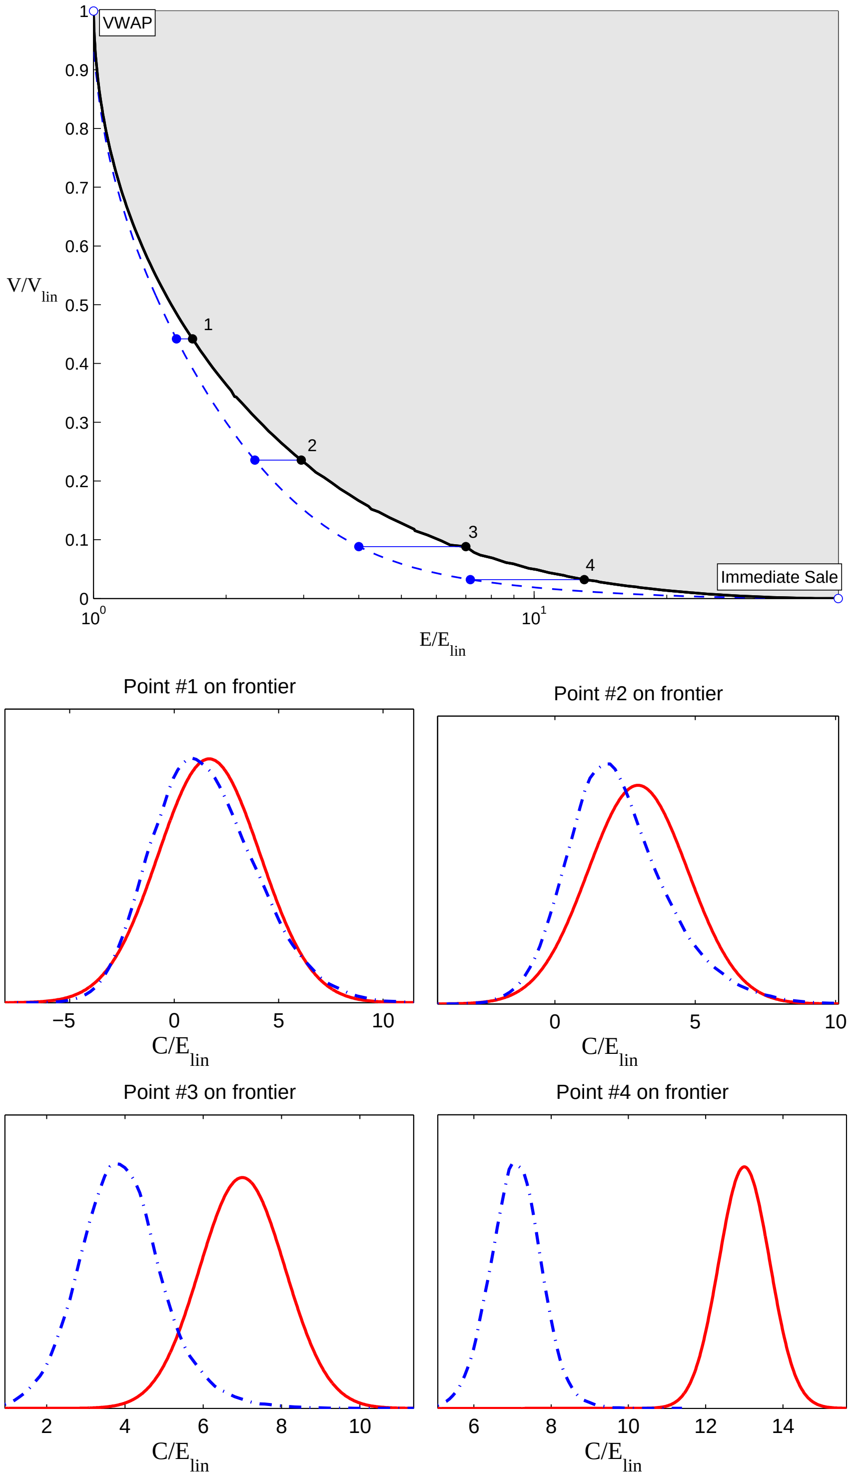

FIGURE 3.5. Distributions of total cost corresponding to four points on the frontier, $\mu = 0.15$ and $N = 50$. The adaptive distributions are reasonably close to Gaussian; recall that the static distribution functions are always exactly Gaussian by construction.

	#1		#2		#3		#4	
	static	adapt	static	adapt	static	adapt	static	adapt
$\mathbb{E}[\cdot]$	1.68	1.52	2.96	2.27	7.00	3.92	13.00	7.09
$\mathrm{Var}[\cdot]$	5.98	5.98	3.19	3.19	1.20	1.20	0.44	0.44
$\mathrm{SVar}[\cdot]$	3.01	3.35	1.61	1.89	0.60	0.68	0.22	0.22
$\mathrm{VaR}_{5.0\%}$	5.70	5.85	5.90	5.58	8.80	5.83	14.09	8.17
$\mathrm{VaR}_{2.5\%}$	6.47	6.77	6.46	6.38	9.14	6.34	14.29	8.41
$\mathrm{VaR}_{1.0\%}$	7.37	7.91	7.12	7.35	9.54	7.02	14.54	8.73
$\mathrm{VaR}_{0.5\%}$	7.98	8.62	7.56	8.00	9.82	7.56	14.70	8.97
$\mathrm{VaR}_{0.1\%}$	9.23	10.18	8.48	9.45	10.38	8.57	15.04	9.43
$\mathrm{CVaR}_{5.0\%}$	6.72	7.09	6.65	6.67	9.25	6.56	14.36	8.51
$\mathrm{CVaR}_{2.5\%}$	7.43	7.91	7.15	7.39	9.57	7.06	14.54	8.75
$\mathrm{CVaR}_{1.0\%}$	8.23	8.87	7.71	8.26	9.93	7.71	14.77	9.06
$\mathrm{CVaR}_{0.5\%}$	8.76	9.52	8.14	8.87	10.20	8.19	14.91	9.28
$\mathrm{CVaR}_{0.1\%}$	9.94	10.96	8.97	10.16	10.69	9.19	15.20	9.80

TABLE 1. Statistics for the adaptive and static cost distribution functions shown in Figure 3.5, obtained by Monte Carlo simulation (10^5 sample paths). For the random variable C, the total cost in units of E_{lin}, the value-at-risk VaR_β is defined by $\mathbb{P}[C \geq \mathrm{VaR}_\beta(C)] = \beta$, and the conditional-value-at-risk $\mathrm{CVaR}_\beta(C) = \mathbb{E}[C \mid C \geq \mathrm{VaR}_\beta(C)]$. Thus, low values for VaR and CVaR are desirable.

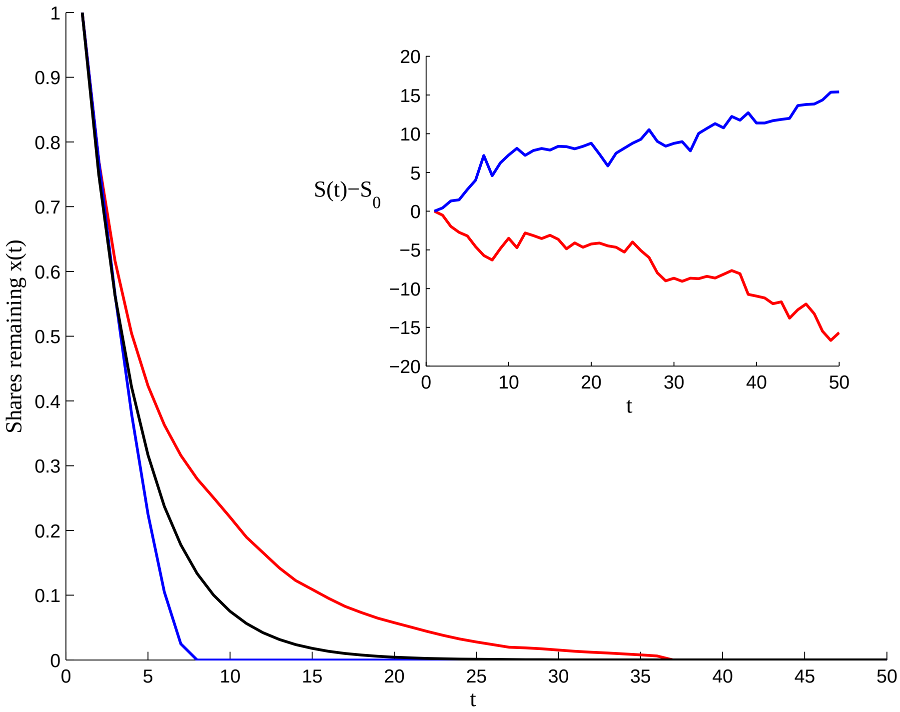

FIGURE 3.6. Optimal adaptive strategy for the point on the efficient frontier in Figure 3.4, having the same variance but lower expected cost than the static trajectory (solid black line), computed using 50 time steps. Specific trading trajectories are shown for two rather extreme realizations of the stock price process. The blue trajectory incurs impact costs that are slightly higher than the static trajectory, but has trading gains because it holds more stock as the price rises. The red trajectory has lower impact costs because of its slower trade rate, but it has trading losses because of the price decline. The mean and variance of the adaptive strategy cannot be seen in this picture, because they are properties of the entire ensemble of possible realizations.

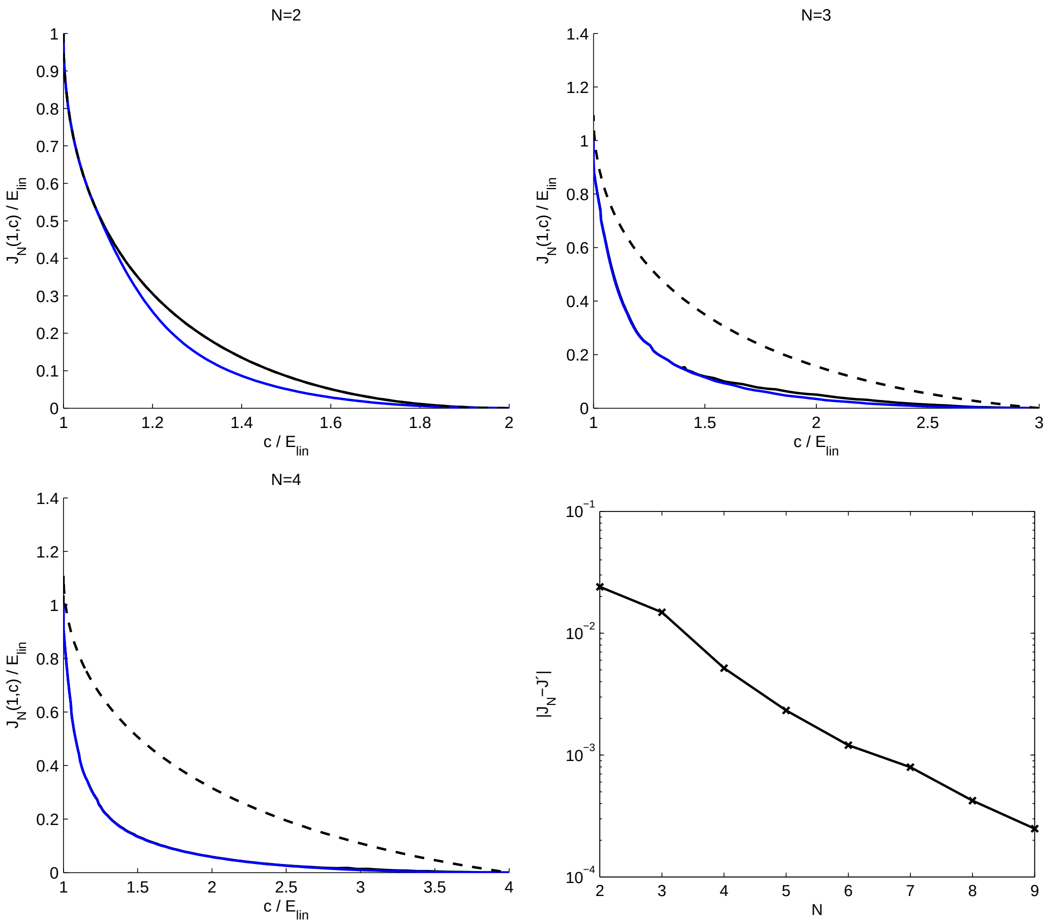

FIGURE 3.7. The first three plots show $J_N(x,c)$ (blue solid line) vs. $J'_N(x,c)$ (black solid line) for $N = 2, 3, 4$ and $\mu = 2$. For $N = 2$, the two curves are clearly separated. For $N = 3$ there is only a very small visible difference for larger values of c (around $c/E_{\text{lin}} \approx 2$), and for $N = 4$ the two curves are practically identical. The dashed line is the static frontier ($\mu = 0$), which coincides for $N = 2$ with J'_2. The last plot shows $\|J_N - J'_N\|_1$ as function of N in semilogarithmic scale.

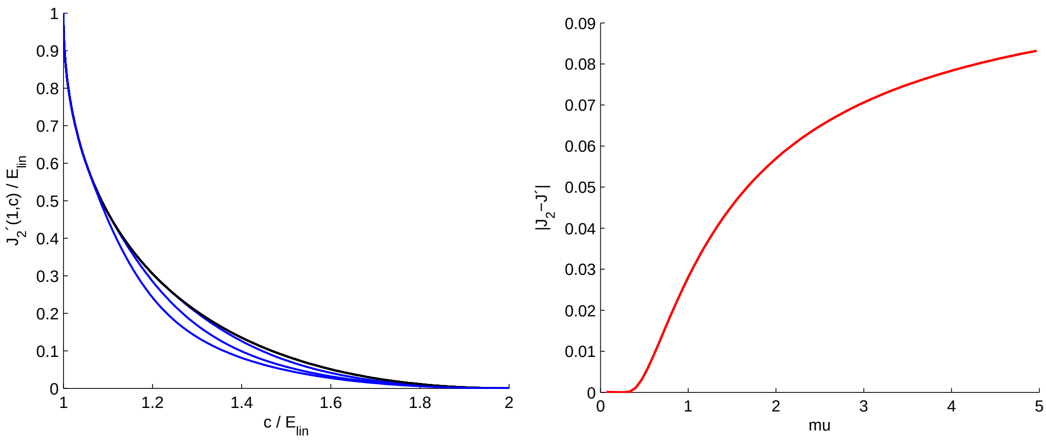

FIGURE 3.8. Left plot shows $J_2(x,c)$ for different values of $\mu \in \{0.25, 0.5, 0.75, 1.0\}$ as a function of c (and $x = 1$). The black line is $J'_2(x,c)$. The right plot shows $\|J_2 - J'_2\|_1$ as a function of μ.

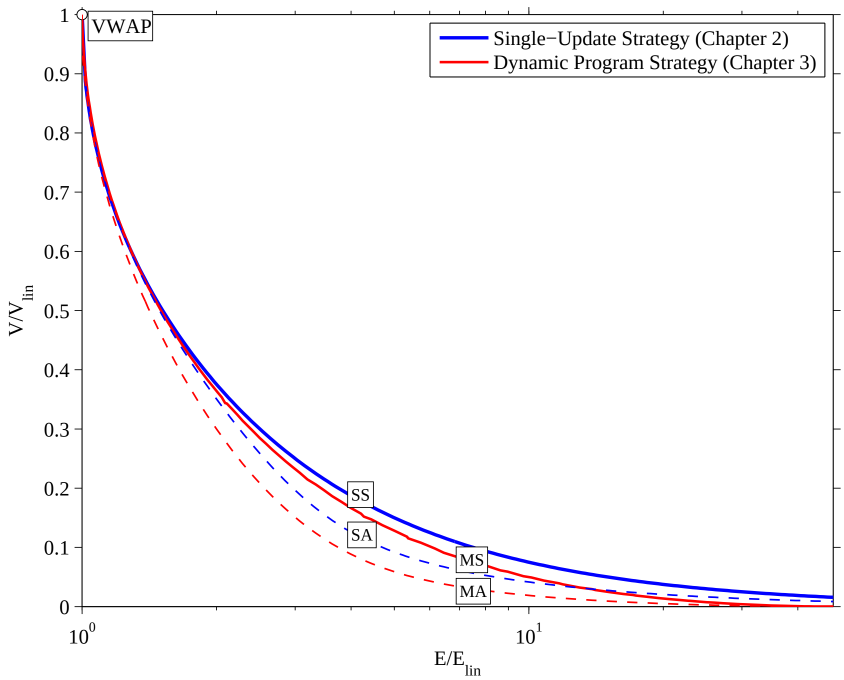

FIGURE 3.9. The blue solid line (labeled SS) is the static frontier (i.e. $\mu = 0$) in the continuous-time model of Chapter 2, and the red solid line (MS) is the static frontier in the discrete-time setting of Chapter 3 (with $N = 50$ time steps). Due to the discretization effect, the two frontiers are not exactly the same. The dashed red curve (SA) and blue curve (MA) are the adaptive frontiers in the two models, both for $\mu = 0.15$. All values are scaled by their values of the corresponding linear strategies (in the continuous-time and discrete-time model, respectively).

CHAPTER 4

Bayesian Adaptive Trading with Price Appreciation

In this chapter, we discuss a model that incorporates another major source of slippage besides market impact, namely *price appreciation.* As proposed in Section 1.3, we consider a model in which the trader uses information from observations of the price evolution during the day to continuously update his estimate of the price momentum resulting from other traders' target sizes and directions. The motivation for this model is the *daily trading cycle*: large institutional traders make investment decisions in the morning and implement them through the trading day. Using this information we determine optimal trade schedules to minimize total expected cost of trading.

4.1. Introduction

Most current models of optimal trading strategies view time as an undifferentiated continuum, and other traders as a collection of random noise sources. Trade decisions are made at random times and trade programs have random durations. Thus, if one observes buy pressure from the market as a whole, one has no reason to believe that this pressure will last more than a short time. From the point of view of optimal trading, price motions are purely random.

Here we present a model for price dynamics and optimal trading that explicitly includes the *daily trading cycle*: large institutional participants make investment decisions overnight and implement them through the following trading day. Within each day, the morning is different from the afternoon, since an intelligent trader will spend the early trading hours collecting information about the targets of other traders, and will use this information to trade in the rest of the day. His main source of information is the observation of the price dynamics.

As in Chapters 2 and 3, this will result in an adaptive execution algorithm. However, in the previous chapter we assumed a pure random walk model with no price momentum or serial correlation, using pure classic mean variance. The adaptive behavior discovered there stems from the trader's risk-aversion. In this chapter we consider a risk-neutral trader, whose motivation for accelerating or slowing down trading is the momentum of the asset price.

The following model may be interpreted as one plausible way to model such price momentum that the trader observes during the course of the day. There is an underlying drift factor, caused by the net positions being executed by other institutional investors. This factor is approximately constant throughout the day because other traders execute across the entire day. Thus price increases in the early part of the day suggest that this factor is positive, which suggests that prices will continue to increase throughout the day. This is different from a short-term momentum model in which the price change across one short period of time is correlated with the price change across a preceding period; most empirical evidence shows that such correlation is weak if it exists at all. Our strategies will exploit this momentum to minimize the expected value of trading costs, somewhat in the spirit of Bertsimas and Lo (1998), except that because we focus on long-term momentum, our results can obtain higher gains.

Using techniques from dynamic programming as well as the calculus of variations, we determine trading strategies that minimize the expectation of total cost. Surprisingly, optimal strategies can be determined by computing a "static" optimal trajectory at each moment, assuming that the best parameter estimates of the unknown price momentum at that time will persist until the end of the day. Loosely speaking, this will be because the expected value of future updates is zero, and thus they do not change the strategy of a risk-neutral trader. In fact, the parameter estimates will change as new price information is observed. The actual optimal strategy will use only the initial instantaneous trade rate of this trajectory, continuously responding to price information. This is equivalent to following the strategy only for a very small time interval Δt, then recomputing. Hence, the optimal strategy is highly dynamic.

In Section 4.2 below, we present a model of Brownian motion with a drift whose distribution is continuously updated using Bayesian inference. In Section 4.3 we present optimal trading strategies, which we illustrate in Section 4.4 with numerical examples.

4.2. Trading Model Including Bayesian Update

Let $(\Omega, \mathcal{F}, P)$ be a probability space with a filtration $\{\mathcal{F}_t\}$ that satisfies the usual conditions, i.e. $(\Omega, \mathcal{F}, P)$ is complete, $\mathcal{F}_0$ contains all P-null sets in $\mathcal{F}$, and $\mathcal{F}_t$ is right-continuous (see for instance Yong and Zhou (1999, p. 63)).

As in Chapters 2 and 3, we model stock prices with an arithmetic instead of the more traditional geometric model. As in the previous chapters, also here our focus is on intraday time scales, where the difference is negligible. We consider trading in a single asset whose price is given by the $\{\mathcal{F}_t\}$-adapted stock price process $S(t)$,

$$S(t) = S_0 + \alpha t + \sigma B(t) \qquad \text{for } t \geq 0 \quad , \tag{4.1}$$

where $B(t)$ is a standard Brownian motion, σ an absolute volatility and α a drift. The reader is referred to standard textbooks for background information on stochastic processes and stochastic calculus, e.g. Oksendal (2003). In the presence of intraday seasonality, we interpret t as a volume time relative to a historical profile.

Our interpretation is that volatility comes from the activity of the "uninformed" traders, whose average behavior can be predicted reasonably well. Mathematically, we assume that the value of σ is known precisely (for a Brownian process, if σ is known to be constant, it can be estimated arbitrarily precisely from an arbitrarily short observation of the process).

We interpret the drift as coming from the activity of other institutional traders, who have made trade decisions before the market opens, and who expect to execute these trades throughout the day. If these decisions are in the aggregate weighted to buys, then this will cause positive price pressure and an upwards drift; conversely for overall net selling. We do not know the net direction of these trades but we can infer it by observing prices. We implicitly assume that these traders are using VWAP-like strategies rather than arrival price, so that their trading is not "front-loaded". This assumption is questionable; if the strategies are front-loaded then the drift coefficient would vary through the day.

Thus we assume that the drift α is constant throughout the day, but we do not know its value. At the beginning of the day, we have a prior belief

$$\alpha \sim \mathcal{N}\big(\bar{\alpha},\, \nu^2\big) \ , \tag{4.2}$$

which will be updated using price observations during the day. There are thus two sources of randomness in the problem: the continuous Brownian motion representing the uninformed traders, and the single drift coefficient representing the constant trading of the large traders.

4.2.1. Bayesian Inference. Intuitively, as the trader observes prices from the beginning of the day onwards, he starts to get a feeling for the day's price momentum.

More precisely, at time t we know the stock price trajectory $S(\tau)$ for $0 \le \tau \le t$. We have the following result to update our estimate of the drift, given the observation $\{S(\tau)\,|\,0 \le \tau \le t\}$.

LEMMA 4.1. *Given the observation* $\{S(\tau)\,|0 \le \tau \le t\}$ *up to time* $t \ge 0$, *we have the posterior distribution for* α,

$$\alpha \sim \mathcal{N}\big(\hat{\alpha}, \hat{\nu}^2\big) \tag{4.3}$$

with

$$\hat{\alpha} = \frac{\bar{\alpha}\sigma^2 + \nu^2\big(S(t) - S_0\big)}{\sigma^2 + \nu^2 t} \quad \textit{and} \quad \hat{\nu}^2 = \frac{\sigma^2}{\sigma^2 + \nu^2 t}\nu^2 \ . \tag{4.4}$$

PROOF. In differential form, (4.1) reads

$$dS = \alpha dt + \sigma dB, \quad S(0) = S_0 \ . \tag{4.5}$$

It follows from the work of Liptser and Shiryaev (2001, chap. 17.7) that for this diffusion process the changes in the conditional expectation $\hat{\alpha}(\tau) = \mathbb{E}_\tau[\alpha]$ and conditional variance $v(\tau) = \hat{\nu}(\tau)^2 = \mathrm{Var}_\tau[\alpha]$ are given by the dynamics

$$d\hat{\alpha} = \frac{v(\tau)}{\sigma^2}\big(dS - \hat{\alpha}(\tau)\, d\tau\big), \quad \hat{\alpha}(0) = \bar{\alpha} \tag{4.6}$$

$$dv = -\frac{v(\tau)^2}{\sigma^2} d\tau, \quad v(0) = \nu^2 \ . \tag{4.7}$$

From (4.7) we immediately obtain

$$v(\tau) = \frac{\nu^2\sigma^2}{\sigma^2 + \nu^2\tau} \ ,$$

and since $\hat{\nu}^2 = v(t)$, we have shown the second part of (4.4).

The solution to the SDE (4.6) for $\hat{\alpha}$ is given by

$$\hat{\alpha} = \frac{\bar{\alpha}\sigma^2 + \nu^2\big(S(\tau) - S_0\big)}{\sigma^2 + \nu^2\tau} \ , \tag{4.8}$$

which is easily verified by Ito's Lemma: For $\hat{\alpha} = f(\tau, S) = \frac{\bar{\alpha}\sigma^2 + \nu^2(S - S_0)}{\sigma^2 + \nu^2\tau}$, Ito's Lemma yields

$$\begin{aligned} d\hat{\alpha} &= \left(\frac{\partial f}{\partial \tau} + \alpha\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2\frac{\partial^2 f}{\partial S^2}\right) d\tau + \sigma\frac{\partial f}{\partial S} dB \\ &\overset{(4.5)}{=} \left(\frac{\partial f}{\partial \tau} + \frac{1}{2}\sigma^2\frac{\partial^2 f}{\partial S^2}\right) d\tau + \frac{\partial f}{\partial S} dS = \frac{\partial f}{\partial \tau} d\tau + \frac{\partial f}{\partial S} dS \ . \end{aligned} \tag{4.9}$$

Straightforward differentiation shows

$$\partial f/\partial S = \frac{v(\tau)}{\sigma^2} \quad \text{and} \quad \partial f/\partial \tau = -\frac{v(\tau)}{\sigma^2} f(\tau, S) \ ,$$

and thus the dynamics (4.6) and (4.9) are indeed the same, completing the proof of (4.4). □

As Lemma 4.1 shows, in fact all of our information about the drift comes from the *final value* $S(t)$ of our observation period $0 \le \tau \le t$ (i.e. the distribution $\mathbb{P}[\alpha|S(t)]$ of α conditional on the single random value $S(t)$ is standard normal with mean and variance (4.4), as also an application of Bayes rule shows[1]). This is also apparent in (4.6), which implies that changes in the drift estimate $\hat{\alpha}$ are perfectly positively correlated with the instantaneous stock price change: we raise our estimate of the mean whenever the security return is above our current best estimate, and vice versa. This is a consequence of the assumption that the true mean α is constant.

Equation (4.4) represents our best estimate of the true drift α, as well as our uncertainty in this estimate, based on combination of our prior belief with price information observed to time t. This formulation accommodates a wide variety of belief structures. If we believe our initial information is perfect, then we set $\nu = 0$ and our updated belief is always just the prior $\alpha = \bar{\alpha}$ with no updating. If we believe we have no reliable prior information, then we take $\nu^2 \to \infty$ and our estimate is $\alpha \sim \mathcal{N}\big((S(t) - S_0)/t, \sigma^2/t\big)$, coming entirely from the intraday observations. For $t = 0$, we will have $S(0) = S_0$, and our belief is just our prior. As $t \to \infty$, our estimate becomes $\alpha \sim \mathcal{N}\big((S - S_0)/t, 0\big)$: we have accumulated so much information that our prior belief becomes irrelevant.

4.2.2. Trading and Price Impact. The trader has an order of $X > 0$ shares, which begins at time $t = 0$ and must be completed by time $t = T < \infty$. Unlike in Chapters 2 and 3, we interpret this as a buy order in this chapter. The definitions and results for sell programs are completely analogous. However, since the "natural" price trend for a stock is upwards (*price appreciation*), the application of our model in terms of a buy program seems more adequate: the trader wants to buy shares of a rising stock. In terms of a sell program, the model would assume an expected decline of the stock price (downward trend, price depreciation).

A *trading strategy* is a real-valued function $x(t)$ with $x(0) = X$ and $x(T) = 0$, representing the number of shares remaining to buy at time t. We require $x(t)$ to be differentiable, and define the corresponding *trading rate* $v(t) = -dx/dt$.

A trading strategy can be any non-anticipating (i.e. adapted to the filtration $\{\mathcal{F}_t\}$) random functional of B. Let $\mathfrak{A}$ denote the set of all such non-anticipating, differentiable random functionals.

[1] By (4.1) and (4.2), we have $\Delta S = S(t) - S(0) = (\bar{\alpha} + \nu\xi_1)t + \sigma\sqrt{t}\xi_2$ with $\xi_1, \xi_2 \sim \mathcal{N}(0, 1)$ i.i.d. Hence, $\Delta S \sim \mathcal{N}(\bar{\alpha}t, (\sigma^2 + \nu^2 t)t)$. We then use Bayes' rule $\mathbb{P}[\alpha|S(t)] = \mathbb{P}[S(t)|\alpha]\,\mathbb{P}[\alpha] \,/\, \mathbb{P}[S(t)]$ and straightforward calculation yields

$$\mathbb{P}[\alpha = a \mid S(t) = s] = \frac{\sqrt{\sigma^2 + \nu^2 t}}{\sigma\nu\sqrt{2\pi}} \, \frac{\varphi\Big(\frac{s - \alpha t}{\sigma\sqrt{t}}\Big)\varphi\Big(\frac{a - \bar{\alpha}}{\nu}\Big)}{\varphi\Big(\frac{s - \bar{\alpha}t}{\sqrt{t}\sqrt{\sigma^2 + \nu^2 t}}\Big)} = \frac{1}{\sqrt{2\pi}\hat{\nu}} \exp\Big(- \frac{(a - \hat{\alpha})^2}{2\hat{\nu}^2} \Big) \ ,$$

with $\hat{\alpha}$ and $\hat{\nu}$ given by (4.4), and $\varphi(x)$ denotes the standard normal distribution. Thus conditional on $S(t)$, $\alpha \sim \mathcal{N}(\hat{\alpha}, \hat{\nu}^2)$.

We use a linear market impact function for simplicity, although empirical work by Almgren et al. (2005) suggests a concave function. Thus the actual execution price is

$$\tilde{S}(t) \;=\; S(t) \,+\, \eta\, v(t) \;, \tag{4.10}$$

where $\eta > 0$ is the coefficient of temporary market impact.

The *implementation shortfall* C is the total cost of executing the buy program relative to the initial value, for which we have the following Lemma, similar to Lemma 2.1 in Chapter 2.

LEMMA 4.2. *For a trading policy* $\{x(\tau)\}$*, a realization* $\{B(\tau)\}$ *of the Brownian motion and a realization* α *of the random drift, the implementation shortfall is*

$$C\big(\{x(\tau)\},\{B(\tau)\},\alpha\big) \;= \sigma \int_0^T x(t)\,dB(t) + \eta \int_0^T v(t)^2\,dt + \alpha \int_0^T x(t)\,dt \;, \tag{4.11}$$

with $v(t) = -dx/dt$*.*

PROOF. To shorten notation, let $C = C\big(\{x(\tau)\},\{B(\tau)\},\alpha\big)$. By the definition of the implementation shortfall C and because of (4.10),

$$\begin{aligned} C \;&=\; \int_0^T \tilde{S}(t)\,v(t)\,dt \;-\; X\,S_0 \\ &= \int_0^T S(t)\,v(t)\,dt \;+\eta \int_0^T v(t)^2\,dt -\; X\,S_0 \;. \end{aligned}$$

Integration by parts for the first integral yields

$$\begin{aligned} C \;&=\; -\big[S(t)x(t)\big]_0^T + \sigma \int_0^T x(t)\,dB(t) + \eta \int_0^T v(t)^2 + \alpha x(t)\,dt \;-\; X\,S_0 \\ &= \sigma \int_0^T x(t)\,dB(t) + \eta \int_0^T v(t)^2\,dt + \alpha \int_0^T x(t)\,dt \;. \end{aligned}$$

□

$C = C\big(\{x(\tau)\},\{B(\tau)\},\alpha\big)$ is a random variable, because the price $S(t)$ is random, the drift α is random and the trading strategy $v(t)$ may be adapted to B.

4.3. Optimal Trading Strategies

We now address the question of what trading strategies are optimal, given the above model for price evolution and market impact. In the "classic" arrival price problem discussed in Chapters 2 and 3, trajectories are determined as a tradeoff between market impact and aversion to risk caused by volatility. The trader wants to complete the trade quickly to eliminate exposure to price volatility; he wants to trade slowly to minimize the costs of market impact. The optimal trajectory is determined as a balance between these two effects, parameterized by a coefficient of risk aversion.

To focus on the drift, which is the most important new aspect of this problem, here we neglect risk aversion; we consider the situation faced by a risk-neutral trader and seek to minimize only the expectation of trading cost. That is, we assume that the pressure to complete the trade rapidly comes primarily from a desire to capture the price motion expressed by the drift α, and it is this effect that must be balanced against the desire to reduce impact costs by trading slowly.

To support this description, we shall generally suppose that the original buy decision was made because the trader's belief has $\bar{\alpha} > 0$. We then expect $\alpha > 0$ in (4.11), and the term $\int \alpha\, x(t)\, dt$ is a positive cost. It may be that the true value has $\alpha < 0$, or that intermediate price movements cause us to form a negative estimate. Because our point of view is that of a broker/dealer executing an agency trade, we shall always require that the trade be completed by $t = T$, unless the instructions are modified.

Our trading goal is to determine a strategy that minimizes the expectation of the implementation shortfall (4.11): determine $x(\tau)$ for $0 \le \tau \le T$ so that

$$\min_{x \in \mathfrak{A}} \mathbb{E}\left[C\big(\{x(\tau)\}, \{B(\tau)\}, \alpha\big) \right] \ . \tag{4.12}$$

This corresponds to determining optimal strategies for a risk-neutral trader.

4.3.1. Static Trajectories. Suppose at some intermediate time $0 \le t_* < T$ we have $x(t_*)$ shares left to buy, and we decide to fix the trading strategy for the remaining time $t_* \le \tau \le T$ *independently* of $\{B(\tau) \,|\, t_* \le \tau \le T\}$; regardless of the stock price process, we follow the same trading trajectory. We will refer to such trading policies as *static*.

Lemma 4.3 gives optimal static trading strategies for a risk-neutral trader, i.e who optimizes (4.12). The proof uses standard methods of the calculus of variations.

LEMMA 4.3. *Let $0 \le t_* < T$. Let $\hat{\alpha} = \hat{\alpha}(t_*, S(t_*))$ be our best estimate (4.4) of α at time t_*, and let $x(t_*)$ be the shares remaining to buy. Then for a risk-neutral trader the optimal static strategy $\{x(\tau) \,|\, t_* \le \tau \le T\}$ specified at $t = t_*$ is given by*

$$x(\tau) \;=\; \frac{T-\tau}{T-t_*}\, x(t_*) \;-\; \frac{\hat{\alpha}}{4\eta}\big(\tau - t_*\big)\big(T-\tau\big), \qquad t_* \le \tau \le T, \tag{4.13}$$

and the instantaneous trade rate at time t_* is

$$v(t_*) \;=\; -x'(\tau)\big|_{\tau = t_*} \;=\; \frac{x(t_*)}{T-t_*} \;+\; \frac{\hat{\alpha}}{4\eta}\big(T - t_*\big) \ . \tag{4.14}$$

The expectation of the remaining cost $\mathbb{E}\left[C\big(\{x(\tau)\}, \{B(\tau)\}, \alpha\big)\right]$ of the optimal static strategy $\{x(\tau) \,|\, t_* \le \tau \le T\}$ is

$$\mathbb{E}\left[C\big(\{x(\tau)\}, \{B(\tau)\}, \alpha\big)\right] = \frac{x(t_*)^2 \eta}{T-t_*} + \frac{x(t_*)\hat{\alpha}(T-t_*)}{2} - \frac{(T-t_*)^3 \hat{\alpha}^2}{48\eta} \ . \tag{4.15}$$

PROOF. Let $x_* = x(t_*)$ and let $\hat{\alpha} = \hat{\alpha}(t_*, S(t_*))$ be given by (4.4). For a static strategy $\{x(\tau) \,|\, t_* \le \tau \le T\}$ specified at t_*, the final total cost $C = C\big(\{x(\tau)\}, \{B(\tau)\}, \alpha\big)$ is a Gaussian variable by Lemma 4.2. We want to determine a trajectory function $x : [t_*, T] \to \mathbb{R}$ with endpoint conditions $x(t_*) = x_*$ and $x(T) = 0$ to optimize

$$\begin{aligned} J[x] \;&=\; \mathbb{E}\left[\sigma \int_{t_*}^{T} x(t)\, dB(t) + \int_{t_*}^{T} \eta v(t)^2 + \alpha x(t)\, dt \;\middle|\; t_* \right] \\ &=\; \int_{t_*}^{T} \eta x'(t)^2 + \hat{\alpha} x(t)\, dt \ , \end{aligned} \tag{4.16}$$

since $\mathbb{E}\left[\alpha \,|\, t_*\right] = \hat{\alpha}$. The Euler-Lagrange equation (see for instance Wan (1995)) for the variational problem $\min\big\{J[x] \,|\, x(t_*) = x_*,\ x(T) = 0\big\}$ is the ordinary differential equation

$$x''(\tau) \;=\; \frac{\hat{\alpha}}{2\eta}, \qquad t_* \le \tau \le T \ . \tag{4.17}$$

The solution to this equation that satisfies the boundary conditions is the trajectory given by (4.13). Its derivative $x'(\tau)$ is

$$x'(\tau) \;=\; -\frac{x(t_*)}{T-t_*} - \frac{(T+t_*-2\tau)\hat{\alpha}}{4\eta}, \qquad t_* \le \tau \le T \;, \tag{4.18}$$

and hence the instantaneous trade rate at time t_* is $v(t_*) \;=\; -x'(\tau)\big|_{\tau=t_*}$ given by (4.14). The expected cost (4.15) follows from straightforward evaluation of the integral in (4.16),

$$\mathbb{E}\left[C\big(\{x(\tau)\},\{B(\tau)\},\alpha\big)\right] = \int_{t_*}^{T} \eta x'(t)^2 + \hat{\alpha}x(t)\,dt$$

for the optimal trajectory (4.13). □

This trajectory (4.13) is the sum of two pieces. The first piece is proportional to $x(t_*)$ and represents the linear (VWAP) liquidation of the current position; it is the optimal strategy to reduce expected impact costs with no risk aversion. The second piece is independent of $x(t)$ and would therefore exist even if the trader had no initial position. Just as in the solutions of Bertsimas and Lo (1998), this second piece is effectively a proprietary trading strategy superimposed on the liquidation. The magnitude of this strategy, and hence the possible gains, are determined by the ratio between the expected drift $\hat{\alpha}$ and the liquidity coefficient η.

As we can see from Lemma 4.3, the optimal path-independent strategy specified at any intermediate t_* for the remainder $[t_*, T]$ uses the best estimate $\hat{\alpha} = \mathbb{E}\left[\alpha \,|\, t_*\right]$ available at time t_*, and assumes that this estimate does not change anymore until time T. Certainly, this strategy is not the optimal strategy for $[t_*, T]$.

Interestingly, though, in the next section we will prove by dynamic programming that we indeed obtain an optimal policy if we trade with the instantaneous trade rate (4.14) at every $0 \le t \le T$ using the *current best drift estimate* $\hat{\alpha} = \mathbb{E}\left[\alpha \,|\, t\right]$ at each moment.

4.3.2. Optimal Strategy. We shall now show that trading with the instantaneous trade rate (4.14) at every $0 \le t \le T$ using the drift estimate $\hat{\alpha} = \mathbb{E}\left[\alpha \,|\, t\right]$ at each moment is the true optimum strategy. Apparently, this strategy is "locally" optimal in the sense that at all times we use all the new information available. However, we need to prove that this leads to a globally optimal strategy.

Our proof uses techniques from dynamic programming; the reader is referred to standard textbooks for background information on this theory (Yong and Zhou, 1999; Fleming and Rishel, 1975; Korn, 1997).

THEOREM 4.4. *The optimal dynamic strategy* $\{x_*(\tau)\}$ *for* (4.12) *is given by the optimal instantaneous trade velocity* $x_*'(t) = -v_*(t)$*,*

$$v_*(t) = \frac{x(t)}{T-t} + \frac{\hat{\alpha}(t,S(t))\cdot(T-t)}{4\eta} \qquad 0 \le t \le T \;, \tag{4.19}$$

and $x(0) = X$*.* $S(t)$ *is the stock price at time* t*,* $\hat{\alpha}(t,S(t))$ *denotes the estimate (4.4) of* α *at time* t*, and* $x(t)$ *is the current number of shares remaining to buy at time* t*.*

The expected implementation shortfall of this strategy is

$$\mathbb{E}\left[C\big(\{x_*(\tau)\},\{B(\tau)\},\alpha\big)\right] = \frac{X^2\eta}{T} \;+\; \frac{X\,\bar{\alpha}\,T}{2} - \;\frac{T^3\bar{\alpha}^2}{48\eta} \;-\; \Delta \;, \tag{4.20}$$

with

$$\Delta \;=\; \frac{\sigma^2 T^2}{48\eta} \int_0^1 \frac{(1-\delta)^3}{(\delta+\rho)^2}\, d\delta, \qquad \textit{with} \quad \rho = \frac{\sigma^2}{\nu^2 T} \;. \tag{4.21}$$

PROOF. We formulate the problem in a dynamic programming framework. The control, the state variables, and the stochastic differential equations of problem (4.12) are given by

$$\begin{aligned}
v(t) &= \text{rate of buying at time } t \\
x(t) &= \text{shares remaining to buy at time } t & dx &= -v\,dt \\
y(t) &= \text{dollars spent up to time } t & dy &= (s+\eta v)\,v\,dt \\
s(t) &= \text{stock price at time } t & ds &= \alpha\,dt + \sigma\,dB
\end{aligned}$$

where $\alpha \sim N(\bar{\alpha}, \nu^2)$, chosen randomly at $t=0$. We begin at $t=0$ with shares $x(0)=X$, cash $y(0)=0$, and initial stock price $s(0)=S$. The strategy $v(t)$ must be adapted to the filtration of B.

To address the endpoint constraint $x(T)=0$, we assume that all shares remaining to buy at time $t=T-\epsilon$, for some $\epsilon>0$, will be purchased using an optimal static trajectory (as in Lemma 4.3) specified at $T-\epsilon$, i.e. using the drift estimate $\hat{\alpha}(T-\epsilon, s(T-\epsilon))$. By (4.15) this final purchase of $x(T-\epsilon)$ shares will cost an expected dollar amount of

$$s(T-\epsilon)\; x(T-\epsilon) + \frac{\eta\, x(T-\epsilon)^2}{\epsilon} + \frac{\hat{\alpha}\;\epsilon\; x(T-\epsilon)}{2} - \frac{\hat{\alpha}^2\;\epsilon^3}{48\eta} \tag{4.22}$$

with $\hat{\alpha} = \hat{\alpha}(T-\epsilon, s(T-\epsilon))$. For every $\epsilon>0$, we will determine an optimal policy for $0 \le t \le T-\epsilon$. Combined with the static trajectory for $T-\epsilon \le t \le T$, we obtain an admissible trading strategy that satisfies $x(T)=0$, and as $\epsilon \to 0$ we end up with an optimal admissible trading strategy for $0 \le t \le T$.

For $\epsilon>0$, conditional on the information available at $T-\epsilon$ the expected overall final dollar amount spent is

$$y(T) = y(T-\epsilon) + \underbrace{s(T-\epsilon)\; x(T-\epsilon) + \frac{\eta\, x(T-\epsilon)^2}{\epsilon} + \frac{\hat{\alpha}\;\epsilon\; x(T-\epsilon)}{2} - \frac{\hat{\alpha}^2\;\epsilon^3}{48\eta}}_{=\,(4.22)} \tag{4.23}$$

with $\hat{\alpha} = \hat{\alpha}(T-\epsilon, s(T-\epsilon))$. We want to determine a control function $\{v(\tau) \mid 0 \le \tau \le T-\epsilon\}$ to minimize the expected value of (4.23),

$$\min_{\{v(\tau)\,|\,0\le\tau\le T-\epsilon\}} \mathbb{E}\left[y(T)\right] \;,$$

where $\{v(\tau) \mid 0 \le \tau \le T-\epsilon\}$ now is unconstrained. Standard techniques lead to the Hamilton-Jacobi-Bellman (HJB) partial differential equation

$$0 \;=\; J_t \;+\; \frac{1}{2}\sigma^2\, J_{ss} \;+\; \hat{\alpha}\, J_s \;\;+\; \min_v \Big(\big(s\, J_y \;-\; J_x\big)\, v \;+\; \eta J_y\, v^2 \Big) \tag{4.24}$$

for the value function $(0 \le t \le T-\epsilon)$

$$J(t,x,y,s) = \min_{\{v(\tau),\, t\le\tau\le T-\epsilon\}} \mathbb{E}\left[y(T)\right] \;,$$

with $y(T)$ given by (4.23). $\hat{\alpha} = \hat{\alpha}(t,s)$ denotes the estimate of α at time t as computed in (4.4). In (4.24) and in the following, subscripts denote partial derivatives. The optimal trade velocity in (4.24) is found as

$$v_*(t,x,y,s) = \frac{J_x \; - \; s\,J_y}{2\eta\,J_y} \tag{4.25}$$

and we have the final HJB partial differential equation

$$0 \;=\; J_t \;+\; \frac{1}{2}\sigma^2\,J_{ss} \;+\; \hat{\alpha}\,J_s \;-\; \frac{\left(s\,J_y - J_x\right)^2}{4\eta\,J_y} \tag{4.26}$$

for $J(t,x,y,s)$ with the terminal condition

$$J(T-\epsilon,x,y,s) = y + sx + \frac{\eta x^2}{\epsilon} + \frac{\hat{\alpha}(T-\epsilon,s)x\epsilon}{2} - \frac{\hat{\alpha}(T-\epsilon,s)^2\,\epsilon^3}{48\eta} \tag{4.27}$$

for all x, y, s.

For $0 \le t \le T-\epsilon$, we define the function

$$\begin{aligned} H(t,x,y,s) = y \;+\; sx \;+\; \frac{\eta x^2}{T-t} \;+\; \frac{(T-t)\;\hat{\alpha}(t,s)\;x}{2} \\ -\;\frac{(T-t)^3\hat{\alpha}(t,s)^2}{48\eta} \;-\; \int_t^{T-\epsilon}\frac{\sigma^2\nu^4(T-\tau)^3}{48\eta(\tau\nu^2+\sigma^2)^2}\,d\tau \;. \end{aligned} \tag{4.28}$$

Obviously, $H(t,x,y,s)$ satisfies the terminal condition (4.27). We will now show that $H(t,x,y,s)$ also satisfies the PDE (4.26).

We have

$$\begin{aligned} \hat{\alpha}_s \;&=\; \frac{\partial}{\partial s}\hat{\alpha}(t,s) = \frac{\nu^2}{t\nu^2+\sigma^2} \\ \hat{\alpha}_t \;&=\; \frac{\partial}{\partial t}\hat{\alpha}(t,s) = -\nu^2\frac{\sigma^2\bar{\alpha}+\nu^2(s-S_0)}{(t\nu^2+\sigma^2)^2} = -\hat{\alpha}\,\hat{\alpha}_s \;. \end{aligned}$$

Furthermore

$$H_t = \frac{x^2\eta}{(T-t)^2} - \frac{\hat{\alpha}x[(T-t)\hat{\alpha}_s+1]}{2} + \frac{(T-t)^2\hat{\alpha}^2}{16\eta} + \frac{(T-t)^3\hat{\alpha}^2\hat{\alpha}_s}{24\eta} + \frac{\sigma^2(T-t)^3\hat{\alpha}_s^2}{48\eta}$$

and

$$H_x = s + \frac{2x\eta}{T-t} + \frac{(T-t)\hat{\alpha}}{2}\,, \qquad H_s = x + \frac{(T-t)\hat{\alpha}_s x}{2} - \frac{(T-t)^3\hat{\alpha}\hat{\alpha}_s}{24\eta}$$

$$H_{ss} = -\frac{(T-t)^3\hat{\alpha}_s^2}{24\eta}\,, \qquad H_y = 1\;.$$

Straightforward calculation shows that $H(t,x,y,s)$ indeed satisfies the PDE (4.26). By (4.25), the corresponding optimal trade velocity $\{v_*(\tau) \mid 0 \le \tau \le T-\epsilon\}$ is

$$v_*(t,x,y,s) = \frac{H_x \; - \; s\,H_y}{2\eta\,H_y} = \frac{x}{T-t} + \frac{\hat{\alpha}(t,s)\cdot(T-t)}{4\eta}\;, \tag{4.29}$$

for all $\epsilon > 0$. From (4.28), we see that the expected total dollar amount spent

$$\mathbb{E}\left[y(T) \mid t, x(t)=x, y(t)=y, s(t)=s\right] = H(t,x,y,s)$$

is decreasing in ϵ. Hence, the optimal strategy $\{x_*(t) \mid 0 \le t \le T\}$ for (4.12) is obtained for $\epsilon \to 0$, and given by the trade rate (4.29). Its expected cost is

$$\begin{aligned} \mathbb{E}\left[C\right] \;&=\; \lim_{\epsilon\to 0} H(0,X,0,S) - XS \\ &=\; \frac{X^2\eta}{T} + \frac{X\bar{\alpha}T}{2} - \frac{T^3\bar{\alpha}^2}{48\eta} - \int_0^T \frac{\sigma^2\nu^4(T-\tau)^3}{48\eta(\tau\nu^2+\sigma^2)^2}\,d\tau\;. \end{aligned}$$

Substituting $\delta = t/T$ in the integral, and defining $\rho = \sigma^2/(\nu^2 T)$ we obtain (4.21). □

We may explicitly determine the gains due to adaptivity, i.e. the extra profit that we make by using all new information available at all times by continuously updating our drift estimate.

COROLLARY 4.5. *Let E_{stat} denote the expected cost of the non-adaptive strategy determined at $t = 0$ using the prior expected drift $\bar{\alpha}$ (given in Lemma 4.3, (4.15)), and E_{dyn} the expected cost of the optimal strategy given in Theorem 4.4,* (4.20). *Then*

$$E_{dyn} \;=\; E_{stat} \;-\; \Delta \;,$$

with Δ given by (4.21).

PROOF. From Lemma 4.3, the static strategy determined at $t_* = 0$ has expected cost

$$E_{\text{stat}} \;=\; \frac{X^2\eta}{T} + \frac{X\bar{\alpha}T}{2} - \frac{T^3\bar{\alpha}^2}{48\eta} \;.$$

□

The gain (4.21)

$$\Delta \;=\; \frac{\sigma^2T^2}{48\eta}\int_0^1 \frac{(1-\delta)^3}{(\delta+\rho)^2}\,d\delta, \qquad \rho = \frac{\sigma^2}{\nu^2 T}, \quad \delta = \frac{t}{T}$$

is the reduction in expected cost obtained by using the Bayesian adaptive strategy (note that $\Delta > 0$).

As discussed in Section 4.3.1, the optimal static trajectory (4.13) contains a proprietary trading part, which leads to the cost reduction $T^3\bar{\alpha}^2/(48\eta)$ in (4.20) independent of initial portfolio size X.

Similarly, the gain Δ in (4.21) is also independent of X. Thus, it represents the gains from another proprietary trading strategy superimposed on the static risk-neutral liquidation profile. The following Lemma gives the asymptotic magnitude of this gain for $T \to 0$ and $T \to \infty$.

LEMMA 4.6. *The gain Δ,* (4.21), *is*

i) $\Delta \sim T^4$ when T is small and
ii) $\Delta \sim T^3$ when T is large.

PROOF. We have

$$\frac{\Delta}{T^4} = \frac{\sigma^2}{48\,\eta}\int_0^1 \frac{(1-\delta)^3}{(T\delta+\sigma^2/\nu^2)^2}\,d\delta \;\xrightarrow{T\to 0}\; \frac{\nu^4}{48\,\sigma^2\eta}\int_0^1 (1-\delta)^3\,d\delta = \frac{\nu^4}{192\,\sigma^2\eta} \;,$$

which shows (i). Evaluating the integral in the expression (4.21) of the gain Δ yields

$$\Delta = \frac{\sigma^6}{48\eta\nu^4}\cdot\frac{1}{\rho^2}\cdot\Big(3(1+\rho)^2\ln\frac{\rho}{1+\rho} + \frac{1}{\rho} + 3\rho + \frac{9}{2}\Big)$$

Thus, since $x\ln x \to 0$ for $x \to 0$,

$$\Delta\cdot\rho^3 = \frac{\sigma^6}{48\eta\nu^4}\cdot\Big(3(1+\rho)^3\frac{\rho}{1+\rho}\ln\frac{\rho}{1+\rho} + 1 + 3\rho^2 + \frac{9}{2}\rho\Big) \;\xrightarrow{\rho\to 0}\; \frac{\sigma^6}{48\eta\nu^4} \;.$$

Hence,

$$\frac{\Delta}{T^3} \;\xrightarrow{T\to\infty}\; \frac{\nu^2}{48\eta} \;,$$

which proves (ii). □

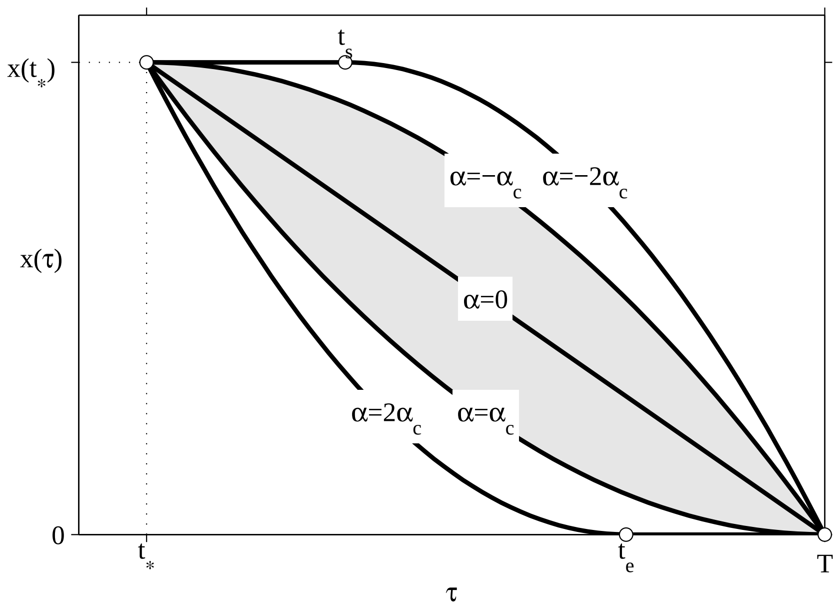

FIGURE 4.1. Constrained solutions $x(\tau)$, starting at time t_* with shares $x(t_*)$ and drift estimate α. For $\alpha > 0$, the trajectories go below the linear profile to reduce expected purchase cost. For $|\alpha| \leq \alpha_c$, the constraint is not binding (shaded region). At $\alpha = \alpha_c$ the solutions become tangent to the line $x = 0$ at $\tau = T$, and for larger values they hit $x = 0$ with zero slope at $\tau = t_e < T$. For $\alpha < -\alpha_c$, trading does not begin until $\tau = t_s > t_*$.

Thus, the adaptivity adds only little value if applied to short-term correlation. This corresponds to the small gains obtained by Bertsimas and Lo (1998).

4.3.3. Constrained Trajectories. In some situations, we may want to constrain the trade direction, and require that the trading strategy must never sell as part of a buy program, even if this would yield lower expected costs (or give an expected profit) because of anticipated negative drift in the price. This is for two reasons: First, we take the point of view of a broker/dealer executing an agency trade for a client, and the client might object bidirectional trading on his account. Second, we neglect the bid/offer spread and other fixed costs, which may reduce the profitability of such reversing strategies.

Let us consider the static trajectories given in Lemma 4.3. The instantaneous trade velocity (4.14) may indeed violate the constraints: if $\hat{\alpha}$ is large then the quadratic term in (4.13) may cause $x(\tau)$ to dip below zero, which would cause the trading rate at some point to become negative, since we require $x(T) = 0$.

The condition $v(t) = x'(t) \geq 0$ constitutes a *nonholonomic inequality constraint* to the variational problem (4.16). Following Wan (1995, chap. 12), $x(t)$ must satisfy the ODE (4.17) in the region of $x'(t) > 0$, and a differentiable solution $x(t)$ must meet the constraint $x'(t) = 0$ smoothly; we cannot simply clip a trajectory that crosses the axis $x = 0$ to satisfy $x \geq 0$, then the derivative $x'(t)$ will be discontinuous. Solutions are obtained by combining the ODE (4.17) in regions of smoothness, with this “smooth pasting” condition at the boundary points.

From (4.18) we see that there is in fact a critical drift value

$$\alpha_c\big(x(t_*),\, T - t_*\big) \;=\; \frac{4\eta\, x(t_*)}{(T-t_*)^2}\;.$$

For $|\hat{\alpha}| \le \alpha_c$, the trade rate $v(\tau) = -x'(\tau)$, (4.18), of the trajectory (4.13) is positive for all $t_* \le \tau \le T$, and the constraint $v(\tau) \ge 0$ never becomes binding.

For high negative drift $\hat{\alpha} < -\alpha_c$ the unconstrained trajectory (4.13) would increase over $x(t_*)$ right after its start at t_*, i.e. the strategy short-sells in the beginning in order to capture additional profit by buying back those extra shares later for lower prices; on the other hand, the constrained trajectory does not start buying shares until some starting time $t_s > t_*$. As already mentioned, by the "smooth pasting" condition the trajectory has to start trading at t_s with zero slope $v(t_s) = 0$.

Conversely, for high positive drift $\hat{\alpha} > \alpha_c$ the unconstrained trajectory dips below zero in the course of the trading in order to sell those additional shares again at the end of the day; contrary, the constrained trajectory stops trading at a shortened end time t_e. Again, the trajectory has to meet t_e with zero slope $v(t_e) = 0$.

The times t_s and t_e are determined by the "smooth pasting" condition. The shortened end time t_e is determined so that $x'(t_e) = x(t_e) = 0$, and is found as

$$t_e = t_* + \sqrt{\frac{4\eta\, x(t_*)}{\hat{\alpha}}}\;.$$

Note that since $\hat{\alpha} > \alpha_c = 4\eta x(t_*)/(T-t_*)^2$, we always have $t_e < T$.

The deferred start time t_s is determined so that $x'(t_s) = 0$ and $x(t_s) = x(t_*)$ and given by

$$t_s = T - \sqrt{\frac{4\eta\, x(t_*)}{-\hat{\alpha}}}\;.$$

Note that since $\hat{\alpha} < -\alpha_c = -4\eta x(t_*)/(T-t_*)^2$, we always have $t_s > t_*$.

Figure 4.1 illustrates these solutions. The overall instantaneous trade rate at time $0 \le t_* \le T$ of a constrained optimal statc trajectory $\{x(\tau) \mid t_* \le \tau \le T\}$ specified at t_* may be summarized as

$$v(t_*) \;=\; \begin{cases} 0, & \hat{\alpha} < -\alpha_c \\ \dfrac{x(t_*)}{T-t_*} \;+\; \dfrac{\hat{\alpha}}{4\eta}\,(T-t_*), & |\hat{\alpha}| < \alpha_c \\ \dfrac{x(t_*)}{t_e-t_*} \;+\; \dfrac{\hat{\alpha}}{4\eta}\,\big(t_e-t_*\big) \;=\; \sqrt{\hat{\alpha}x(t_*)/\eta}, & \hat{\alpha} > \alpha_c \end{cases} \tag{4.30}$$

where $\hat{\alpha} = \hat{\alpha}\big(t_*, S(t_*)\big)$ is the best drift estimate at time t_* given by Lemma 4.1.

From this discussion about the constrained static trajectories, it is conceivable that the optimal adaptive constrained strategy is – similarly to the unconstrained case in Theorem 4.4 – obtained by trading with the instantaneous trade rate (4.30) at every $0 \le t_* \le T$ using the *current best drift estimate* $\hat{\alpha} = \mathbb{E}\left[\alpha \mid t_*, s(t_*)\right]$ at each moment. The dynamic programming approach to prove the optimality of this strategy is still possible in principle, and leads to

$$0 \;=\; J_t \;+\; \frac{1}{2}\sigma^2\, J_{ss} \;+\; \hat{\alpha}\, J_s \;+\; \min_{v\ge 0}\Big(\big(s\, J_y \;-\; J_x\big)\, v \;+\; \eta J_y\, v^2\Big)\;, \tag{4.31}$$

as in (4.24) – but with the crucial difference that now we have the additional constraint $v \geq 0$ in "$\min_{v\geq 0} \ldots$". With this sign constraint the optimal velocity becomes

$$v_*(t,x,y,s) = \max\left\{ \frac{J_x \; - \; s\,J_y}{2\eta\,J_y},\; 0 \right\}$$

instead of (4.25). Substituting this v_* back in (4.31), instead of the PDE (4.26) we obtain a PDE that is more highly nonlinear, and no explicit closed form solution seems to be available. The imposition of the constraint should not change the relation between the static and the dynamic solutions, but we leave this as an open question.

4.4. Examples

Figures 4.2 and 4.3 show examples of the strategies computed by this method. To produce these pictures, we began with a prior belief for α having mean $\bar{\alpha} = 0.7$ and standard deviation $\nu = 1$. For each trajectory, we generated a random value of α from this distribution, and then generated a price path $S(t)$ for $0 \leq t \leq 1$ with volatility $\sigma = 1.5$. For example, on a stock whose price is \$100 per share, these would correspond to 1.5% daily volatility, and an initial drift estimate of +70 basis points with a substantial degree of uncertainty.

We set the impact coefficient $\eta = 0.07$ and the initial shares $X = 1$, meaning that liquidating the holdings using VWAP across one day will incur realized price impact of 7 basis points[2] for a stock with price \$100 per share.

Then for each sample path, we evaluate the share holdings $x(t)$ using the Bayesian update strategy (4.30) and plot the trajectories. For comparison, we also show the optimal static trajectory using only the initial estimate of the drift. In Figure 4.2, to illustrate the features of the solution, we show a rather extreme collection of paths, having very high realized drifts. In Figure 4.3 we show a completely representative selection.

4.5. Conclusion

We have presented a simple model for momentum in price motion based on daily trading cycles, and derived optimal risk-neutral adaptive trading strategies. The momentum is understood to arise from the correlated trade targets being executed by large institutional investors. The trader begins with a belief about the direction of this imbalance, and expresses a level of confidence in this belief that may range anywhere from perfect knowledge to no knowledge. This belief is then updated using observations of the price process during trading. Under the assumptions of the model, our solutions deliver better performance than non-adaptive strategies.

It is natural to ask whether this model can be justified by empirical data. In the model, the random daily drift is superimposed on the price volatility caused by small random traders. In theory, these two sources of randomness can be disentangled by measuring volatility on an intraday time scale and comparing it to daily volatility. If daily volatility is higher than intraday, then the difference can be attributed to serial correlation of the type considered here. In practice, because real price processes are far from Gaussian, it might be difficult to do this comparison even if one restricts attention to days when there is large institutional flow.

[2] One basis point equals 0.01%.

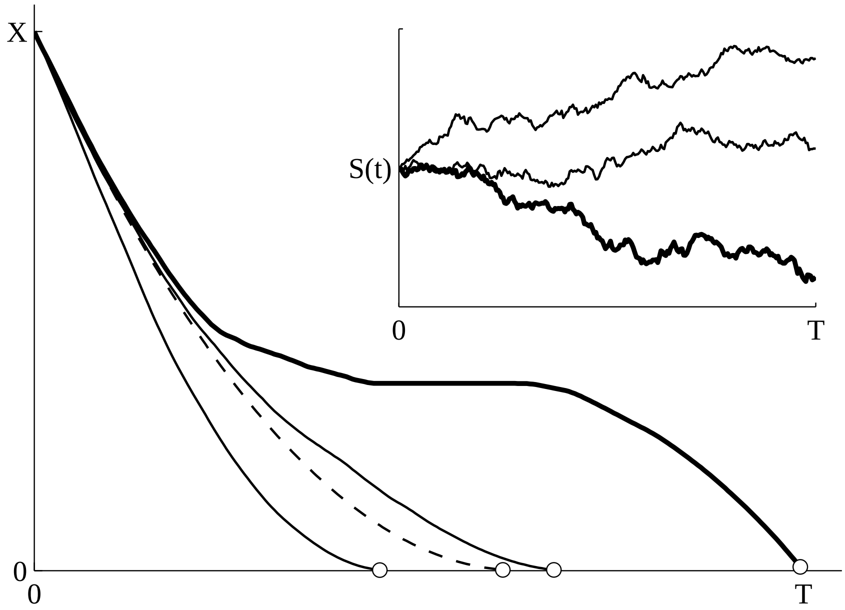

FIGURE 4.2. Sample trajectories of the optimal strategy (4.30). The inset shows price processes, and the main figure shows trade trajectories $x(t)$. The dashed line is the static trajectory using the prior belief for the drift value. The thick line corresponds to the falling stock price path. In that realization, the trader has to adjust his prior belief (that is, a positive drift) substantially; the new drift estimate is negative, and since we are considering a buy program, the trader slows down trading in order to buy at cheaper prices later on. In this figure, we have selected realizations with very high drift to highlight the solution behavior, including temporary interruption of trading.

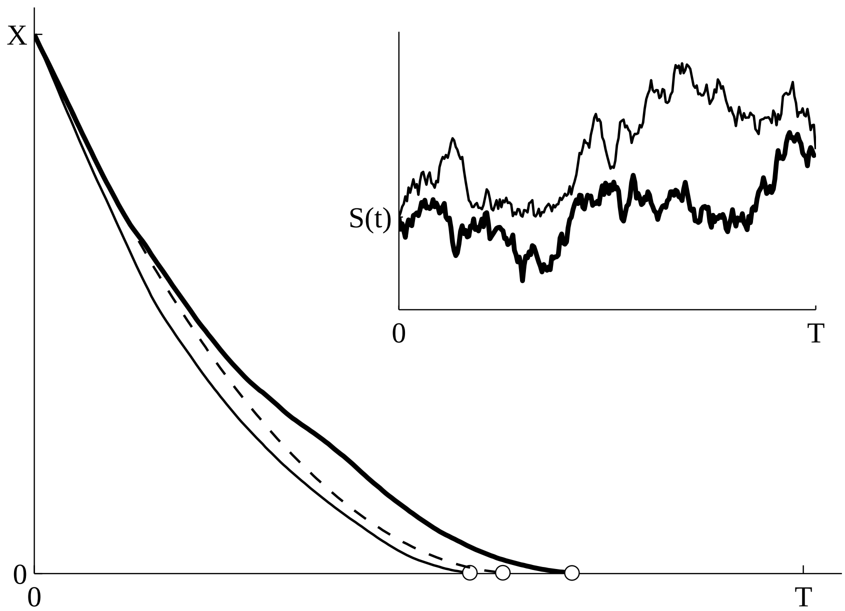

FIGURE 4.3. Same layout as Figure 4.2, but with more realistic sample paths. The thick path causes the trader to estimate a lower drift value than for the light path, and he comparatively trades slower.

CHAPTER 5

Multiperiod Portfolio Optimization

In this chapter we show how the dynamic programming principle introduced in Chapter 3 can be applied to solve the discrete-time multiperiod mean-variance portfolio problem. Our solution coincides with the optimal strategy given by Li and Ng (2000).

5.1. Introduction

Research on portfolio theory was pioneered in the 1950s with the seminal work of Markowitz (1952, 1959) on mean-variance efficient portfolios for a single period investment. While it is natural to extend Markowitz's work to multiperiod settings, difficulties arise from the definition of the mean-variance objective $\mathbb{E}[X] - \lambda \operatorname{Var}[X]$, which does not allow a direct application of the dynamic programming principle (Bellman, 1957) due to the square of the expectation in the variance term.

Work on multiperiod and continuous-time formulations of portfolio selection took a different approach and employed expected utility criteria (von Neumann and Morgenstern, 1953). Here, the expected terminal wealth $\mathbb{E}[u(w_T)]$ for a utility function $u(\cdot)$ is optimized, see for instance Mossin (1968) and Samuelson (1969). Typical choices for $u(\cdot)$ are logarithmic, exponential, power or quadratic functions. One of the most important breakthroughs in this area is due to Merton (1969). Using Ito's lemma and Bellman's dynamic programming principle, he solved the problem of optimal consumption and investment in continuous-time where the prices of the risky assets follow diffusion processes and was able to derive closed-form optimal dynamic portfolio strategies for some particular cases of utility functions $u(\cdot)$.

In the original mean-variance formulation, Dantzig and Infanger (1993) show how multiperiod portfolio problems can be efficiently solved as multistage stochastic linear programs, using scenario generation and decomposition techniques. However, no analytical closed-form solution is obtained.

In a rather influential paper, Li and Ng (2000) give a closed-form solution for the classical mean-variance problem in discrete time. They introduce an embedding of the original mean-variance problem into a tractable auxiliary problem with a utility function of quadratic type. Their embedding technique provides a general framework of stochastic linear-quadratic (LQ) optimal control in discrete time. Zhou and Li (2000) generalize this theory to continuous-time settings, and also obtain a closed-form solution for the efficient frontier. In fact, prior to this work White (1998) studied a similar approach to characterizes mean-variance efficient solution sets by a parametric solution sets of first and second moment.

The embedding technique of Li and Ng (2000) and Zhou and Li (2000) set off a series of papers on dynamic mean-variance optimization. Lim and Zhou (2002) consider stock price processes with random interest rate, appreciation rates, and volatility coefficients. Zhou and Yin (2003) discuss mean-variance portfolio selection in a model that features regime switching of the market parameters. Li, Zhou, and Lim (2002) add a constraint on short selling, and give an explicit optimal investment strategy. Jin and Zhou (2007) consider risk measures different from variance, for instance weighted mean-variance problems where the risk has different weights on upside and downside risk. Leippold et al. (2004) present a geometric approach and decompose the set of multiperiod mean-variance optimal strategies in an orthogonal set of basis strategies. Li and Zhou (2006) prove that a mean-variance efficient portfolio realizes the (discounted) targeted return on or before the terminal date with a probability greater than 80%, *irrespective* of the (deterministic) market parameters, targeted return and investment horizon.

Parallel to this line of work, Bielecki, Jin, Pliska, and Zhou (2005) employ a second approach to solve multiperiod and continuous-time mean-variance problems, namely using a decomposition technique that was first introduced by Pliska (1986). They reduce the problem of continuous-time mean-variance portfolio selection to solving two subproblems: first find the optimal attainable wealth X^* (which is a random variable), and then determine the trading strategy that replicates X^*. Using such a technique, Richardson (1989) also tackled the mean-variance problem in a continuous-time setting. Duffie and Richardson (1991) and Schweizer (1995) study the related mean-variance hedging problem, where optimal dynamic strategies are determined to hedge contingent claims in an imperfect market.

In this chapter, we show how the dynamic programming approach used in Chapter 3 for the problem of risk-averse optimal execution of portfolio transactions can be applied to solve the discrete-time multiperiod mean-variance portfolio problem. As in Chapter 3, after introducing an additional state variable, a suitable application of the dynamic programming principle reduces the multiperiod problem to a series of optimization problems. Whereas we had to resort to numerical treatment in Chapter 3, for the portfolio selection problem without shortsale constraint we obtain analytical solutions in each step and inductively derive an explicit optimal dynamic investment strategy. This optimal strategy and the efficient frontier indeed coincide with the results of Li and Ng (2000).

The efficient frontier for a universe of stocks with expected excess return vector $\boldsymbol{\mu}$ and covariance matrix Σ may be summarized as follows: for an investor with an excess return target (over an investment in the riskless asset) of α and T periods of investment, the optimal dynamic (*path-dependent*) policy yields a variance of $V_{dyn} = \alpha^2/((\theta^2+1)^T - 1)$, where $\theta = \boldsymbol{\mu}'\Sigma^{-1}\boldsymbol{\mu}$. Contrary, if the investor uses an optimal static (*path-independent*) policy, his strategy has a variance of $V_{stat} = \alpha^2/(T\theta^2)$. For all $\theta > 0$ (i.e. excess returns are possible), $V_{dyn} < V_{stat}$.

The key fact that lets us determine a closed-form solution is a correlation observation in each step: the portfolio return in each time period is *perfectly anticorrelated* to the targeted return for the remaining time: after a fortunate period profit the investor will try to conserve his realized gains and put less capital at risk in the remainder, i.e. aim for a lower expected return. That is, the investor's risk aversion changes in response to past performance, making him *more risk-averse* after *positive portfolio performance*. We saw a similar behavior of the dynamic execution strategy in Chapter 3.

This change in risk-aversion of the optimal multiperiod investment strategy is very apparent in the direct dynamic programming approach that we propose. Hence, we hope that our technique contributes to the understanding of the behavior of mean-variance optimal dynamic strategies for multiperiod problems.

As it is well known, mean-variance based preference criteria have received theoretical criticism (see for instance Maccheroni et al., 2004). However, mean-variance portfolio optimization retains great practical importance due to its clear intuitive meaning. Compared to utility based criteria, mean-variance optimization has the advantage that no assumptions on the investor's utility function are needed. From the perspective of an investment company, who does not know about its clients' utility or about their other investment activities, one solution is to offer the efficient frontier of mean-variance optimal portfolio strategies, and let the clients pick according to their needs. Another advantage is that the trade-off between risk and return is explicitly given, and not only implicitly contained in the utility function. The relationship between mean-variance optimization and the expected utility framework is discussed, for instance, by Kroll et al. (1984) and Markowitz (1991).

The remainder of this chapter is organized as follows. We formulate the discrete-time multiperiod mean-variance portfolio selection problem in Section 5.2. In Section 5.3, we show how mean-variance optimization becomes amenable to dynamic programming by a suitable choice of the value function. Using this framework, we derive the explicit solution for the multiperiod portfolio problem in Section 5.4.

5.2. Problem Formulation

An investor has an initial wealth of w_0, and wants to invest across T periods with no consumption, generating intermediate wealth $w_1, \ldots, w_T$. His only concern is to maximize the final wealth w_T, while also minimizing its riskiness (drawdowns, *etc,* are of no concern).

There are n risky assets with prices $S_t = (S_t^1, \ldots, S_t^n)'$ and one riskless asset with price S_t^0 (all vectors are column vectors, and $'$ denotes transpose). The riskless asset has a fixed return $r = S_{t+1}^0 / S_t^0 > 0$. The excess returns of the risky assets at time t are $\mathbf{e}_t = (\xi_t^1, \ldots, \xi_t^n)'$, with

$$\mathbf{e}_t^i \;=\; \frac{S_{t+1}^i}{S_t^i} - r \;.$$

The excess return vectors $\mathbf{e}_0, \ldots, \mathbf{e}_{T-1}$ are i.i.d. random variables with known mean vector $\boldsymbol{\mu} = \mathbb{E}\left[\mathbf{e}_t\right]$ and covariance matrix $\Sigma = \mathbb{E}\left[(\mathbf{e}_t - \boldsymbol{\mu})(\mathbf{e}_t - \boldsymbol{\mu})'\right]$. Since we assume the returns $\mathbf{e}_t$ identically distributed, $\boldsymbol{\mu}$ and Σ are time-independent. In principle, the framework in this chapter also works for non-identically distributed excess returns $\mathbf{e}_t$ as well, but closed-form solutions of the final efficient frontier and optimal strategies will be much more difficult to obtain (if at all).

At least one component of $\boldsymbol{\mu}$ must be nonzero so that excess expected profits are possible—the sign is not important since we allow arbitrary short sales—and Σ must be strictly positive definite so that there are no non-risky combinations of the risky assets.

We denote by $\Omega \subset \mathbb{R}^n$ the set of possible values of each $\mathbf{e}_t$, and we make no assumptions about the structure of Ω. The returns may be multivariate normal, a multinomial tree taking m distinct values at each step ($m > n$ so that Σ may be positive definite), or anything else. All we need are the mean and variance as stated.

Let $\mathcal{I}_t$ be the σ-algebra generated by $\{\mathbf{e}_0, \ldots, \mathbf{e}_{t-1}\}$, for $t = 0, \ldots, T$. For $t = 0, \ldots, T-1$, $\mathcal{I}_t$ is the information available to the investor before he makes his investment at time t; since w_0 is nonrandom, $\mathcal{I}_t$ also includes the wealth $w_0, \ldots, w_t$. $\mathcal{I}_T$ is the information set *after* the end of the last investment period, when he simply collects the final value w_T.

An investment plan for the t-th period is an $\mathcal{I}_t$-measurable vector $\mathbf{y}_t = (y_t^1, \ldots, y_t^n)'$, where y_t^i is the amount invested in the i-th risky asset. The amount invested in the riskless asset is $y_t^0 = w_t - \sum_{i=1}^n y_t^i$, and the system dynamics can be summarized in the discrete time stochastic system

$$w_{t+1} = f\big(w_t, \mathbf{y}_t, \mathbf{e}_t\big) \tag{5.1}$$

for $t = 0, 1, \ldots, T-1$, with

$$f(w_t, \mathbf{y}_t, \mathbf{e}_t) = r w_t + \mathbf{e}_t' \mathbf{y}_t \ . \tag{5.2}$$

The series of these investment decisions $(\mathbf{y}_0, \ldots, \mathbf{y}_{T-1})$ constitutes a nonanticipating portfolio policy π. We let $\mathcal{A}_0$ be the set of all such policies; we do not consider portfolio constraints other than mean and variance limits, but such restrictions could easily be included in the definition of $\mathcal{A}_0$. Also, our overall approach can be extended to general systems with structure (5.1), independently of the specific form (5.2).

The problem of the mean-variance investor can be stated in any one of the following three forms:

P1(α) Minimize the variance of wealth for a given level of expected wealth:

$$\min_{\pi \in \mathcal{A}_0} \mathrm{Var}\,[w_T] \quad \text{such that} \quad \mathbb{E}\,[w_T] = \alpha \ .$$

P2(σ) Maximize the expected value of wealth for a given level of variance:

$$\max_{\pi \in \mathcal{A}_0} \mathbb{E}\,[w_T] \quad \text{such that} \quad \mathrm{Var}\,[w_T] = \sigma^2 \ .$$

P3(λ) Maximize a combination of expectation and variance for a given risk aversion coefficient $\lambda \geq 0$:

$$\max_{\pi \in \mathcal{A}_0} \big(\ \mathbb{E}\,[w_T] - \lambda\, \mathrm{Var}\,[w_T] \ \big) \ .$$

These formulations are essentially equivalent (Li and Ng, 2000). In practice one might prefer to meet or exceed a return target, or to incur not more than a maximum variance, rather than the equality constraints as stated above. For technical reasons in Section 5.4 we prefer the equality targets.

In stating P1–P3, it is implicit that expectation and variance are evaluated at $t = 0$, before any information is revealed. The portfolio policy π lays down a rule for the entire investment process, for every possible series of "states of nature" until time T (all possible stock price paths). We can understand this as programing a computer at time $t = 0$, in order to optimize the expectation and variance of final wealth w_t *measured at time* $t = 0$. This program might contradict the risk-reward preference of the "future self" of the investor at some later time $t = k$: if the investor were to re-program the computer at $t = k$ in order to optimize mean and variance of w_T – *measured at time* $t = k$ – according to his risk-reward preference, he might want to do differently as he sees gains or losses made between $t = 0$ and $t = k$ as *sunk cost.*

An example of this type of optimization in the financial industry is trading a portfolio on behalf of a client. At the time when the client entrusts his bank (or other institutional fund manager) with his money, he communicates his risk-reward preference. He is then only given reports of the portfolio performance at the end of the year, leaving the daily trading activities to the bank. Consequently, at the beginning of each year the bank has to determine a portfolio strategy that optimizes the client's wealth at the end of the year, since that is how the bank is evaluated by the client.

5.3. Dynamic Programming

Of the three formulations above, the third P3(λ) is the most general, often appearing as a Lagrange multiplier approach to either of the first two. It is also appealing because it is strongly reminiscent of a utility function $u(w_T)$, and it is known that multi-step problems involving utility maximization $\max \mathbb{E}\left[u(w_T)\right]$ can be solved by dynamic programming, reducing them to a sequence of single-period problems.

But despite this apparent similarity, problem P3(λ) cannot be used directly as a utility function on the stochastic system (5.1,5.2), since the variance $\operatorname{Var}\left[w_T\right] = \mathbb{E}\left[w_T^2\right] - \mathbb{E}\left[w_T\right]^2$ involves the square of the expectation, and hence does not fit into the framework of expected utility maximization. The dynamic programming approach relies on the "smoothing property" $\mathbb{E}\left[\mathbb{E}\left[u(w_T) \mid \mathcal{I}_s\right] \mid \mathcal{I}_t\right] = \mathbb{E}\left[u(w_T) \mid \mathcal{I}_t\right]$ for $s > t$, which does not hold for the variance term $\operatorname{Var}\left[u(w_T)\right]$. Zhou and Li (2000) overcome this obstacle by using an embedding of $\mathbb{E}\left[w_T\right] - \lambda \operatorname{Var}\left[w_T\right]$ into a family of truely quadratic utility functions $\mathbb{E}\left[\rho w_T - \lambda w_T^2\right]$ with an additional parameter $\rho \in \mathbb{R}$ and relating the solution set for (λ, ρ) to the original problem P3(λ). Unfortunately, the introduction of the additional parameter ρ greatly complicates the problem and makes it hard to obtain explicit solutions.

As in Chapter 3, we now show how we can apply dynamic programming directly to the multi-period mean-variance optimization problem, with a suitable choice of the value function. In the following, we will solve the problem in formulation P1(α), and derive the solutions for P2(σ) and P3(λ) from that. In fact, the approach could as well adopted to solve instead P2(σ) directly.

Recall that $\mathbf{y}_t$ is the investor's investment decision at time $t = 0, \ldots, T-1$, and that the excess return of this investment from t to $t+1$ is $\mathbf{e}_t' \mathbf{y}_t$. Let

$$g_t(\mathbf{y}_t, \mathbf{e}_t) = r^{T-t-1} \mathbf{e}_t' \mathbf{y}_t \tag{5.3}$$

be the eventual value at time T of this excess gain, reinvested for the remaining periods at the riskless rate. Note that

$$w_T - r^{T-t} w_t = \sum_{s=t}^{T-1} g_s(\mathbf{y}_s, \mathbf{e}_s) \quad \text{for } t = 0, 1, \ldots, T-1 \quad .$$

For $t = 0, \ldots, T-1$, let $\mathcal{A}_t$ be the set of all nonanticipating portfolio policies that start at time t: the investor makes his first investment at t, and the last at $T-1$. Let $\mathcal{G}_t(w_t, \hat{\alpha}) \subset \mathcal{A}_t$ be the set of all non-anticipating portfolio policies that start at time t with wealth w_t, for which the expected excess gain of final wealth over the riskless investment is exactly $\hat{\alpha}$:

$$\mathcal{G}_t(w_t, \hat{\alpha}) = \left\{ \pi \in \mathcal{A}_t \;\middle|\; \mathbb{E}\left[w_T - r^{T-t} w_t \;\middle|\; \mathcal{I}_t \right] = \hat{\alpha} \right\} \quad . \tag{5.4}$$

It is easy to see that $\mathcal{G}_t$ is nonempty for all w_t and $\hat{\alpha}$, for nonzero $\boldsymbol{\mu}$ and with no investment limits; for example, the constant investment $y_t = \cdots = y_{T-1} = y$ yields an expected excess gain of $\boldsymbol{\mu}' y \cdot (1 + r + \cdots + r^{T-t-1})$, which may take any desired value $\hat{\alpha}$ by choice of $\mathbf{y}$.

For $t = 0, \ldots, T-1$, consider the value function

$$J_t(w_t, \hat{\alpha}) = \min_{\pi \in \mathcal{G}_t(w_t, \hat{\alpha})} \operatorname{Var}\left[w_T \mid \mathcal{I}_t \right] = \min_{\pi \in \mathcal{G}_t(w_t, \hat{\alpha})} \operatorname{Var}\left[\sum_{s=t}^{T-1} g_s(\mathbf{y}_s, \mathbf{e}_s) \,\middle|\, \mathcal{I}_t \right] , \tag{5.5}$$

which is the minimum variance of final wealth achievable by an investment policy for $t, \ldots, T-1$ such that the expected final excess gain is $\hat{\alpha}$, starting with current wealth w_t. The solution to problem P1(α) is $J_0\big(w_0, \alpha - r^T w_0\big)$, together with the corresponding optimal policy $\pi^* \in \mathcal{A}_0$.

As in Chapter 3 we use the law of total variance (see Lemma 3.1), the counterpart to the smoothing property of the expectation operator mentioned above. In particular, for $s = t+1$,

$$\operatorname{Var}\left[w_T \mid \mathcal{I}_t \right] = \operatorname{Var}\left[\mathbb{E}\left[w_T \mid \mathcal{I}_{t+1} \right] \mid \mathcal{I}_t \right] + \mathbb{E}\left[\operatorname{Var}\left[w_T \mid \mathcal{I}_{t+1} \right] \mid \mathcal{I}_t \right] . \tag{5.6}$$

In this expression, $\mathbb{E}\left[w_T \mid \mathcal{I}_{t+1} \right]$ is still uncertain at time t, since it depends on the outcome of the period return $\mathbf{e}_t$. The first term can thus be thought of the risk resulting from the uncertainty between t and $t+1$.

$\operatorname{Var}\left[w_T \mid \mathcal{I}_{t+1} \right]$ is the risk that we will face in the remaining periods from $t+1$ and T. The second term is therefore the mean risk for the remainder after the next period.

As in Chapter 3, the key idea is that the expected gain for the remaining periods between $t+1$ and T does not have to be the same for all possible outcomes of $\mathbf{e}_t$. At time t, we specify target expected returns for the remaining time, depending on the portfolio return in the period to come. Thus, we make our investment strategy for $t+2, \ldots, T$ dependent on the performance of our portfolio over the next period. For instance, we might plan that after a fortunate windfall profit we will conserve the realized gain and be content with less expected return for the remainder, putting less capital at risk. Crucially, we choose this rule to optimize mean and variance measured at the current time t; it is this that makes dynamic programming possible.

We specify $\mathbb{E}\left[w_T \mid \mathcal{I}_{t+1} \right]$ in terms of a $\mathcal{I}_{t+1}$-measurable random variable

$$z_t \;=\; \mathbb{E}\left[w_T \mid \mathcal{I}_{t+1} \right] \;-\; r^{T-t-1} w_{t+1} \;.$$

The value of z_t may depend on $\mathbf{e}_t$ (asset prices and wealth at $t+1$) in addition to $w_0, \ldots, w_t$ and $\mathbf{e}_0, \ldots, \mathbf{e}_{t-1}$. Thus, z_t may be interpreted as a function $\Omega \to \mathbb{R}$, measurable in the probability measure of $\mathbf{e}_t$. Let $\mathcal{L}$ be the set of all such functions.

The function z_t is determined at time t, and specifies for each possible outcome of $\mathbf{e}_t$ what excess return target we will set for the remaining periods. Recall that $\mathbf{y}_t$ is $\mathcal{I}_t$-measurable and hence may not depend on $\mathbf{e}_t$. Whereas the dimension of $\mathbf{y}_t$ is the number of assets n, the dimension of z_t is the cardinality of Ω, which may be a finite number m or infinity.

In terms of z_t and the dynamics (5.1,5.2), we have

$$\mathbb{E}\left[w_T \mid \mathcal{I}_{t+1} \right] = r^{T-t-1} w_{t+1} + z_t = r^{T-t} w_t + r^{T-t-1} \mathbf{e}_t' \mathbf{y}_t + z_t \tag{5.7}$$

and, assuming that we make optimal decisions from $t+1$ to T,

$$\operatorname{Var}\left[w_T \mid \mathcal{I}_{t+1} \right] = J_{t+1}(w_{t+1}, z_t) = J_{t+1}\big(r w_t + \mathbf{e}_t' \mathbf{y}_t, z_t \big) \;. \tag{5.8}$$

Not all functions z_t are compatible with our current excess return target $\hat{\alpha}$. Using (5.7), we may write the eventual expected excess return as

$$\mathbb{E}\left[w_T - r^{T-t} w_t \,\middle|\, \mathcal{I}_t \right] \;=\; r^{T-t-1} \boldsymbol{\mu}' \mathbf{y}_t \;+\; \mathbb{E}\left[z_t \right] \;\; .$$

Thus the set of $(\mathbf{y}, z)$ pairs that yield the desired final excess profit of $\hat{\alpha}$, as required by (5.4), is

$$\tilde{\mathcal{G}}_t(\hat{\alpha}) \;=\; \left\{ \mathbf{y} \in \mathbb{R}^n,\; z \in \mathcal{L} \;\middle|\; r^{T-t-1} \boldsymbol{\mu}' \mathbf{y} + \mathbb{E}\left[z \right] = \hat{\alpha} \right\} \; .$$

This is simply a linear condition on the pair (y, z): if we choose an investment $\mathbf{y}$ that yields a low expected profit this period, then we must compensate by setting higher targets z for the remaining periods. The above expressions are valid for $t = 0, \ldots, T-2$; for $t = T-1$, it is easy to see that the same expressions hold without the terms including z.

Suppose we have already determined $J_{t+1}(w_{t+1}, \hat{\alpha})$. By the dynamic programming principle with (5.6,5.7,5.8), we obtain

$$\begin{aligned} J_t(w_t, \hat{\alpha}) &= \min_{\mathcal{G}_t(w,\hat{\alpha})} \operatorname{Var}\left[w_T - r^{T-t} w_t \,\middle|\, \mathcal{I}_t \right] \\ &= \min_{(y,z) \in \tilde{\mathcal{G}}_t(\hat{\alpha})} \left\{ \operatorname{Var}\left[r^{T-t-1} \mathbf{e}_t' \mathbf{y} + z \right] + \mathbb{E}\left[J_{t+1}(r w_t + \mathbf{e}_t' \mathbf{y}, z) \right] \right\} \; . \end{aligned} \tag{5.9}$$

In this expression, expectation and variance of $\mathbf{e}$ and $z = z(\mathbf{e})$ are taken using the density of $\mathbf{e} = \mathbf{e}_t$ on its sample space Ω.

Equation (5.9) completely describes the dynamic program to find optimal Markovian policies for $\mathrm{P1}(\alpha)$.

As seen in Chapter 3, even for more complicated system dynamics $f(\cdot)$ than (5.2), a dynamic program for multiperiod mean-variance optimization problems in discrete time can be derived. For the portfolio transaction problem in Chapter 3, we had to resort to numerical treatment to solve the dynamic program. Contrary, in the following we will show that for the portfolio selection problem we can solve (5.9) analytically, and inductively derive a closed-form solution for the value function at time $t = 0$.

5.4. Explicit Solution

Before we solve the optimization problem (5.9) in the backwards step, we need to find the solution for the single-period problem $J_{T-1}(w, \hat{\alpha})$ first, which is easily obtained as the solution of a constrained quadratic program:

LEMMA 5.1.

$$J_{T-1}(w, \hat{\alpha}) = \hat{\alpha}^2 / \theta^2 \;\; , \qquad \theta^2 = \boldsymbol{\mu}' \Sigma^{-1} \boldsymbol{\mu} \;\; . \tag{5.10}$$

PROOF. For current wealth w and required expected excess gain $\hat{\alpha}$, the single-period mean-variance problem reads

$$\min_{\mathbf{y} \in \mathbb{R}^n} \mathbf{y}' \Sigma \mathbf{y} \quad \text{s.t.} \quad \boldsymbol{\mu}' \mathbf{y} = \hat{\alpha} \;\; . \tag{5.11}$$

The solution is easily found as $\mathbf{y} = \hat{\alpha} \Sigma^{-1} \boldsymbol{\mu} / \theta^2$, which yields the $t = T-1$ value function $J_{T-1}(w, \hat{\alpha}) = \hat{\alpha}^2 / \theta^2$. □

Now we are ready to derive the explicit solution of the portfolio selection problem by the dynamic program stated in Sect. 5.3.

LEMMA 5.2. *The optimal investment decision at time t for a mean-variance investor with current wealth $w_t = w$, who is optimizing P1($\hat{\alpha} + r^{T-t} w_t$), i.e.*

$$\min \ \mathrm{Var}\left[w_T \mid \mathcal{I}_t \right] \quad s.t. \quad \mathbb{E}\left[w_T - r^{T-t} w_t \,\middle|\, \mathcal{I}_t \right] = \hat{\alpha} \ ,$$

is the control $(\mathbf{y}_t, z_t)$ given by

$$\mathbf{y}_t^* = \frac{\hat{\alpha}}{r^{\tau-1}} \frac{1+\theta_{\tau-1}^2}{\theta_\tau^2} \Sigma^{-1} \boldsymbol{\mu} \tag{5.12}$$

$$z_t^* = \hat{\alpha} \frac{\theta_{\tau-1}^2}{\theta_\tau^2} \Big(1 - \boldsymbol{\mu}' \Sigma^{-1} (\mathbf{e}_t - \boldsymbol{\mu}) \Big) \tag{5.13}$$

with $\tau = T - t$ and

$$\theta_\tau^2 = (\theta^2 + 1)^\tau - 1 \quad \text{for } 0 \le \tau \le T, \tag{5.14}$$

with $\theta^2 = \boldsymbol{\mu}' \Sigma^{-1} \boldsymbol{\mu}$. Here, $\mathbf{y}_t^*$ is the optimal investment vector and z_t^* gives the planned expected gain the investor commits himself for the remaining periods $t+1, \ldots, T$ as a function of the realized outcome of the return $\mathbf{e}_t$ over the period to come. The value function (5.5) *is*

$$J_t(w_t, \hat{\alpha}) = \inf \Big\{ \mathrm{Var}\left[w_T \mid \mathcal{I}_t \right] \ \Big| \ \mathbb{E}\left[w_T - r^{T-t} w_t \,\middle|\, \mathcal{I}_t \right] = \hat{\alpha} \Big\} = \frac{\hat{\alpha}^2}{\theta_\tau^2} \ . \tag{5.15}$$

PROOF. We prove by induction. Since $\theta_0 = 0$ and $\theta_1 = \theta$, Lemma 5.1 proves the result for $t = T-1$. For $0 \le t \le T-2$, assume that $J_{t+1}(w, \hat{\alpha}) = \hat{\alpha}^2 / \theta_{\tau-1}^2$, and then the dynamic programming step (5.9) is

$$J_t(w, \hat{\alpha}) \ = \inf_{\mathbf{y}, z} \Big\{ V(y, z) \ \Big| \ E(y, z) = \hat{\alpha} \Big\} \tag{5.16}$$

with

$$\begin{aligned} E(\mathbf{y}, z) \ &= \ r^{\tau-1} \boldsymbol{\mu}' \mathbf{y} \ + \ \mathbb{E}\left[z \right] \\ V(\mathbf{y}, z) \ &= \ \mathrm{Var}\left[r^{\tau-1} \mathbf{e}' \mathbf{y} + z \right] \ + \ \frac{1}{\theta_{\tau-1}^2} \ \mathbb{E}\left[z^2 \right] \ . \end{aligned}$$

Here and in the following we use $\mathbf{e} = \mathbf{e}_t$ instead of $\mathbf{e}_t$ to shorten notation.

Our first observation is that, since

$$\mathrm{Var}\left[r^{\tau-1} \mathbf{e}' \mathbf{y} + z \right] \ = \ \mathrm{Var}\left[r^{\tau-1} \mathbf{e}' \mathbf{y} \right] \ + \ \mathrm{Var}\left[z \right] \ + \ 2 r^{\tau-1} \, \mathrm{Cov}\left[\mathbf{e}' \mathbf{y}, z \right] \ ,$$

we can generally reduce $V(y, z)$ by introducing negative correlation between $\mathbf{e}'\mathbf{y}$ and z, where correlation between two random variables X and Y with positive variance is the standard $\rho(X, Y) = \mathrm{Cov}\left[X, Y \right] / \sqrt{\mathrm{Var}\left[X \right] \mathrm{Var}\left[Y \right]}$. In fact, we now show that in the optimal solution, these two terms have precisely minimal correlation $\rho(\mathbf{e}'\mathbf{y}, z) = -1$.

To see that, suppose $\rho\big(\mathbf{e}'\mathbf{y}^*, z^*\big) > -1$. Consider the control $(\mathbf{y}^*, \tilde{z})$ with $\tilde{z} = a\,\mathbf{e}'\mathbf{y}^* + b$, where $a = -\sqrt{\mathrm{Var}\left[z^* \right] / \mathrm{Var}\left[\mathbf{e}'\mathbf{y}^* \right]}$ and $b = \mathbb{E}\left[z^* \right] - a\,\mathbb{E}\left[\mathbf{e}'\mathbf{y}^* \right]$. Then z^* and $\tilde{z}$ have identical first and second moments, but $\rho(\mathbf{e}'\mathbf{y}^*, \tilde{z}) = -1 < \rho(\mathbf{e}'\mathbf{y}^*, z^*)$ implies $\mathrm{Cov}\left[\mathbf{e}'\mathbf{y}^*, \tilde{z} \right] < \mathrm{Cov}\left[\mathbf{e}'\mathbf{y}^*, z^* \right]$, and thus the control $(\mathbf{y}^*, \tilde{z})$ yields a lower value of V for the same E. This correlation observation is the key fact that lets us determine the optimal solution independently of the structure of Ω.

Thus, up to a set of zero measure, we must have

$$z = ar^{\tau-1}(\mathbf{e} - \boldsymbol{\mu})'\mathbf{y} + b \ , \tag{5.17}$$

for some $a < 0$ and $b \in \mathbb{R}$ (note that $\boldsymbol{\mu}'\mathbf{y}$ is a constant). Then

$$\begin{aligned} \mathbb{E}[z] &= b \\ \mathbb{E}\left[z^2\right] &= b^2 + a^2 r^{2(\tau-1)}\mathbf{y}'\Sigma\mathbf{y} \\ \mathrm{Var}\left[r^{\tau-1}\mathbf{e}'\mathbf{y} + z\right] &= (a+1)^2 r^{2(\tau-1)}\mathbf{y}'\Sigma\mathbf{y} \ , \end{aligned}$$

and we may write $E(y,z)$ and $V(y,z)$ in terms of y, a, and b as

$$E(\mathbf{y},a,b) = r^{\tau-1}\boldsymbol{\mu}'\mathbf{y} + b \ , \tag{5.18}$$

$$V(\mathbf{y},a,b) = \frac{b^2}{\theta_{\tau-1}^2} + \left((a+1)^2 + \frac{a^2}{\theta_{\tau-1}^2}\right) r^{2(\tau-1)}\mathbf{y}'\Sigma\mathbf{y} \ . \tag{5.19}$$

Instead of solving (5.16) directly, we optimize the trade-off function

$$H(w,\lambda) = \sup_{\mathbf{y}\in\mathbb{R}^n, a<0, b\in\mathbb{R}} L(\mathbf{y},a,b,\lambda) \qquad \text{for } \lambda > 0 \ , \tag{5.20}$$

with

$$L(\mathbf{y},a,b,\lambda) = E(\mathbf{y},a,b) - \lambda V(\mathbf{y},a,b) \ . \tag{5.21}$$

Since $L(\mathbf{y},a,b,\lambda)$ is a convex quadratic function in $(\mathbf{y},a,b)$, there exists a unique optimum, which we can find by first-order conditions.

The first-order conditions $\partial L/\partial a = 0$ and $\partial L/\partial b = 0$ immediately yield

$$a^* = -\frac{\theta_{\tau-1}^2}{1+\theta_{\tau-1}^2} \quad \text{and} \quad b^* = \frac{\theta_{\tau-1}^2}{2\lambda} \ .$$

Plugging a^* and b^* into the first-order condition for $\mathbf{y}$, we obtain

$$\mathbf{y}^*(\lambda) = \frac{1+\theta_{\tau-1}^2}{2\lambda r^{\tau-1}}\,\Sigma^{-1}\boldsymbol{\mu} \ ,$$

then

$$z^*(\lambda) = \frac{\theta_{\tau-1}^2}{2\lambda}\left(1 - \boldsymbol{\mu}'\Sigma^{-1}(\mathbf{e}-\boldsymbol{\mu})\right)$$

and finally

$$\begin{aligned} E^*(\lambda) &= \frac{\theta_{\tau-1}^2(\theta^2+1)+\theta^2}{2\lambda} = \frac{\theta_\tau^2}{2\lambda} \\ V^*(\lambda) &= \frac{\theta_{\tau-1}^2(\theta^2+1)+\theta^2}{4\lambda^2} = \frac{\theta_\tau^2}{4\lambda^2} \ , \end{aligned}$$

where we use the definition (5.14) of θ_t. Solving $E^*(\lambda) = \hat{\alpha}$ gives $\lambda = \theta_\tau^2/2\hat{\alpha}$, and we obtain (5.12,5.13) by substituting λ above, and finally

$$J_t(w,\hat{\alpha}) = V^*(\lambda) = \hat{\alpha}^2/\theta_\tau^2 \ ,$$

which completes the inductive proof. □

We can give the optimal policy $(\mathbf{y}_0, \ldots, \mathbf{y}_{T-1})$ explicitly in terms of $\mathbf{e}_i$. Let

$$A_t = \prod_{s=0}^{t-1} \Big(1 - \boldsymbol{\mu}' \Sigma^{-1} (\mathbf{e}_s - \boldsymbol{\mu}) \Big) \tag{5.22}$$

$(A_0 = 1)$. Then

THEOREM 5.3. *For $T \geq 1$, the adaptive solution to P1(α) is the mean-variance efficient frontier, $E = \mathbb{E}[w_T]$ and $V = \mathrm{Var}[w_T]$,*

$$E = \alpha, \quad V(\alpha) = \big(\alpha - r^T w_0\big)^2 / \theta_T^2 \ , \tag{5.23}$$

and the optimal portfolio strategy $(\mathbf{y}_0, \ldots, \mathbf{y}_{T-1})$ in terms of $\mathbf{e}_i$ is

$$\mathbf{y}_t = \big(\alpha - r^T w\big) \frac{1 + \theta_{T-t-1}^2}{r^{T-t-1} \theta_T^2} A_t \Sigma^{-1} \boldsymbol{\mu} \ , \tag{5.24}$$

with A_t from (5.22) and θ_T from (5.14).

PROOF. Equation (5.23) follows directly from Lemma 5.2. By induction, we will now prove

$$\mathbf{y}_t = \big(\alpha - r^T w\big) \frac{1 + \theta_{T-t-1}^2}{r^{T-t-1} \theta_T^2} A_t \Sigma^{-1} \boldsymbol{\mu} \ , \tag{5.25}$$

$$z_t = \big(\alpha - r^T w\big) \frac{\theta_{T-t-1}^2}{\theta_T^2} A_{t+1} \ . \tag{5.26}$$

For $t = 0$, (5.25, 5.26) follows from Lemma 5.2. Suppose (5.25, 5.26) hold at time $t-1$. Then, again by Lemma 5.2

$$\begin{aligned} \mathbf{y}_t &= z_{t-1} \frac{1 + \theta_{T-t-1}^2}{r^{T-t-1} \theta_{T-t}^2} \Sigma^{-1} \boldsymbol{\mu} \\ &= \big(\alpha - r^T w\big) \frac{\theta_{T-t}^2}{\theta_T^2} A_t \cdot \frac{1 + \theta_{T-t-1}^2}{r^{T-t-1} \theta_{T-t}^2} \Sigma^{-1} \boldsymbol{\mu} \end{aligned}$$

and

$$\begin{aligned} z_t &= z_{t-1} \frac{\theta_{T-t-1}^2}{\theta_T^2} \Big(1 - \boldsymbol{\mu}' \Sigma^{-1} (\mathbf{e}_t - \boldsymbol{\mu}) \Big) \\ &= \big(\alpha - r^T w\big) \frac{\theta_{T-t}^2}{\theta_T^2} A_t \cdot \frac{\theta_{T-t-1}^2}{\theta_{T-t}^2} \Big(1 - \boldsymbol{\mu}' \Sigma^{-1} (\mathbf{e}_t - \boldsymbol{\mu}) \Big) \ , \end{aligned}$$

completing the inductive proof. □

From Theorem 5.3, we can easily derive the solutions for P2(σ) and P3(λ).

COROLLARY 5.4. *Let $T \geq 1$. Let θ_T and A_t defined by* (5.14) *and* (5.22).

(i) The adaptive solution to P2(σ) is given by the mean-variance efficient frontier

$$\big\{ (E(\sigma), \sigma^2) \mid E(\sigma) = w_0 r^T + \sigma \theta_T \text{ and } \sigma \geq 0 \big\} \ .$$

(ii) The solution to P3(λ), i.e. the value of $U(\lambda) = \mathbb{E}[w_T] - \lambda - \mathrm{Var}[w_T]$ under an optimal adaptive strategy, is

$$U(\lambda) = w_0 r^T + \theta_T^2 / (4\lambda) \ .$$

Furthermore, the optimal portfolio strategy $(\mathbf{y}_0, \ldots, \mathbf{y}_{T-1})$ *is given by*

$$\mathbf{y}_t = \gamma \cdot \frac{1+\theta_{T-t-1}^2}{r^{T-t-1}\theta_T^2} A_t \Sigma^{-1} \boldsymbol{\mu} \ , \tag{5.27}$$

where $\gamma = \theta_T \sigma$ *in the case of P2(*σ*) and* $\gamma = \theta_T^2/(2\lambda)$ *in the case of P3(*λ*).*

Proof. From (5.23), for given $V(\alpha^*) = \sigma^2$ we first obtain $\alpha^* = wr^T + \theta_T \sigma$. The optimal solution for P2(σ) is then given by the optimal solution (5.24) for P1(α) with $\alpha = \alpha^* = wr^T + \theta_T \sigma$.

For P3(λ), note that the optimal solution for $\max_\pi \{\mathbb{E}[w_T] - \lambda \operatorname{Var}[w_T]\}$ for given $\lambda \geq 0$ can only lie on the E-V-frontier given by (5.23). Thus, it is equivalent to solve $\max_{\alpha \geq 0}\{\alpha - \lambda(\alpha - w_0 r^T)^2/\theta_T^2\}$, which is optimized by $\alpha^* = \theta_T^2/(2\lambda) + w_0 r^T$ with optimal value $U(\lambda) = w_0 r^T + \theta_T^2/(4\lambda)$, and the optimal solution for P3(α) is given by the optimal solution for P1(α) with $\alpha = \alpha^* = \theta_T^2/(2\lambda) + w_0 r^T$. □

As can be seen from (5.24), similar to the well-known fund-separation theorem the proportions of the assets in the risky part of the portfolio in an optimal adaptive investment policy are constant throughout, given by $\Sigma^{-1}\boldsymbol{\mu}$. What varies is the total amount invested in risky assets.

The expression for the mean-variance efficient frontier coincide with the expression given by Li and Ng (2000). However, the optimal policy is specified differently. Li and Ng (2000) give the optimal investment decision at time t as a function of the current wealth w_t. In Theorem 5.3 and Corollary 5.4 the optimal investment decision is given as a function of the sequence of the returns $\mathbf{e}_t$.

We can also rewrite the optimal strategy as a function of the wealth process w_t.

Theorem 5.5. *Let* $\hat{\alpha} = \alpha - w_0 r^T$ *be the expected excess return. The adaptive solution to P1(*α*), given in Theorem 5.3, can be rewritten as a function of the wealth process* w_t *as*

$$\mathbf{y}_t r^{-t} = \frac{\Sigma^{-1}\boldsymbol{\mu} r}{(\theta^2+1)} \left(\hat{\alpha} \frac{1+\theta_T^2}{\theta_T^2} - \left(w_t r^{-t} - w_0\right) \right) \ . \tag{5.28}$$

Proof. We prove by induction. By Theorem 5.3, we have

$$\mathbf{y}_0 = \hat{\alpha} \frac{1+\theta_{T-1}^2}{r^{T-1}\theta_T^2} \Sigma^{-1} \boldsymbol{\mu} \ .$$

Since $(1+\theta_T^2)/(1+\theta^2) = 1 + \theta_{T-1}^2$, this shows the base case $t = 0$.

For the inductive step, note that by Theorem 5.3,

$$\begin{aligned} \mathbf{y}_{t+1} &= \hat{\alpha} \frac{1+\theta_{T-t-2}^2}{r^{T-t-2}\theta_T^2} A_t \big(1 - \boldsymbol{\mu}'\Sigma^{-1}(\mathbf{e}_t - \boldsymbol{\mu})\big) \Sigma^{-1} \boldsymbol{\mu} \\ &= \mathbf{y}_t \frac{r}{\theta^2+1} \big(1 + \theta^2 - \boldsymbol{\mu}'\Sigma^{-1}\mathbf{e}_t\big) = \mathbf{y}_t r \left(1 - \frac{\boldsymbol{\mu}'\Sigma^{-1}\mathbf{e}_t}{\theta^2+1}\right) \ . \end{aligned}$$

Hence, with the inductive hypothesis, we conclude

$$\mathbf{y}_{t+1} - \mathbf{y}_t r = -\mathbf{y}_t r \frac{\mathbf{e}_t'\Sigma^{-1}\boldsymbol{\mu}}{1+\theta^2} \overset{(5.28)}{=} -r\Sigma^{-1}\boldsymbol{\mu} \frac{\mathbf{e}_t'\mathbf{y}_t}{\theta^2+1} \ .$$

By the dynamics (5.1, 5.2), $w_{t+1} = rw_t + \mathbf{e}_t'\mathbf{y}_t$. Thus, again using the inductive hypothesis (5.28),

$$\begin{aligned}
\mathbf{y}_{t+1} &= \mathbf{y}_t r - r\frac{\Sigma^{-1}\boldsymbol{\mu}(w_{t+1} - rw_t)}{\theta^2+1} \\
&= \frac{\Sigma^{-1}\boldsymbol{\mu}}{(\theta^2+1)r^{T-t-2}}\left(\hat{\alpha}\frac{1+\theta_T^2}{\theta_T^2} - \left(w_t r^{T-t} - w_0 r^T\right)\right) \\
&\quad - \frac{r\Sigma^{-1}\boldsymbol{\mu}(w_{t+1} - rw_t)}{\theta^2+1} \\
&= \frac{\Sigma^{-1}\boldsymbol{\mu}}{(\theta^2+1)r^{T-(t+1)-1}}\left(\hat{\alpha}\frac{1+\theta_T^2}{\theta_T^2} - \left(w_t r^{T-(t+1)} - w_0 r^T\right)\right) \quad ,
\end{aligned}$$

completing the inductive proof. □

The formulation of the strategy in Theorem 5.28 resembles the mean-variance optimal strategy given by Richardson (1989) in a continuous-time setting: the discounted amount invested in the risky assets at each time is linear in the discounted gain. From a practical point of view, this may be considered unrealistic, since it assumes the ability to continue trading regardless how high the cumulated losses. To address this issue, Bielecki et al. (2005) add the additional restriction that bankruptcy is prohibited. In fact, such constraints can also be incorporated into the dynamic programming principle for mean-variance that we introduced in Section 5.3, namely as constraints on the set $\mathcal{G}_t(w, \hat{\alpha})$ in (5.9). However, closed-form solutions for the optimal strategy might no longer be available. For instance, for the optimal execution problem in Chapter 3 we required pure sell/buy programs (i.e. a no shorting constraint).

In the remainder of this section, we show how much the optimal dynamic, *path-dependent* strategy in Theorem 5.3 improves over a *path-independent* investment strategy, which is not allowed to respond to changes in the asset price or previous portfolio performance.

Improvement over Path-Independent Strategy. Suppose we specify a path-independent ("static") investment strategy at time $t = 0$, given by the T real-valued vectors $\mathbf{y}_t \in \mathbb{R}^n$ ($t = 0 \dots T-1$).

Lemma 5.6. *Let $T \geq 1$. The static mean-variance efficient frontier, $E = \mathbb{E}[w_T]$ and $V = \mathrm{Var}[w_T]$, is given by*

$$E = \alpha, \quad V(\alpha) = \left(\alpha - w_0 r^T\right)^2/(T\theta^2) \quad . \tag{5.29}$$

The optimal static path-independent investment strategy, specified at time $t = 0$, reads

$$\mathbf{y}_t = \gamma\frac{\Sigma^{-1}\boldsymbol{\mu}}{r^{T-t-1}}, \quad t = 0, 1, \dots, T-1 \quad , \tag{5.30}$$

with $\gamma = (\alpha - w_0 r^T)/(T\theta^2)$ for P1(α), $\gamma = \sigma/(\theta\sqrt{T})$ for P2(σ) and $\gamma = 1/(2\lambda)$ for P3(λ).

Proof. Let $\tilde{\mathbf{y}}_t = r^{T-t-1}\mathbf{y}_t$. The final wealth reads $w_T = w_0 r^T + \sum_{t=0}^{T-1} \mathbf{e}_t'\tilde{\mathbf{y}}_t$, and since $\tilde{\mathbf{y}}_t$ are deterministic, we have

$$\mathbb{E}[w_T] = w_0 r^T + \sum_{t=0}^{T-1} \boldsymbol{\mu}'\tilde{\mathbf{y}}_t \quad \text{and} \quad \mathrm{Var}[w_T] = \sum_{t=0}^{T-1} \tilde{\mathbf{y}}_t'\Sigma\tilde{\mathbf{y}}_t \quad .$$

We solve P3(λ) by maximizing the trade-off function

$$U(\tilde{\mathbf{y}}_0, \dots, \tilde{\mathbf{y}}_{T-1}, \lambda) = \mathbb{E}[w_T] - \lambda\,\mathrm{Var}[w_T], \quad \lambda \geq 0 \tag{5.31}$$

for $\tilde{\mathbf{y}}_0, \ldots, \tilde{\mathbf{y}}_{T-1} \in \mathbb{R}^n$. For fixed $\lambda \geq 0$, the optimal investment decisions $\tilde{\mathbf{y}}_t$ are given by $\tilde{\mathbf{y}}_t = \Sigma^{-1}\boldsymbol{\mu}/(2\lambda)$, $t = 0, 1, \ldots, T-1$, yielding $\mathbb{E}[w_T] = w_0 r^T + T\theta^2/(2\lambda)$ and $\mathrm{Var}[w_T] = T\theta^2/(4\lambda^2)$. From that, we immediately obtain the static mean-variance frontier (5.29), as well as the trading strategies $\tilde{\mathbf{y}}_t$ (and thus $\mathbf{y}_t$) for the three problems P1(α), P2(σ) and P3(λ). □

As noted by Bajeux-Besnainou and Portait (2002), the optimal static T period policy given in Lemma 5.6 already improves over the classical one-period buy-and-hold policy in the original setup of Markowitz (1952, 1959). For $T = 1$ the optimal adaptive policy and the optimal static policy coincide, recovering the one-period solution (Merton, 1972).

From (5.23) and (5.29) we obtain that for given level of expected return α, the optimal dynamic and optimal static strategies have a variance of

$$V_{dyn} = \big(\alpha - w_0 r^T\big)^2/\theta_T^2 \quad \text{and} \quad V_{stat} = \big(\alpha - w_0 r^T\big)^2/(T\theta^2)\ ,$$

respectively. Thus, the optimal dynamic strategy indeed strictly improves over the static, path-independent trading schedule:

Corollary 5.7. *For all $\theta > 0$, we have $\theta_T = \sqrt{(\theta^2+1)^T - 1} > \sqrt{T}\theta$, and $\theta_T = \sqrt{T}\theta + \mathcal{O}\big(\theta^2\big)$. That is, the optimal adaptive investment policy given in Theorem 5.3 strictly improves the optimal static strategy, and for $\theta \to 0$, the two mean-variance frontiers coincide.*

Proof. We have $\theta_T^2 = (\theta^2+1)^T - 1 = \sum_{i=1}^T \binom{T}{i}\theta^{2i}$, thus $\theta_T = \sqrt{T}\theta + \mathcal{O}\big(\theta^2\big)$, and $\theta_T > \sqrt{T}\theta$ for all $\theta > 0$. □

CHAPTER 6

Optimal k-Search and Pricing of Lookback Options

In this chapter, we study the *k-search* problem, in which an online player is searching for the k highest (respectively, lowest) prices in a series of price quotations which are sequentially revealed to him. We analyze the problem using competitive analysis, and we give deterministic and randomized algorithms as well as lower bounds for the competitive ratio. We also show the usefulness of these algorithms to derive bounds for the price of "lookback options".

6.1. Introduction

In this chapter, we consider the following online search problem: an online player wants to sell (respectively, buy) $k \geq 1$ units of an asset with the goal of maximizing his profit (minimizing his cost). At time points $i = 1, \ldots, n$, the player is presented a price quotation p_i, and must *immediately* decide whether or not to sell (buy) one unit of the asset for that price. The player is required to complete the transaction by some point in time n. We ensure that by assuming that if at time $n - j$ he has still j units left to sell (respectively, buy), he is compelled to do so in the remaining j periods. We shall refer to the profit maximization version (selling k units) as k-max-search, and to the cost minimization version (purchasing k units) as k-min-search.

We shall make no modeling assumptions on the price path except that it has finite support, which is known to the player. That is, the prices are chosen from the real interval $\mathcal{I} = \{x \mid m \leq x \leq M\}$, where $0 < m < M$ (see Figure 6.1). We define the *fluctuation ratio* $\varphi = M/m$. Let $\mathcal{P} = \bigcup_{n \geq k} \mathcal{I}^n$ be the set of all price sequences of length at least k. Moreover, the length of the sequence is known to the player at the beginning of the game.

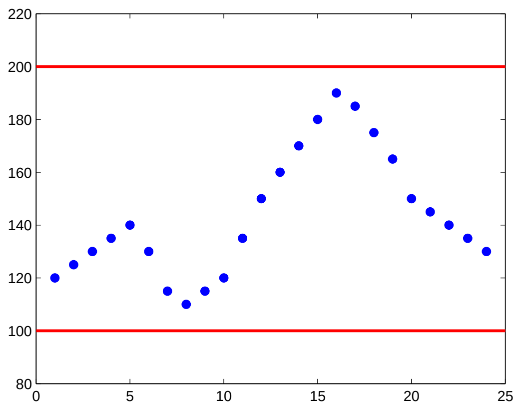

FIGURE 6.1. Example of a bounded price path with $m = 100, \varphi = 2$.

6.1.1. Online Algorithms and Competitive Analysis. In this section we briefly sketch the concepts of online algorithms and competitive analysis. The reader is referred to the textbook of Borodin and El-Yaniv (1998) for more details.

An *online algorithm* is an algorithm that processes its input piece-by-piece, without having the entire input available from the start. Many problems in disciplines such as computer science, economics or operations research are intrinsically online in that they require immediate decisions to be made in real time (for instance, trading). From a broader perspective, online algorithms fall into the field of *decision-making in the absence of complete information.*

Competitive analysis is a method for analyzing online algorithms, in which the performance of an online algorithm is compared to the performance of an optimal *offline algorithm* that has access to the entire input in advance. The optimal offline algorithm is in a sense hypothetical for an online problem since it requires clairvoyant abilities and is thus not realizable. The *competitive ratio* of an algorithm is understood as the *worst-case* ratio of its performance (measured by a problem specific value) and the optimal offline algorithm's performance. One imagines an "adversary" that deliberately chooses difficult inputs in order to maximize this performance ratio. Competitive analysis thus falls within the framework of *worst-case analysis.*

More precisely, let ALG be an online algorithm, and let OPT be the offline optimum algorithm (which knows the entire input sequence in advance) for the same problem. In the case of *k-search*, the input sequence is a price sequence σ, chosen by an adversary out of the set $\mathcal{P}$ of admissible sequences. Let $\mathsf{ALG}(\sigma)$ and $\mathsf{OPT}(\sigma)$ denote the objective values of ALG and OPT when executed on $\sigma \in \mathcal{P}$. The *competitive ratio* of ALG is defined for maximization problems as

$$\mathrm{CR}(\mathsf{ALG}) = \max\left\{ \frac{\mathsf{OPT}(\sigma)}{\mathsf{ALG}(\sigma)} \;\middle|\; \sigma \in \mathcal{P} \right\} \quad ,$$

and similarly, for minimization problems

$$\mathrm{CR}(\mathsf{ALG}) = \max\left\{ \frac{\mathsf{ALG}(\sigma)}{\mathsf{OPT}(\sigma)} \;\middle|\; \sigma \in \mathcal{P} \right\} \quad .$$

We say that ALG is *c-competitive* if it achieves a competitive ratio not larger than c. For randomized algorithms, we substitute the *expected* objective value $\mathbb{E}[\mathsf{ALG}]$ for ALG in the definitions above.

The roots of competitive analysis lie in classical combinatorial optimization theory. Although not named "competitive analysis", Yao (1980) applied it to study how well online bin packing algorithms perform. Later it was first explicitly advocated by Sleator and Tarjan (1985) in the context of list updating and paging problems. Following that new approach, the term "competitive algorithm" was introduced by Karlin, Manasse, Rudolph, and Sleator (1988).

Competitive analysis can also be viewed from a game theoretical perspective as a zero-sum game between an online player and the adversary (Borodin and El-Yaniv, 1998): the online player chooses an algorithm in order to minimize the competitive ratio, and the adversary constructs inputs so as to maximize the competitive ratio. Using results from game theory interesting implications for the design and analysis of competitive online algorithms can be derived. For instance, one application of the well-known Minimax-Theorem is *Yao's principle* (Yao, 1977) for proving lower bounds, which we will use and explain in Section 6.3.

The competitive analysis approach differs from other analysis frameworks that use distributional assumptions of the input and study the distribution of an algorithm's performance, typically the expected performance (or possibly also dispersion measures, e.g. variance). This approach is very common in mathematical finance, economics and operations research, and we did follow such a framework in Chapters 2 to 5 to analyze portfolio and trading strategies.

Any conceptual model has advantages and disadvantages. Competitive analysis has the disadvantage of being too pessimistic by assuming that inputs are generated by some evil fiend who would only pick the worst possible, whereas in practical situations those "bad" instances rarely if ever occur. In fact, an algorithm that performs well under the competitive ratio measure might be just mediocre in the most common situations – since it is only designed to perform well in the bad ones.

On the other hand, frameworks that use distributional assumptions and assess online algorithms for instance by its expected performance have the disadvantage that their implications are highly dependent on the underlying distributional assumptions – which are sometimes not precisely known in practical applications. Also, the desire for analytic tractability often leads to very simplified models.

The advantage of competitive analysis lies in its typically weak modeling assumptions. In that sense, they provide results that are robust and – albeit sometimes pessimistic – hold even when an algorithm's input distributions are not perfectly known. In mathematical finance, where conventionally distributional models are used, recently there have been applications of such "worst-case" performance guarantees (Korn, 2005; Epstein and Wilmott, 1998; DeMarzo et al., 2006).

6.1.2. Related Work. El-Yaniv et al. (2001) study the case $k = 1$, i.e. *1-max-search* and the closely related *one-way trading* problem with the competitive ratio (as defined in Section 6.1.1) as performance measure. In the latter, a player wants to exchange some initial wealth to some other asset, and is again given price quotations one-by-one. However, the player may exchange an *arbitrary fraction* of his wealth for each price. Hence, the k-max-search problem for general $k \geq 1$ can be understood as a natural bridge between the two problems considered by El-Yaniv et al. (2001), with $k \to \infty$ corresponding to the one-way trading problem. This connection will be made more explicit later.

Recently, Fujiwara and Sekiguchi (2007) provide an average-case competitive analysis for the one-way trading problem, and give optimal average-case threat-based algorithms for different distributions of the maximum price.

Several variants of search problems, have been extensively studied in operations research and mathematical economics. Most of the work follows a *Bayesian* approach: optimal algorithms are developed under the assumption that the prices are generated by a known distribution. Naturally, such algorithms heavily depend on the underlying model. Lippmann and McCall (1976, 1981) give an excellent survey on search problems with various assumptions on the price process. More specifically, they study the problem of job and employee search and the economics of uncertainty, which are two classical applications of series search problems. Rosenfield and Shapiro (1981) study the situation where the price follows a random process, but some of its parameters may be random variables with known prior distribution. Hence, Rosenfield and Shapiro (1981) try to get rid of the assumption of the Bayesian search models that the underlying price process is *fully* known to the player. Ajtai, Megiddo, and Waarts (2001) study the classical *secretary* problem. Here, n objects from an ordered set are presented in random order, and the player has to accept k of them so that the final decision about each object is made only on the basis of its rank relative to the ones already seen. They consider the problems of maximizing the probability of accepting the best k objects, or minimizing the expected sum of the ranks (or powers of ranks) of the accepted objects. In this context, Kleinberg (2005) designes an $(1-\mathcal{O}(1/\sqrt{k}))$-competitive algorithm for the problem of maximizing the sum of the k chosen elements.

6.1.3. Results and Discussion. In contrast to the Bayesian approaches mentioned above, El-Yaniv et al. (2001) circumvent almost all distributional assumptions by resorting to competitive analysis and the minimal assumption of a known finite price interval. Here we also follow this approach. The goal is to provide a generic search strategy that works with any price evolution, rather than to retrench to a specific stochastic price process. In many applications, where it is not clear how the generating price process should be modeled, this provides an attractive alternative to classical Bayesian search models. In fact, in Section 6.4 we give an application of k-max-search and k-min-search to *robust option pricing* in finance, where relaxing typically made assumptions on the (stochastic) price evolution to the minimal assumption of a *price interval* yields remarkably good bounds.

Before we proceed with stating our results, let us introduce some notation. For $\sigma \in \mathcal{P}, \sigma = (p_1, \ldots, p_n)$, let $p_{max}(\sigma) = \max_{1 \le i \le n} p_i$ denote the maximum price, and $p_{min}(\sigma) = \min_{1 \le i \le n} p_i$ the minimum price. Let W denote *Lambert's W-function*, i.e., the inverse of $f(w) = w\exp(w)$. For brevity we shall write $f(x) \sim g(x)$, if $\lim_{x\to\infty} f(x)/g(x) = 1$. It is well-known that $W(x) \sim \ln x$.

Our results for deterministic k-max-search are summarized in Theorem 6.1.

THEOREM 6.1. *Let $k \in \mathbb{N}$, $\varphi > 1$. There is a r^*-competitive deterministic algorithm for k-max-search, where $r^* = r^*(k, \varphi)$ is the unique solution of*

$$\frac{\varphi - 1}{r^* - 1} = \left(1 + \frac{r^*}{k}\right)^k \quad , \tag{6.1}$$

and there exists no deterministic algorithm with smaller competitive ratio. Furthermore, we have

(i) $r^*(k, \varphi) \sim \sqrt[k+1]{k^k \varphi}$ for fixed $k \ge 1$ and $\varphi \to \infty$,
(ii) $r^*(k, \varphi) \sim 1 + W(\frac{\varphi-1}{e})$ for fixed $\varphi > 1$ and $k \to \infty$.

The algorithm in the theorem above is given explicitly in Section 6.2. Interestingly, the optimal competitive deterministic algorithm for the one-way trading problem studied by El-Yaniv et al. (2001) has competitive ratio exactly $1 + W(\frac{\varphi-1}{e})$ (for $n \to \infty$), which coincides with the ratio of our algorithm given by the theorem above for $k \to \infty$. Hence, k-max-search can indeed be understood as a natural bridge between the 1-max-search problem and the one-way trading problem.

For deterministic k-min-search we obtain the following statement.

Theorem 6.2. *Let $k \in \mathbb{N}$, $\varphi > 1$. There is a s^*-competitive deterministic algorithm for k-min-search, where $s^* = s^*(k, \varphi)$ is the unique solution of*

$$\frac{1 - 1/\varphi}{1 - 1/s^*} = \left(1 + \frac{1}{ks^*}\right)^k \ , \tag{6.2}$$

and there exists no deterministic algorithm with smaller competitive ratio. Furthermore, we have

(i) $s^*(k, \varphi) \sim \sqrt{\frac{k+1}{2k}\varphi}$ *for fixed $k \geq 1$ and $\varphi \to \infty$,*
(ii) $s^*(k, \varphi) \sim (W(-\frac{\varphi-1}{e\varphi}) + 1)^{-1}$ *for fixed $\varphi > 1$ and $k \to \infty$.*

The algorithm in the theorem above is also given explicitly in Section 6.2. Surprisingly, although one might think that k-max-search and k-min-search should behave similarly with respect to competitive analysis, Theorem 6.2 states that this is in fact *not* the case. Indeed, according to Theorems 6.1 and 6.2, for large φ, the best algorithm for k-max-search achieves a competitive ratio of roughly $k \sqrt[k]{\varphi}$, while the best algorithm for k-min-search is at best $\sqrt{\varphi/2}$-competitive. Similarly, when k is large, the competitive ratio of a best algorithm for k-max-search behaves like $\ln \varphi$, in contrast to k-min-search, where a straightforward analysis (i.e. series expansion of the W function around its pole) shows that the best algorithm achieves a ratio of $\Theta(\sqrt{\varphi})$. Hence, algorithms for k-min-search perform in the worst-case rather poorly compared to algorithms for k-max-search.

Furthermore, we investigate the performance of *randomized* algorithms for the problems in question. El-Yaniv et al. (2001) gave a $\mathcal{O}(\ln \varphi)$-competitive randomized algorithm for 1-max-search, but did not provide a lower bound.

Theorem 6.3. *Let $k \in \mathbb{N}$, $\varphi > 1$. For every randomized k-max-search algorithm* RALG *we have*

$$\mathrm{CR}(\mathsf{RALG}) \geq (\ln \varphi)/2 \ . \tag{6.3}$$

Furthermore, there is a $2 \ln \varphi$-competitive randomized algorithm for $\varphi > 3$.

Note that the lower bound above is independent of k, i.e., randomized algorithms cannot improve their performance when k increases. In contrast to that, by considering Theorem 6.1, as k grows the performance of the best *deterministic* algorithm improves, and approaches $\ln \varphi$, which is only a multiplicative factor away from the best ratio that a randomized algorithm can achieve.

Our next result is about randomized algorithms for k-min-search.

Theorem 6.4. *Let $k \in \mathbb{N}, \varphi > 1$. For every randomized k-min-search algorithm* RALG *we have*

$$\mathrm{CR}(\mathsf{RALG}) \geq (1 + \sqrt{\varphi})/2 \ . \tag{6.4}$$

Again, the lower bound is independent of k. Furthermore, combined with Theorem 6.2, the theorem above states that for *all* $k \in \mathbb{N}$, randomization *does not* improve the performance (up to a multiplicative constant) of algorithms for k-min-search, compared to deterministic algorithms. This is again a difference between k-max-search and k-min-search.

6.1.4. Application to Robust Valuation of Lookback Options. In Section 6.4 we will use competitive k-search algorithms to derive upper bounds for the price of *lookback options*. Recall from Chapter 1 that a lookback *call* allows the holder to buy the underlying stock at time T from the option writer at the historical minimum price observed over $[0, T]$, and a lookback *put* to sell at the historical maximum.

A key argument in the famous Black-Scholes option pricing model (Black and Scholes, 1973) is a *no arbitrage condition*. Loosely speaking, an arbitrage is a zero-risk, zero net investment strategy that still generates profit. If such an opportunity came about, market participants would immediately start exploiting it, pushing prices until the arbitrage opportunity ceases to exist. Black and Scholes essentially give a dynamic trading strategy in the underlying stock by which an option writer can risklessly hedge an option position. Thus, the no arbitrage condition implies that the cost of the trading strategy must equal the price of the option to date.

In the model of Black and Scholes trading is possible continuously in time and in arbitrarily small portions of shares. Moreover, a central underlying assumption is that the stock price follows a geometric Brownian motion (Shreve, 2004, for instance), which then became the standard model for option pricing. While it certainly shows many features that resemble reality fairly, the behavior of stock prices in practice is not fully consistent with this assumption. For instance, the distribution observed for the returns of stock price processes are non-Gaussian and typically heavy-tailed (Cont, 2001), leading to underestimation of extreme price movements. Furthermore, in practice trading is discrete, price paths include price jumps and stock price volatility is not constant. As a response, numerous modifications of the original Black-Scholes setting have been proposed, examining different stochastic processes for the stock price (Merton, 1976; Cont and Tankov, 2004; Heston, 1993, for instance).

In light of the persistent difficulties of finding and formulating the "right" model for the stock price dynamic, there have also been a number of attempts to price financial instruments by *relaxing the Black-Scholes assumptions* instead. The idea is to provide *robust* bounds that work with (almost) any evolution of the stock price rather than focusing on a specific formulation of the stochastic process. In this fashion, DeMarzo et al. (2006) derive both upper and lower bounds for option prices in a model of bounded quadratic variation, using competitive online trading algorithms. In the mathematical finance community, Epstein and Wilmott (1998) propose non-probabilistic models for pricing interest rate securities in a framework of "worst-case scenarios". Korn (2005) combines the random walk assumption with a worst-case analysis to tackle optimal asset allocation under the threat of a crash.

In this spirit, using the deterministic k-search algorithms from Section 6.2 we derive in Section 6.4 upper bounds for the price of lookback calls and puts, under the assumption of *bounded stock price paths* and non-existence of arbitrage opportunities. Interestingly, the resulting bounds show similar qualitative properties and quantitative values as pricing in the standard Black-Scholes model. Note that the assumption of a bounded stock price is indeed very minimal, since without *any* assumption about the magnitude of the stock price fluctuation in fact *no* upper bounds for the option price apply.

6.2. Deterministic Search

Consider the following *reservation price policy* $\mathsf{RPP_{max}}$ for k-max-search. Prior to the start of the game, we choose *reservation prices* p_i^* $(i = 1 \dots k)$. As the prices are sequentially revealed, $\mathsf{RPP_{max}}$ accepts the first price that is at least p_1^* and sells one unit. It then waits for the first price that is at least p_2^*, and subsequently continues with all reservation prices. $\mathsf{RPP_{max}}$ works through the reservation prices in a strictly sequential manner. Note that $\mathsf{RPP_{max}}$ may be forced to sell at the last prices of the sequence, which may be lower than the remaining reservations, to meet the constraint of completing the sale.

The proof of the lemma below generalizes ideas used for *1-max-search* by El-Yaniv et al. (2001).

LEMMA 6.5. *Let* $k \in \mathbb{N}$, $\varphi > 1$. *Let* $r^* = r^*(k, \varphi)$ *be defined as in* (6.1). *Then the reservation price policy* $\mathsf{RPP_{max}}$ *with reservation prices given by*

$$p_i^* = m \left[1 + (r^* - 1) \left(1 + \frac{r^*}{k} \right)^{i-1} \right] \quad , \tag{6.5}$$

satisfies $k\, p_{max}(\sigma) \leq r^* \cdot \mathsf{RPP_{max}}(\sigma)$ *for all* $\sigma \in \mathcal{P}$. *In particular,* $\mathsf{RPP_{max}}$ *is a* r^**-competitive algorithm for the* k*-max-search problem.*

PROOF. For $0 \leq j \leq k$, let $\mathcal{P}_j \subseteq \mathcal{P}$ be the sets of price sequences for which $\mathsf{RPP_{max}}$ accepts *exactly* j prices, excluding the forced sale at the end. Then $\mathcal{P}$ is the disjoint union of the $\mathcal{P}_j$'s. To shorten notation, let us write $p_{k+1}^* = M$. Let $\varepsilon > 0$ be fixed and define the price sequences

$$\forall 0 \leq i \leq k \ : \quad \sigma_i \ = \ p_1^*, p_2^*, \dots, p_i^*, \underbrace{p_{i+1}^* - \varepsilon, \dots, p_{i+1}^* - \varepsilon}_{k}, \underbrace{m, m, \dots, m}_{k} \ .$$

Observe that as $\varepsilon \to 0$, each σ_j is a sequence yielding the worst-case ratio in $\mathcal{P}_j$, in the sense that for all $\sigma \in \mathcal{P}_j$

$$\frac{\mathsf{OPT}(\sigma)}{\mathsf{RPP_{max}}(\sigma)} \leq \frac{kp_{\max}(\sigma)}{\mathsf{RPP_{max}}(\sigma)} \ \leq \ \frac{kp_{j+1}^*}{\mathsf{RPP_{max}}(\sigma_j)} \ . \tag{6.6}$$

Thus, to prove the statement we show that for $0 \leq j \leq k$ it holds $kp_{j+1}^* \leq r^* \cdot \mathsf{RPP_{max}}(\sigma_j)$. A straightforward calculation shows that for all $0 \leq j \leq k$

$$\sum_{i=1}^{j} p_i^* = m \left[j + k(1 - 1/r^*) \left((1 + r^*/k)^j - 1 \right) \right] \ .$$

But then we have for $\varepsilon \to 0$, the competitive ratio is arbitrarily close to

$$\forall\, 0 \leq j \leq k : \quad \frac{kp_{j+1}^*}{\mathsf{RPP_{max}}(\sigma_j)} = \frac{kp_{j+1}^*}{\sum_{i=1}^{j} p_i^* + (k-j)m} = r^* \ .$$

Thus, from (6.6) the r^*-competitiveness of $\mathsf{RPP_{max}}$ follows immediately. □

The sequence (6.5) is illustrated in Figure 6.2.

REMARK 6.6. While the proof above shows that the reservation prices in (6.5) are in fact the optimal choice, let us also briefly give an intuition on how to construct them. First, note that we have to choose the p_i^*'s such that

$$\frac{kp_1^*}{km} \stackrel{!}{=} \frac{kp_2^*}{p_1^* + (k-1)m} \stackrel{!}{=} \cdots \stackrel{!}{=} \frac{kM}{\sum_{i=1}^{k} p_i^*} \stackrel{!}{=} r^*. \tag{6.7}$$

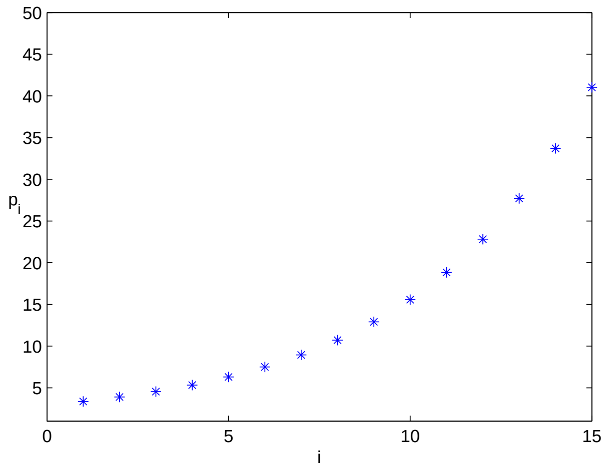

FIGURE 6.2. Sequence of reservation prices (6.5) for k-max-search. The parameter values are $\varphi = 50, m = 1, k = 15$.

(The nominator is the objective value of OPT on σ_i as $\varepsilon \to 0$, whereas the denominator is the value of $\mathsf{RPP}_{\mathsf{max}}$ on the same sequence.) For $0 \le i \le k$, let $r_i = p_i^*/p_1^*$. By comparing adjacent terms in (6.7), it is easy to see that r_i satisfies the simple recurrence

$$r_i = r_{i-1}\big(1 + p_1^*/(km)\big) - 1/k, \quad \text{and} \quad r_1 = 1 \enspace ,$$

and standard methods readily yield a closed formula for p_i^* in terms of p_1^*. Furthermore, using (6.7) we obtain the explicit expression for p_1^*.

From the choice of reservation prices in Lemma 6.5, we see that in fact no deterministic algorithm will be able to do better than $\mathsf{RPP}_{\mathsf{max}}$ in the worst-case.

LEMMA 6.7. *Let $k \ge 1$, $\varphi > 1$. Then $r^*(k, \varphi)$ given by (6.1) is the lowest possible competitive ratio that a deterministic k-max-search algorithm can achieve.*

PROOF. Let ALG be any deterministic algorithm. We shall show that ALG cannot achieve a ratio lower than $r^*(k, \varphi)$. Let $p_1^*, \dots, p_k^*$ be the price sequence defined by (6.5). We start by presenting p_1^* to ALG, at most k times or until ALG accepts it. If ALG never accepts p_1^*, we drop the price to m for the remainder, and ALG achieves a competitive ratio of $p_1^*/m = r^*(k, \varphi)$. If ALG accepts p_1^*, we continue the price sequence by presenting p_2^* to ALG at most k times. Again, if ALG never accepts p_2^* before we presented it k times, we drop to m for the remainder and ALG achieves a ratio not lower than $kp_2^*/(p_1^* + (k-1)m) = r^*(k, \varphi)$. We continue in that fashion by presenting each p_i^* at most k times (or until ALG accepts it). Whenever ALG doesn't accept a p_i^* after presenting it k times, we drop the price to m. If ALG subsequentially accepts all $p_1^*, \dots, p_k^*$, we conclude with k times M. In any case, ALG achieves only a ratio of at most $r^*(k, \varphi)$. □

With the above preparations we are ready to prove Theorem 6.1.

PROOF OF THEOREM 6.1. The first statement follows from Lemma 6.5 and Lemma 6.7. To show (i), first observe that for $k \ge 1$ fixed, $r^* = r^*(\varphi)$ must satisfy $r^* \to \infty$ as $\varphi \to \infty$, and r^*

is an increasing function of φ. Let $r_+ = k^{\frac{k}{k+1}} \sqrt[k+1]{\varphi}$. Then, for $\varphi \to \infty$, we have

$$(r_+ - 1)\left(1 + \frac{r_+}{k}\right)^k = (1 + o(1))\left(k^{\frac{k}{k+1}} \sqrt[k+1]{\varphi} \cdot \left(k^{-\frac{1}{k+1}} \sqrt[k+1]{\varphi}\right)^k\right) = (1 + o(1))\varphi \ .$$

Furthermore, let $\varepsilon > 0$ and set $r_- = (1 - \varepsilon)k^{\frac{k}{k+1}} \sqrt[k+1]{\varphi}$. A similar calculation as above shows that for sufficiently large φ we have

$$(r_- - 1)\left(1 + \frac{r_-}{k}\right)^k \geq (1 - 3k\varepsilon)\,\varphi \ .$$

Thus, $r = (1 + o(1))\, k^{\frac{k}{k+1}} \sqrt[k+1]{\varphi}$ indeed satisfies (6.1) for $\varphi \to \infty$. For the proof of (ii), note that for $k \to \infty$ and φ fixed, equation (6.1) becomes $(\varphi - 1)/(r^* - 1) = e^{r^*}$, and thus

$$(\varphi - 1)/e = (r^* - 1)\, e^{r^* - 1} \ .$$

The claim follows from the definition of the W-function. □

Similarly, we can construct a reservation price policy $\mathsf{RPP}_{\mathsf{min}}$ for k-min-search. Naturally, $\mathsf{RPP}_{\mathsf{min}}$ is modified such that it accepts the first price *lower* than the current reservation price.

Lemma 6.8. *Let $k \in \mathbb{N}$, $\varphi > 1$. Let $s^* = s^*(k, \varphi)$ be defined as in (6.2). Then the reservation price policy $\mathsf{RPP}_{\mathsf{min}}$ with reservation prices $p_1^* > \cdots > p_k^*$,*

$$p_i^* = M\left[1 - \left(1 - \frac{1}{s^*}\right)\left(1 + \frac{1}{ks^*}\right)^{i-1}\right] \ , \tag{6.8}$$

satisfies $\mathsf{RPP}_{\mathsf{min}}(\sigma) \leq s^*(k, \varphi) \cdot k\, p_{min}(\sigma)$, and is a $s^*(k, \varphi)$-competitive deterministic algorithm for k-min-search.

Proof. The proof is analogous to the proof of Lemma 6.5. Again, for $0 \leq j \leq k$, let $\mathcal{P}_j \subseteq \mathcal{P}$ be the sets of price sequences for which $\mathsf{RPP}_{\mathsf{min}}$ accepts *exactly* j prices, excluding the forced sale at the end. To shorten notation, define $p_{k+1}^* = m$. Let $\varepsilon > 0$ be fixed and define the price sequences

$$\sigma_i = p_1^*, p_2^*, \ldots, p_i^*, \underbrace{p_{i+1}^* + \varepsilon, \ldots, p_{i+1}^* + \varepsilon}_{k}, \underbrace{M, \ldots, M}_{k} \ , \quad \text{for } 0 \leq i \leq k \ .$$

As $\varepsilon \to 0$, each σ_j is a sequence yielding the worst-case ratio in $\mathcal{P}_j$, in the sense that for all $\sigma \in \mathcal{P}_j$, $0 \leq j \leq k$,

$$\frac{\mathsf{RPP}_{\mathsf{min}}(\sigma)}{\mathsf{OPT}(\sigma)} \leq \frac{\mathsf{RPP}_{\mathsf{min}}(\sigma)}{kp_{\min}(\sigma)} \leq \frac{\mathsf{RPP}_{\mathsf{min}}(\sigma_j)}{kp_{j+1}^*} \ . \tag{6.9}$$

Straightforward calculation shows that for $\varepsilon \to 0$

$$\forall\, 0 \leq j \leq k: \quad \frac{\mathsf{RPP}_{\mathsf{min}}(\sigma_j)}{kp_{j+1}^*} = \frac{\sum_{i=1}^{j} p_i^* + (k - j)M}{kp_{j+1}^*} = s^* \ ,$$

and hence

$$\forall\, \sigma \in \mathcal{P}: \quad \frac{\mathsf{RPP}_{\mathsf{min}}(\sigma)}{kp_{\min}(\sigma)} \leq s^* \ .$$

Since $\mathsf{OPT}(\sigma) \geq kp_{\min}(\sigma)$ for all $\sigma \in \mathcal{P}$, this also implies that $\mathsf{RPP}_{\mathsf{min}}$ is s^*-competitive. □

The sequence (6.8) is illustrated in Figure 6.3.

Again, no deterministic algorithm can do better than $\mathsf{RPP}_{\mathsf{min}}$ in Lemma 6.8.

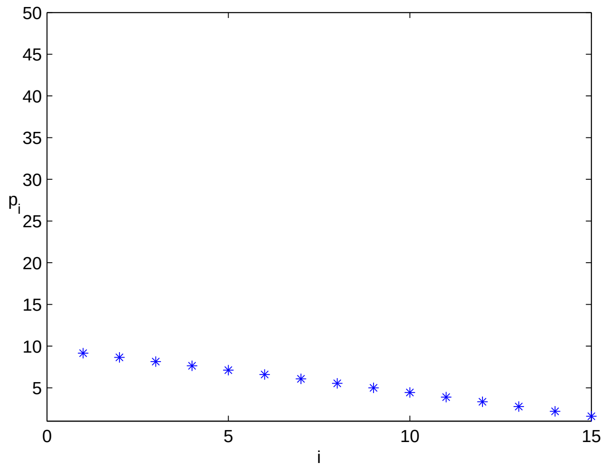

FIGURE 6.3. Sequence of reservation prices (6.8) for k-min-search. The parameter values are $\varphi = 50, m = 1, k = 15$.

LEMMA 6.9. *Let* $k \geq 1$*,* $\varphi > 1$*. Then* $s^*(k, \varphi)$ *given by* (6.2) *is the lowest possible competitive ratio that a deterministic* k*-min-search algorithm can achieve.*

The proof of Lemma 6.9 is identical to the proof of Lemma 6.7, except that the adversary now uses the prices defined by (6.8) and the roles of m and M are interchanged.

Using Lemma 6.8 and Lemma 6.9 we can now prove Theorem 6.2.

PROOF OF THEOREM 6.2. The first part of the Theorem follows directly from Lemma 6.8 and Lemma 6.9. To show (i), first observe that for $k \geq 1$ fixed, $s^* = s^*(\varphi)$ must satisfy $s^* \to \infty$ as $\varphi \to \infty$, and s^* is an increasing function of φ. With this assumption we can expand the right-hand side of (6.2) with the binomial theorem to obtain

$$\frac{1 - 1/\varphi}{1 - 1/s^*} = 1 + \frac{1}{s^*} + \frac{k-1}{2k\,(s^*)^2} + \Theta\left((s^*)^{-3}\right) \implies \frac{1}{\varphi} = \frac{k+1}{2k\,(s^*)^2} + \Theta\left((s^*)^{-3}\right) \ .$$

By solving this equation for s^* we obtain (i). For the proof of (ii), first observe that for $\varphi \geq 1$ fixed, $s^* = s^*(k)$ must satisfy $s^*(k) \leq C$, for some constant C which may depend on φ. Indeed, if $s^*(k) \to \infty$ with $k \to \infty$, then by taking limits on both sides of (6.2) yields

$$1 - \frac{1}{\varphi} = \lim_{k\to\infty} \left(1 + \frac{1}{ks^*(k)}\right)^k = 1 \ ,$$

which is a contradiction. Thus, $s^* = \Theta(1)$ and we obtain from (6.2) by taking limits

$$\frac{1 - 1/\varphi}{1 - 1/s^*} = \lim_{k\to\infty} \left(1 + \frac{1}{ks^*}\right)^k = e^{1/s^*} \ ,$$

and (ii) follows immediately by the definition of the W-function. □

6.3. Randomized Search

6.3.1. Lower Bound for Randomized k-max-search. We consider $k = 1$ first. The optimal deterministic online algorithm achieves a competitive ratio of $r^*(1, \varphi) = \sqrt{\varphi}$. As shown by El-Yaniv et al. (2001), randomization can dramatically improve this. Assume for simplicity that

$\varphi = 2^\ell$ for some integer ℓ. For $0 \le j < \ell$ let $\mathsf{RPP}_{\mathsf{max}}(j)$ be the reservation price policy with reservation $m2^j$, and define EXPO to be a uniform probability mixture over $\{\mathsf{RPP}_{\mathsf{max}}(j)\}_{j=0}^{\ell-1}$.

LEMMA 6.10 (Levin, see El-Yaniv et al. (2001)). *Algorithm* EXPO *is* $\mathcal{O}(\ln \varphi)$*-competitive.*

We shall prove that EXPO is in fact the optimal randomized online algorithm for 1-max-search. We will use the following version of *Yao's principle* (Yao, 1977).

THEOREM 6.11 (Yao's principle). *For an online maximization problem denote by* $\mathcal{S}$ *the set of possible input sequences, and by* $\mathcal{A}$ *the set of deterministic algorithms, and assume that* $\mathcal{S}$ *and* $\mathcal{A}$ *are finite. Fix any probability distribution* $y(\sigma)$ *on* $\mathcal{S}$*, and let* S *be a random sequence according to this distribution. Let* RALG *be any mixed strategy, given by a probability distribution on* $\mathcal{A}$*. Then,*

$$\mathrm{CR}(\mathsf{RALG}) = \max_{\sigma \in \mathcal{S}} \frac{\mathsf{OPT}(\sigma)}{\mathbb{E}\left[\mathsf{RALG}(\sigma)\right]} \ge \left(\max_{\mathsf{ALG} \in \mathcal{A}} \mathbb{E}\left[\frac{\mathsf{ALG}(\mathrm{S})}{\mathsf{OPT}(\mathrm{S})} \right] \right)^{-1} . \tag{6.10}$$

Note that the first expectation is taken with respect to the randomization of the algorithm RALG, whereas the second expectation is taken with respect to the input distribution $y(\sigma)$.

The reader is referred to standard textbooks for a proof (e.g. Chapter 6 and 8 of Borodin and El-Yaniv (1998)). In words, Yao's principle states that we obtain a lower bound on the competitive ratio of the best randomized algorithm by calculating the performance of the best deterministic algorithm for a chosen probability distribution of input sequences. Note that (6.10) gives a lower bound for *arbitrary* chosen input distributions. However, only for well-chosen y's we will obtain strong lower bounds.

We first need to establish the following lemma on the representation of an arbitrary randomized algorithm for k-search.

LEMMA 6.12. *Let* RALG *be a randomized algorithm for the* k*-max-search problem. Then* RALG *can be represented by a probability distribution on the set of all deterministic algorithms for the* k*-max-search problem.*

PROOF. The proof of the statement is along the lines of the proof of Theorem 1 in El-Yaniv et al. (2001). Here we only sketch the proof idea. Using game-theoretic terminology, RALG may be either a mixed strategy (a distribution on deterministic algorithms, from which one is randomly chosen prior to the start of the game) or a behavioral strategy (where an independent random choice may be made at each point during the game). As we have perfect recall in k-search (player has no memory restrictions), k-search is a linear game. For linear games, every behavioral strategy has an equivalent mixed algorithm. Thus, we can always assume that RALG is a mixed strategy given by a probability distribution on the set of all deterministic algorithms. □

The next lemma yields the desired lower bound.

LEMMA 6.13. *Let* $\varphi > 1$*. Every randomized* 1*-max-search algorithm* RALG *satisfies*

$$\mathrm{CR}(\mathsf{RALG}) \ge (\ln \varphi)/2 \ .$$

Proof. Let $b > 1$ and $\ell = \log_b \varphi$. We define a finite approximation of $\mathcal{I}$ by $\mathcal{I}_b = \{mb^i \mid i = 0 \dots \ell\}$, and let $\mathcal{P}_b = \bigcup_{n \geq k} \mathcal{I}_b^n$. We consider the 1-max-search problem on $\mathcal{P}_b$. As $\mathcal{P}_b$ is finite, also the set of deterministic algorithms $\mathcal{A}_b$ is finite. For $0 \leq i \leq \ell - 1$, define sequences of length ℓ by

$$\sigma_i = mb^0, \dots, mb^i, m, \dots, m \ . \tag{6.11}$$

Let $\mathcal{S}_b = \{\sigma_i \mid 0 \leq i \leq \ell - 1\}$ and define the probability distribution y on $\mathcal{P}_b$ by

$$y(\sigma) = \begin{cases} 1/\ell & \text{for } \sigma \in \mathcal{S}_b \ , \\ 0 & \text{otherwise} \ . \end{cases}$$

Let $\mathsf{ALG} \in \mathcal{A}_b$. Note that for all $1 \leq i \leq \ell$, the first i prices of the sequences σ_j with $j \geq i - 1$ coincide, and ALG cannot distinguish them up to time i. As ALG is deterministic, it follows that if ALG accepts the i-th price in $\sigma_{\ell-1}$, it will accept the i-th price in all σ_j with $j \geq i - 1$. Thus, for every ALG, let $0 \leq \chi(\mathsf{ALG}) \leq \ell - 1$ be such that ALG accepts the $(\chi(\mathsf{ALG}) + 1)$-th price, i.e. $mb^{\chi(\mathsf{ALG})}$, in $\sigma_{\ell-1}$. ALG will then earn $m\, b^{\chi(\mathsf{ALG})}$ on all σ_j with $j \geq \chi(\mathsf{ALG})$, and m on all σ_j with $j < \chi(\mathsf{ALG})$. To shorten notation, we write χ instead of $\chi(\mathsf{ALG})$ in the following. Thus, we have

$$\mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] = \frac{1}{\ell}\left[\sum_{j=0}^{\chi-1} \frac{m}{mb^j} + \sum_{j=\chi}^{\ell-1} \frac{mb^\chi}{mb^j}\right] = \frac{1}{\ell}\left[\frac{1 - b^{-\chi}}{1 - b^{-1}} + \frac{1 - b^{-(\ell-\chi)}}{1 - b^{-1}}\right] ,$$

where the expectation $\mathbb{E}[\cdot]$ is with respect to the probability distribution $y(\sigma)$. If we consider the above term as a function of χ, then it is easily verified that it attains its maximum at $\chi = \ell/2$. Thus,

$$\max_{\mathsf{ALG} \in \mathcal{A}_b} \mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] \leq \frac{1}{\ell}\left(1 - \frac{1}{\sqrt{\varphi}}\right)\frac{2b}{b-1} \leq \frac{1}{\ln \varphi} \cdot \frac{2b \ln b}{b - 1} \ . \tag{6.12}$$

Let Υ_b be the set of all randomized algorithms for 1-max-search with possible price sequences $\mathcal{P}_b$. By Lemma 6.12, each $\mathsf{RALG}_b \in \Upsilon_b$ may be given as a probability distribution on $\mathcal{A}_b$. Since $\mathcal{A}_b$ and $\mathcal{S}_b$ are both finite, we can apply Theorem 6.11. Thus, for all $b > 1$ and all $\mathsf{RALG}_b \in \Upsilon_b$, we have

$$\mathrm{CR}(\mathsf{RALG}_b) \geq \left(\max_{\mathsf{ALG} \in \mathcal{A}_b} \mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right]\right)^{-1} \geq \ln \varphi \, \frac{b - 1}{2b \ln b} \ .$$

Let Υ be the set of all randomized algorithms for 1-max-search on $\mathcal{P}$. Since for $b \to 1$, we have $\mathcal{A}_b \to \mathcal{A}$, $\Upsilon_b \to \Upsilon$ and $(b - 1)/(2b \ln b) \to \frac{1}{2}$, the proof is completed. □

In fact, Lemma 6.13 can be generalized to arbitrary $k \geq 1$.

Lemma 6.14. *Let $k \in \mathbb{N}$, $b > 1$ and $\varphi > 1$. Let* RALG *be any randomized algorithm for k-max-search. Then, we have*

$$\mathrm{CR}(\mathsf{RALG}) \geq (\ln \varphi)/2 \ .$$

Proof. Let $1 < b < \varphi$ and $\ell = \log_b \varphi$. We define $\mathcal{P}_b$ and $\mathcal{A}_b$ as in the proof of Lemma 6.13. For $0 \leq i \leq \ell - 1$, define

$$\sigma_i = \underbrace{mb^0, \dots, mb^0}_{k}, \dots, \underbrace{mb^i, \dots, mb^i}_{k}, \underbrace{m, \dots, m}_{k(\ell-i-1)} \ . \tag{6.13}$$

Let $\mathcal{S}_b = \bigcup_{0 \le i \le \ell-1} \sigma_i$ and define the probability distribution y on $\mathcal{P}_b$ by

$$y(\sigma) = \begin{cases} 1/\ell & \text{for } \sigma \in \mathcal{S}_b \ , \\ 0 & \text{otherwise} \ . \end{cases}$$

Similarly as in the proof of Lemma 6.13, we characterize every algorithm $\mathsf{ALG} \in \mathcal{A}_b$ by a vector $(\chi_i)_{1 \le i \le k}$ where mb^{χ_i} is the price for which ALG sells the i-th unit on $\sigma_{\ell-1}$. By construction, we have $\chi_1 \le \cdots \le \chi_k$. (Recall that $\sigma_{\ell-1}$ is the sequence that is increasing until the very end.) Note that for all $1 \le i \le \ell$, the sequences $\{\sigma_j \mid j \ge i-1\}$ are not distinguishable up to time ik, since the first ik prices of those sequences are identical. Let $0 \le j \le \ell - 1$ and $t = \max\{i \,|\, \chi_i \le j\}$. When presented σ_j, ALG accepts all prices mb^{χ_i} for which $\chi_i \le j$. Hence, we have $\mathsf{OPT}(\sigma_j) = kmb^j$ and $\mathsf{ALG}(\sigma_j) = (k-t)m + \sum_{s=1}^{t} mb^{\chi_s}$, i.e. ALG can successfully sell for its first t reservation prices. To abbreviate notation, let $\chi_0 = 0$ and $\chi_{k+1} = \ell$, and define $\delta_t = \chi_{t+1} - \chi_t$. Taking expectation with respect to $y(\sigma)$, we have

$$\begin{aligned}
\mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] &= \frac{1}{\ell}\sum_{t=0}^{k}\sum_{j=\chi_t}^{\chi_{t+1}-1} \frac{(k-t)m + \sum_{s=1}^{t} mb^{\chi_s}}{kmb^j} \\
&= \frac{1}{\ell}\sum_{t=0}^{k} \frac{\left(k-t+\sum_{s=1}^{t} b^{\chi_s}\right)\sum_{j=0}^{\delta_t-1} b^{-j}}{kb^{\chi_t}} \\
&= \frac{1}{\ell}\sum_{t=0}^{k} \frac{\left(k-t+\sum_{s=1}^{t} b^{\chi_s}\right)(1-b^{-\delta_t})}{kb^{\chi_t}(1-b^{-1})} \ .
\end{aligned}$$

Straightforward yet tedious algebra simplifies this expression to

$$\mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] = \frac{\sum_{t=1}^{k} 1 - b^{-\chi_t} + \sum_{t=1}^{k} 1 - b^{-(\ell-\chi_t)}}{\ell k(1-b^{-1})} \ ,$$

and the maximum over $\{\chi_1, \ldots, \chi_k\}$ is attained at $\chi_1 = \cdots = \chi_k = \ell/2$. Thus, defining $\chi = \ell/2$ we have

$$\max_{\mathsf{ALG} \in \mathcal{A}_b} \mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] \le \frac{1}{\ell}\left[\frac{1-b^{-\chi}}{1-b^{-1}} + \frac{1-b^{-(\ell-\chi)}}{1-b^{-1}}\right] = \frac{2b}{\ell\,(b-1)}\left(1 - \frac{1}{\sqrt{\varphi}}\right) \ ,$$

which is exactly (6.12) in the proof of Lemma 6.13. Thus, we can argue as in the remainder of the proof of Lemma 6.13, and let again $b \to 1$ to conclude that $\mathrm{CR}(\mathsf{RALG}) \ge (\ln \varphi)/2$ for all randomized algorithms RALG for k-max-search. □

Giving an optimal randomized algorithm for k-max-search is straightforward. For $1 < b < \varphi$ and $\ell = \log_b \varphi$, EXPO_k chooses j uniformly at random from $\{0, \ldots, \ell-1\}$, and sets all its k reservation prices to mb^j.

LEMMA 6.15. *Let $k \in \mathbb{N}$. EXPO_k is an asymptotically optimal randomized algorithm for the k-max-search problem with $\mathrm{CR}(\mathsf{EXPO}_k) = \ln\varphi \cdot \frac{(b-1)}{\ln b}\frac{\varphi}{\varphi-1}$.*

PROOF. We want to determine

$$\mathrm{CR}(\mathsf{EXPO}_k) = \max_{\sigma \in \mathcal{P}} \mathsf{R}(\sigma), \quad \text{where} \quad \mathsf{R}(\sigma) = \frac{\mathsf{OPT}(\sigma)}{\mathbb{E}\left[\mathsf{EXPO}_k(\sigma)\right]} \ . \tag{6.14}$$

Obviously, a sequence σ that maximizes R is non-decreasing, since rearranging σ does not change the objective value of OPT, but may increase the objective value of EXPO_k. Let

$$\hat{\sigma} = \underbrace{m, \ldots, m}_{k}, \underbrace{mb^1, \ldots, mb^1}_{k}, \underbrace{mb^2, \ldots, mb^2}_{k}, \ldots, \underbrace{mb^{\ell-1}, \ldots, mb^{\ell-1}}_{k}, \underbrace{M, \ldots, M}_{k} \ .$$

In the following, we will prove that $\hat{\sigma}$ is a worst-case sequence for (6.14). We will first show that for every non-decreasing sequence $\sigma = p_1, p_2, \ldots, p_N$ it holds

$$\mathsf{R}(\sigma) \le \mathsf{R}(\sigma \circ \hat{\sigma}) \ , \tag{6.15}$$

where $\sigma \circ \hat{\sigma}$ is the concatenation of σ and $\hat{\sigma}$. Let P be the reservation price of EXPO_k, and let $\mathsf{EXPO}_k^j(\sigma) = \mathbb{E}\left[\mathsf{EXPO}_k(\sigma) \mid P = mb^j\right]$. To see the first inequality we shall show that for all $0 \le j < \ell$,

$$\frac{\mathsf{EXPO}_k^j(\sigma \circ \hat{\sigma})}{\mathsf{OPT}(\sigma \circ \hat{\sigma})} \le \frac{\mathsf{EXPO}_k^j(\sigma)}{\mathsf{OPT}(\sigma)} \ , \tag{6.16}$$

which yields

$$\frac{1}{\mathsf{R}(\sigma)} = \mathbb{E}\left[\frac{\mathsf{EXPO}_k(\sigma)}{\mathsf{OPT}(\sigma)}\right] \ge \mathbb{E}\left[\frac{\mathsf{EXPO}_k(\sigma \circ \hat{\sigma})}{\mathsf{OPT}(\sigma \circ \hat{\sigma})}\right] = \frac{1}{\mathsf{R}(\sigma \circ \hat{\sigma})} \ ,$$

i.e., $\mathsf{R}(\sigma) \le \mathsf{R}(\sigma \circ \hat{\sigma})$. To see (6.16), note that if p_l is the first price accepted by EXPO_k^j in σ, then EXPO_k^j will also accept $p_{l+1}, \ldots, p_{l+k-1}$. This follows from the property of σ being non-decreasing and from the fact that all reservation prices of EXPO_k^j are identical. Now we distinguish two cases: either $l = N - k + 1$ (i.e. EXPO_k^j accepts the last k prices in σ, possibly forced by the constraint to finish the sale by the end of the sequence σ) or $l < N - k + 1$ (i.e. EXPO_k^j can successfully sell all k units for prices of at least its reservation price mb^j). In the first case, $\mathsf{EXPO}_k^j(\sigma) = \mathsf{OPT}(\sigma)$ and (6.16) follows trivially, since we always have $\mathsf{EXPO}_k^j(\sigma \circ \hat{\sigma}) \le \mathsf{OPT}(\sigma \circ \hat{\sigma})$. In the second case, $\mathsf{OPT}(\sigma \circ \hat{\sigma}) = kM \ge \mathsf{OPT}(\sigma)$ and $\mathsf{EXPO}_k^j(\sigma \circ \hat{\sigma}) = \mathsf{EXPO}_k^j(\sigma)$, since EXPO_k^j accepted k prices before the end of σ was reached, and it cannot accept any prices in the second part of $\sigma \circ \hat{\sigma}$. Hence, (6.16) also holds in this case. This shows (6.15).

Now observe that for any non-decreasing σ we have

$$\mathbb{E}\left[\mathsf{EXPO}_k(\sigma \circ \hat{\sigma})\right] \ge \mathbb{E}\left[\mathsf{EXPO}_k(\hat{\sigma})\right] \ ,$$

since for every j algorithm EXPO_k^j accepts k prices in $\sigma \circ \hat{\sigma}$ that are at least mb^j, but in $\hat{\sigma}$ it accepts k times exactly its reservation price mb^j. Combined with the fact that $\mathsf{OPT}(\sigma \circ \hat{\sigma}) = \mathsf{OPT}(\hat{\sigma}) = kM$, this yields

$$\mathsf{R}(\sigma \circ \hat{\sigma}) \le \mathsf{R}(\hat{\sigma}) \ .$$

With (6.15), this implies that $\hat{\sigma}$ is a worst-case sequence for (6.14). Therefore, we have

$$\mathrm{CR}(\mathsf{EXPO}_k) = \mathsf{R}(\hat{\sigma}) = \frac{kM}{\frac{1}{\ell}\sum_{j=0}^{\ell-1} kmb^j} = \ell\frac{\varphi\,(b-1)}{\varphi-1} = \ln\varphi\,\frac{\varphi}{\varphi-1}\cdot\frac{(b-1)}{\ln b} \ ,$$

since $M = \varphi m$ and $b^\ell = \varphi$. □

For $\varphi > 3$ and $b < 3/2$, we have $\frac{\varphi}{\varphi-1}\frac{(b-1)}{\ln b} < 2$, and hence combining Lemma 6.14 and 6.15 we immediately obtain Theorem 6.3.

6.3.2. Randomized k-Min-Search. The proof of the lower bound for k-min-search, Theorem 6.4, uses an analogous version of Yao's principle (see for instance Theorem 8.5 in Borodin and El-Yaniv (1998)).

THEOREM 6.16 (Yao's principle for cost minimization problems). *For an online cost minimization problem Π, let the set of possible input sequences $\mathcal{S}$ and the set of deterministic algorithms $\mathcal{A}$ be finite, and given by $\mathcal{S} = \{\sigma_1, \dots, \sigma_n\}$ and $\mathcal{A} = \{\mathsf{ALG}_1, \dots, \mathsf{ALG}_m\}$. Fix any probability distribution $y(\sigma)$ on $\mathcal{S}$. Let* RALG *be any mixed strategy, given by a probability distribution $x(a)$ on $\mathcal{A}$. Then,*

$$\mathrm{CR}(\mathsf{RALG}) = \max_{\sigma \in \mathcal{S}} \frac{\mathbb{E}\left[\mathsf{RALG}(\sigma)\right]}{\mathsf{OPT}(\sigma)} \geq \min_{\mathsf{ALG} \in \mathcal{A}} \mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] .$$

We are now ready to prove Theorem 6.4.

PROOF OF THEOREM 6.4. We shall consider first the case $k = 1$. Let $\mathcal{S} = \{\sigma_1, \sigma_2\}$ with

$$\sigma_1 = m\sqrt{\varphi}, M, \dots, M \quad \text{and} \quad \sigma_2 = m\sqrt{\varphi}, m, M, \dots, M \ ,$$

and let $y(\sigma)$ be the uniform distribution on $\mathcal{S}$. For $i \in \{1, 2\}$, let ALG_i be the reservation price policy with reservation prices $p_1^* = m\sqrt{\varphi}$ and $p_2^* = m$, respectively. Obviously, the best deterministic algorithm against the randomized input given by the distribution $y(\sigma)$ behaves either like ALG_1 or ALG_2. Since

$$\mathbb{E}\left[\frac{\mathsf{ALG}_i}{\mathsf{OPT}}\right] = (1 + \sqrt{\varphi})/2, \quad i \in \{1, 2\} \ ,$$

the desired lower bound follows from the min-cost version of Yao's principle. For general $k \geq 1$, we repeat the prices $m\sqrt{\varphi}$ and m in σ_1 and σ_2 k times each. Observe that in that case we can partition the set of all deterministic algorithms into $k+1$ equivalence classes, according to the number price quotations accepted from the first k prices $m\sqrt{\varphi}, \dots, m\sqrt{\varphi}$, as σ_1 and σ_2 are not distinguishable until the $(k+1)$th price. Suppose ALG accepts j times the price $m\sqrt{\varphi}$. Then we have

$$\mathbb{E}\left[\frac{\mathsf{ALG}}{\mathsf{OPT}}\right] = \frac{1}{2}\left(\frac{jm\sqrt{\varphi} + (k-j)M}{km\sqrt{\varphi}} + \frac{jm\sqrt{\varphi} + (k-j)m}{km}\right) = (1 + \sqrt{\varphi})/2$$

for all $0 \leq j \leq k$, and the lower bound follows again from Yao's principle. □

6.4. Robust Valuation of Lookback Options

In this section, we use the deterministic k-search algorithms from Section 6.2 to derive upper bounds for the price of lookback options under the assumption of *bounded stock price paths* and non-existence of arbitrage opportunities. We consider a discrete-time model of trading. For simplicity we assume that the interest rate is zero. The price of the stock at time $t \in \{0, 1, \dots, T\}$ is given by S_t, with S_0 being the price when seller and buyer enter the option contract.

Recall that the holder of a lookback call has the right to buy shares from the option writer for the price $S_{min} = \min\{S_t \,|\, 0 \leq t \leq T\}$. We shall assume that the lookback call is on $k \geq 1$ units of the underlying stock.

Note that since $S_{min} \leq S_T$, the option holder is never worse off executing the option at the expiry date T (and then immediately selling the shares for S_T) rather than to forgo his option. Hence, a lookback call option can always be considered as executed at expiry. This is in contrast to a European call option, where the option holder is not interested in executing his option if the price S_T at expiry is below the pre-specified strike price K.

Neglecting stock price appreciation, upwards and downwards movement of the stock price is equally likely. Consequently, we will assume a symmetric trading range $[\varphi^{-1/2}S_0, \varphi^{1/2}\, S_0]$ with $\varphi > 1$. We shall refer to a stock price path that satisfies $S_t \in [\varphi^{-1/2}S_0, \varphi^{1/2}\, S_0]$ for all $1 \leq t \leq T$ as a (S_0, φ) price path.

THEOREM 6.17. *Assume $(S_t)_{0\leq t\leq T}$ is a (S_0, φ) stock price path. Let $s^*(k, \varphi)$ be given by* (6.2)*, and let*

$$V^*_{Call}(S_0, \varphi) = S_0(s^*(k, \varphi) - 1)/\sqrt{\varphi}\ . \tag{6.17}$$

Let V be the option premium paid at time $t = 0$ for a lookback call option on k shares expiring at time T. Suppose we have $V > V^*_{Call}(k, S_0, \varphi)$. Then there exists an arbitrage opportunity for the option writer, i.e., there is a zero-net-investment strategy which yields a profit for all (S_0, φ) stock price paths.

PROOF. In the following, let C_t denote the money in the option writer's cash account at time t. At time $t = 0$, the option writer receives V from the option buyer, and we have $C_0 = V$. The option writer then successively buys k shares, one-by-one, applying $\mathsf{RPP}_{\mathsf{min}}$ for k-min-search with reservation prices as given by (6.8). Let H be the total sum of money spent for purchasing k units of stock. By Lemma 6.8 we have $H \leq ks^*(k, \varphi)S_{min}$. At time T the option holder purchases k shares from the option writer for kS_{min} in cash. As noted above, a lookback call option can always be considered executed at the expiration time T; if the option holder does not execute his option, the option writer simply sells the k shares again on the market for $kS_T \geq kS_{min}$.

After everything has been settled, we have

$$C_T = V - H + kS_{min} \geq V + kS_{min}(1 - s^*(k, \varphi))\ .$$

Because of $S_{min} \geq S_0/\sqrt{\varphi}$ and $V > V^*_{\mathrm{Call}}(S_0, \varphi)$, we conclude that $C_T > 0$ for all possible (S_0, φ) stock price paths. Hence, this is indeed a zero net investment profit for the option writer on all (S_0, φ) stock price paths. □

Under the no-arbitrage assumption, we immediately obtain an upper bound for the value of a lookback call option.

COROLLARY 6.18. *Under the no-arbitrage assumption, $V \leq V^*_{Call}(k, S_0, \varphi)$, with V^*_{Call} defined in* (6.17).

Using Lemma 6.5 and similar no-arbitrage arguments, also an upper bound for the price of a lookback put option can be established.

Note that it is not possible to derive non-trivial *lower bounds* for lookback options in the bounded stock price model, as a (S_0, φ)-price path may stay at S_0 throughout, making the option mature worthless for the holder. To derive lower bounds, there must be a promised fluctuation of the stock. In the classical Black-Scholes pricing model, this is given by the volatility of the Brownian motion.

We shall remark that in practice there is of course no trading range in which the stock price stays with certainty; what we rather can give are trading ranges in which the price stays with sufficiently high probability. V^*_{Call} is then to be understood as a bound for the option price up to a certain residual risk. Note that whereas the Black-Scholes-type price (6.18), which shall be in the next paragraph, has no such residual risk within the Black-Scholes model, it does certainly have significant *model risk* due to the fact that the Black-Scholes assumptions might turn out to be incorrect in the first place.

Comparison to Pricing in Black-Scholes Model. In the classical Black-Scholes setting, Goldman, Sosin, and Gatto (1979) give closed form solutions for the value of lookback puts and calls. Let σ be the volatility of the stock price, modeled by a geometric Brownian motion, $S(t) = S_0 \exp(-\sigma^2 t/2 + \sigma B(t))$, where $B(t)$ is a standard Brownian motion. Let $\Phi(x)$ denote the cumulative distribution function of the standard normal distribution. Then, for zero interest rate, at time $t = 0$ the value of a lookback call on one share of stock, expiring at time T, is given by

$$V^{\text{BS}}_{\text{Call}}(S_0, T, \sigma) = S_0\big(2\,\Phi(\sigma\sqrt{T}/2) - 1\big) \ . \tag{6.18}$$

The hedging strategy is a certain roll-over replication strategy of a series of European call options. Everytime the stock price hits a new all-time low, the hedger has to "roll-over" his position in the call to one with a new strike. Interestingly, this kind of behavior to act only when a new all-time low is reached resembles the behavior of $\mathsf{RPP}_{\mathsf{min}}$ for k-min-search.

For a numerical comparison of the bound $V^*_{\text{Call}}(k, S_0, \varphi, T)$ with the Black-Scholes-type pricing formula (6.18), we choose the fluctuation rate $\varphi = \varphi(T)$ such that the expected trading range $\big[\mathbb{E}(\min_{0\le t\le T} S_t),\, \mathbb{E}(\max_{0\le t\le T} S_t)\big]$ of a geometric Brownian motion starting at S_0 with volatility σ is the interval $\big[\varphi^{-1/2}S_0, \varphi^{1/2}S_0\big]$.

Figure 6.4 shows the results for $\sigma = 0.2$, $S_0 = 20$ and $k = 10$. As can be seen from the graph, the bound V^*_{Call} is qualitatively and quantitatively similar to the Black-Scholes-type pricing (6.18). However, it is important to note that the two pricing formulas rely on two very different models. In the Black-Scholes model, (6.18) *is* the correct price for a lookback option. On the other hand, the key advantage of the price bound (6.17) are its weak modeling assumptions on the price dynamics, and hence the price bound holds even in situations where the Black-Scholes assumptions might break down. Certainly, both concepts have strengths and weaknesses, and a good analysis consults both.

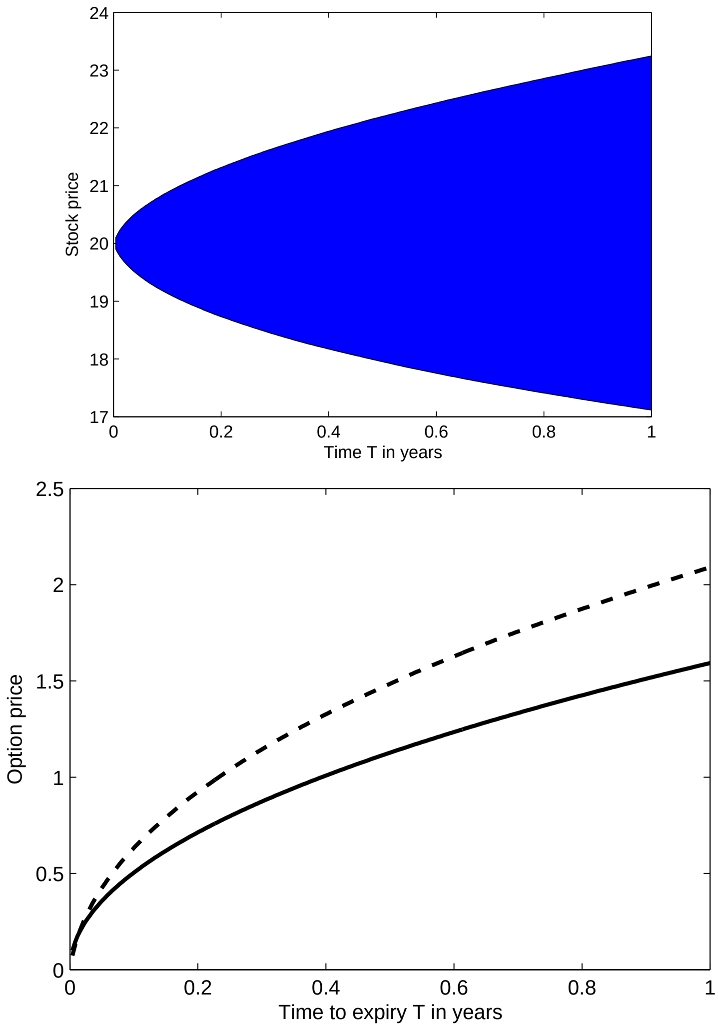

FIGURE 6.4. The left plot shows the expected trading range of a geometric Brownian motion with volatility $\sigma = 0.2$ and $S(0) = 20$. The right plot shows the price of a lookback call with maturity T in the Black-Scholes model (solid line) and the bound V^*_{Call} (dashed line), with $k = 100$ and $\varphi(T)$ chosen to match the expected trading range.

APPENDIX A

Trading with Market Impact: Optimization of Expected Utility

In this section, we shall briefly outline expected utility optimization for the continuous-time trading model with market impact introduced in Chapter 2, Section 2.2. Instead of the mean-variance framework as in Chapters 2 and 3, a risk-reward balance is sought by optimizing the expected *utility* of the final total amount captured. Different types of utility functions $u(\cdot)$ can be used, and varying levels of risk aversion can be captured by parameters in $u(\cdot)$.

As discussed in Chapters 1–3, mean-variance optimization retains practical advantages compared to the optimization of expected utility. Mean and variance are very easily understood, and correspond to how trade results are reported in practice. The broker/dealer can provide a family of efficient strategies without assumptions on the clients' utility and his other investment activities. Furthermore, another problem with utility functions is that the risk-aversion depends on the initial wealth, and in an institutional context it is very difficult to assign a value to this wealth: is it the size of the individual order, the size of the day's orders, or the client's total wealth under management?

For expected utility optimization, it is convenient to formulate the problem as a set of stochastic differential equations, which incorporate a control function that we seek to optimize. As in Chapter 2, suppose we are doing a sell program to complete at time T. Then the control, the state variables, and the dynamic equations are

$$\begin{aligned} v(t) &= \text{rate of selling} & & \\ x(t) &= \text{shares remaining to sell} & dx &= -v\,dt \\ y(t) &= \text{dollars in bank account} & dy &= (s - \eta v)\,v\,dt \\ s(t) &= \text{stock price} & ds &= \sigma\,dB \end{aligned}$$

where dB is the increment of a Brownian motion, η is the coefficient of linear market impact, and σ is absolute volatility. We begin at $t = 0$ with shares $x(0) = X$, cash $y(0) = 0$, and initial stock price $s(0) = S$. The strategy $v(t)$ must be adapted to the filtration of B.

We define the value function at time t in state (x, y, s) as the maximal achievable expected utility of the final dollar amount in the bank account,

$$J(t,x,y,s) \;=\; \max_{\substack{\{v(t')\,|\,t\le t'\le T\}\\ \text{s.t. } x(T)=0}} \mathbb{E}\left[u\big(y(T)\big)\,\middle|\,t,x,y,s\right]$$

for a utility function $u(\cdot)$. The utility function is increasing ("more is always better than less"). Its curvature reflects whether decision makers are risk averse (concave utility functions) or risk seeking (convex utility functions).

Standard techniques (see e.g. textbooks by Yong and Zhou (1999); Fleming and Rishel (1975); Korn (1997)) lead to the Hamilton-Jacobi Bellman (HJB) partial differential equation

$$0 \;=\; J_t \;+\; \frac{1}{2}\sigma^2 J_{ss} \;+\; \max_v\bigl(\,(sJ_y - J_x)v \;-\; \eta J_y v^2\,\bigr) \quad ,$$

where subscripts denote partial derivatives. The quadratic is easily maximized explicitly, giving the optimal trade rate

$$v_* \;=\; \frac{sJ_y - J_x}{2\eta\, J_y}$$

and the final PDE for $J(\tau, x, y, s)$

$$J_\tau \;=\; \frac{1}{2}\sigma^2 J_{ss} \;+\; \frac{(sJ_y - J_x)^2}{4\eta\, J_y} \tag{A.1}$$

with $\tau = T - t$. This is to be solved on $\tau > 0$, with initial condition

$$J(0, x, y, s) \;=\; u(y) \quad \text{for } x = 0 \;. \tag{A.2}$$

It is crucial to handle the constraint $x(T) = 0$ correctly. One way to do so is by adding the expected cost of selling any remaining shares at time T with a linear liquidation strategy in an extra time period $[T, T + \epsilon]$, and then let $\epsilon \to 0$ (see Chapter 4, proof of Theorem 4.4).

While the PDE (A.1) is the same for all utility functions $u(\cdot)$, different solutions arise for different choices of $u(\cdot)$ due to the initial condition (A.2). Schied and Schöneborn (2007) consider exponential utility, $u(y) = -\exp(-\alpha y)$ with risk-aversion parameter $\alpha > 0$. This utility function has constant absolute risk-aversion (CARA), and it can be shown that the static trading trajectories of Almgren and Chriss (2000), see Lemma 2.2, are the optimal strategies; that is, the optimal control $v(t, x, y, s) = v(t, x)$ is deterministic and does not depend on the stock price s or the dollar amount in the bank account y, which are driven by the realization of the Brownian motion. For other utility functions with non-constant absolute risk-aversion (for instance, power utility functions $u(y) = (y^{1-\gamma} - 1)/(1 - \gamma)$, $0 < \gamma \neq 1$, or logarithmic utility $u(y) = \log y$), optimal strategies will in general be dynamic and may depend on s and y as well. As discussed by Li and Ng (2000) and Zhou and Li (2000), the family of quadratic utility functions $u(y) = \rho y - \lambda y^2$ may be used to derive mean-variance optimal strategies.

Problems in solving the problem for utility functions other than exponential utility arise from the high nonlinearity of the PDE (A.1), especially the term J_y in the denominator; in the case of exponential utility, J_y turns out to be constant (as a result of the constant risk-aversion), which significantly simplifies the PDE. We leave the solution of the PDE for other utility functions and the interpretation of these results as an open question for future research.

Notation

$\mathbb{N}, \mathbb{R}$	natural and real numbers, respectively
$\log n$, $\ln n$	logarithms of n to base 2 and e, respectively
$o(f(n))\,, \mathcal{O}(f(n))$	Landau symbols
$\mathbf{dom}\, f$	domain of the function f
f_x	partial derivative $\frac{\partial f}{\partial x}$ of function f with respect to x
$\mathcal{N}(\mu, \sigma^2)$	normal distribution with mean μ and variance σ^2
Φ	cumulative distribution function of $\mathcal{N}(0, 1)$
$B(t)$	Brownian motion
$\mathbb{P}\left[\mathcal{E}\right]$	probability of event $\mathcal{E}$
$\mathbb{P}\left[\mathcal{E} \mid \mathcal{E}'\right]$	probability of event $\mathcal{E}$ conditioned on $\mathcal{E}'$
$\mathbb{E}\left[X\right]$	expectation of random variable X
$\mathbb{E}\left[X \mid \mathcal{E}\right]$	expectation of the random variable X conditioned on $\mathcal{E}$
$\mathrm{Var}\left[X\right]$	variance of X, $\mathrm{Var}\left[X\right] = \mathbb{E}\left[X^2\right] - \mathbb{E}\left[X\right]^2$
$\mathrm{Var}\left[X \mid \mathcal{E}\right]$	variance of the random variable X conditioned on $\mathcal{E}$
$\mathrm{Cov}\left[X, Y\right]$	covariance of the random variables X and Y,
	i.e. $\mathrm{Cov}\left[X, Y\right] = \mathbb{E}\left[XY\right] - \mathbb{E}\left[X\right]\mathbb{E}\left[Y\right]$
$\rho(X, Y)$	correlation coefficient of the random variables X and Y,
	i.e. $\rho(X, Y) = \mathrm{Cov}\left[X, Y\right] / \sqrt{\mathrm{Var}\left[X\right]\mathrm{Var}\left[Y\right]}$
$\mathrm{Skew}\left[X\right]$	skewness of the random variable X,
	i.e. $\mathrm{Skew}\left[X\right] = \mathbb{E}\left[(X - \mathbb{E}\left[X\right])^3\right] / \sqrt{\mathrm{Var}\left[X\right]}^3$
$W(x)$	Lambert's W-function 88
$\mathrm{CR}(\mathsf{ALG})$	competitive ratio of algorithm ALG 86
OPT	offline optimum algorithm 86

List of Figures

Bibliography

Ajtai, M., N. Megiddo, and O. Waarts (2001). Improved algorithms and analysis for secretary problems and generalizations. *SIAM Journal on Disc. Math. 14*(1), 1–27.

Almgren, R. and N. Chriss (1999). Value under liquidation. *Risk 12*(12), 61–63.

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

Almgren, R. and N. Chriss (2003, June). Bidding principles. *Risk*.

Almgren, R. and J. Lorenz (2006). Bayesian adaptive trading with a daily cycle. *J. Trading 1*(4), 38–46.

Almgren, R. and J. Lorenz (2007). Adaptive arrival price. *Algorithmic Trading III*, 59–66.

Almgren, R., C. Thum, E. Hauptmann, and H. Li (2005). Equity market impact. *Risk 18*(7, July), 57–62.

Almgren, R. F. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. *Appl. Math. Fin. 10*, 1–18.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999). Coherent measures of risk. *Math. Finance 9*(3), 203–228.

Artzner, P., F. Delbaen, J.-M. Eber, D. Heath, and H. Ku (2007). Coherent multiperiod risk adjusted values and bellman's principle. *Annals of Operations Research 152*(1), 5–22.

Bajeux-Besnainou, I. and R. Portait (2002). Dynamic, deterministic and static optimal portfolio strategies in a mean-variance framework under stochastic interest rates. In P. Willmott (Ed.), *New Directions in Mathematical Finance*, pp. 101–114. New York: Wiley.

Bellman, R. (1957). *Dynamic programming.* Princeton Univeristy Press, Princeton, N. J.

Bertsimas, D., G. J. Lauprete, and A. Samarov (2004). Shortfall as a risk measure: properties, optimization and applications. *Journal of Economic Dynamics and Control 28*(7), 1353–1381.

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

Bielecki, T. R., H. Jin, S. R. Pliska, and X. Y. Zhou (2005). Continuous-time mean-variance portfolio selection with bankruptcy prohibition. *Math. Finance 15*(2), 213–244.

Black, F. and M. S. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy 81*(3), 637–54.

Borodin, A. and R. El-Yaniv (1998). *Online computation and competitive analysis.* New York: Cambridge University Press.

Boyd, S. and L. Vandenberghe (2004). *Convex optimization.* Cambridge: Cambridge University Press.

Butenko, S., A. Golodnikov, and S. Uryasev (2005). Optimal security liquidation algorithms. *Computational Optimization and Applications 32*(1), 9–27.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance 1*, 223–236.

Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes.* CRC Press.

Dantzig, G. B. and G. Infanger (1993). Multi-stage stochastic linear programs for portfolio optimization. *Ann. Oper. Res. 45*(1-4), 59–76.

DeMarzo, P., I. Kremer, and Y. Mansour (2006). Online trading algorithms and robust option pricing. In *Proc. of the ACM Symp. on Theory of Comp., STOC*, pp. 477–486.

Densing, M. (2007). *Hydro-Electric Power Plant Dispatch-Planning – Multi-Stage Stochastic Programming with Time-Consistent Constraints on Risk.* Ph. D. thesis, ETH Zurich. DISS. ETH Nr. 17244.

Duffie, D. and H. R. Richardson (1991). Mean-variance hedging in continuous time. *The Annals of Applied Probability 1*(1), 1–15.

El-Yaniv, R., A. Fiat, R. M. Karp, and G. Turpin (2001). Optimal search and one-way trading online algorithms. *Algorithmica 30*(1), 101–139.

Engle, R. and R. Ferstenberg (2007). Execution risk. *Journal of Portfolio Management* (Winter).

Epstein, D. and P. Wilmott (1998). A new model for interest rates. *International Journal of Theoretical and Applied Finance 1*(2), 195–226.

Fleming, W. H. and R. W. Rishel (1975). *Deterministic and Stochastic Optimal Control.* Springer-Verlag.

Fujiwara, H. and Y. Sekiguchi (2007). Average-Case Threat-Based Strategies for One-Way-Trading.

Gandhi, D. K. and A. Saunders (1981). The superiority of stochastic dominance over mean variance efficiency criteria: Some clarifications. *Journal of Business Finance and Accounting 8*(1), p51 – 59.

Goldman, M. B., H. B. Sosin, and M. A. Gatto (1979). Path dependent options: "buy at the low, sell at the high". *The Journal of Finance 34*(5), 1111–1127.

He, H. and H. Mamaysky (2005). Dynamic trading policies with price impact. *Journal of Economic Dynamics and Control 29*(5), 891–930.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Rev. of Fin. Stud. 6*(2), 327–343.

Huberman, G. and W. Stanzl (2005). Optimal liquidity trading. *Review of Finance 9*(2), 165–200.

Hull, J. C. (2002). *Options, Futures, and Other Derivatives.* Prentice Hall.

Jin, H. and X. Zhou (2007). *Stochastic Analysis and Applications*, Chapter Continuous-Time Markowitz's Problems in an Incomplete Market, with No-Shorting Portfolios, pp. 435–459.

Springer-Verlag.

Karlin, A. R., M. S. Manasse, L. Rudolph, and D. D. Sleator (1988). Competitive snoopy caching. *Algorithmica 3*(1), 79–119.

Kissell, R. and M. Glantz (2003). *Optimal Trading Strategies.* Amacom.

Kissell, R. and R. Malamut (2005). Understanding the profit and loss distribution of trading algorithms. In B. R. Bruce (Ed.), *Algorithmic Trading*, pp. 41–49. Institutional Investor.

Kleinberg, R. D. (2005). A multiple-choice secretary algorithm with applications to online auctions. In *SODA*, pp. 630–631. SIAM.

Konishi, H. and N. Makimoto (2001). Optimal slice of a block trade. *Journal of Risk 3*(4), 33–51.

Korn, R. (1997). *Optimal Portfolios: Stochastic Models for Optimal Investment and Risk Management in Continuous Time.* World Scientific.

Korn, R. (2005). Worst-case scenario investment for insurers. *Insurance Math. Econom. 36*(1), 1–11.

Krokhmal, P. and S. Uryasev (2007). A sample-path approach to optimal position liquidation. *Annals of Operations Research 152*(1), 193–225.

Kroll, Y., H. Levy, and H. M. Markowitz (1984). Mean-variance versus direct utility maximization. *The Journal of Finance 39*(1), 47–61.

Leippold, M., F. Trojani, and P. Vanini (2004). A geometric approach to multiperiod mean variance optimization of assets and liabilities. *J. Econom. Dynam. Control 28*(6), 1079–1113.

Levy, H. (1992). Stochastic dominance and expected utility: Survey and analysis. *Management Science 38*(4), 555–593.

Levy, H. (2006). *Stochastic dominance* (Second ed.), Volume 12 of *Studies in Risk and Uncertainty.* New York: Springer. Investment decision making under uncertainty.

Li, D. and W.-L. Ng (2000). Optimal dynamic portfolio selection: multiperiod mean-variance formulation. *Math. Finance 10*(3), 387–406.

Li, X. and X. Y. Zhou (2006). Continuous-time mean-variance efficiency: the 80% rule. *Annals of Applied Probability 16*, 1751.

Li, X., X. Y. Zhou, and A. E. B. Lim (2002). Dynamic mean-variance portfolio selection with no-shorting constraints. *SIAM J. Control Optim. 40*(5), 1540–1555.

Lim, A. E. B. and X. Y. Zhou (2002). Mean-Variance Portfolio Selection with Random Parameters in a Complete Market. *Mathematics of Operations Research 27*(1), 101–120.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics 47*(1), 13–37.

Lippmann, S. A. and J. J. McCall (1976). The economics of job search: a survey. *Economic Inquiry XIV*, 155–189.

Lippmann, S. A. and J. J. McCall (1981). The economics of uncertainty: selected topics and probabilistic methods. *Handbook of mathematical economics 1*, 211–284.

Liptser, R. S. and A. N. Shiryaev (2001). *Statistics of random processes. II* (expanded ed.), Volume 6 of *Applications of Mathematics (New York)*. Berlin: Springer-Verlag. Applications, Translated from the 1974 Russian original by A. B. Aries, Stochastic Modelling and Applied Probability.

Lorenz, J., K. Panagiotou, and A. Steger (2007). Optimal algorithms for k-search with application in option pricing. In *Proc. of the 15th European Symp. on Algorithms (ESA 2007).*, pp. 275–286.

Lorenz, J., K. Panagiotou, and A. Steger (2008). Optimal algorithms for k-search with application in option pricing. *Algorithmica. To appear*.

Maccheroni, F., M. Marinacci, A. Rustichini, and M. Taboga (2004). Portfolio selection with monotone mean-variance preferences. ICER Working Papers - Applied Mathematics Series 27-2004, ICER - International Centre for Economic Research.

Madhavan, A. (2002). Vwap strategies. *Transaction Performance: The Changing Face of Trading Investment Guides Series. Institutional Investor Inc.*, 32–38.

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

Markowitz, H. M. (1959). *Portfolio selection: Efficient diversification of investments.* Cowles Foundation for Research in Economics at Yale University, Monograph 16. New York: John Wiley & Sons Inc.

Markowitz, H. M. (1991). Foundations of portfolio theory. *Journal of Finance 46*(2), 469–77.

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

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

Merton, R. C. (1972). An analytic derivation of the efficient portfolio frontier. *The Journal of Financial and Quantitative Analysis 7*(4), 1851–1872.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics 3*(1-2), 125–144.

Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica 34*(4), 768–783.

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

Oksendal, B. (2003). *Stochastic differential equations* (Sixth ed.). Universitext. Berlin: Springer-Verlag. An introduction with applications.

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

Pliska, S. R. (1986). A stochastic calculus model of continuous trading: optimal portfolios. *Math. Oper. Res. 11*(2), 370–382.

Richardson, H. R. (1989). A minimum variance result in continuous trading portfolio optimization. *Management Sci. 35*(9), 1045–1055.

Rosenfield, D. B. and R. D. Shapiro (1981). Optimal adaptive price search. *Journal of Economic Theory 25*(1), 1–20.

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

Schied, A. and T. Schöneborn (2007). Optimal portfolio liquidation for cara investors.

Schweizer, M. (1995). Variance-optimal hedging in discrete time. *Math. Oper. Res. 20*(1), 1–32.

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance 19*(3), 425–442.

Shreve, S. E. (2004). *Stochastic calculus for finance. II.* Springer Finance. New York: Springer-Verlag. Continuous-time models.

Sleator, D. D. and R. E. Tarjan (1985). Amortized efficiency of list update and paging rules. *Comm. ACM 28*(2), 202–208.

Steinbach, M. C. (2001). Markowitz revisited: mean-variance models in financial portfolio analysis. *SIAM Rev. 43*(1), 31–85.

The Economist (2005). The march of the robo-traders. In *The Economist Technology Quarterly, September 17th*, pp. 13. The Economist.

The Economist (2007). Algorithmic trading: Ahead of the tape. In *The Economist, June 21st*, pp. 85. The Economist.

von Neumann, J. and O. Morgenstern (1953). *Theory of Games and Economic Behavior, 3rd ed.* Princeton: Princeton University Press.

Walia, N. (2006). Optimal trading: Dynamic stock liquidation strategies. Senior thesis, Princeton University.

Wan, F. Y. M. (1995). *Introduction to the calculus of variations and its applications.* Chapman and Hall Mathematics Series. New York: Chapman & Hall.

White, D. J. (1998). A parametric characterization of mean-variance efficient solutions for general feasible action sets. *Journal of Multi-Criteria Decision Analysis 7*(2), 109–119.

Yao, A. C. C. (1977). Probabilistic computations: toward a unified measure of complexity. In *18th Symp. on Foundations of Comp. Sci. (1977)*, pp. 222–227. IEEE Computer Society.

Yao, A. C. C. (1980). New algorithms for bin packing. *J. Assoc. Comput. Mach. 27*(2), 207–227.

Yong, J. and X. Y. Zhou (1999). *Stochastic controls: Hamiltonian systems and HJB equations.* Springer-Verlag.

Zhou, X. Y. and D. Li (2000). Continuous-time mean-variance portfolio selection: a stochastic LQ framework. *Appl. Math. Optim. 42*(1), 19–33.

Zhou, X. Y. and G. Yin (2003). Markowitz's mean-variance portfolio selection with regime switching: A continuous-time model. *SIAM Journal on Control and Optimization 42*(4), 1466–1482.

Curriculum Vitae

Julian Michael Lorenz

born on December 19, 1978 in Pfaffenhofen a.d. Ilm, Germany

1987 – 1998	Secondary Education Schyrengymnasium Pfaffenhofen, Germany
1998 – 1999	Mandatory Military / Civilian Service BRK Pfaffenhofen
1999 – 2004	Studies in Mathematics at TU München, Germany Degree: Diplom-Mathematiker
	Scholarship of German National Academic Foundation (Studienstiftung des Deutschen Volkes)
2003	Studies at National University of Singapore, Singapore LAOTSE Exchange Program
since March 2005	Ph.D. student at ETH Zürich, Switzerland