Normative Approach to Market Microstructure Analysis

Yuriy Nevmyvaka

CMU-RI-TR-05-32

Submitted in partial fulfillment of the requirements

for the degree of Doctor of Philosophy

School of Computer Science

Carnegie Mellon University

Pittsburgh, PA 15213

July 2005

Thesis Committee:

Katia Sycara, RI, Chair

Duane Seppi, GSIA, Chair

Matt Mason, RI

Steve Smith, RI

Michael Kearns, University of Pennsylvania

In this thesis we propose a normative approach to market microstructure analysis. We study, model, and quantify low-level high-frequency interactions among agents in financial markets. This is an environment where electronic agents are much better positioned to both make decisions and take actions because of the amount of information and the rapid pace of activity, which overwhelm humans. Unlike previous work in this area, we are not only interested in explaining why microstructure variables (prices, volumes, spreads, order flow, etc) behave in a certain way, but also in determining optimal policies for agents interacting in this environment. Our prescriptive – as opposed to explanatory – method treats market interactions as a stochastic control problem. We suggest a quantitative framework for solving this problem, describe a reinforcement learning algorithm tailored to this domain, and conduct empirical studies on very large datasets of high-frequency data. We hope that our research will lead not just to automation of market activities, but to more orderly and efficient financial markets.

Acknowledgements

While I got to do all the typing, this thesis is first and foremost a product of collaboration with a lot of great people both in the academia and in the industry. I would like to thank the members of my thesis committee, who contributed their considerable knowledge, resources, and patience to my graduate education: Katia Sycara, Duane Seppi, Michael Kearns, Matt Mason, and Steve Smith. Ph.D. student from University of Pennsylvania Elliot Feng worked very closely with me and helped me tremendously with the empirical studies. I would like to acknowledge my colleagues at Lehman Brothers, who contributed their many years of experience to my research: Amy Papandreou, Andy Ellner, Stephano Risa, Dimitry Nesterov, Andy Westhead, Bill Warren, Tim Genirs, and others. And finally, I want to thank my friends and family for their love and support through the years.

Table of Content

Introduction

In this thesis we propose a normative approach to market microstructure analysis. It is our ambition to study, model, and quantify low-level high-frequency interactions among agents in financial markets. Unlike previous work in this area, however, we are not only interested in explaining why microstructure variables (prices, volumes, spreads, order flow, etc) behave in a certain way, but also in determining optimal policies for agents interacting in this environment. Our normative – as opposed to explanatory – method treats market interactions as a stochastic control problem. We suggest a quantitative framework for solving such problem, describe a reinforcement learning algorithm tailored to this domain, and conduct empirical studies on massive datasets of high-frequency data (1.5 years on millisecond scale). Trade execution optimization is the centerpiece of this thesis, but we also demonstrate the flexibility of the proposed framework by applying it to other microstructure-based trading problems without major modifications.

Market Microstructure is a discipline that studies how securities prices are determined under different market mechanisms. From the point of view of Economics, price is a point of intersection of supply and demand schedules; in Neoclassical Finance, price is again a single number – an intrinsic value of a security on which all market participants agree and thus can trade an arbitrary number of shares at this one price. Neither concept is faithful to the real world: in financial markets prices are in constant flux, agents have divergent expectations about values, information is never fully accessible, investors act under different constraints, etc. Therefore, it is said that financial markets are engaged in a continuous, on-going process of *price discovery* – through a series of repeating negotiations among themselves, market participants are trying to determine true values of traded securities.

Every single market actor should be aware of this process, since these short-term market fluctuation affect transaction prices, and ultimately the investors' wealth. Let's say, an investor decides (either through research or some information that he believes nobody else has) that a certain stock is undervalued. To take advantage of such opportunity, this investor would like to buy the stock, wait for its value to converge to its true level, and then sell for a profit. If the current market consensus about the stock's value could be summarized by a single number, then this investor would buy as many shares as his risk preference and capital constraints allow at this consensus price. However, in the light of the on-going price discovery process, his decision to "buy stock A" is underspecified. Investor can act impatiently, and demand the best price available immediately – this will guarantee that his trade takes place, but will result in an inferior price; or he can wait for a price improvement, but risk not transacting and thus not taking advantage of the investment opportunity. This example shows that the implementation of an investment objective affects actual prices at which securities are bought and sold and thus cash flows to investors. To execute trades intelligently, one must account for the so-called *market frictions* – the bid-ask spread, price and volume volatility, order book composition, market resiliency, and many others. It is our goal to quantify these short-term market microstructure effects and derive optimal actions for any investment strategy.

Current state-of-the-art market microstructure research can be divided into two currents: theoretical and empirical. Theoretical works model individual market participants: dealers, informed traders, noise traders, and so on. Given a population of traders and their utility functions, theoretical models derive market equilibriums, and each trader's optimal equilibrium strategy. The benefit of this analysis is that, given a market setting, a theoretical model can tell an agent exactly what actions to take (whether buy or sell, and what prices and volumes). However, there are several problems with this method: first, equilibrium determination is analytically intractable and computationally expensive for all but simplest settings; second, theoretical models rely on many unobservable factors, such as utility functions, risk preference, inventory holdings, and so on. This makes it difficult or even impossible to apply such models to the real world data.

On the other end of the spectrum are empirical models: they concentrate on the variables that can be readily observed – prices, volumes, spreads, etc – instead of modeling the behavior of actors who originate these variables. In treating market microstructure factors as time series, we can model their behavior without understanding the underlying processes. Empirical methods can tell us, for example, that spreads tend to revert to their mean values, that extreme price movements tend to reverse, that volume misbalances are indicative of future price movements, and so on, but they give no recommendations about how agents should behave under such conditions.

We propose a Reinforcement Learning-based method that delivers the best of the two approaches: it is an empirical method driven by observable market variables, but it is created specifically to derive optimal actions (optimal order prices, more specifically) based on the uncovered relationships among microstructure variables. We implement a novel RL algorithm, which takes advantage of the problem's structure to reduce the dimensionality of the search space and to use the available data in the most efficient way. We conducted large scale empirical studies on very large datasets to gain an insight into the information content of various microstructure factors and their influence on the market participant's optimal actions. We also show how our normative method can be applied to a variety of problems.

This document is structured as follows. We start by introducing the environment where we operate: we describe the mechanisms that govern modern equity markets, concentrating on the aspects that induce microstructure effects. This section should give the reader an appreciation for the problem we are trying to solve. We then survey the state-of-the-art research in market microstructure, simulated markets, and AI. In the following chapter we describe the experimental setup in which our models are tested – we emphasize the realism of our simulator built around the historical activity records from an actual electronic exchange, but we also point out its inherent limitations.

In the Market Microstructure chapter, we build a theoretical foundation that supports the validity of our research. We start with the Neoclassical Finance Theory, which tells us how prices are supposed to behave (as a random walk) and how securities should be valued (according to their risk measure) in frictionless markets. However, the real world markets are not frictionless, and we show how neoclassical theories no longer hold on the microstructure level. We explain sources of market frictions and point out positive and negative autocorrelations in transaction prices that these frictions induce. We then present a systematic way to reason about market microstructure effects: we start with the simplest model, and build it out by adding layers of complexity. All discussion is geared toward the automation of microstructure strategies.

In the remainder of the thesis, we present our empirical results. We start by discussing two simplistic implementations of current empirical models. In analyzing the obtained results, we point out the shortcomings of available methods and provide a motivation for a normative approach. We present our new method in the context of a specific microstructure problem – efficient trade execution. Executing trades efficiently means minimizing trading costs, so we explain how these costs can be quantified. We start with a simple single-period model and demonstrate that combining limit and market orders results in lower trading costs than when using market orders only. We develop an explanatory model, which shows how we diminish market impact by distributing it over time.

We then move on to a complex multi-period problem. We explain how we translate the real world information into a lower-dimensional state space. To find an optimal policy, we need to perform a search over this space. We propose a reinforcement learning algorithm, which exploits the structure of the execution problem to shrink the number of strategies that must be considered and which re-uses available data for added efficiency. We then present our results – an investigation of how various microstructure factors influence transaction prices and optimal actions. We also describe how the same methodology can be applied to other trading strategies.

We believe that the main contributions of this thesis are the following:

(1) *market simulator*, which uses records from an actual electronic exchange and includes every event that happens (submissions, cancellations, hidden orders, etc) and not just transaction prices and quotes

(2) *theoretical foundation* for automated trading strategies on the microstructure level

(3) *normative approach* to empirical microstructure research

(4) *novel RL algorithm* tailored to microstructure interactions

(5) *empirical study* of market microstructure factors

Research Background: Modern Financial Markets

Since our approach is largely empirical, and we are confident that our results are practical and transferable into the real world, we first explain where our research fits in modern financial markets. In order to appreciate the importance of our contribution, the reader ought to understand the role financial markets play in our society, who market participants are, how markets are organized and regulated, and finally what the most important trading venues are. All this will be briefly explained in this chapter. The goal of this chapter is to give the reader enough background on financial markets to make him or her understand and appreciate our contribution to the field.

A. Why Do We Need Markets?

While the concept of free markets may be the main ideological pillar of the capitalist society, role, functioning, and influence of the modern financial markets is obscured from most people. Centuries-old metaphor of Adam Smith's "invisible hand" [Smith, 1776] still holds true today: we cannot see financial markets themselves, but their impact is manifested everywhere. While it's market crises, financial crimes, and outsized losses that make front pages, markets exist to provide an orderly way to efficiently distribute goods, services, and capital in the economy, and have been successful in doing so for centuries.

The *macro-economic* role of financial markets is, first, the efficient capital allocation among competing projects and, second, enabling of risk sharing for market participants. Companies need to raise money to finance their projects that produce real economic growth, and markets serve as voting mechanisms that award more capital to the most promising (i.e. likely to succeed, high expected value) projects. Securities' trading also represents a continuous process of re-evaluation of companies' worth and financial well-being, which in turn leads to a more efficient economy.

Through trading investors can transfer securities – and thereby associated values and risks – from one to another. This ability to quickly move assets and liabilities to the entities that are best positioned to hold and manage them is called liquidity and is one of the main reasons for success of capitalist societies. If financial securities were non-transferable, then the initial investors (in the case of stocks) or lenders (bonds) would have been committed to carrying the issuing company's risk indefinitely. Such inflexibility would make investors less willing to commit their money, and make it difficult for companies to raise capital to expand their activities, which, in turn, would result in a general economic slowdown and a decrease in the nation's wealth.

This non-stop flow of money and securities between companies and investors (primary market) and among the investors (secondary market) breeds a class of facilitators or intermediaries, such as exchanges, brokers, dealers, and others. They constitute an important middle-layer, which contributes to the overall efficiency, and their activities and impact will be examined closely in this thesis.

The second, *micro-economic* function of financial markets, which is of the greatest interest to us, is the process of price discovery. One can think of trading as an on-going process of determining the value of a security through continuous negotiations among market participants. Every trade that takes place means that two parties have agreed on the worth of the traded stock or bond at that particular point in time. However, prices are changing constantly as a consequence of difference of opinions among investors (*divergent expectations*), new information arriving to the market, temporary fluctuations of supply and demand for available shares (*liquidity shocks*). This search for the true value of a security, know as *price discovery*, is the main subject of our analysis. What are the dynamics of this dynamic process? What is the optimal behavior in this constantly evolving environment? These are some of the questions we will try to answer.

While we normally think about financial market as a single concept, investors actually participate in a number of different markets: the market for securities, the market for information, and the market for transaction services. The *market for securities* deals with the determination of the intrinsic (or long-term) worth of a stock or a bond: what is the present value of future cash flows that the owner of a security is entitled to? Asset pricing and fundamental valuation fall into this category. This market is often analyzed under the assumption of frictionless and costless trading. The *market for information* deals with the supply and demand for information: transactions and price movements are oftentimes the ways in which private information gets disseminated to all the market participants. Existence of this market can explain the prolonged interest in technical analysis, incentives for arbitrageurs and stock analysts, etc.

The *market for transaction services* is the one we are the most interested in. This market prices the services of financial intermediaries: brokers and dealers. This pricing is reflected in predominant spreads and commissions. This market is also called *market for liquidity* – traders who wish to trade at a given moment in time demand liquidity at that moment, and this liquidity is supplied either by the dealer's quotes, or limit orders, or a combination of the two. It is very important to be aware of this multi-layered nature of financial markets, since the three components are highly interrelated. For example, if a trader prices an illiquid asset without taking this reality into consideration, acquiring a position in such asset may lead to disastrous consequences. Similarly, liquidity providers must take the intrinsic value and information effects into consideration when posting their quotes and submitting their limit orders.

B. Market Actors

1. Issuers

The three main groups of market participants have been already mentioned: to raise money for their projects, *issuing companies* sell securities to *investors*, who in turn exchange these securities with other investors to balance their holdings, and *intermediaries* facilitate all flows of money and securities.

All market activities originate from issuers. A simple act of transferring money and paper among investors creates no economic value by itself, and is more akin to gambling. Markets exist to transfer capital into those hands that will employ it in the most productive way, which means creating tangible benefits by producing goods or offering services. Companies turn to investors for capital to finance their projects. They may choose to either borrow (issue bonds – a promise to repay borrowed money with percentages in the future) or offer to share the proceeds from their projects (issue stock – an offer of partial ownership of the company). Investors then get to asses the merits of proposed projects and determine how much they are willing to pay for the securities offered. In the great majority of cases, in order to sell securities to investors in exchange for cash, issuers employ intermediaries: investment banks and their staff of corporate finance specialists and salesforce.

It is important to keep in mind this notion of fundamental value that is attached to a stock (or almost any other security). A common share is a claim on a [very small] portion of all future cash flows generated by a corporation, and the value of these cash flows – properly discounted – should be equal to the price of the stock. This is the basis for fundamental analysis of financial securities. However, there is so much uncertainty associated with the future projects of any company, that in practice this "true value" of a stock is anybody's guess. This is why equity prices change so rapidly, as investors are trying to discover this true price through negotiations and assimilation of new information. And this is why we have to study market microstructure and understand the dynamic properties of this process. While our analysis of market microstructure issues is predominantly technical (as opposed to fundamental) in its nature, this link to the real world is always present. Example: while in many models, we treat stock prices as random walks (stochastic processes that can wonder up and down without pre-set limits), in reality stock holders have a limited liability. This means that equity prices cannot go below zero even if the values of underlying projects/companies can. We have to use logarithms of prices in order to address this issue in mathematical modeling.

One microstructure aspect that is tied to the issuers is information dissemination: public companies are obligated by law to provide all the important information to *all* investors on a timely basis. Publishing news periodic financial reports is supposed to help the investors to better assess the company's value. We are interested in examining how information gets impounded into prices once it becomes public, and how the investors anticipate information arrival – i.e. why prices change ahead of the news.

2. Investors

Investors (sometimes referred to as *traders*) are the primary source of market activity. Profit-motivated price discovery is largely their responsibility, and we spend most of the time analyzing their behavior. There are many ways to classify traders.

From the organizational point of view, we can separate investors into *individual* investors and *institutions*. Institutions can be mutual funds, pension funds, hedge funds, banks, insurance companies, trading desks of large corporations, etc. The main aspect that separates institutional investors is that they usually prefer to trade in much larger sizes than individuals. According to some studies, institutional trading accounts for 70 to 80 percent of all equity volume in the US [Schwartz et al., 2004]. Aside from the average order size, the difference between an institutional and an individual trader has become largely blurred. With the recent advances in communications technologies individuals can execute their orders in real time, receive news and analysis as quickly as institutions, and a vast quantity of market research along with advanced analytical tools is available to them. One can argue that the long-standing competitive advantage of institutions over individuals has all but disappeared.

Investors (both individual and institutional) can pursue a large number of strategies. We will name a few most common types here. Investors who use fundamental analysis to determine the "true price" of a stock are known as *value investors*. They buy securities believed to be significantly undervalued They tend to accumulate relatively large positions and hold them for extended periods, since it can take a long time for market prices to converge to their intrinsic value. Those investors who believe that stock prices indeed resemble random walks and are therefore impossible to predict may choose to invest into *index funds* – pools of money designed to track broad market indexes (S&P500, DJIA, NAS, etc). The idea is that while "picking winners" may be difficult, the broad market is moving up over time, reflecting the real value created by the economy as a whole; thus the goal is to replicate the payouts of the broad market as cheaply as possible.

Although looked down upon by the academia, *technical investing* is alive and well. These strategies are based on an assumption that stock prices exhibit persistent patterns, which once discovered can be exploited. Technical trading involves searching for price "trends", "support and resistance level", etc. The critics of the technical analysis are quick to point out that such patterns cannot persist because anticipation of future price movements will make them disappear. For example, if a price of a certain stock always goes up in January, then investors will try to accumulate a position in this stock sometime in December, which will push the price up before January and the pattern exists no longer. Recent research, however, suggests that there may be a link between market microstructure and technical analysis: price trends may be the consequence of dynamic (i.e. gradual) adjustment of prices to the incoming information (see [Schwartz and Francioni, 2004]) and support levels may reflect excess liquidity available at a certain price – if a large order is submitted at a pre-specified price, it can take a long time for transaction prices to get through that level (see [Kavajecz and Odders-White, 2002]).

Another popular class of trading strategies is *arbitrage*. The idea is that if two financial instruments promise the same flow of money to their holders, they should trade at the same price (since they have the same intrinsic value). If this isn't the case, then a trader buys the undervalued instrument, sells the overvalued one, and guarantees himself a riskless profit when the two prices converge. The simplest example of arbitrage: price of gold in the US must be the same as the price of gold in Great Britain; otherwise a trader can buy the gold where it is relatively cheaper, transfer it to the other country and sell it for a profit (minus the cost of transfer of course). The reality is many times more complex, and the matter of fact is that "free profits" are difficult to come by.

For modeling purposes, we can think of traders as being *active* or *passive*. Active traders trade frequently, demand liquidity, and employ market orders or aggressive limit orders. Passive traders do not initiate transactions themselves, but rather supply liquidity for the active traders via limit orders. Arbitrageurs are an example of active traders, while value investors can be passive traders.

Another distinction can be drawn between *liquidity* (or *uninformed*) and *informed* traders. Liquidity traders transact to rebalance their portfolios. This can be caused by a change in their risk preference, need for cash, or some other exogenous factors. Informed traders, on the other hand, have some private information about the value of a security (or at least they think that they do), and trade to acquire a position in such security before this private information becomes public and gets reflected in the security's price. As a general rule, liquidity traders tend to lose to the informed traders. We can think of index fund managers as "uninformed" – they buy stock not because they believe that it will go up, but because it belongs to a market index they are replicating. An informed investor can be someone with an exceptional knowledge of a particular company or industry (an ex-executive, for example) or someone with superior analytical tools.

3. Intermediaries

The first type of financial middle-men is a *market facility* or an *exchange*, within which the trading takes place. If the markets as institutions did not exist – i.e. if there were no notion of financial intermediation – then an investor who wanted to buy an asset would have to find another investor willing to sell this same asset. Furthermore, the first trader would have to contact every single possible seller in order to find the best price. Obviously, such process would be prohibitively expensive if at all feasible. This is why financial exchanges were established, and in this context, we treat them as intermediaries that bring buyers and sellers together at the same location. This location can be either physical – the floor of the New York Stock Exchange (NYSE), for example – or virtual, as in the case of Electronic Communications Networks (ECNs). We will present a more detailed market typology later in this chapter. While we say that exchanges facilitate transactions in space, this, however, may not be sufficient.

Imagine an investor who wishes to sell a certain amount of stock. He arrives to the market place, but finds nobody who is willing to buy at that very moment, or at least not at an acceptable price. This trader can either exit the market or wait for a buyer to arrive. In either case, the investor incurs some undesirable cost – i.e. his utility (satisfaction) is clearly lower than if he could trade immediately. If the seller chooses to leave the market place after not immediately finding a counterparty for his trade, he also hurts the eventually arriving buyer, who now has no seller to transact with. While this example with only two traders is admittedly contrived, we can think of many other situations when the overall system utility is reduced because the traders "missed" one another – for example, the two traders that could have agreed on a mutually acceptable price are forced to trade on disadvantageous terms because they did not arrive to the market at the same time.

This problem can be resolved with the introduction of intermediaries (middle agents): the first trader transacts with an intermediary and leaves; the intermediary waits for the second trader to arrive and enters into an offsetting transaction. He charges both traders a fee, such that all three of them are better off than before. Thus both traders ultimately transact with one another even if they are not present in the marketplace at the same time. Just like financial exchanges bring investors together in *space*, middle agents (market makers) bring investors together in *time*. The diagram in Figure 1 should help visualize this concept: in (a), both the buyer and the seller arrive at the same time and transact; in (b), the traders are separated in time, which results in either no trading or the seller having to wait; and in (c), both traders transact immediately through a middle agent.

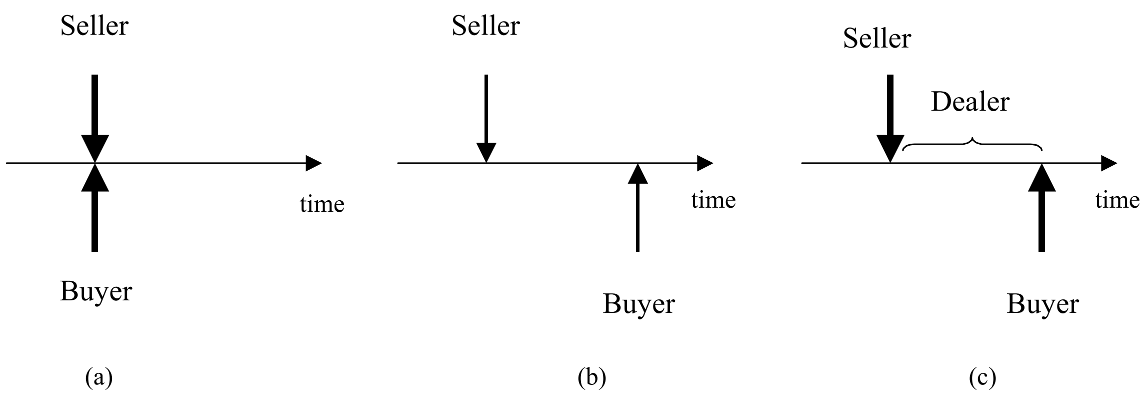

Figure 1.

We classify financial intermediaries into *dealers* (or *market makers*) and *brokers*. The situation above – when a middle agent buys securities from one trader and sells them to another trader at a later point in time – is a classic example of a dealer behavior. Dealers are the designated counterparties for the rest of market participants – when a trader wants to transact, he can trade immediately with a market maker at prices that the latter quotes. The key concept to understand about the dealer's role is that by committing to trading at all times, the dealer participates in the market as a *principal* – i.e. he trades for his own account. For example, when the market maker buys stock from the first trader, he has to hold on to this stock for some time. During this time, if the price of the stock falls, the dealer loses money. Conversely, if the price rises, he makes money. Then there is a possibility that the buyer never arrives, and the dealer gets stuck with an unbalanced portfolio (holding a stock he doesn't want). One of the marker maker's main goals is to manage these uncertainties and to be reimbursed for his troubles by other market participants.

Dealers are the central players in many organized markets. One of our goals is to understand the dealer's behavior and to explore the ways to automate his valuable services. We can find dealers in almost all major securities markets: NYSE (specialists), Nasdaq (OTC market makers), bond markets, currencies, etc. Electronic Communication Networks are a notable exception: they have no explicitly designated market makers and all trading happens by matching investors' orders directly with one another. This does not mean, however, that the dealers' services are not needed on ECNs – ad-hoc liquidity providers can make trading over ECNs cheaper, thus benefiting the entire system. We will look closer at these issues later in this thesis.

The main feature that differentiates *brokers* from dealers is that brokers do not take on risk. They do not trade on their own account and get paid by commission for arranging transactions on behalf of their clients. Instead of taking a position from one trader and maybe transferring it to another trader in the future, as in the example above, a broker would help one of his clients to find a counterparty for his trade out of a pool of his other clients, available dealers, electronic networks, etc. In most cases, brokers help find the *best execution venue* for those actors who do not have direct access to the market. Individual investors often transact through brokers. This thesis is mostly concerned with institutional trading, and thus broker's activities are of little interest to us.

C. Market Organization

Now that we have reviewed the importance of financial markets to our society and have introduced the main participants, we need to lay out the rules and mechanisms through which market participants interact.

1. Order Types

An *order* is a way for an investor to initiate a transaction. If a trader wishes to make a trade, he submits a directive to either a broker or directly to the market (if he has such access) to either buy or sell a specified number of shares. A *market order* directs the broker to transact immediately at the best price he can find or executes immediately against the best price quoted in the market in the case of direct access. When submitting a *limit order*, a trader specifies a reservation price at which he is still willing to transact. For example, when a trader submits a buy limit order, he is willing to buy a specified quantity of stock at a limit price or lower. For a limit sell, it's the opposite: limit price or higher.

…
25.56 – 300
25.55 – 1000 Sell Orders
25.35 – 200
25.30 – 150

25.21 – 200
25.19 – 300
25.15 – 785 Buy Orders
25.10 – 170
…

Figure 2.

Incoming limit orders that cannot execute immediately at their specified prices, get entered into *order books*, where they are arranged by price. They transact against the incoming market orders. See Figure 2 for an order book snapshot: every order in the book consists of the specified volume, reservation price, and an indication of whether it's an instruction to buy or to sell. The way prices are specified right now (the lowest sell price is higher than the highest buy price), no immediate transaction is possible, but an incoming market buy will transact $25.30, market sell – at $25.21. There is also a number of more complex conditional orders, "good till cancelled", "fill-or-kill", etc., but their behavior can be replicated using these basic order types, and we never use them explicitly in our models or implementation.

2. Market Mechanisms

Generally, we classify market mechanisms into two broad categories: *quote-driven* and *order-driven*. In a quote-driven (or *dealer*) market, dealers post their quotes, and all transactions take place at those quotes – investors cannot transact directly with one another, all trade flow goes through market makers. Bond market and currency market are pure dealer markets, Nasdaq also started as a dealer market. Under this mechanism, the dealer has to buy every time an investor wishes to sell, and has to sell when an investor seeks to make a purchase. The dealer participates in every single transaction and is the designate counterparty for the rest of market participants. Dealer's buy quote (called *bid*) is always lower than his sell quote (*ask*), which creates a so-called *bid-ask spread* enabling the market maker to sell at higher prices than he buys. The spread reimburses the dealer for providing liquidity to the market.

It is essential to understand that even though all trades go through the dealer he does not set prices in the market: the buyers ultimately transact with sellers, and sellers with buyers. The dealer just interposes himself between the two groups of traders and guarantees that there will always be the "other side" of the trade when a trader of either type arrives to the market place. Quote-driven markets can feature multiple dealers in the same securities, competing for the order flow among themselves – public investors will only transact with the dealer who has the highest bid or the lowest ask in the market. Besides striving to attract the order flow by posting competitive quotes, market makers have to control their inventory. Indeed, if the dealer only receives sell orders (thus accumulating a large positive inventory), he never gets to "buy low and sell high" in order to profit from the spread. Because of these realities, quote-driven markets are best suited for somewhat illiquid securities where his presence can add some structure and support to the process of price discovery.

In a pure order-driven (or *auction*) market, investors transact directly with one another by submitting limit and market orders; there are no designated market makers. ECNs are a canonical example of such market mechanism. As described above, limit orders form order books, thereby establishing prices and volumes at which incoming market orders can transact. These books are visible to all market participants. It is said that limit orders *provide liquidity* to the market (essentially performing the function of the market maker in the quote-driven market), and that market orders *demand or take liquidity*. This process is also known as a *continuous double-auction*.

Note that the traders who use market orders still end up paying the bid-ask spread in the auction market just as the do in the dealer market. Now the spread is the distance between the highest buy and the lowest sell, and these two orders were not necessarily submitted by the same trader. Market orders traders pay the spread, but receive immediate execution, while limit order traders receive price improvement, but face execution uncertainty (their orders may never execute). In order for the pure auction market to exist, both types of traders must be present in the market place.

A variation of the continuous double-auction is the *periodic call auction*. Trading under this mechanism is no longer continuous, but happens at pre-determined times in batches: at the beginning of the time interval all traders submit their bids – prices and quantities – and then the central exchange mechanism (an auctioneer) clears all trades at a single price, using some well-defined clearing algorithm. Here there is no notion of the bid-ask spread, which makes trading through market orders cheaper, but traders give up their ability to transact any time the want – they have to wait until the next auction is called.

In reality, almost all modern financial markets are *hybrid* markets: they combine features of all mechanisms just described. Markets can open and close with a call auction, maintain an open order book visible to the public, and employ designated market makers at the same time. We will take a closer look at several US markets at the end of this chapter.

3. Rules, Regulations, and Frictions

Other aspects of modern securities markets useful for the understanding of models and applications presented in this thesis are the rules of precedence, tick size, and a realistic example of the trading process. The *rules of precedence*, while complex in reality, for our purposes can be summed up as follows: order with the most advantageous (highest for the buys, lowest for the sells) price transacts first; for the orders at the same price, the one submitted the earliest transacts first.

The rules of precedence are closely tied to the *tick size* – the minimum allowable distance between posted prices. For example, until 1997, the tick size on the NYSE was 1/8th of a dollar – i.e. allowable prices were $1, $1 1/8, $1 1/4, $1 3/8, etc. Since then the tick size has been reduced to 1/16th of a dollar, and, eventually, to a penny. $0.01 is the mandated tick size for all US equity markets today. The effects of tick size reduction are not clear-cut. On one hand, the proponents of tick reduction say that smaller tick size increases the competitiveness/aggressiveness of dealer's quotes and limit orders, which leads to tighter spreads, and reduces the cost of trading. On the other hand, smaller "price steps" make it easier to "undercut" or "step in front of" a resting order. If the tick size is 1/8 or $0.125, and the top order in the book is at $10, then to step in front of it one will have to bid $10.125, which is a significant increase. Whereas, if the tick is $0.01, then stepping in front of an order costs virtually nothing, which renders time priority rule essentially useless and makes limit orders more vulnerable, thereby discouraging the limit order submission. This means that while the spreads are in fact tighter, the pre-committed liquidity or *depth* at those quotes is much smaller than before when spreads were larger, which makes trading more expensive. Which of the two effects – tighter spreads or reduced liquidity – is dominant has been a topic of multiple studies, but no consensus has been reached thus far. For a more in-depth discussion of this issue see [Goldstein and Kavajecz, 2000] and [Chordia and Subrahmaniyam, 1995].

We capture in Figure 2 how the trading process is generally modeled. Either two traders meet in some trading environment, and exchange shares for cash, or trade in the same fashion with the market maker. While this abstraction gives a fair representation of what happens in the real world, it is nonetheless beneficial to be aware of the process that really takes place and be conscious of the *market frictions* that transpire.

In [Stoll, 2001], the process is broken up into several stages. First, the market provides *information* about past and current prices and quotes, which allows the investor to make his trading decision. While this seems obvious, not that long ago the NYSE kept this information private, disseminating it only to its members. Today, pricing data is available in real time for all of the US equity markets. The situation is not the same with the limit order books. While the books are open to the public on NYSE and ECNs, this is not the case in many other markets, such as bond market.

The second step, after the investor decides to transact, his broker goes through a process of *order routing*. The broker examines different options for the client's order execution and sends the order to an appropriate venue. He can send it on the floor of an exchange, to the "in-house" trader, over to an ECN, etc. The broker's obligation to his client is to find the best price for the client's order, which is not always evident because of market fragmentation, payment for order flow, and other frictions. Order gets routed either through an electronic system or via a phone.

The third stage of the trading process is the *execution* proper. In the electronic market, a computer simply crosses the incoming order with an outstanding limit order or a quote, whereas when dealing with human market makers it may get more complicated. The dealers, out of fear of being "picked off" by better informed traders, naturally try to delay the execution – even if by a few seconds – to observe the order flow or information arrival. While it appears that electronic systems offers a better deal from the customers' standpoint, their volumes can be thinner and such systems can potentially be exploited. We will attempt to address these shortcomings in our models and implementations.

Finally, the fourth phase of the process is the *trade settlement*. Since after the trade cash and securities do not physically change hands, this stage amounts to making entries in the computer and comparing records of all sides involved in the trade – some trades can be erroneous, there can be glitches in recordkeeping in various systems, human error aspect is introduced, etc. After all possible disagreements are resolved, the trade is cleared and is irreversible, the ownership of securities is transferred, and cash is put in the correct bank account. This clearing/settlement process takes three days to complete after the trade takes place.

As it must be evident at this point, actual trading process is fairly involved and complicated, which suggests that a simple random walk – prices wondering around and instantaneously adjusting to new information – may not be all that realistic of an assumption. In order to act optimally in the real-world markets, we must study these market frictions and account for their effects in our investing process.

D. US Equity Markets

In this section, we briefly examine the organization of the most important US markets: New York Stock Exchange and Nasdaq. The descriptions below are still oversimplified because of the myriad rules and regulations that govern these entities, but the reader should get a fair picture of how these markets function.

1. The New York Stock Exchange

The NYSE is the oldest equity market in the US, tracing its heritage to 1792 (see [New York Stock Exchange]). It has about 1,400 *seats* (memberships) and 2,600 listed companies with a total market value of about 12 trillion, according to [Schwartz and Francioni, 2004]. All transactions happen on the physical *trading floor* located in downtown Manhattan, and the majority of trades involve human traders reaching an agreement through negotiations.

The exchange employs market makers called *specialists*. Each stock has a specialist assigned to it, and most specialists cover multiple stocks. Orders arrive to the specialist in two ways – through an electronic system (SuperDOT) and through the floor traders (sometimes referred to as "trading crowd" in analytical models). Specialists participate in the market as both brokers and dealers: they can help arrange a trade between two floor traders, or they can use their own inventory to provide liquidity.

When an investor wants to make a trade in an NYSE-listed stock, he contacts his broker. The broker's firm can trade on its own account and use its inventory to complete the transaction. Otherwise, the broker can enter the order into the electronic system or forward it along to a floor trader who works for the broker's firm. Small orders entered into the electronic system execute automatically, while larger ones show up in the specialist's order book, which is also visible to other traders. After that, the specialist executes orders in his book against incoming market orders or his own quotes, depending on the order type and price changes.

In the case when the order is routed to the floor trader, the trader approaches the specialist's post (the dealer's physical location) where other traders in the same security congregate. This trader can negotiate the price of the transaction with his counterparts from other firms. If the agreement is reached, the specialist gets notified and the transaction gets recorded. If no transaction within the trading crowd is possible, then the trader transacts with the specialist. If his order is a limit order, it gets entered into the specialist's book, where it awaits execution. If it's a market order, then it will transact with one of the outstanding limit orders or the *specialist's quote*. At all times, the specialist quotes two prices – the bid and the ask – at which he is willing to respectively buy and sell shares. The ask is higher than the bid, which creates the bid-ask spread. To each price the dealer attaches *depth* – a number of shares up to which he is committed to his posted price. For example, if the bid is $20.35, and the depth at the bid is 1,000 shares, this means that the dealer is willing to buy up to 1,000 shares at $20.35 per share. As far as the rules of precedence go, the dealer's quotes are treated as competing limit orders – i.e. the specialist cannot trade for his own account if he has outstanding limit orders at more advantageous prices; these limit order have to get crossed first. For example, let's say that the dealer's bid is $20.35, ask – $20.45, and there is an outstanding buy order in his order book for $20.38. Then if a new market sell arrives, it will transact with the outstanding buy at $20.38, and not with the dealer's quote at $20.35. There are also many other regulations constraining specialist's activities: he cannot step in front of client's orders, he cannot interfere with transactions that can happen without his participation, and so on and so forth.

This detailed description of the NYSE specialist's role exemplifies the decision problem any market maker faces. Essentially, the specialist observes the trade flow and the state of his order book and then makes a decision what prices to quote and at what depth. An automated specialist would either make these same decisions autonomously, or at least serve as a tool for a human dealer, facilitating the decision process.

2. Nasdaq

Nasdaq started in 1971 as an electronic quote-dissemination system for the over-the-counter (OTC) stocks – essentially all the equity not traded on NYSE. Now it is the second-largest exchange in the US, best-known for many technology companies that rose to prominence over the past two decades. There are about 4,000 companies listed on the

Nasdaq national market (the main exchange) with a total capitalization of about 4 trillion dollars, according to [Harris, 2003].

In the non-centralized Nasdaq market we have multiple competing market makers for each stock, ranging anywhere from 3 to 50 [Schwartz and Francioni, 2004]. They are geographically dispersed (there is no physical floor), and connected by telephone and computer networks. While NYSE relies on explicit rules and regulations that obligate the specialists to maintain fair and orderly markets, Nasdaq relies more on the competitive forces of having multiple dealers competing for the order flow by posting attractive quotes.

A broker, looking for the best execution for his client, will give a call to several dealers and ask for a quote, without revealing either the direction (buy or sell) or the size of his order. Each market maker – who is usually a trader at an investment bank – will communicate his bid, ask, and depth at both quotes, after which the broker will pick the best price for his client, call that dealer back and finalize the trade. Actually, today all the dealer's quotes are consolidated in an electronic system, which the broker can access via a terminal; there is also automatic execution of small order just like on the NYSE, and all outstanding limit orders are visible to the public. While Nasdaq started as a pure dealer market, now it is a hybrid dealer-auction market.

Market makers have the same obligation as their NYSE counterparts to execute incoming orders against outstanding orders if this results in a more advantageous price. All orders are stored in and displayed through central Nasdaq computers. There is no competing trading crowd in this setting, but standing orders in the open limit order book provide additional competition to the dealers' quotes. The above description is somewhat incomplete, but it captures the essence of the Nasdaq market maker's decision process: looking at all the outstanding limit orders, and all other dealers' quotes, but only seeing his own order flow, he has to set bid, ask and depth.

3. Electronic Communication Networks

Electronic Communication Networks rose to prominence during the technology boom of the late 1990s. The biggest names are Instinet (founded in 1969), Archipelago, and Island (in the process of forming INET through a merger with Instinet). They have attracted about 40% of volume in Nasdaq stock, but are having a harder time diverting the order flow from the NYSE. They can be viewed as legitimate competitors to the traditional exchanges – they attract investors by offering instantaneous and anonymous execution. As a reflection of these competitive pressures, mergers between NYSE and Archipelago and between Nasdaq and Instinet are currently pending.

ECNs do not have designated market makers. The absence of an official position does not signify the absence of demand for the dealer's services (additional liquidity, to be precise). Because of a multitude of factors, not all the investors willing to trade maintain outstanding orders at all times, lots of orders get cancelled before execution, and many submitted orders are for small quantities of stock. Therefore, for many stocks, order books are rather *thin*, which makes trading expensive. A trader, who fills the book with additional limit orders on both sides, thereby providing liquidity to other traders, is effectively a market maker on the ECN. Since his orders exist on both sides of the market, there is the same notion of the bid-ask spread and its updating from the trader's perspective. We will take a much closer look at this particular scenario in the next chapter.

Conclusion

The goal of this chapter was to impress upon the reader all the complexities of the real-world financial markets. In current academic (and even industrial) research, these fine details are regularly disregarded in favor of obtaining elegant and tractable solutions to the problems of securities pricing and event modeling. This can be a very dangerous practice because when these simplified models are transferred to the real world and the end users (investors and intermediaries) become overly reliant on them, market frictions tend to manifest themselves in the most destructive way. For example, the list of derivative disasters runs dozens of pages long (see most notably the description of the Long-Term Capital Management debacle that threatened the entire US financial system in [Lowenstein, 2000]), and most of them can be explained by the failure to account for market microstructure effects.

The main reason to study market microstructure is because of its impact on the bottom line: we need to know *actual prices* at which we can transact, and how do they depend on our timing, volume, prior actions, general market conditions, and so on. We will present a more analytical examination of these issues in the chapter titled "Market Microstructure".

Review of the Related Literature

The goal of this chapter is to provide a brief overview of literature that was used as a foundation for the work presented in this thesis or in some way related to the subject of market microstructure. The following is more of a catalogue designed to give the reader a general idea about the research in related areas. Here we will limit ourselves to classification of various sources and brief descriptions of their main ideas. The reader should also get a historical perspective of how the field of Market Microstructure research has evolved over the three decades of its existence. Specific approaches and findings that are especially pertinent to microstructure-based automated trading and optimization are presented in greater detail in the "Market Microstructure" chapter of this thesis.

We first present several fundamental books that can help better understand both the larger area of Quantitative Finance as well as specific methods and tools used in current research. Second, we will discuss several Market Microstructure surveys and a large number of papers examining optimal dealer behavior under different market mechanisms. Then we will review an important body of literature on modeling limit order submission and survey the most widely-used microstructure-related trading strategies. Finally, we will look at some specific contributions from the Computer Science community: simulated markets and learning/optimization techniques.

Throughout this chapter, we emphasize the uniqueness of our approach to market microstructure research. We start with canonical econometric models that explain behavior of market participants and then use advanced AI techniques to apply these theoretical intuitions to the real world data. Furthermore, our goal is to not only explain the interplay among many market variables, but also to prescribe optimal actions based on microstructure relationships.

A. Fundamentals of Financial Markets.

While automated trading is of a significant practical importance, and a lot of work in this area has been done in the brokerage industry over the past two decades, published results are few and far between for fairly obvious reasons. To the best of our knowledge, there are no comprehensive books on the subject. There are no written guidelines or "how-to" rulebooks on what is the optimal action for a trader in a given situation. [Kissell and Glantz, 2003] can be considered the lone exception – they offer a comprehensive top-down approach to trading cost minimization.

In the spirit of this thesis' goal – to establish an analytical foundation for microstructure trading – our approach will be based on general economic principles applied to financial markets: supply and demand for a security, strategic decisions based on private information, expectations maximization, etc. The branch of Economics, which aims to quantify these aspects is called *Econometrics*, and [Campbell et al., 1997] serves as an excellent, even if rather broad, primer on this subject.

To model securities prices and other events, while incorporating the general uncertainty of financial markets, we use a number of stochastic (random) processes. Knowing the properties of such processes helps us build not overly complex and yet accurate simulations of financial markets. The classic textbook in this area is [Merton, 1990]. While we emphasize throughout this thesis that the real-world transaction prices can deviate from information-efficient random walk prices, classical finance theory of efficient markets is a good starting point. It provides a tractable analytical framework that helps us reason about the risk-returns trade-offs and optimal portfolio composition. [Markowitz, 1952] and [Sharpe, 1964] are generally considered to be the intellectual foundation of the modern finance theory. Markowitz is credited with creating the portfolio theory and showing the benefits of diversification. Sharpe proposes an analytical approach to quantifying the relative risk of different securities and extends this concept to portfolio selection – how to select a combination of securities that achieves the highest return for the lowest level of risk.

We are still mostly interested in market mechanisms on the lowest level – i.e. individual transactions and not end-of-the-day prices – which falls within the domain of Market Microstructure. [O'Hara, 1995] is a comprehensive overview of leading academic theories in this area. Two other complete and yet more practically-oriented surveys are [Harris, 2003] and [Schwartz and Francioni, 2004].

The number of individual academic papers on financial modeling and market organization is truly enormous, so we will mention only a handful of more pertinent ones. [Madhavan, 1992] and [Stoll and Whaley, 1990] develop theoretical frameworks for the process of price formation in different markets. [Black, 1971] is a visionary paper describing the advantages of automated securities exchanges, which are becoming more and more of a standard nowadays. The evolution of this trend of automated trading and the related challenges are documented in [Becker et al., 1992]. And, finally, [Domowitz, 2001] looks ahead at potential future developments in electronic exchanges, concentrating on liquidity provision and the role of financial intermediaries.

Moving from concepts to methods, the two broad categories of tools that we employ are Machine Learning, used to tune models' parameters to past data, and Time-Series Analysis, used to investigate relationships between different processes (i.e. transaction price and past prices, or transaction price and bid-ask spread, etc.). The two excellent textbooks in these disciplines are respectively [Mitchell, 1997] and [Yafee and McGee, 2000]. Papers devoted to specific techniques are innumerable and will be left out in this chapter, since all the leading techniques are covered sufficiently for our purposes in the above books that also contain extensive references.

We emphasize our reliance on Market Microstructure theory in our work, and thus we find surveys on this topic extremely helpful. They can potentially explain to readers some important concepts not covered in either Fundamentals or Theory and Models chapters. The earliest progress in the area is documented in [Cohen et al., 1979]. [Stoll, 2001] is probably the most accessible and comprehensive review, which includes the discussion of the trading process, bid-ask spread, market organization, and implications for other areas of Finance. [Madhavan, 2000], while similar in its nature and structure, provides a complimentary reading to the previous work, since the author manages to present every topic from a slightly different perspective. [Madhavan, 2002] is a guide tailored specifically for market professionals, and thus is the most accessible even if slightly simplified. [Bias et al., 2004] is the most up-to-date survey from this category.

All the Market Microstructure theories that are particularly relevant to market makers are re-hashed in [Stoll, 1999], where the author presents a number of ways that the dealer's role can be interpreted: auctioneer, price stabilizer, information aggregator, and liquidity provider. In reality, it is likely the mixture of all of the above. Another interesting take on the same problem is presented in [Stoll, 2000], where all the microstructure effects (discrete prices, bid-ask spread, transparency, reporting, etc.) are called "frictions" – as in "in the frictionless markets none of this would have mattered and prices would have been a perfect reflection of available information". The underlying theories are essentially the same, but this paper helps to better understand the nature of the issues we are dealing with. [Hasbrouck, 2004] is a microstructure tutorial chiefly concerned with the time-series techniques applied to the subject.

B. Dealer Models under Different Market Mechanisms.

In this section we will look at the most prominent theories that examine dealers activities, and most importantly the sources and behavior of the bid-ask spread. The papers in this section are mostly set in the quote-driven markets – where all transactions go through the market makers and limit orders are not modeled explicitly. The main findings are universally applicable nonetheless to many other settings.

Since the bid-ask spread is the central aspect in the study of dealers' activities, it is fitting to start with [Cohen et al., 1981], which proves that the existence of the spread is inevitable in all but perfect markets. [Amihud and Mendelson, 1980] and [Ho and Stoll, 1981] are the two founding papers in the bid-ask spread literature. They essentially postulate that the market maker's quotes are primarily functions of his inventory holdings, and that the dealer adjusts them to balance his inventory – i.e. to prevent the accumulation of large positive or negative position in the underlying securities. In other words, using his quotes, the dealer induces such transactions from the rest of market participants that move his inventory towards some desired levels. A more recent work, which also adopts inventory as the central explanatory variable is [Madhavan and Smidt, 1993]. Its main finding is that inventory effects are indeed important, at least in the middle to long term. Another empirical study, which attempts to establish the importance of inventory effects is [Hasbrouck and Sofianos, 1993], which confirms that inventory adjustments in response to trades can be delayed significantly – up to one or two months.

The other, perhaps alternative, approach to explaining the presence and behavior of the bid-ask spread is offered by [Glosten and Milgrom, 1985]. Since the market maker is a committed counterparty to any trade, he is bound to loose out on transactions with traders who know more than him. The bid-ask spread exists to compensate the dealer for these loses, since it is being charged indiscriminately to both informed and uninformed traders. This same idea is expanded and cast in a more complex framework in [Easley and O'Hara, 1987] and [Easley and O'Hara, 1992], where more periods and more choices (higher branching factor) are introduced into the model.

The truth is that both effects – inventory and information – influence the dealer's decision making, and therefore must both be incorporated into theories. [O'Hara and Oldfield, 1986] is one of the first publications to recognize this and to develop a joint model. A number of empirical studies sprung out of this "debate" trying to determine which effect is responsible for what portion of the bid-ask spread. [Hasbrouck, 1988] and [Stoll, 1989] are two prominent examples of such efforts. [Huang and Stoll, 1994] and [Huang and Stoll, 1997] go even further by introducing other explanatory variables (a futures index and quotes covariance, for example) to model the bid-ask spread evolution. [Chau, 2002] is the most recent publications on this subject, which challenges some of the established concepts.

This brings us to a somewhat different class of models, which do not try to explain the underlying processes that affect price formation, but simply look at time series of variables and try to determine how these variables influence each other, without making specific assumptions. [Roll, 1984], for example, suggests that the spread is simply a square root function of a covariance of stock price changes. This line of reasoning is extended in [Glosten, 1987], [Choi et al., 1988], and [Hasbrouck, 1991], all of the authors developing more complex "relaxed" (as opposed to "structured") models.

C. Extended Models of Quote-Driven Markets.

Whereas the work presented in Section B deals with modeling of a single market maker, such setting needs not be true in the real world. To incorporate a significant number of practical "frictions" other models were developed.

One possible extension is shown in [Ho and Stoll, 1983], which applies the findings from [Ho and Stoll, 1981] to a market with multiple dealers. While it is generally accepted that competition between market makers leads to tighter spreads and lower costs of trading for investors, [Dennet, 1993] presents a multi-dealer framework that arrives at a rather simple yet interesting result: as more market makers compete in the same market, the risk of posting aggressive quotes for each individual dealer increases, which leads to wider spreads and higher trading costs. In other words, in a scenario where too many dealers compete in the same market, investors may actually be better off having a single monopolistic market maker. Another model of competing dealers is developed in [Dutta and Madhavan, 1997], where competing market makers make strategic decisions to compete or collude – this paper is based on game theoretic principles.

The dealers can face competition not only from other dealers, but also from competing traders. [Seppi, 1997] models an NYSE-type specialist, who competes with the "trading crowd" for incoming orders. The basic conclusion is that each category of liquidity providers in this setting is better equipped to handle orders of certain sizes. Other papers that examine competition between dealers and open limit order books in hybrid markets are [Peterson and Sirri, 2002], [Ellul et al., 2003], and [Boehmer et al., 2004]. They look at the dealer-book cohabitation from the two angles: can a dealer exploit the information conveyed by the book and his order flow, and what are optimal actions of different trader types in a market with a strategic dealer.

Another real-world "wrinkle" that can be introduced into the market making equation is payment for order flow. Essentially, dealers can compete for brokers routing their clients' orders to a particular dealer. A market maker can agree to pay a small rebate (fraction of a penny usually) per share to a broker who sends him an order to execute. [Chordia and Subrahmanyam, 1995] develop a model, which incorporates this reality and perform some empirical tests as well. They examine the practice of dealers trying to lure customers in NYSE stocks off the exchange by paying for the order flow. The article concludes that such practice results in sub-optimal prices obtained for customers, which the authors explain with the lack of transparency in quotes dissemination. A game-theoretic article [Parlour and Rajan, 2002] arrives to a similar conclusion that payment for order flow results in wider spreads, which ultimately works to the investors' disadvantage. Finally, the competition for order flow between exchanges is studied in [Parlour and Seppi, 2003]. This work shows the possibility of co-existence of competing markets with different execution systems, which can compete for investors by offering either better prices or higher liquidity, and thereby attracting investors with different preferences.

D. Limit Order As a Source of Liquidity.

As Electronic Communications Networks are gaining in popularity and the NYSE and Nasdaq has opened up their order books to the public, the interest in limit order trading has risen considerably in the research community and the order book-level data has become available for scientific analysis. As a result, a large body of literature has sprung out recently on the order submission strategies and other aspects of trading in limit order markets.

The two seminal papers in this area are [Glosten, 1994] and [Bias et al., 1995]. The first paper details the advantages offered by an open limit order book: tighter spreads, convenience for small orders, sufficient liquidity, etc. The second work is a comprehensive empirical study of the order book behavior on the Paris Bourse (Euronext Paris currently). The authors investigate order submission strategies in that specific setting. Among their findings: "thin" books induce limit order submission, whereas "thick" books make market orders more likely; large spreads induce limit order submissions, tight spreads – market orders, and so on.

The most pertinent limit order theories for our research are those that provide us with models of optimal order submission: given the trader's preferences and the state of the order book, should he submit a market order, a limit order, or refrain from trading altogether? And if he chooses to submit a limit order, then how "aggressive" should this order be – i.e. how close to the opposite side of the market should it be positioned? These are not easy questions to answer, and we are not aware of an all-in-one comprehensive model, which can produce a recipe for the optimal strategy.

Most publications tackle some portion of this large decision problem. One of the main trade-offs inherent in limit order trading is time-to-execution vs. price improvement: the further from the inside market is the submitted order, the better is its price compared to the current price level, but the longer is the expected time for this order to get filled. A probabilistic approach to quantifying this trade-off is presented in [Foucault et al., 2003a] and [Harris, 1997]. An alternative method is suggested in [Lo et al., 2000], where the price is modeled as a stochastic process, and time-to-execution becomes the first time this process passes through the specified price level. The information effect of limit order trading (similar to the idea from the dealers' market setting) is captured in [Hollifield et al., 2003a] and [Hollifield et al., 2003b] – this model outputs the optimal order placement strategy, given the trader's preferences. This approach can be particularly useful for an automated market maker, since it helps the dealer decide if the book needs additional liquidity. [Sandas, 2001] goes even further by deriving not only the optimal price level, but also the optimal quantity, for a given state of the book and trader's preferences. Another paper, which incorporates the depth of quote into the dealer's decision process, is [Kavajecz, 1998].

A game-theoretic take on order submission and price dynamics is presented in [Parlour, 1998], [Goettler et al., 2003], and [Goettler et al., 2004] – authors derive an equilibrium for strategic market participant and examine the implications of their solutions. [Rosu,2004] also proves an existence of Markov equilibrium and uses his model to derive an expected shape of the limit order book and to explain the book's behavior following a large buy or sell order. [Smith, 2001] and [Vayanos, 2001] offer somewhat different approaches to the order submission problem: the earlier paper incorporates stock-specific factors and uses data from actual traders to calibrate the model, whereas the later work models the market as a group of strategic traders with private endowments, who are trying to efficiently re-distribute the risk among themselves. [Back and Baruch, 2004] compare a dealer market, a limit order market, and a uniform-price auction from the utility/surplus point of view, and conclude that the limit-order market and a uniform-price auction are actually equivalent.

Another broad category of limit order markets research is the empirical search for the optimal limit order price. These studies are the closest in spirit to this thesis. A number of papers attempt to derive the optimal order submission strategy in different markets: [Chordia et al., 2002], [Hasbrouck and Harris, 1996] and [Kaniel and Liu, 2004] examine traders' behavior on the NYSE, [Griffith et al., 2000] look at the Toronto Stock Exchange, while [Barclay et al., 2002] and [Ronaldo, 2004] study the ECNs. All these papers are similar to our work, since they set out to determine what orders do (or should) traders submit under different conditions: buys tend to follow buys, sells follow sells, small spreads induce market orders, large spreads – limit orders, transaction prices tend to revert after large moves, and so on. [Bloomfield et al., 2003] is a different look at the same problem: they use a simulated order-driven market to determine what kind of orders informed and uninformed traders submit. While in the most of canonical microstructure theories it has always been assumed that informed traders use only market orders, this paper shows that informed traders are actually more likely to use limit orders – i.e. they become market makers and provide liquidity to the uninformed traders, since they know the "right" price.

As it was the case in the dealership markets, it is always tempting to apply basic statistical methods to vast amounts of information generated by limit order markets to look for trends and dependencies. [Bouchaud et al., 2002], [Coppejans and Domowitz, 1999], and [Wang, 2002] are some examples of this approach. Some of their findings are useful for automated trading in order-driven markets – i.e. the shape of the order book, the effects of "peaks" and "holes" in the book, etc. While searching for patterns in prices is certainly not the goal of this research, an interesting study – [Kavajecz and Odders-White, 2002] – suggests that there can be a meaningful link between the technical analysis and market microstructure theories. For example, statistical analysis can help locate liquidity in order books, and market microstructure can explain what this means for likely future price movements.

Since limit order trading is a very rich domain, there are many other individual topics that warrant our attention. Because liquidity provision is pivotal to the limit order trading, quantifying liquidity becomes an important issue. [Irvine et al., 2000] suggest a simple yet effective measure of the "cost of roundtrip trade" – i.e. how expensive is it to buy and immediately sell a certain number of securities in a given order book. Limit orders have to be monitored, or face the risk of being "picked off" as soon as new information reaches the market. The way to quantify the cost of monitoring is suggested in [Foucault et al., 2003b]. [Handa and Schwartz, 1996] and [Handa et al., 2002] examine the profitability of limit order trading strategies and price formation in these markets. [Ahn et al., 2001] and [Degryse et al., 2001] study the effects of respectively volatility and large market orders on the behavior of the order book.

One more branch of the limit order trading research is the comparison the role of limit orders in different real-world markets: [Chung et al., 1999] look at the limit order's influence on the NYSE bid-ask spreads, and find out that this type of liquidity is responsible for the large portion of quotes. [Hasbrouck and Saar, 2002] conduct a similar study for the Island ECN, paying particular attention to volatility.

E. Liquidation and Other Automated Strategies.

The central part of our research and the topic of the majority of publications surveyed above are the actions of any single trader "against" the market as a whole. The inherent problem with this approach is that trading conditions – order arrival rates, proportion of "informed" traders, volatility, supply and demand, etc. – do not stay constant. Therefore, if a very good model is calibrated using a certain set of past data, and if trading environment changes dramatically in the future (imagine a period of a market-wide sell-off, for example), then an automated trader will do poorly, potentially losing large sums of money. This reality, on one hand, can be interpreted as an indication that electronic trading agents are destined to be tools for human traders rather than their replacements. On the other hand, however, it may be possible to make an automated trader recognize the change in market conditions and adjust his strategy accordingly.

While the universal solution is far from being evident, a solid first step is to study a trader's performance within a mix of trading strategies.

We are mostly interested in strategies that can be easily automated to facilitate testing and parameter selection. The most popular strategies of this kind are those dealing with either acquisition or liquidation of large blocks of securities. Buying or selling a large quantity in one shot will move the price adversely, making trading prohibitively expensive. Therefore large blocks are often broken up into small pieces and are bought or sold over longer periods of time (several days, for example). What is the optimal way to break up a block and what is the best timing for transactions is a subject of a significant number of publications. We will frequently refer to this research area as *optimal execution*; [Kissell and Glantz, 2003] and [Best Execution, 2002] provide a comprehensive overview of the issues involved. [Bertsimas and Lo, 1998] use dynamic programming to construct the most favorable schedule for large transactions. Their model can also be expanded to a portfolio of securities – see [Bertsimas et al., 2000]. [Huberman and Stanzl, 2001] expand this work by introducing transaction costs, more complex price impact functions, and continuous trading. [Harris, 1996] analyzes the problem of best execution from the agent-principal perspective: what are the broker's incentives to really find the best available price in all markets for their clients.

Two papers – [Almgren and Chriss, 1999] and [Almgren and Chriss, 2000] – bring together the concepts of liquidity and a more widely accepted notion of price risk. The basic idea is when trying to calculate the risk of holding a particular security (trying to estimate possible loss of value over some time horizon) or simply performing a valuation, practitioners often disregard the market microstructure aspect of securities markets. For example, valuing a significant holding at a single observed or expected price is simply not realistic, since the liquidation of this holding will inevitably move this price adversely. This reality must be included in all related calculations. Two other strategies for asset liquidations are presented in [Dubil, 2002] and [Krokhmal and Urysaev, 2003]; both use theoretical models and quantitative methods to determine the optimal trading schedule. Finally, [Weber and Rosenow, 2003] examine the problem of efficient execution within a pure limit order market: they look at the price impact of trades and at how a limit order book "re-fills" itself following a large market impact.

Another group of strategies, very similar in spirit to those mentioned above, are so-called VWAP (Volume-Weighted Average Price) strategies. The idea is to complete a large transaction over a given period of time in such a way that its average price will be equal or more favorable than the average price in the entire market over the same period. A simple example – if you break a large block into a large number of small pieces and sell them off at regular intervals over one trading day, you are virtually guaranteed to achieve market VWAP over that day. In practice, however, it's much more difficult, since the more pieces you create, the more will you have to pay in commissions, among other complications. For an overview of these strategies, consult [Madhavan, 2002], and for a theoretical analysis of competitive VWAP algorithms see [Kakade et al., 2003].

F. Computer Science Community Contribution.

In this thesis, our goal is to pull together past findings and useful tools from both Finance and Computer Science to combine and leverage the strengths of these two disciplines. Up to this point, however, the overwhelming majority of literature reviewed came from the Finance side. This does reflect the reality that Computer Science publications on in this area are scarce, maybe because this particular domain is considered "too applied". Indeed, many of the CS papers use the electronic trading domain as a setting to test some algorithm rather than a problem that requires a solution on its own (the way it's treated in Finance). We would like to reconcile the two approaches.

The first publication on automated market making – [Hakansson et al., 1985] – is a testimony that the interest in creating an electronic trader has been around for two decades. The authors created a simple agent with a single goal of "demand smoothing" – acting as a counterparty when misbalances arise on either buy or sell side. They tried an array of "rules" that a market maker should follow, and some of their experiments are fairly insightful. The bottom lime of this work is not very surprising – any "hard-coded" (non-adaptive) set of rules is bound to fail sooner or later on a set of new data that it has not been specifically tailored to.

A number of very interesting recent publications on electronic market making came out of the MIT AI Lab – particularly [Chan and Shelton, 2001], [Kim and Shelton, 2002], and [Das, 2003]. While all three papers describe an implementation of an electronic dealer, there are significant differences between their approaches and ours. They test their strategies in simulated markets, as opposed to using a real-world price feed, and they only use market orders, not employing limit order books. They rely mostly on the information-based models, similar to those in [Glosten and Milgrom, 1985], which are described in detail in Section B of Market Microstructure chapter. These studies implement both the analytical solutions of the financial theories and reinforcement learning algorithms that can be trained on past data. The later approach is similar in spirit to our own use of Reinforcement Learning, but our ultimate goals differ significantly.

Simulated markets is probably the strongest contribution of the CS community to date. Pitting trading agents in a controlled competition has always been an exciting event and a great motivation for research advancement. One of the first steps was taken by the Santa Fe Institute Double Auction Tournament [Rust et al., 1994]. Trading Agent Competition (TAC) – [Wellman et al., 2003] – is a more recent and still extremely popular event. While these markets are highly stylized and based on a variety of auction mechanisms, the simulation of financial markets has also been attempted: see [Poggio et al., 1999] and [Kearns and Ortiz, 2003]. We have used the later simulated environment during the initial stages of our research.

Finally, we should mention a significant overlap between the Computer Science and Statistical work. For example, [Papageorgiou, 1997] is very similar to statistical time series studies discussed earlier, even though it's a product of MIT AI Lab. Another example is [Thomas, 2003], where the author combines technical analysis with the news-based trading. While we are not explicitly looking for price patterns in our work, there may be a link between the technical analysis and market microstructure (as described in [Kavajecz and Odders-White, 2002]), plus the effect of unexpected news on stock prices is undeniable.

Similar to classical statistical techniques, like time-series analysis, we employ several Machine Learning algorithms to derive optimal trading strategies based on available data. There are volumes written on the subject, but we use only very generic ideas (as described in [Steward and Russell, 1998] and [Mitchell, 1997]) and adapt them to our specific problems.

To conclude, we would like to combine the wealth of theoretical models produced by financiers with the powerful tools and experimental platforms created by computer scientists to conduct a thorough investigation of automated trading domain and, hopefully, produce some applications that can be effective in real-world financial markets.

Experimental Setup

This chapter is based primarily on material from [Kearns and Ortiz, 2003], [Nevmyvaka et al., 2003], and [Nevmyvaka et al., 2005]. It is dedicated to a more detailed description of the limit order trading over ECNs, the functioning of the exchange simulator that we use for our experiments, and the specificities of the automated trading in this environment.

A. ECN Trading.

Although we have already touched upon the trading process on Electronic Communication Networks in the Research Background chapter, we take a closer look at it here. To recap: a market order is an instruction from a client to the dealer (or to the exchange) to buy or sell a certain quantity of stock at the *best available* price – i.e. "buy 100 shares of MSFT at the best price available right now" – whereas a limit order is an instruction to buy or sell a specified quantity of stock at a *specified or more advantageous* price: "sell 100 shares of MSFT at $25.53 or higher". Therefore, market orders guarantee execution of customer's transaction, but not the price at which such transaction will occur, and limit orders guarantee a certain price, but transaction may never happen.

Island ECN is a purely electronic market, which only uses limit orders and employs no designated middlemen. All liquidity comes from customers' limit orders that are arranged in order books (essentially two priority queues ordered by price) as shown in Figure 1a (limit price – number of shares). As we can see, the orders are arranged from most aggressive to less aggressive – increasing in price in the sell book, and decreasing in the buy book. Top of the buy book (the highest priced buy order) and top of the sell book (the lowest sell) are called *bid* and *ask* respectively, and together they form the *inside* *market*, and the distance between the two is called the *bid-ask spread*. In our case, the bid is \$25.21, ask – \$25.30, and the spread is \$0.09. The *depth* at the bid is 200 shares, depth at the ask – 150 shares. The *tick size* is \$0.01 – a penny.

(a)		(b)	(c)
…		…	…
25.56 – 300		25.56 – 300	25.56 – 300
25.55 – 1000	Sell Orders	25.55 – 1000	25.55 – 1000
25.35 – 200		25.35 – 200	**25.35 – 100**
25.30 – 150	← Top sell	25.30 – 150	
Inside Market			
25.21 – 200	← Top Buy	25.21 – 200	25.21 – 200
25.19 – 300		**25.20 – 1000**	25.20 – 1000
25.15 – 785	Buy Orders	25.19 – 300	25.19 – 300
25.10 – 170		25.15 – 785	25.15 – 785
…		25.10 – 170	25.10 – 170
		…	…

Figure 1.

Now let's look at how the orders execute. If a new limit order arrives, and there are no orders on the opposite side of the market that can satisfy the limit price, then such order is being entered into the appropriate order book. For example, in Figure 1b, a new buy order for 1000 shares at \$25.20 or less has arrived, but the best sell order is for \$25.30 or more; thus no transaction is possible at the moment, and the new order gets entered into the buy queue according to its price. Now, another buy order arrives for 250 shares at \$25.40 or less. This order gets transacted (or *crossed*) with the outstanding orders in the sell queue: 150 shares are bought at \$25.30 and another 100 shares are bought at \$25.35. The resulting order book is shown in Figure 1c – the partially transacted sell order at \$25.35 is now at the top of the book. This shows that even though there are no explicit market orders in pure electronic markets, immediate and guaranteed execution is still possible by specifying a limit price that falls inside the opposite order book. Such orders are also referred to as *marketable limit orders*.

Note that he first order in (b) added liquidity to the buy book for incoming sell orders to cross with; whereas the second order in (c) removed liquidity by transacting with outstanding sell orders. Also, we can see that a market order, while removing liquidity, also increases the bid-as spread – in (c), the spread increases from \$0.09 to \$0.14. We will discuss the implications of this property later on.

All crossing on ECNs is performed electronically by a computer respecting the price and time priority, without intervention of any intermediaries. Since all the orders are being processed electronically, it is easy to establish the exact sequence of order arrival. As a marketable order arrives, it gets executed “instantaneously” and the execution price is being publicly reported as the *last price*. It is important to remember that such price can be fairly “stale” (outdated) in less liquid stocks, where transactions don’t happen very often.

These are the founding principals on which our market simulator is built.

B. The INET ECN.

The main idea behind the simulator is to merge the real-world orders from the INET ECN (a product of a merger between Instinet and Island ECNs) with the artificial orders produced by electronic simulated agents. There are essentially two ways to trade Nasdaq stocks: first, a trader can send his order to a designated dealer and trade either at the dealer’s quote, or let the dealer cross the order with a more competitive outstanding limit order (this process is described in the Research Background chapter). Second, a trader can trade directly with other traders over an Electronic Communications Network. ECNs cover the same Nasdaq stocks as human dealers at investment banks, but display only orders that were submitted to that ECN. It is apt to think of ECNs as separate electronic exchanges that bring together traders who want to transact directly with one another in Nasdaq stocks. In fact, several ECNs have applied for exchange status.

On the positive (for a trader) side, traders pay less commissions by bypassing the brokers and transacting directly with other traders, but on the downside, if there is a better quote posted by another dealer, or a more competitive limit order submitted through a dealer, they will not be displayed on the network, which makes trading more expensive by not showing the best prices to traders. This does not work the same the other way – the dealers can see *both* all the quotes and orders in the dealers network, and all the activity in all ECNs. This gives them a certain advantage, since they can always transact at best prices (human dealers can trade via ECNs).

Which one of the two effects – lower commissions vs. not always seeing the best prices – dominates the other is not clear, and probably depends on the trader's preferences, but it is a fact that ECNs are successfully competing with the Nasdaq OTC dealer network, by capturing a significant volume of all Nasdaq trading. For example, INET executes one out of five trades in Nasdaq stocks ([Harris, 2003]). The other major ECN is Archipelago.

There is one important aspect of using the INET data that we have to be mindful of: since INET is just an alternative trading system, and not the main exchange for Nasdaq stocks, it is possible that price discovery does not take place on INET, but in the larger inter-dealer Nasdaq. We will discuss in detail what price discovery means in the Market Microstructure chapter, and will further discuss the practical implications of this issue for our empirical results.

From the technology standpoint, INET accepts and disseminates orders through a number of common brokerage routing systems and also provides two APIs to its non-brokerage clients. These two protocols – called ITCH and OUTCH – provide data feeds and allow order placement respectively. One INET's service, which traditional brokerages find difficult to match, is the ability to execute automated trading strategies through the OUTCH API. Island charges its customers a relatively small subscription fee, and applies per-trade *fees* and *rebates*. If a trader submits a limit order, which executes against an outstanding limit order (a marketable limit order essentially) he is charged a small fee for each share transacted (currently the fee is $0.003 per share). And if a trader submits an order, which gets entered into the order book first, and then transacts with an incoming marketable order, such trader gets a rebate (currently $0.002 per share). For more details, see [INET ECN]. Simply put, a trader who removes liquidity from the book pays a fee, and the trader who provides liquidity gets a rebate when this liquidity is used.

This discussion is not solely aimed at familiarizing the reader with the functioning of the INET ECN, but also has important implications for automated trading analysis and development. Traditional dealership models treat trading profits as the only source of trader's income. The reality is such, however, that limit order traders can be liquidity providers and thus are in great position to take advantage of these rebates offered by ECNs. Depending on the volume that flows through such "ad-hoc" dealer, the cash flow from rebates can be substantial to the point that the trader can take trading losses (or make zero trading profit) and still come out ahead by providing a lot of liquidity to the market and being remunerated for this service. In other words, under the standard interpretation of the market making activity, the dealer is getting paid for his services by the impatient trades via the bid-ask spread, but in the real world, the exchanges also value these services and reward market makers through the above rebate system. While the later source of profits is deterministic (number of shares transacted multiplied by the rebate rate), and thus less interesting from the academic standpoint, it also influences the non-deterministic trading component, and so it must be taken into account. One of the obvious implications is the one just mentioned – an electronic market maker can follow an inherently money-losing strategy and still remain in business by being supported by rebates.

The other INET protocol – ITCH – provides transaction-level data at the nominal cost, which allows its users to reconstruct the state of the limit order book at any point in time. This service is necessary for those traders who want to analyze the past behavior of prices, order flow, spreads, etc. This feed, and more importantly, the historical data generated by it, is the foundation for our simulator, which we describe in greater detail below.

C. Data.

The data available to us is the output of the ITCH protocol: every transaction (with a timestamp) that happens on INET in a particular stock. Here is a small sample of our data after some pre-processing to make it easier to handle:

Code	Ticker	Price	Size	UID	Tstamap	Telapsed
A	AMZN	20.2	100	3845588	49239515	0
X	AMZN	20.18	1000	3842201	49240018	12356
P	AMZN	20.209	-21	3846349	49243136	5842
X	AMZN	20.21	-33800	3840077	49243861	24499
A	AMZN	20.25	-500	3846581	49244016	0
E	AMZN	20.2	100	3845588	49244872	5357

Table 1.

Code signifies the action to be taken:

A – add a new order to the book

X – cancel a previously submitted order

E – execute a standing (previously submitted) order

P – execute "hidden" volume

All these actions should be self-explanatory, with an exception of P: traders on INET are allowed to submit "hidden" orders, which get entered into the order book, but are not displayed to other market participants (and thus are not included into the ITCH feed). The presence of such orders is manifested only when an incoming marketable order gets executed against them. Hidden orders also lose their time priority – i.e. for a given price level, hidden orders can execute only after all visible orders have been executed.

As described earlier, every order comes with a limit price and a corresponding size. For cancellations and executions, the limit price of a standing order is specified along with the number of shares cancelled or executed, since partial transactions are allowed. Every new order (code 'A') comes with a unique id, while ids for execution and cancellations reference the standing order that is being modified (see first and last lines in Table 1). Every action is time-stamped, and we also keep track of time that has elapsed between an order being submitted and executed or cancelled.

While historical ITCH files report all events in all stocks for a single trading day in one very large file, we break out stocks of interest into separate text files in a format of Table 1. These single-stock plain-text files are still very large – an average of about 50,000 lines and size of 2MB. This once again underscores the fact that events take place at a very high frequency, and we routinely see multiple orders being submitted, executed, and cancelled within 1 second.

While we have records of all stocks traded on INET, we break out and use for our research only a small number of representative equities, which vary across industries, activity levels and trading characteristics. A sample of stocks we look at is summarized in Table 2 (trading volume as of July 2004).

Symbol	**Company**	**Daily Volume (shares)**
MSFT	Microsoft Corp.	14 mil.
QCOM	Qualcomm	4 mil
NVDA	NVIDIA Corp.	1.3 mil
MERQ	Mercury Interactive	0.6 mil

Table 2.

For the stocks we investigate, we use daily records from January 2003 to July 2004, which amounts to about 360 trading days.

D. The Simulator.

Having described the data available to us, we can move on to describing the simulator we have created in order to run controlled experiments on historical transaction records.

Since in our files every action is time-stamped, and all transactions are recorded electronically, the exact sequence of order flow is unambiguous. These features allow us to precisely reconstruct buy and sell limit order books at any point in time. We start at the very beginning of a trading day with two empty books and then add new orders into appropriate books, as we read commands from the file line-by-line. In constructing the limit order books, we maintain price and time priority just like in the real-world markets. We follow execution and cancellation directives just as they appear in the ITCH files. Through this continuous update process we always maintain order books correctly just as they looked like to actual traders on that particular trading day.

At designated times, we insert into the order flow some “artificial” orders that represent various trades we set out to investigate. We then perform all the executions and maintain orders priorities in the way described above. Such setup allows us to run “what-if" simulations in historical order book. The three panels of Figure 1 are representative of the kind of electronic matching process that we are performing.

In addition to book reconstruction and order matching, our simulator also performs record keeping: for every artificial order, we keep track of when it was submitted, at what price and size, with which other orders did it transact, when, at what prices, and what sizes. This allows us to keep track with precision of such microstructure dimensions as time-to-execution, expected executed volume as a function of price, average execution price, etc. We can then compare these measurements to some pre-determined benchmark and thus assign a "score" to a given strategy, which in turn allows for a principled way to distinguish "good" and "bad" results.

We investigate optimal order submission strategies by examining the pre-committed liquidity, which is essentially what limit order books represent. When the trader submits a limit order, he gives an option to the rest of the market participants to transact at a pre-specified price up to the order's size – i.e. other traders can "lift off" liquidity, while knowing ex-ante how much they are paying. This is crucial for our analysis – transaction costs in our model come from two sources: the bid-ask spread and price concessions as payment for liquidity. When a trader submits a market order, he has to first "step over" the spread and then pay increasingly disadvantageous prices the deeper in the opposing book he needs to reach to satisfy his liquidity demands. On the other hand, when he submits a limit order, he risks having the price move away from his order and then being forced to demand liquidity at the end of the trading interval. Historical limit order books allow us to quantify and compare these two dimensions of order submission.

Let's summarize the advantages that this setup gives us. This environment is realistic – it lets us test real-world trading strategies without putting real money at risk. Using full reconstructed limit order books adds extra complexity and realism when compared to simpler trades-and-quote datasets used in previous studies: first, there is statistical evidence that there is additional information contained in the limit order books, which can help predict future prices (see [Ranaldo, 2004], [Griffith et al., 2000], [Kaniel and Liu, 2004] among others). Second, we can investigate the order book dynamics and study how the liquidity in the market evolves with time and events. This setup is also relatively novel – very few studies use fully-reconstructed limit order books from real-world transaction data ([Weber and Rosenow, 2003] operate in a setting closest to ours technologically). Finally, this is a controlled environment – we can change traders' behavior and see the effects under the same market conditions for all cases.

All this being said, our chosen simulated environment unfortunately has its shortcomings. First, activity that we see in our data represents only the Island order flow and not the market for a given stock in its entirety (there is a number of other trading venues for Nasdaq stocks). We have mentioned earlier that this reality may distort our view of the price formation process. Fortunately, this problem is not particularly serious, since it does not invalidate our analysis. Maybe some particular number we output – optimal limit order price, for instance – can be biased due to the incompleteness of our data, but it is not the goal of this thesis to provide exact and definitive numerical answers. We propose a general method of quantifying and optimizing trade execution parameters, and our conclusions are unaffected by the kind of data that we use. When our methods are implemented in practice, however, it is certainly more advantageous to have access to (and thus base your decisions on) a consolidated order book as opposed to some fragment of the market.

A similar data-related hurdle is the fact that we cannot see hidden orders in the records that we use. By the same token, while this may bias the numerical output of our experiments, our models remain valid. By not including the hidden volume in the analysis, we are strictly overestimating transaction costs, since there can be some potential inexpensive liquidity that we don't see. Estimating the proportion of hidden orders to visible orders, and correcting our models' outputs by a single factor can alleviate this issue.

The most serious inadequacy of our simulator is that this market, based on historical order flow, cannot react to the arrival of artificial orders. For example, in the real world when a large limit order appears in the book close to the inside market, this may incite the other side of the market to transact more aggressively (by taking advantage of this "pocket of liquidity"); or the traders on the same side may be tempted to "step in front" of the resting large order and thus guarantee price protection for their own orders.

We do not see these types of behavior in our simulator. Therefore, we should be cautious when transferring our finding from paper to the real markets. Throughout our analysis, we will point out which aspects can be affected by this lack of realism. Ultimately, it may be possible to model these strategic games around the bid-ask spread within a simulated market, but this will add another layer of complexity, and will still fall short of the real-world interplay.

From the implementation standpoint, there are also difficulties. Dealing with full limit order books is more cumbersome, significantly slower, and requires much larger quantities of data than a more conventional setup that relies solely on transaction prices and prevailing market quotes.

Overall, we strongly believe that the strong points of our simulator greatly outweigh its shortcomings. Our artificial environment built around actual limit order books is very well suited for investigating optimal trading strategies in financial markets.

Market Microstructure: Process of Price Formation

In this chapter, we explain in great detail the concept of market microstructure – the ideological foundation of this thesis. We will start by introducing the Efficient Market Hypothesis (EMH) – a theory that explains how perfectly frictionless markets functions. It is also the linchpin of all modern Finance (more or less). Understanding how efficient markets behave is a very useful starting point, but we will show that empirical tests do not fully support EMH or CAPM (Capital Asset Pricing Model – another fundamental theory). This brings us to the existence of market frictions – realities that explain why and how actual markets' behavior deviates from the behavior of perfectly efficient markets. Market microstructure, as a branch of Finance, aims to understand, model, and quantify these frictions. We will enumerate them, explain where they come from, and then present a number of increasingly-complex models that capture this microstructure behavior. We will describe these models from the point of view of automation: how can they be applied to the development of electronic trading agents. While we had to mention some of these concepts in the Related Work chapter, here we will go into much more detail, and only present the issues of direct relevance to optimized electronic trading. Unless otherwise specified, the general discussion of market microstructure here is inspired by broad surveys of equity markets organization and dealers' behavior, such as [Stoll, 2000], [Stoll, 2001], [Madhavan, 2000], [Bias et al., 2004], [Harris, 2003], [Schwartz and Francioni, 2004] and others.

A. Efficient Markets: Myth or Reality?

1. Frictionless Environment

First, let us introduce a concept of *frictionless markets*, which is the environment in which the Capital Asset Pricing Model is derived. We assume that there are

- no taxes

- no transaction costs (i.e. commissions)
- unlimited short-selling (selling shares one doesn't own)
- unlimited borrowing and lending at some risk-free rate
- markets are perfectly liquid (buying and selling happens at a single price in unlimited quantities)
- all investors know the *distribution* of future returns of all assets (it is the same as saying that they have homogenous expectations)
- covariance among assets is constant and known to all investors

The above assumptions imply that all information about every security can be summarized by the following numbers: stock's expected return, variance of this expected return (a measure of uncertainty or riskiness), and the covariance of expected returns with other stocks. These numbers are also known as *return* (E[r] or μ), *volatility* (σ(r) or σ), and *beta* (β). Beta is actually a correlation between the stock's return and the return of a portfolio of all stocks in the market.

While the above view is undeniably oversimplified, it still provides us with several key insights that largely hold true even when above assumptions are relaxed:

- the nature of risk: it is the variance of the distribution of future returns
- the relationship between risk and return: in order to achieve higher return, one must assume more risk
- benefits of diversification: idiosyncratic risk of holding any individual stock can be reduced through holding a well-diversified portfolio (since correlations are assumed to be stable)

2. Portfolio Selection

One fundamental theory, which is derived in this environment and will be very useful in our market microstructure study, is the *Optimal Portfolio Theory* developed by Henry Markowitz (see [Markowitz, 1952]).

It is easy to agree that a rational (risk-averse) investor seeks to maximize his returns and minimize his risk. Figure 1 depicts a number of portfolios – we can think of them as individual stocks for now – and the investor tries to decide which one to chose. What can we say about choosing stock A? It is a suboptimal decision because each stock

B and C are strictly better than A: B offers a higher expected return than A for the same level of risk, while C offers the same return as A, but with much less risk involved.

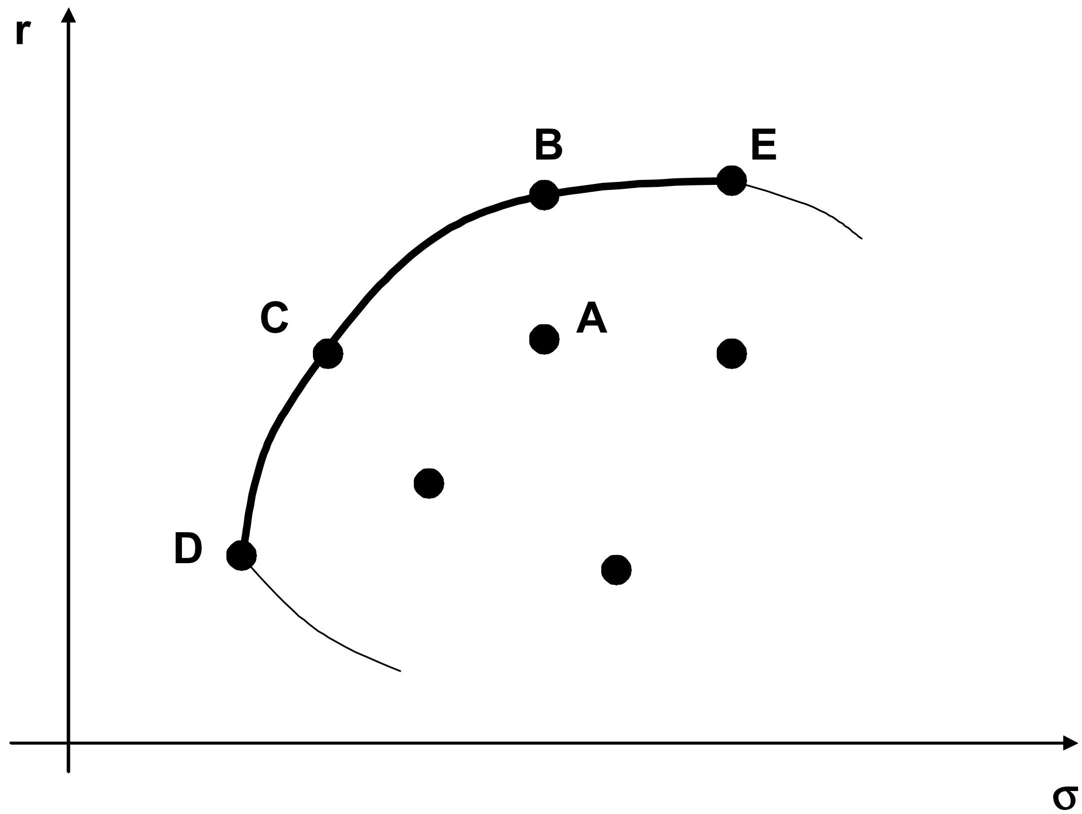

Figure 1.

It is more difficult to decide between B and C: B promises a higher return, but with more uncertainty, while C is less risky, but also returns less. It is the investor's risk preference (as embedded in his utility function) that allows him to choose between these two alternatives.

If we take this line of reasoning to other points (stocks) in Figure 1, we can see that it is suboptimal to select points not on the D-C-B-E arc. For any point inside the arc, we can always find another point with either higher return or lower risk; the later point will lie on the D-C-B-E arc. Points that comprise this outer arc share the following property: to move from one point to another you have to give up something in one category to improve in the other category. If you want to achieve higher return (moving to the right along the curve), you will have to assume more risk; and if you want to decrease your risk (moving to the left), you will have to settle for a lower return. Point D represents a stock with the lowest risk, and point E is a stock with the highest return. Rational investor will only select stocks along the D-C-B-E arc, which is appropriately called the *efficient frontier*.

(Note: in the above discussion, we assumed for simplicity that points on Figure 1 are stocks. In reality, an investor does not pick a single stock, but a *portfolio* of stocks. So, individuals stocks may be rationally situated inside the efficient frontier, but all optimal portfolios must lie on the frontier.)

3. CAPM

The second major advancement in Finance theory is the capital asset pricing model (CAPM). Sharpe, Lintner, and Mossin are credited with its creation – see [Sharpe, 1964], [Lintner, 1965], and [Mossin, 1966]. Derived from the efficient portfolio theory we just described, CAPM can tell us what a price should be for an individual security.

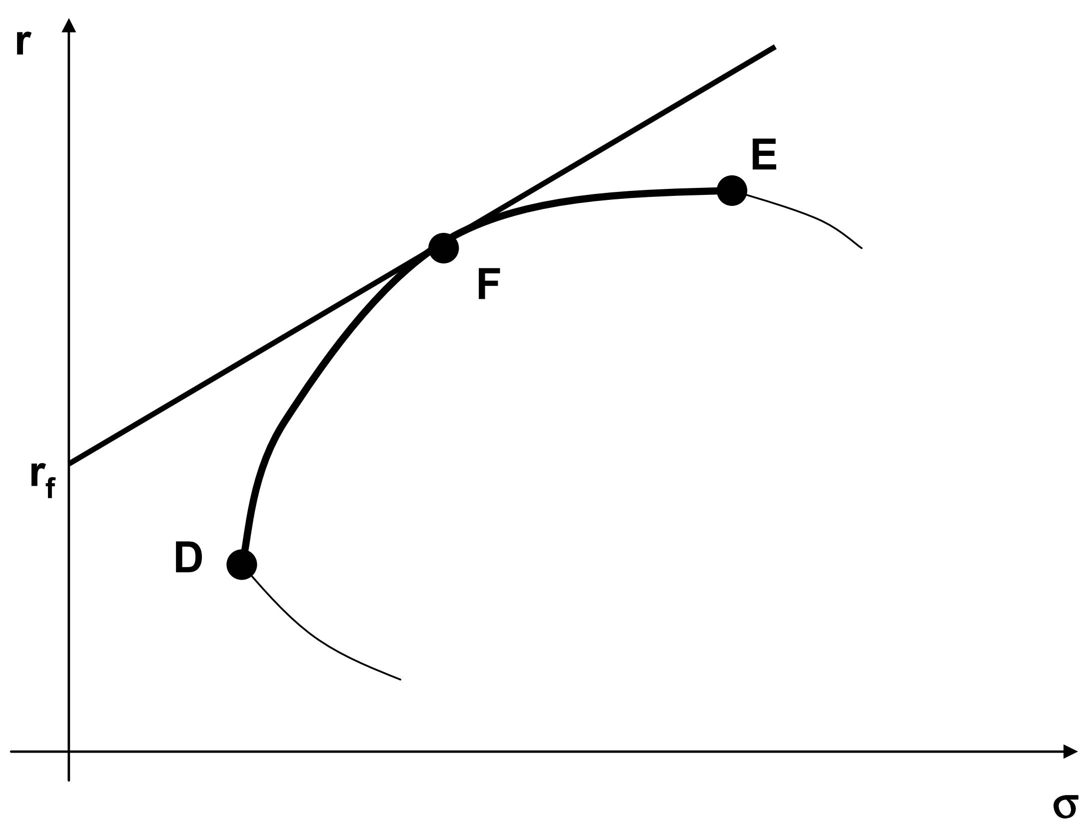

Figure 2.

In Figure 2, we introduce into our simple economy a risk-free rate r_f at which limitless lending and borrowing can be done. Now we have a risk-free asset, which provides some guaranteed return for no risk (think US Treasuries in the real markets). In our coordinate system, this asset is plotted on the vertical axis. The line r_f–F shows what kind of performance can be achieved by combining the risk-free security and portfolio F. On the extreme left, we are fully invested in the riskless asset (we are lending money at the risk-free rate); between r_f and F, we lend some money and invest the rest in the portfolio F; to the right of F, we are *leveraged* – we borrow extra money and invest everything in F.

We can observe from Figure 2 that the risk-return profile of the r_f-F line dominates the profile of the D-F-E arc – since r_f-F is a tangent – and thus all rational investors (we assume that everyone has homogenous expectations) will make their portfolio decisions along the r_f-F line, according to their risk preference. What do we know about portfolio F? When market is in equilibrium, all shares of all stocks must be held by investors, which makes F the portfolio of all stocks in the market. Line r_f-F is known as the *capital market line*. It can be described by the following equation:

$$r = r_f + [(r_{market} - r_f)/\sigma_{market}]\sigma \qquad (1)$$

This relationship shows that investor must be reimbursed with a return beyond the risk-free rate for assuming risk.

The relationship between risk and return for a single asset is derived following the same reasoning:

$$r_i = r_f + \beta_i\,(r_{market} - r_f), \qquad (2)$$

where $\beta = \sigma_i/\sigma^2_{market}$ is a measure of variability of the ith asset as compared to variability of the market portfolio. Beta of the risk-free security is zero; beta of the market portfolio is 1. Under the CAPM assumptions, all assets and all portfolios (including the market portfolio) must lie on the *security market line*, as shown in Figure 3.

One of the main insights of the CAPM theory is that a return on any asset is related to the market return, which allows us to separate risk into *systematic* and *unsystematic*. Unsystematic risk can be eliminated by a diversified portfolio, whereas systematic (or market) risk cannot. Also, because the price of every stock depends only on how its return covaries with the market portfolio, demand for an individual stock is horizontal – i.e. each stock has an *intrinsic* value above which everybody wants to buy, and below everybody wants to sell.

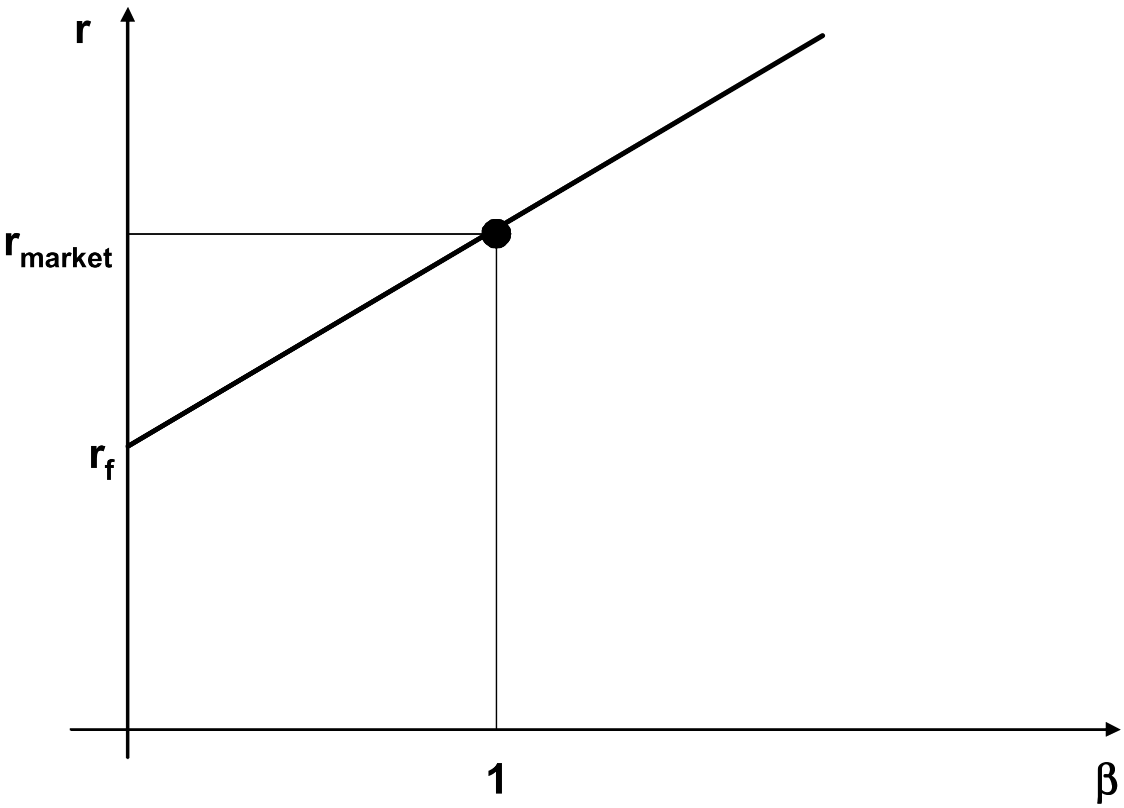

Figure 3.

4. Information Efficiency

Perhaps the most important (far-reaching) assumption that CAPM makes is the assumption that all investors have the same expectations about future return distributions and their variances. We will not discuss here whether it is reasonable to assume that investors have homogenous expectations, but we will look instead at one important dimension of the investors' decision-making process: how information gets reflected in securities prices.

To test the information efficiency of a market, we have to study if traders can make excess returns by trading on information. Markets are said to be efficient if excess profits cannot be made from the following information:

- past prices (weak-form efficiency)
- publicly-available information (semi-strong-form efficiency)
- all information – public, private, and past prices (strong-form efficiency)

The above statements constitute the *Efficient Markets Hypothesis* (EMH).

The weak-form market efficiency postulates that future prices cannot be predicted from past prices. If successive price changes are statistically independent and identically distributed, then we call such price process a *random walk*. [Fama, 1970] is a comprehensive survey of academic studies on random walks in capital markets. Here is a basic argument for why stock prices should follow a random walk – i.e. why they must not exhibit any persistent patterns. Imagine if a stock consistently went up in price during the month of January. In that case, investors would want to buy this stock in December to profit from the impending rise. However, the increased demand for the stock in December will drive its price higher, and the January appreciation will not happen. It is said that such pattern has been *arbitraged away* – it disappeared because many investors tried to profit from it all at once. In efficient markets, this is what happens to all patterns: they are supposed to disappear "instantaneously" because there are lots of people out there looking for them.

The above example means that all expectations about the future movements in the stock price are fully reflected in the current market price; therefore, in the future the price is as likely to go up as to go down, which makes the price process a random walk. The EMH is a very important insight into financial markets and a starting point for our analysis: EMH tells us how prices must behave in perfectly frictionless markets, any deviation from this behavior is an evidence of operational inefficiencies in the real-word markets, which is precisely what we have set out to explore.

Very similar arguments can be developed for both semi-strong and strong-form efficiencies. Is it profitable to trade on the newly-released information? Can one make outsized profits by buying a stock following positive news, and selling it following negative news? In efficient markets, the answer is a definite "no" – most news are anticipated by market participants and are already impounded into the stock price when the announcement is being made; when the development is unanticipated, the price will adjust to it "instantaneously", making consistent profits impossible. Under the strongest form of EMH, even the non-public information gets reflected in stock prices through anticipation and continuous competitive financial analysis.

5. Implications and Tests

So, what does the Modern Finance Theory – frictionless markets, optimal portfolio selection, capital asset pricing model, rational expectations, efficient market hypothesis, etc. – tell us about trading and investing? It tells us that it is hopeless to try to anticipate future prices, "beat the market", or "time the market", which means that technical analysis doesn't work and fundamental analysis should yield "normal" profits at best (since very many people are engaging in it at the same time). Furthermore, stocks should be priced just like equation (2) tells us. In practical terms, if an investor believes that the above conclusions are true, he should invest in broad-market mutual funds (index-funds) and low-risk bonds – i.e. select a point on the market line in Figure 3, according to his risk preference, and then get an appropriate mixture of the market portfolio and the risk-free asset. The only decision variable for traders is how much to buy or to sell – all transactions happen at prices determined by the "invisible hand" of market forces.

But is there support for these theories in the real world? In a word: yes. *A Random Walk Down Wall Street* [Malkiel, 1999] is the best-known book that defends the ideas presented above. It appears that over a long run (months and years) stocks indeed tend to behave like random walks, not exhibiting any repeating statistically-distinguishable patterns. Since pattern-finding is the domain of technical analysis, numerous academic studies were performed to see if following technical rules can produce outsized profits. The overwhelming majority of technical analysis techniques cannot withstand academic scrutiny – see [Malkiel, 1999] for a complete survey. This means that there is strong evidence that it is very difficult (read impossible) to predict future prices from past records, which is consistent with the weak-form market efficiency.

There is also support for the stronger forms of information efficiency. A seminal paper [Fama et al., 1969] shows that investors correctly anticipate positive news: stocks are as likely to go up as down following stock splits. There are also numerous studies showing that active managers of mutual funds tend to underperform passive managers (i.e. index funds) – see [Sharpe, 1966] and [Jensen, 1969].

But not all the evidence is on the side of efficient markets. CAPM, for example does not do so well in the face of empirical evidence. [Fama and French, 1992] investigate if stocks' beta is indeed predictive of future returns (remember that high-beta stocks are expected to return more because they are also riskier). Their findings look similar to Figure 4. This is markedly different from the market line in Figure 3.

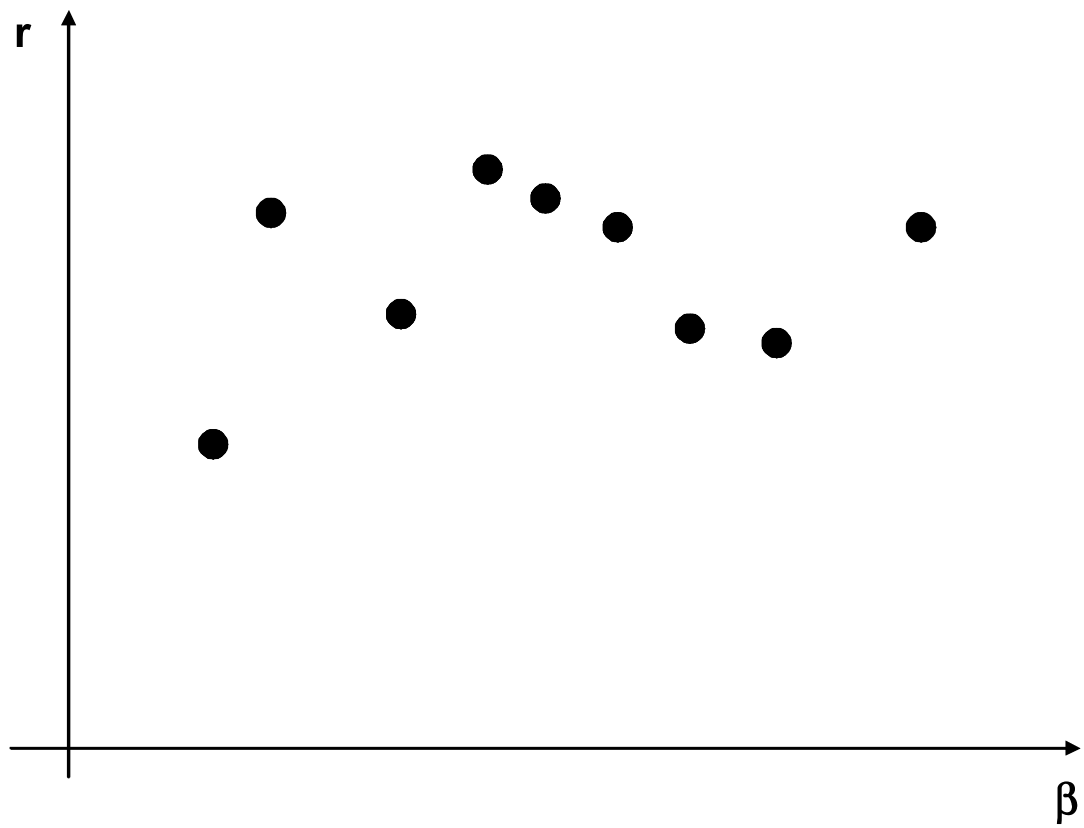

Figure 4.

There is no visible correlation between beta and returns, despite what the theory predicts. There are several other studies that arrive to the same conclusion.

There are also limits on the extent of random walk behavior in stock prices. An extensive critique of the EMH is provided in [Lo and MacKinlay, 2001]. The authors perform extensive tests to check if prices indeed follow random walk, and they reject this hypothesis in many cases. They also suggest specific cases when EMH does not hold (index arbitrage, for example). There is a very large body of evidence that suggest that prices are not random over short time periods – i.e. intraday. Other prominent works that question efficient pricing in markets are [Modigliani and Cohn, 1979] and [Schiler, 1981]. Furthermore, even though it is looked down upon by the academics, the technical analysis is alive and well, and practitioners continue claiming that there is money to be made through charting and other techniques.

Finally, it is commonly accepted that new information does not get impounded into market prices instantaneously – it is a dynamic process. For example, [Dann et al., 1977] show that when large block trades are broken up into smaller pieces and are transacted sequentially, trading rules can be formulated to take advantage of predictable price impacts. There are more recent studies with the same conclusion. For obvious reasons, predictability in this environment is very short-lived.

6. Third Dimension.

The bottom line of the above discussion is the following: over long horizons, markets do appear efficient and investors maybe better off investing in index funds and government bonds instead of trying to pick "winning" stock or attempting to "time the market". But the assumptions of rational frictionless markets tend to break down in the short run: investors have divergent expectations and market mechanics come into play. Also, the proponents of efficient markets claim that mispricings will be immediately arbitraged away by the profit-seeking traders, but they fail to specify how to identify such mispricings.

In our view, the main reason why classical theories tend to break down is because they are concentrated solely on two dimensions of financial markets: risk and return. They correctly conclude that in order to achieve higher returns, one must assume more risk. But there is also the third very important dimension: liquidity. It deals with prices at which trades can actually take place, as opposed to reported transaction prices from the past. This is the missing link that can reconcile the risk-return methodology with the empirical evidence.

Expected return and variance of an asset can help assess the investment's performance, but liquidity must be analyzed as well, since it affects prices at which this investment can be made, and they, in turn, affect expected returns and variances. Because of this connection, in this thesis we try to separate the trading process from the investment process – during the trading phase, the goal is to attain the best prices under the high-level investment strategy. While this is a helpful concept, which allows us to concentrate only on liquidity and leave out other considerations (portfolio selection, for example), in the real world trading and investment decisions should be a tightly-coupled project, since transaction prices (trading) influence stocks' returns and risk (investing).

B. Market Frictions

In this section, we will specifically point out real-world realities and aspects of markets operations that prevent prices from being perfectly efficient. We call these "wrinkles" *market frictions*, and the area of Market Microstructure is devoted to understanding, analyzing, and modeling them.

1. Dynamic Price Discovery

In frictionless markets, all participants know the future distributions of returns and their variances. The market is said to have homogenous expectations. Since everyone agrees on possible future prices, there is a single current price, at which any number of shares can be transacted, and investors trade only to accommodate their risk preferences. In the real world, investors have *divergent expectations* – they do not always agree on future returns distributions. There are many reasons why market participants can disagree: the number of factors that can affect stock prices is extremely large, and different people can assign different weights to these factors and can generally have different interpretations of causes and effects. In other words, there are very few situations, where there is a single conjuncture that "all rational individuals" can agree on, instead we have a number of competing logically coherent hypothesis, and the market becomes a voting mechanism that decides which scenario is perceived as the most likely one.

In a market with divergent expectations, there is not a single price that everybody agrees upon and that can accommodate any number of shares. It is said that the true underlying price has to be "discovered". Indeed, the process of *price discovery* is the central topic in Market Microstructure, and it is mentioned often in this thesis. By its nature, price discovery is a dynamic process where traders with divergent valuations negotiate in the marketplace. This is the most fundamental reason why short-term prices differ from what is predicted by EMH.

One of the visible manifestations of this on-going dynamic negotiation process is the heightened short-term volatility. Simply put, volatility observed in equity markets appears to be much higher than can be explained by factors that affect the "true value" of the traded security. Prices bounce around literally non-stop, while new information relevant to any particular stock arrives only periodically and much less frequently then price changes. This effect is especially pronounced over short time periods (intraday), and volatility settles down to expected levels over longer horizons (months, years). See [Ozenbas et al., 2002] for specific numbers. There are other explanations for this phenomenon, as we discuss later, but divergent expectations are the main contributor.

Second, not only do traders have heterogeneous expectations, but they are also differently informed. It is a fact that some investors can have a better-quality information about the underlying value of a stock. This could happen because of a personal experience (industry executives, company insiders), superior analysis (access to private comprehensive databases, superior analytical tools, etc), or private information (un-released news). Note that most of the sources of superior information above are not considered illegal and thus can be acted upon. This reality turns financial markets into yet another source of information dissemination: better-informed traders will try to take advantage of their knowledge by buying or selling stocks, which will move prices closer to their equilibrium level. Every trade can be interpreted as a probabilistic signal about the value of a security, and uninformed traders can make inferences about private information by observing the order flow. Informed traders, in turn, will be strategic in their trading in order to keep their edge for as long as possible. And again, we end up with a dynamic process which does not fit into the EMH framework, and makes prices diverge from the random walk behavior.

2. Bid-Ask Spread and Other Payments for Liquidity

We mostly talk about the *bid-ask spread* as the dealer's livelihood – a compensation for providing liquidity to the market. While there is no place for the spread in frictionless markets, [Cohen et al., 1981] show that any market (short of perfect) must have a spread with or without a dealer. Spreads exist as a payment for immediacy that a market order receives. Here is a basic intuition why a spread must exist: a market order has a zero risk of non-execution, but a limit order submitted infinitesimally close to the other side of the market has a risk of non-execution that is not infinitesimal, since market can always move away from such order. Therefore, a trader submitting a market order must pay more than a limit order trader, and it's the bid-ask spread that captures this payment for immediacy.

How does the presence of the bid-ask spread influence transaction prices and volatility? Most of the seller-initiated trades happen at the bid or close to the bid, most of the buyer-initiated trades happen close to the ask. Since both buyers and seller are present in the market simultaneously, this creates a so-called *bid-ask bounce*: trades happen alternatively at the bid and at the ask, which induces negative autocorrelation in transaction prices (uptick is more likely to follow a downtick than an uptick) and accentuates short-term volatility (prices appear more volatile then if all trades were to happen at the mid-spread). Such behavior is yet another source of deviation from random walk.

As we have described in the Background chapter, both quotes – bid and ask – come with a certain volume attached to them. This volume indicates how many shares liquidity providers (dealers or limit-order traders) are willing to buy at the bid and sell at the ask. If the impatient trader wants to trade up to the quoted volume, then the spread is the only cost he has to pay for liquidity, but if he demands more volume than quoted at the best bid or offer (BBO), he must accept further price concessions – i.e. transact at inferior prices to compensate other traders for providing additional volume.

To illustrate this mechanism, in Figure 5 we reproduce a familiar snapshot of a hypothetical limit order book.

…	…
25.56 – 300	25.56 – 300
25.55 – 1000	**25.55 – 1000**
25.35 – 200	**25.35 – 200**
25.30 – 150 ← Best sell	**25.30 – 150** ← Best sell
25.25 - mid-spread	25.25 - mid-spread
25.20 – 200 ← Best buy	25.20 – 200 ← Best buy
25.19 – 300	25.19 – 300
25.15 – 785	25.15 – 785
25.10 – 170	25.10 – 170
…	…
A	**B**

Figure 5.

In case A, an impatient trader wants to buy 100 shares. He can do so at the ask, by paying \$25.30 per share. If we assume the mid-spread price to be a proxy for the "fair value" of the stock, then the impatient trader pays 5 cents per share for liquidity (or 20 bp: (\$25.30-\$25.25)/\$25.25 = 20 bp). In case B, we see the same book, but now a trader needs to buy 1000 shares. The best offer is good for only 150 shares, so the trader has to "*walk the book*" by paying increasingly higher prices for additional volume: \$25.30 for 150 shares, \$25.35 for 200 shares, and \$25.55 for the remaining 650 shares. The average transaction price now is \$25.4725 per share or 88 basis points. This situation is illustrative of general market realities: the more shares trader wants to transact, the higher his trading costs are. It is said that a large trade has *market impact* – post-trade prices move in the direction of the trade. Price impact of trades must be contrasted with the EMH setting where current prices can accommodate arbitrary volumes.

Figure 6 demonstrates the effect that payment for liquidity has on transaction prices. We can observe that even when the underlying value evolves gradually, transaction prices can fluctuate considerably around the "true price". This is how market microstructure manifests itself in financial markets – it induces autocorrelations in prices over short time horizons as well as accentuated short-term volatility.

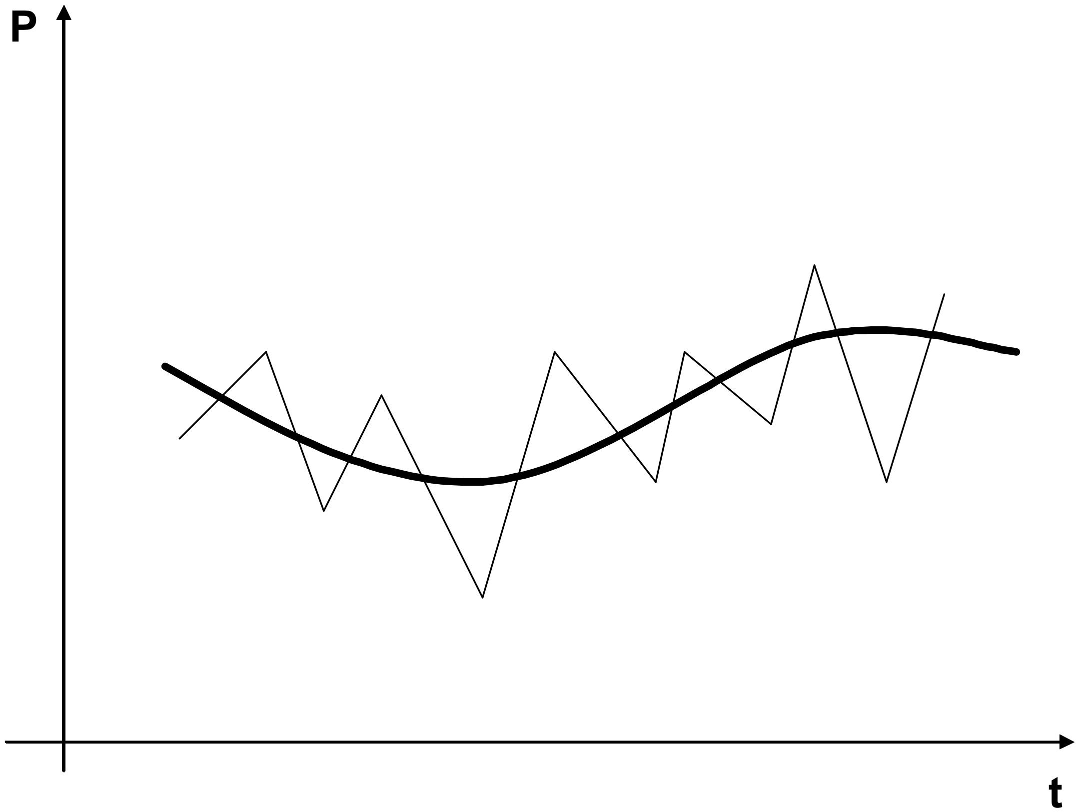

Figure 6.

In light of all market frictions that have been reviewed so far, it is increasingly apt to view trading process as continuous auction process where traders negotiate both prices and volumes for their transactions.

3. Trader Types

We have mentioned that market participants have different investment objectives, trade in different sizes, and use different strategies. Here we discuss trader types whose presence in the markets have particular effect on market microstructure: institutional investors and technical traders.

The main difference about institutional order flow is the trade size. Curiously enough, while institutional order flow accounts for about 70 or 80 percent of transacted volume on US stock exchanges, modern financial markets are designed to better accommodate individuals, not institutions. A telling statistic: average quoted volume and average trade size today is measured in hundreds of shares (maybe low thousands), whereas an average institutional trade is for hundred thousands or even millions of shares. Trading an institutional order in modern markets is being appropriately compared to pushing an elephant through a keyhole. A lot of literature is dedicated to specificities of handling institutional transactions: [Best Execution, 2002], [Schwartz and Francioni, 2004], [Schwartz et al., 2004], [Schwartz et al., 2005], [Dann et al., 1977], [Kraus and Stoll, 1972], and others.

The main influence of institutional trading on market microstructure is the piecemeal sale or acquisition of large blocks – a behavior, which induces positive autocorrelations in transaction prices. It is impossible to bring the entire block to the market at once: first, it is prohibitively expensive, and second, there simply may not be enough shares posted in the limit order book to fill a million-share order. Thus, a standard procedure in the industry is to cut a large block into smaller portions and transact them sequentially over the course of a day or multiple days.

A sequence of one-sided trades creates positive autocorrelations in transaction prices (buys are more likely to follow buys, and sells are more likely to follow sells); this also creates pressure on prices, which makes them deviate from random walk behavior and exhibit trends instead.

If institutional trading indeed induces price trends, other market participants will try to exploit this pattern by trading ahead of the block order. If an institutional investor is trying to acquire a block of shares, other traders can accumulate a position in the stock, wait until institutional trades have driven the price up and then sell for a profit. This trend-following strategy is an example of *technical trading*.

Technical trend followers (a.k.a. momentum traders) wait until they are confident that a price is heading in one direction (presumably being moved by a large institutional trade), and then they either buy (if price is moving up) or sell (if price is going down) the shares of that company; they wait until the trend ends (price either levels out or reverses its direction) and liquidate their position for a profit. Of course, this is easier said then done, since it is difficult to tell if several upticks in a row signal the start of a longer uptrend, or will the price go back down immediately; it is as difficult to tell if the trend has come to its end or just temporarily slowed down. Such “false breakouts” make trend following far from a sure bet. But such strategies are widespread nonetheless and do have an important impact on market microstructure: they accentuate trends, making prices to overshoot their equilibrium levels thus making the dynamic price discovery process more volatile.

When some unexpected positive news arrive to the market, prices must adjust upwards; this is a gradual process, and a trend may develop. When momentum traders see the trend forming, they start buying the stock and thus speeding up the correction. But as multiple traders join the same trend, the price can go past its new equilibrium level, since momentum trading creates trends by itself. When better-informed traders realize that the new level is too high (the news was good, but not that good), they will start selling, driving the price down towards its proper level. This may start a new downtrend, momentum traders will jump in, the price will miss its equilibrium level once again, and so on ad nauseum. This is clearly an exaggerated example, but it shows how trend-followers contribute to positive auto-correlation in prices and again disrupting the random walk behavior.

Not all technical traders are momentum traders. Some technical analysts believe that there are some natural *support and resistance levels* that transaction prices are unlikely to penetrate. Such traders will provide liquidity at those levels believing that at that point prices are likely to reverse their direction. Again, if enough traders follow this same strategy and have homogenous believes about support/resistance levels, price reversion will become a self-fulfilling prophecy. This type of technical trading actually dampens price volatility and induces negative autocorrelation in prices (buys are likely to follow sells, not buys).

We saw that different trader types can make prices to either trend or revert to some mean value, but the point here is that rational market participants can divert prices from their equilibrium levels. This has to be taken into account when studying market microstructure.

4. Summary of Price Correlations

In frictionless markets, price changes are uncorrelated, whereas in the real-world prices can be positively or negatively correlated due to various microstructure factors. Let's firs recap which factors induce particular behavior.

Negative autocorrelation (reversion to mean):

- bid-ask spread (compensation for immediacy)
- market impact (compensation for liquidity)
- dynamic price adjustment (first "overshooting" the equilibrium value and then reverting back to it)
- technical trading (support and resistance levels)

Positive autocorrelation (trending):

- dynamic price adjustment (gradual price changes to reflect new information)
- institutional trading (splitting of block trades and creating one-sided order flow)
- technical trading (momentum trading)

The ultimate question is, of course, what is the aggregate effect of all the above microstructure factors on stock prices? Do prices trend? Do they mean-revert? Or do trending and reversion cancel each other out perfectly, and prices follow random walk, just like prescribed by the EMH? And the ultimate answer is: it depends. Under some circumstances trending factors prevail, under others – mean reversion. Which one it is depends on the individual stock's characteristics, the size of the time window that we are interested in, current market conditions, and so in. In our own empirical investigation, we see evidence of negative autocorrelation, but prices remain hard to predict even over short time horizons.

The most important thing here is to be aware that these microstructure realities exist and they must be taken into account when modeling prices or implementing real-world trading strategies.

Final aside on correlations: while we are exclusively interested in behavior of single stocks, we have to mention *cross-correlations* – how changes in the price of one stock can affect prices of other stocks. Some equities may move in lockstep, while others can go in opposite directions. The same microstructure factors enumerated above must be taken into account when trading a portfolio of securities – see [Bertsimas et al., 2000] for an example – but this issue is beyond the scope of this thesis.

5. Defining Liquidity

In this thesis, we often talk about supplying, demanding, and pricing liquidity. In this sub-section we try to formalize (to an extent) the concept of liquidity.

The textbook definition of liquidity is an ability to sell an asset quickly and in large quantities without substantially affecting the asset's price. From the microstructure point of view however, we are more interested in liquidity as a cost of trading – by transacting a large quantity at once, the trader has to "trade through" a potentially significant number of outstanding limit orders, thereby removing liquidity from the book and moving his average price further away from the consensus value. The smaller is this price movement, the more liquid is the stock.

As important as liquidity is for various market participants, it cannot be summarized in a single variable. In order to assess how liquid the market is, we can look at the size of the bid-ask spread, volume at the bid and at the ask, effective spread, overall volume, short-term volatility, etc. (see [Irvine et al., 2000]). In our research, we focus mainly on the *pre-committed* liquidity – i.e. liquidity that is available before the trade takes place, which rules out hidden limit order, incoming orders, trades arranged in "upstairs" markets, etc.

Liquidity can be aptly viewed as a service traded in a separate marketplace, with certain agents demanding this service and other supplying it. For example, we have discussed a concept of patient and impatient traders: impatient traders want to transact immediately and thus demand liquidity. On the other hand, opportunistic patient traders don't have to transact at specified times or volumes and thus can provide liquidity (serve as counterparties) to impatient (or constraint) market participants at profitable terms. Naturally, market makers fall into this category, but it also includes all other traders that can afford to pursue limit-order trading strategies.

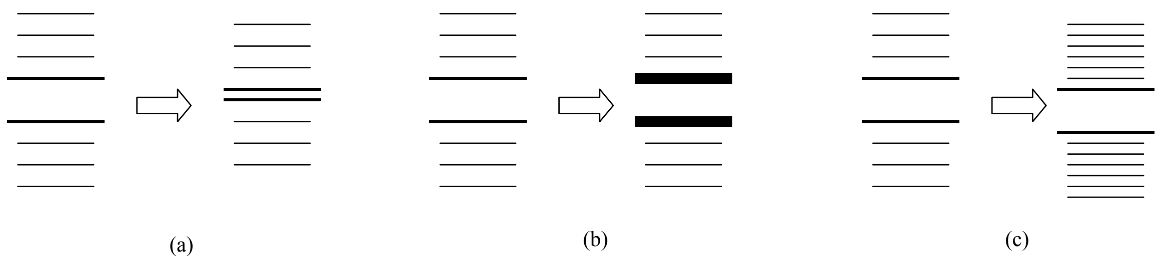

Figure 7.

Let's look at some specific examples of variables that affect liquidity (read cost of trading). Figure 7 illustrates several elementary factors: tighter spreads (a), higher volume at the bid and at the ask (b), and more densely populated book (c) each result in higher liquidity, all other factors held constant. In reality, it's the interplay of these basic factors that determines how expensive it is to trade in a particular market.

One aggregate measure of liquidity was suggested in [Irvine et al., 2000]. It's called the Cost of a Round Trip trade of D dollars, or CRT(D). Let's say that a trader wants to first buy and then immediately sell a certain number of stocks; he seeks instant liquidity on both sides of the market for whatever reason. In a frictionless market, such transaction should cost nothing to the trader, which makes this measure a useful benchmark – the higher is the percentage cost of the round trip trade, the less liquid is the market. Another nice feature of this metric is that it allows us to quantify liquidity beyond the inside spread and the depth of the bid and ask quotes, while incorporating these features as well. While the inside spread and depth provide sufficient information to quantify the cost of trading of small orders, as the size of a transaction increases, these numbers become less meaningful: although a tight spread suggests a liquid market, trading may still turn out very costly if the book is sparsely populated beyond the inside quotes. Computing CRT(D) is trivial: convert D dollars into S shares at the mid-point of the spread, and then calculate how much will it cost to first buy and then sell these shares using prices and quantities in the order book; dividing this cost by D will give you a per-dollar cost of a roundtrip trade.

In our experiments, we are more interested in looking at liquidity on a particular side of the market (either buy or sell) – we essentially assume that liquidity is symmetric given enough observations. This measure can be computed by comparing costs of buying/selling S shares at the mid-quote (as in the perfectly liquid market; such cost is zero) and buying/selling the same quantity against the outstanding limit orders.

Comparing liquidity on both sides of the market can also reveal misbalances and possibly provide clues about future price movements or help decide whether to submit a limit or market order. With all its virtues, the main shortcoming of CRT(D) and similar methods is that it is not capable of summarizing the state of the entire book in a single number – it always depends on the trade size D. Therefore, if we need an aggregate assessment of available liquidity, we will have to compute CRT(D) for many values of D. This measure is very helpful nonetheless when assessing ex-ante profitability of a particular trade.

We emphasize that the state of the order book beyond the inside spread must be taken into account, and that market participants should think about liquidity as a type of commodity, which makes financial transactions less costly. This is especially pertinent for an automated trading, where the goal is to examine the state of a limit order book, evaluate available liquidity, and based on that decide what the optimal action should be.

C. Theoretical Model: Towards Automation of Price Discovery

This section is an overview of some theoretical concepts that provide a rigid analytical foundation for reasoning about market microstructure and electronic trading. The majority of research in this area originates from the branch of Finance called Market Microstructure Theory, which strives to determine how prices are being discovered (or set) in financial markets. While in economic theory price is simply the intersection of supply and demand schedules, in reality such schedules are unobservable, and prices are established through a series of negotiations – a dynamic process we are trying to model. In general, Microstructure research suggests theoretical models of behavior of various market participants – informed traders, liquidity traders, market makers, etc. – and then derives an equilibrium solution for the each participant's optimization (profit maximization) problem. The resulting model can be tested on (or fitted to) some empirical data, which is usually a record of transactions over a certain time period. We aim to present all the proposed models from the point of view of automation and design of electronic trading agents.

It is convenient to separated microstructure models into three categories, depending on what they consider the main source of the bid-ask spread in the marketplace. *Inventory* models suggest that the dealer is reimbursed for holding an unbalanced portfolio and adjusts his quotes as a function of his inventory's deviation from some desired level. *Information* models state that since the liquidity provider is committed to his quotes, he gives away free options to other market participants; therefore, informed traders will trade with the dealer only when it's profitable for them (they will exercise their "free option" when it's in-the-money). The dealer loses on such trades and has to be compensated by the bid-ask spread. *Liquidity* models are the most recent; they incorporate the order book information into their framework. While inventory and information models view the dealer as an institution (a designated intermediary as seen on the NYSE floor or NASDAQ OTC market), liquidity theories see a market maker as any trader willing to supply liquidity to the order book for a fee reflected in the bid-ask spread. Despite the institution vs. trading strategy distinction, all three theoretical approaches are pertinent and interrelated. Even though liquidity models are best suited for our experimental set-up and trading mechanisms of modern financial markets, we will present all three approaches and show where they fit into our main goal – understanding price formation, liquidity provision, and creating an automated trader. In writing this modeling section, we rely on similar surveys, such as [O'Hara, 1995], [Stoll, 2001], [Madhavan, 2000], and [Madhavan, 2002].

1. Inventory Models.

a. Order Flow Based Approach.

The simplest model of the bid-ask spread (as presented, for example in [Amihud and Mendelson, 1980]) is where the quotes are functions of the dealer's stock holdings. The dealer's inventory is modeled as a Markovian state variable, and his decision variables (bid and ask quotes) depend on the value of the state variable (his inventory level). Permissible stock inventory levels are represented by a finite number of states: [-K, -K+1, …, L-1, L], where K represents maximum short position, L – maximum long position. Order arrival rate is assumed to follow a Poisson distribution, but is a function of posted quote: the higher the bid, the more likely next offer to be a buy offer, etc. We define the market maker's revenue and cost functions:

$$R(\mu) = \mu P_a,$$

and

$$C(\lambda) = \lambda P_b,$$

Where C is cash outflow, R – cash inflow, μ – probability of a sell order arriving, λ – probability of a buy order arriving, P_a – ask price, P_b – bid price. μ is a monotone decreasing function of P_a and λ is a monotone increasing function of P_b. The market maker's objective function is defined as follows:

$$g(\underline{\lambda}, \underline{\mu}) = \sum_{k=0}^{M} \phi_k q_k,$$

where $\underline{\lambda} = (\lambda_0, \dots, \lambda_{M-1})$, $\underline{\mu} = (\mu_0, \dots, \mu_{M-1})$, M = L+K. M is a total number of allowable states, q_k is cash flow $R(\mu_k) - C(\lambda_k)$, and ϕ_k is the limiting probability of being in state k. By differentiating the objective function and looking at the relationship between neighboring states, we can show that optimal (λ_k, μ_k) are aligned along the curve defined by

$$[R(\mu) - \mu R'(\mu)] - [C(\lambda) - \lambda C'(\lambda)] = g.$$

This result yields several implications. First, the optimal bid and ask prices are monotone decreasing functions of the market maker's inventory position. The dealer raises both his quotes as his inventory declines, and lowers them as his inventory increases. Second, the dealer has a preferred inventory position – he manipulates quotes to bring his holdings to a desired level. And finally, the bid-ask spread must always be positive.

If we assume linear supply and demand functions, $D(P_a) = \gamma - \delta P_a$ and $S(P_b) = \alpha + \beta P_b$. Then $R(\mu) = (1/\delta)(\gamma\mu - \mu^2)$, and $C(\lambda) = (1/\beta)(\lambda^2 - \lambda\alpha)$. ($D(P_a)$ is essentially μ, and $S(P_b)$ is λ). Optimal (λ_k, μ_k) will lie on an ellipse

$$\mu^2/\delta + \lambda^2/\beta = g.$$

The rates induced by the market maker's quotes at his preferred position are:

$$\lambda^p = \mu^p = \sqrt{g\beta\delta/(\beta+\delta)}\,.$$

The market clearing rates, at which the demand and supply schedules intersect and $P_a = P_b$ (think "true price"), are

$$\lambda^c = \mu^c = (\gamma\beta + \alpha\delta)/(\beta + \delta).$$

It follows that $P_b^p < P^c < P_a^p$ – i.e. preferred bid and ask prices straddle the market-clearing ("true") price. The size of the bid-ask spread can also be readily defined:

$$s = P_a(\mu) - P_b(\lambda) = (\gamma/\delta + \alpha/\beta) - (\mu/\delta + \lambda/\beta).$$

Creation of an automated market maker based on this simple model presents a number of difficulties. First of all, this model does not include information aspects of trading – it assumes that all transactions happen with the liquidity (uninformed) traders. While this limitation comes from the inherent simplicity of this model, there is another conceptual problem, especially damaging from the implementation standpoint. Although this model claims to be (and essentially is) inventory-driven, the inventory never shows up explicitly in any of the equations above. Dealer's quotes are functions of the order arrival rates, not his inventory. In simple terms, this model is saying the following: "establish such quotes that buy and sell orders arrive at the same rate; the resulting stable inventory will be your preferred position". The most suspect part of this model is that this preferred position has no connection with the underlying value of the security traded, or its other characteristics (long-term volatility, for example). If we try to implement this model as described above, we will have to learn a functional relationship between the posted quotes and corresponding arrival processes (supply and demand schedules $D(P_a)$ and $S(P_b)$). While this isn't necessarily impossible, such schedules can be arbitrarily complex and can change with time. A more elegant and tractable solution would be to address this problem backwards – try to establish which level of inventory induces equal arrival rates on both sides of the spread. Reinforcement learning, amongst other techniques, can be employed here. While it seems logical to suggest that preferred inventory level should be zero – and this may indeed be a good starting point – several studies ([Hasbrouck and Sofianos, 1993] and [Madhavan and Smidt, 1993]) suggest that there is a preferred dealer inventory, which is a function of the stock's price, trading volume, and other factors.

b. Classic Inventory Model.

A more complex and probably the best known out of all inventory models is the one developed in [Ho and Stoll, 1981] and extended to multiple dealers in [Ho and Stoll, 1983]. As opposed to an institutionalized monopolist from the first model, this approach treats a market maker as a regular investor who assumes the role of default counterparty for everyone else. With this role come certain risks, and the dealer must be reimbursed for taking on these risks through the bid-ask spread. There are three sources of risk in this model: (1) holding a suboptimal portfolio induced by the dealer's commitment to his quotes (the dealer is assumed to be risk-averse); (2) simple order processing costs – i.e. computers, human labor, taxes, etc.; and (3) the risk of losing to informed traders. While in the first inventory model the market maker is assumed to be risk neutral, in this case the dealer's attitude towards risk becomes one of the variables. The general idea behind this model is that the dealer sets such prices that maximize his expected utility over some fixed time horizon.

This model assumes an existence of an exogenous "true price" p of a stock, which is known to the dealer. Then the dealer sets his quotes at $p_a = p + a$ and $p_b = p + b$. We're interested in determining functional relations that guide a and b. As in the previous models, order follow is a Poisson process and depends on the dealer's quotes. Therefore, the dealer faces transactions (or inventory) uncertainty. In addition, the dealer starts with some pre-existing portfolio X with random returns dX:

$$dX = r_x X dt + X dZ_x,$$

where r_x is the mean return, and Z_x is a Weiner process with mean zero and variance σ_X^2.

The dealer's portfolio consists of cash, stock, and some base wealth. He can lend and borrow at the risk-free rate r. The value of the cash account changes as follows:

$$dF = rF\,dt - (p - b)\,dq_b + (p + a)\,dq_a,$$

where q_b and q_a are quantities bought and sold. The dealer's inventory follows a similar process:

$$dI = r_I I\,dt + p\,dq_b - p\,dq_a + I\,dZ_I.$$

Note that inventory is valued at the "true price" p, while cash is influenced by actual transaction prices. Finally, the base wealth evolves as follows:

$$dY = r_Y Y dt + Y dZ_Y.$$

The dealer's optimization problem is to find bid and ask prices that will maximize his terminal wealth after T time steps:

$$W_T = F_T + I_T + Y_T.$$

This maximized value is given by the value function J(*) such that

$$J(t, F, I, Y) = \max_{a,b} E[U(W_T) \mid t, F, I, Y],$$

where U is the utility function. Since there is no intermediate consumption before time T, the recursion relation implied by the principle of optimality is

$$dJ(t, F, I, Y) = 0 \text{ and } J(t, F, I, Y) = U(W_t). \tag{3}$$

To find a solution to the dealer's problem, we need to find the ask and the bid adjustments a and b that solve (3) for each state.

The solution to this problem requires application of stochastic calculus (Ito's Lemma more specifically). Solving this problem in general terms is not straightforward, so several transformations and simplifications are introduced. First, we assume linear supply and demand functions (similar to the first model):

$$\lambda_a = \lambda(a) = \alpha - \beta a;$$

$$\lambda_b = \lambda(b) = \alpha - \beta b.$$

The two functions are symmetric. We also define buy and sell operators:

$$SJ = S[J(F, I, Y)] = J(F + Q, I - Q, Y),$$

$$BJ = B[J(F, I, Y)] = J(F - Q, I + Q, Y),$$

Where Q is the quantity bought or sold. The sell operator S acting on J describes the dealer's derived utility after a sale by the dealer: utility will increase if the sale moves dealer's inventory closer to the desired position, and will decrease if it moves inventory further away. It works the same way for the buy operator B. It can be shown then that

$$b = \alpha/2\beta + (J - BJ)/2BJ_FQ, \tag{4}$$

$$a = \alpha/2\beta + (J - SJ)/2SJ_FQ, \text{ and} \tag{5}$$

$$s = \alpha/\beta + (J - SJ)/2SJ_FQ + (J - BJ)/2BJ_FQ.$$

In other words, the dealer should set such quotes that maximize expected profits associated with stochastic order arrivals plus a risk premium term. Using Taylor's series expansion and a number of new assumptions (like quadratic utility function, for example) it can be further shown that the size of the spread

$$s = \alpha/\beta - ½Z(Q/W)\sigma_I^2\tau + ½\,Z(Q/W)\sigma_I^2[(r_I - r + G_I) + Z(r_w + 2\Pi/W)]\,\tau^2, \quad (6)$$

where $Z = U''W/U'$, the coefficient of relative risk aversion, τ, the time horizon ($\tau = 1$ corresponds to one year), $G_I = r_I + ½\sigma_I^2$, the instantaneous growth in the variance of I, and $\Pi = \alpha^2 Q/4\beta$, which is the maximum expected profits from either sales or purchases. Although this is only an approximation to the dealer's pricing problem, and this form is somewhat cumbersome, but at least all the variables can be observed in the real world or estimated. For more details regarding mathematical techniques used in this derivation and elsewhere in this chapter refer to [Merton, 1990].

Let's examine some of the implications of this solution. First of all, the size of the spread depends on the dealer's time horizon τ – the longer the horizon, the higher is the uncertainty on the dealer's investment portfolio; the dealer has to be compensated for carrying this risk with a larger spread. When τ is almost zero, the dealer sets a monopolistic spread $s = \alpha/\beta$; the rest of (6) is the risk premium term. The second observation is that this risk premium depends on the dealer's attitude towards risk Z, the size of the transaction Q, and on how volatile the stock is (G). The most important implication, however, is that neither the transaction uncertainty nor the inventory position influence the spread. This means that to counteract possible inventory misbalances, the dealer will not change the size of his spread, but instead will move both bid and ask quotes in the same direction (up or down) to induce the order flow that will return him to the balanced portfolio.

We can define this adjustment as $d = a - b$. Then, given s and d, we can write a and b as

$$a = (s + d)/2, \qquad b = (s - d)/2.$$

From (4) and (5), d can be written as

$$d = (J - SJ)/2SJ_FQ - (J - BJ)/2BJ_FQ,$$

or

$$d = -Z(I/W)\sigma_I^2\tau - Z(I/W)\sigma_I^2[(r_I - r + G_I) + Z(r_w + 2\Pi/W)]\,\tau^2. \quad (7)$$

This expression is very similar to (6) except that 2I has replaced Q. Clearly, this adjustment depend on the inventory that has been acquired. When $d = 0$, both bid and ask quotes are equidistant from the true price p, whereas when $d < 0$ ($I > 0$), both quotes are shifted down to encourage dealer sales, which should eventually drive d back to zero. The inverse happens when $d > 0$. Equations for both s and d can be re-written in a slightly more complex form for the case of asymmetric supply and demand. This model can also be extended to multiple market makers – essentially at any point in time one of the market makers will have the "most misbalanced" position, which will push his quote further then everyone else's, and all public transaction will happen with him until his inventory normalizes; at that point, another market maker's quote will become most aggressive, and so on.

Now that the quote update rule is completely derived and explained, it is necessary to mention some potential limitations of this model. First, the model assumes finite horizon, which implies assets liquidation at a single price at some point. Furthermore, if the time horizon is one day, then the spread should be largest in the morning and diminish through the day, which goes against the empirical U-shaped patterns of spread sizes. Assuming the existence of an exogenous "true price" of a stock can also be problematic, since there is no way to observe it in the real world. Some proxy will have to be used, which may or may not affect the rest of the model. Poisson process assumption for the order flow also needs not be true in the reality. In the derivation of this model, only market orders are used, leaving out the existence of limit orders. Transaction size (quote depth) Q is a constant here, whereas it can be another output variable together with a and b (actually two variables: one for the bid quote, another – for the ask). And, finally, while the existence of informed traders is assumed, it is not modeled explicitly, which certainly is a shortcoming, which will be addressed in great detail in the following section.

To conclude the discussion of the second inventory model, let's examine it from the point of view of an automated market maker. The ease with which this theory can be applied may be its best feature. It is readily parameterizable. Time horizon τ can be set to one day or several weeks. True price p and the liquidation price can be taken as a mid-spread from some external market. Parameter α can be set as an average number of transaction during some time period, and then β can be calibrated so that α/β corresponds to the average observed spread over the same period (i.e. a stock with 2,000 transactions per day and average spread of 0.02 will get $\alpha = 2000$, $\beta = 100000$). Coefficient of risk aversion Z – can be an arbitrary number; risk neutrality may be a good starting point. All σ's can be estimated from transactions records. Initial endowments I and W can be set at arbitrary numbers (zeroes, for example). Almost all of the above parameters can be tweaked either manually or using search/learning techniques to test sensitivity and see which setting achieve the best results – either from the point of view of profitability or fit to the past data.

Overall, here's the basic concept behind all inventory models. As securities exchanges help buyers and sellers meet at one place, market makers help counterparties meet in *time*. In other words, the dealer ends up holding one side of a trade waiting for a counterparty to arrive. During this waiting period, the dealer is exposed to changes in the asset price, and he is compensated for both the service he provides and the risk he is taking by the bid-ask spread. Furthermore, to mitigate his exposure, the dealer posts such quotes that will induce transactions that will drive his inventory towards the desired level: too much inventory – induce public buying, too little – induce selling. Figure 8 shows a graphical representation of this concept.

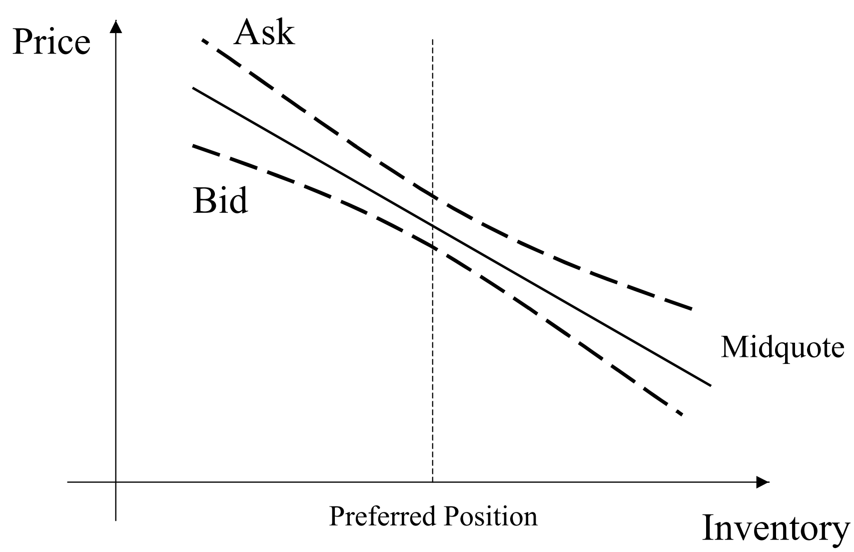

Figure 8.

Empirically, the cornerstones of this theory are the existence of a preferred positions among market makers and serial correlations in trades (a buy should follow a sell, a sell should follow a buy) due to inventory balancing. Practical investigations (([Hasbrouck and Sofianos, 1993], [Madhavan and Smidt, 1993], [Madhavan and Smidt, 1993] and many others), however, yielded mixed results: some researchers find these features in the real world, some don't. The reason for this is because there is more to the story.

2. Information Models.

While everything presented in previous sub-section is very reasonable, those models ignore (or at least do not take into account explicitly) one prominent reality of financial markets – the presence of the so-called *informed* traders. We refer to those traders as "informed" in a sense that the know *more than the dealer does* about the value of the security traded. These traders will transact with the dealer only when they believe the dealer's quotes are posted away from the real price of the security and a profitable trading opportunity is present. This goes back to the concept of a market maker giving a "free option" to the public, which gets exercised by the informed traders when it's in-the-money. By this definition, dealers always loose out to the informed traders. If this was the only type of traders in the market places, the dealers would have been driven out of business through constant loses or by being forced to post spreads so wide that no transactions would take place. Fortunately for everybody, there are other traders as well – they are referred to as the *uninformed* or *liquidity* traders. These are the traders who trade for exogenous reasons (need for cash, end of the investment horizon, portfolio rebalancing, hedging, etc.) – i.e. they do not buy or sell securities based on their expectations (or presumed superior knowledge) of future prices.

The bottom line – the market maker strives to offset his loses from trading with the informed trader by his gains from trading with the uninformed. The latter gains come from the existence of the bid-ask spread. The problem is that prior to the trade, the market maker cannot tell if his counterparty is informed or not. The fact that information also affects prices and is being disseminated through trades is what differentiates the models in this section from those described earlier.

a. Initial Idea.

One of the first treatments of the information-based trading from the dealer's point of view was presented in [Copeland and Galai, 1983]. This is a one-step model where the market maker sets his quotes and then faces a probabilistic arrival of either an informed or an uninformed trader, who can either buy, sell or not trade at all. The stock price P is drawn from a known density distribution f(P). The informed traders know this price, whereas other market participants (dealers and the uninformed) know only the distribution. An informed trader arrives with probability π_I, uninformed one – with probability $(1 - \pi_I)$. Uninformed trader buys with probability π_{BU}, sells with probability π_{SU}, and does nothing with probability $\pi_{NU} = 1 - \pi_{BU} - \pi_{SU}$. The informed trader makes the trade that maximizes his profits. All trades are assumed to be the same size.

The market maker expects to lose

$$\int_{Pa}^{\infty} (P - P_a)\, f(P)dP + \int_{0}^{Pb} (P_b - P)f(P)dP$$

from trading with the informed traders, and to gain

$$\pi_{BU}(P_a - P) + \pi_{SU}(P - P_b) + \pi_{NU}0$$

from trading with the uninformed. Therefore, the dealer's objective function he is

$$J = \pi_I \left[\int_{Pa}^{\infty} (P - P_a)\, f(P)dP + \int_{0}^{Pb} (P_b - P)f(P)dP\right] + (1 - \pi_I)\left[\pi_{BU}(P_a - P) + \pi_{SU}(P - P_b)\right]. \quad (8)$$

The dealer's goal is to pick such P_a and P_b that will maximize J.

This model is not terribly useful in practice, since it's a one-shot framework, but it makes several important points. It shows how to explicitly incorporate the information effects into the liquidity provider's expected profits. It also highlight the importance of knowing the probabilities of informed and uninformed trading and supply-demand elasticities. Another useful conclusion is regarding the source of the bid-ask spread: even with competitive dealers, as long as the probability of informed traders arrival is greater than zero, there will always be a positive spread in the market.

b. Sequential Trades Model.

While the idea from the previous section seems pretty straightforward – the dealer loses to informed traders and charges the uninformed traders the spread – if the new information (the true price) is not instantly revealed after the trade, the loses to the informed traders become difficult to quantify or even to identify. In the more realistic scenario, the size of the loss depends not only on current prices, but also on how quickly the prices reflect new information. This realization implies that first, multiple rounds of this "trading game" must be examined, and, second, trades in themselves can reveal new information and affect price formation.

This idea was developed in [Glosten and Milgrom, 1985]. The main premise is that agents will reflect their private information through trading – they will buy if the stock is undervalued, and sell if it's overvalued. However, they can also buy and sell simply because they need liquidity. Since the market maker cannot tell which type of trade is which (informed or uninformed), he has to continuously adjust his believes about the value of the stock, conditioned on his order flow. As the market maker receives orders, he continuously updates his quotes, which will eventually converge to the expected value of the stock – i.e. the dealer can learn all the information from the informed traders. As opposed to pure inventory models, this approach treats the order flow as an endogenous process by tying it to the value of the asset being traded. In other words, it addresses the way information is being impounded into prices, therefore relating to the efficient markets hypothesis and rational expectations theory. To implement this sequential trades information model, we will use Bayesian learning to represent the market maker's belief update process.

The assumptions made in this model are similar to those made in the one-shot model. The market maker and all the traders are assumed risk neutral and competitive. The asset's eventual "true price" is given by a random variable V, there are no transaction costs and no inventory costs (the dealer can borrow and lend freely). At each iteration, a trader is chosen probabilistically out of the pool of all traders – informed and uninformed – and can either buy one unit (share, round lot, etc.), sell one unit, or do nothing. If the trader wants to transact more than one share, he has to go back into the pool and wait for the next time he is picked. Such system is chosen to prevent all the informed traders from transacting at once, or trading extremely large volumes and instantly revealing their information.

The dealer sets such prices that his expected profit is zero. This is the consequence of the competitive assumption – all profits above normal will be taken away by the competition. The market maker's prices are equal to his expectations of the asset value, given a particular type of trade: the bid is the expected value of the stock, given that a trader wants to sell, and ask is the same value, given that a trader wants to buy. We refer to such prices as "regret-free". Since we assume that each trade reveals some information, the market maker revises his quotes after each trade to keep his prices regret-free.

Let's examine the mechanics of beliefs update in the Bayesian learning framework. Suppose that the stock value can be either low or high – V_L or V_H – and let's denote events of a sale and a purchase as S and B. Market maker's bid and ask prices reflect his expectations of the value of a stock, given a transaction type:

$$a = E[V \mid B] = V_L P[V = V_L \mid B] + V_H P[V = V_H \mid B],$$

$$b = E[V \mid S] = V_L P[V = V_L \mid S] + V_H P[V = V_H \mid S].$$

In order to calculate a and b, we will have to know the following probabilities: $P[V = V_L \mid B]$, $P[V = V_H \mid B]$, $P[V = V_L \mid S]$, and $P[V = V_H \mid S]$. We apply Bayes Rule to define and update these probabilities. For example:

$$P[V = V_L \mid B] = P[V = V_L]\, P[B \mid V = V_L] \,/\, (P[V = V_L]\, P[B \mid V = V_L] + P[V = V_H]\, P[B \mid V = V_H]).$$

Example. To make this more clear, it may be beneficial to have a look at a stylized problem. Let's say that V can be either 0 or 1, and the dealer's initial beliefs are $P[V = 0] = ½$, $P[V = 1] = ½$. The population of traders is split evenly between informed and uninformed: $P[inf] = ½$, $P[uninf] = ½$. Informed traders transact only when it's profitable, while the uninformed sell and buy with probability ½.

$$P[S \mid V = 0] = P[inf]\, P[S \mid inf, V = 0] + P[uninf]\, P[S \mid uninf, V = 0]$$

$$= (½)(1) + (½)(½) = ¾,$$

$$P[S \mid V = 1] = P[inf]\, P[S \mid inf, V = 1] + P[uninf]\, P[S \mid uninf, V = 1]$$

$$= (½)(0) + (½)(½) = ¼,$$

$$P[V = 0 \mid S] = (½)(¾)/[(½)(¾) + (½)(¼)] = ¾,$$

$$P[B \mid V = 0] = P[inf]\ P[B \mid inf, V = 0] + P[uninf]\ P[B \mid uninf, V = 0] = ¼,$$

$$P[B \mid V = 1] = P[inf]\ P[B \mid inf, V = 1] + P[uninf]\ P[B \mid uninf, V = 1] = ¾,$$

$$P[V = 0 \mid B] = ¼,$$

$$a = E[V \mid B] = (1)(¾) + (0)(¼) = ¾,$$

$$b = E[V \mid S] = (1)(¼) + (0)(¾) = ¼.$$

These are regret-free bid and ask prices that the market maker should set, given his initial assumptions. Now, suppose a buy happened. We update our probabilities using the Bayes Rule:

$$P[V = 0 \mid B, B] = P[V = 0 \mid B]\ P[B \mid V = 0] / (P[V = 0 \mid B]\ P[B \mid V = 0] + P[V = 1 \mid B]\ P[B \mid V = 1]) = (¼)(¼) / [(¼)(¼) + (¾)(¾)] = 1/10,$$

$$P[V = 0 \mid B, S] = P[V = 0 \mid B]\ P[S \mid V = 0] / (P[V = 0 \mid B]\ P[S \mid V = 0] + P[V = 1 \mid B]\ P[S \mid V = 1]) = ½,$$

$$a = (0.1)(0) + (0.9)(1) = 0.9,$$

$$b = (0.5)(0) + (0.5)(1) = 0.5.$$

It has been proven that following such updating process, the dealer's quotes will eventually converge to the market's consensus value of the security – i.e. all the private information from the market participants will be impounded into the market maker's prices. For more details on Bayesian learning and the proof of convergence refer to [Mitchell, 1997] and [O'Hara, 1995]. To make all the above formulas more intuitively compelling, it may be helpful to consider the diagram in Figure 9.

The stock's value can be either high or low with $P = \theta$ and $P = 1 - \theta$ correspondingly; informed trader arrives with $P = \mu$, uninformed – with $P = 1 - \mu$; when the value is high, the informed always buys, when it's low – he sells; the uninformed buys with $P = \gamma^B$ and sells with $P = \gamma^S$, regardless of the stock value. The resulting probabilities of buys and sells are on the right side of the diagram. All the market maker needs to do to determine the probability of a purchase or a sale, is to add up the corresponding entries in the tree. In our case, the probability of a sale is $(1 - \mu)\gamma^S + (1 - \theta)\mu$, and probability of a buy is $(1 - \mu)\gamma^B + \theta\mu$. More complex trading games can be created by adding additional nodes to this tree. This process is iterated as new trades come in, and the dealer's quotes and believes will eventually converge to full-information levels.

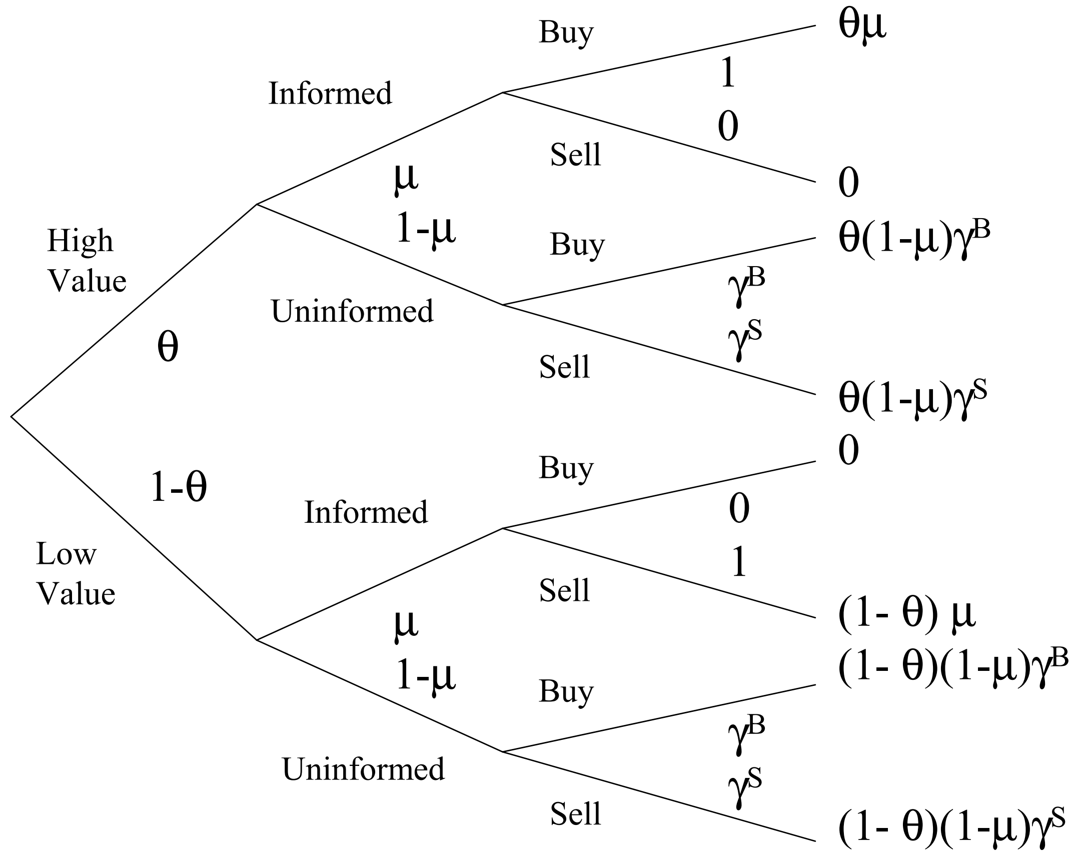

Figure 9.

This model has a number of important implications. It shows that the spread must be present in the market, and is a function of the nature of underlying information, the number of informed traders, and traders' elasticities. The prices form a Martingale – it is impossible to predict future transaction prices from the past records – markets exhibit a weak-form efficiency because private information does not get instantly revealed. The time series of transaction prices should be serially uncorrelated under this model, as opposed to the negative serial correlation in the inventory theories.

This model is easy to adapt to automated market making: the rules for setting and updating the quotes are derived from the Bayes Rule and are completely spelled out above. The only challenge is to determine the initial parameters such as proportion of informed traders, trading probabilities of uninformed traders, the range of values for the stock price, etc. – essentially parameters θ, μ, and γ from the diagram. The solution to this problem is to take a past transaction record, fit our model to it, and then use the calibrated parameters for future transactions.

The simple model like this, however, is unlikely to be robust enough to capture the real world's complexity. It adopts binary treatment of prices, assumes fixed trade size, doesn't specify how long will it take for the quotes to converge. The modeling of information and the trading crowd are also oversimplified. There may or may not be new information, information can arrive in "waves", informed/uninformed ratio is unlikely to stay constant, etc. Informed traders will want to trade more than one unit to maximize their profits. They may also want to behave strategically – i.e. trade in smaller quantities to conceal the fact that they are informed. One example: [Easley and O'Hara, 1987] introduced additional layers of complexity by allowing the information to exist or not exist, and by allowing the traders to trade in large and small quantities. This amount to adding another layer on the left side of Figure 9 (information there or not) and increasing the branching factor on the right side (trade in small or large quantities). Although the problem remains highly stylized, its complexity increases manifold, and we end up with multiple possible equilibria. The same happens in numerous models with strategic traders: to come up with tractable solutions, multiple simplifications and generalizations have to be made, and we end up with a slew of complex equations that are impossible to transfer into the real world. Therefore we will not examine such models in this section, since they are of little interest to an automated market maker.

3. Combined Models.

The two preceding sections described two different explanations of the bid-ask spread in the securities markets. The key point to retain is that these theories are not mutually exclusive – the presence of the spread can be attributed to *both* inventory and information effects. In this section we present a simple model of the short-term behavior of stock prices, similar to those presented in [Stoll, 1989], [Huang and Stoll, 1994], [Huang and Stoll, 1997], and [Stoll, 2001].

First, let's define the change in the quote's midpoint (midpoint $M = (P_a + P_b)/2$):

$$dM = M_t - M_{t-1} = \lambda(s/2)Q_{t-1} + \varepsilon_t, \quad (9)$$

where Q is the trade indicator (1 – buy, -1 – sell), s is the bid-ask spread, λ is the fraction of the half-spread by which quotes respond to the trade at t-1 (this includes both inventory and information effects), and ε is public information shock (idiosyncratic change in the value of the stock). According to this model, quotes change either because of the new public information or because the previous trade has induced this change (either through conveying some new information, or through misbalancing the dealer's inventory, or both). Trades take place either at the ask or at the bid, half-spread above or below midpoint:

$$P_t = M_t + (s/2)Q_t. \tag{10}$$

Combining (9) and (10), yields

$$dP = (s/2)(Q_t - Q_{t-1}) + \lambda(s/2)Q_{t-1} + \varepsilon_t. \tag{11}$$

To evaluate the market maker's performance, we need to define the *realized* as opposed to *quoted* spread. It has been stated before that the dealer profits from the bid-ask spread. It is important to understand, however, that the profits from the roundtrip trade (i.e. acquiring a position from the public, and then liquidating it at a later point it time by entering into an offsetting transaction) may be smaller then the quoted spread. For example, following a public sale the dealer's realized half-spread is the price change, as defined in (11), conditional on a purchase at the bid. Since the quotes change, the realized profit will be less then if they stayed constant. This difference between the realized and quoted spread can help to determine the relative importance of inventory and information factors.

In term of our model, the expected realized half-spread conditional on a transaction at the bid is

$$E[dP_t \mid Q_{t-1} = -1] = (s/2)(E(Q_t) + 1) + \lambda(s/2)(-1). \tag{12}$$

It depends on the expected sign of the next spread $E(Q_t)$, and on λ. Let's define π as a probability of a trade reversal – probability that a sale will follow a purchase and vice versa. Then,

$$E(Q_t \mid Q_{t-1} = -1) = \pi + (1 - \pi)(-1).$$

If purchases and sales are equally likely, then $\pi = 0.5$ and $E(Q_t) = 0$. The value of λ depends on the presence of inventory and information effects.

If neither the information nor his inventory holdings influenced the dealer's quotes, then $\pi = 0.5$ and $\lambda = 0$ (trading doesn't affect quotes). By substituting these values in (12), we get the realized half-spread of s/2. This is a highly unrealistic scenario, where the only costs that the dealer is facing are the fixed material costs of being a dealer (salary, computer, communications, etc.). Such trivial setting is referred to as "*order processing*" model, where the dealer's sole responsibility is to clear trades, and transaction prices would simply tend to "bounce" between the bid and the ask.

If the information were the only source of the spread, then $\lambda = 1$, and, by construction, the quotes are such that buys and sells are equally likely – i.e. $\pi = 0.5$. This results in the realized spread being zero, which is consistent with our information model, where the spread is used only to cover loses from trading with informed traders.

Finally, if the spread is influenced just by inventory considerations, $\lambda = 1$. When a trade introduces a misbalance into the dealer's inventory, the dealer moves both of his quotes up or down to induce a trade in the opposite direction. This means that transactions at the bid and at the ask no longer have the same probabilities. Therefore, $E(Q_t) = 2\pi - 1$, and the realized half-spread is $(2\pi - 1)(s/2)$. For example, if $\pi = 0.75$, then $E(Q_t) = 0.5$, and the realized half-spread is s/4. The quotes should eventually adjust to their initial level, and the full half-spread will be earned on the initial transaction, but the complete reversal in one trade is unlikely. The values of the serial covariances for these alternative models are summarized in Table 1.

Source of Spread	**Price Covariance**	**Quote Covariance**
Order Processing ($\lambda = 0$, $\pi = 0.5$)	$-0.25s^2$	0
Information ($\lambda = 1$, $\pi = 0.5$)	0	0
Inventory ($\lambda = 1$, $0.5 < \pi < 1$)	$-0.25s^2 < cov_p < 0$	$-0.25s^2 < cov_q < 0$

Table 1.

Wile the above discussion considers only the polar cases, this simple model can help us assess the composition of the bid-ask spread. The most straightforward approach is the following: if the average quoted half-spread is 5 cents, and the average realized half-spread is 2 cents, we can conclude that the information portion of the spread is 3 cents, while the inventory and order processing component is 2 cents. From the automation point of view, we can take a record of past transactions, calculate average quoted and realized spreads, calculate proportional influences of inventory and information, and assign appropriate coefficients to the outputs of models from sections A and B. [Stoll, 1989] takes the model from Table 1 and applies it to a number of NASDAQ stocks. He comes up with the following decomposition of the spread:

Component	Contribution
Order Processing	47%
Information	43%
Inventory	10%

Table 2.

More formally, we can describe the "combined" update model with the following set of equations:

$$P_a = k_1 E[V_t \mid Q_t = 1] - k_2\phi(I_t + I^*) + \varepsilon_t,$$

$$P_b = k_1 E[V_t \mid Q_t = -1] - k_2\phi(I_t + I^*) + \varepsilon_t,$$

where I^* is optimal inventory level, $\phi(.)$ is a function of the inventory misbalance, which will bias both quotes to induce a "balance-restoring" trade, ε_t is public information shock, and k_1 and k_2 are the respective weights of information and inventory influences on the spread.

4. Limit Order Models

All the previously described models aim to describe the dealer's behavior. But all the insights gained from those models apply directly to limit orders, since limit order serve the same purpose as designated market makers – they provide liquidity to the market. It is correct to view limit order traders as dealers' competition. In this sub-section, we describe models that analyze limit-order submission strategies directly.

a. Limit Orders Formalized

First, it may be beneficial to present a slightly more formal treatment of limit orders in order to highlight several important trade-offs in this domain: probability of execution vs. price improvement, and probability of execution vs. information content of the trade. [Glosten, 1994] is considered the founding work in this domain. Let's suppose that the limit order at the top of the sell queue has a price P_a and quantity Q_a, and the next lowest sell order is at the price $P_a + \delta$ and quantity Q_L. This second limit order executes only if the quantity Q of the incoming market buy is higher than Q_a. Conditional on the market order being uninformed, the expected profit from this limit order is:

$$(P_a + \delta - E[v \mid uninf])\ \min(Q_L, Q - Q_a). \qquad (13)$$

This expression is positive, since the top of the sell book (P_a) is higher than the unconditional expectation of the security's value. If the incoming market order is informed, then the expected profit is:

$$-\ (E[v \mid uninf] - P_a - \delta)\ \min(Q_L, Q - Q_a). \qquad (14)$$

This term is negative because an informed trader only buys when $v > P_a + \delta$. The overall expected profit is the weighted average of (13) and (14), each expression weighted by the probability that an observed trade is either informed or uninformed given its size and the state of the order books.

At higher prices, the probability of transacting against an uninformed order is lower, but profits from such transaction are higher. Secondly, (13) and (14) describe the expected profit given that a transaction takes place, which may or may not happen. This brings us to the second trade-off: the further the price is from the inside market, the lower is the probability of execution, and higher is the expected profit. It is also essential to remember that each limit order in the book includes in itself a free option for an incoming market order to execute at a set price when the value of the security moves up or down. This means that placing a limit order must be accompanied by its monitoring, which can render such strategies costly and can be another reason for the market maker's presence in this setting.

b. Quantifying the Trade-offs.

Here we will develop a general theory applicable to limit order markets: how do we quantify the trade-offs that are inherent in limit order trading (see above), how do we choose between limit and market orders, among other aspects.

In this model, based largely on [Foucault et al., 2003a], we assume that buyers and sellers arrive sequentially and want to transact one share at a time. All markets participants are willing to trade regardless of prevailing price, which makes them liquidity traders. However, they have different preferences concerning the time of execution: we partition them into "patient" and "impatient". These execution constraints can be driven by private information, but for the purposes of this model, we consider these influences external. Upon arrival to the market, a trader decides to place either a market or a limit order to minimize his total trading costs. Such decision is contingent on the state of the book.

Limit prices are meaningful only within some interval [A, B], which straddles the value of the security traded – it can be assumed that there is an infinite supply of shares at prices above A, and infinite demand at prices below B. Traders arrive to the market according to a Poisson process with a parameter $\lambda > 0$. The number of traders within interval τ is distributed proportionally to $\lambda\tau$, expected time between trader arrivals is therefore $1/\lambda$. Traders can either submit a market order to get immediate execution or submit a limit order, which will improve the price, but delay execution. Traders waiting costs are proportional to the time it takes for an order to execute; such dependency highlights the main trade-off: immediacy versus price improvement. Patient traders incur a cost δ_1 per unit of time, impatient – δ_2 ($\delta_2 > \delta_1$). Δ is the smallest tick ($0.01, for example), a is the lowest ask, b – highest bid (A > a > b > B). s is the inside spread a – b. We express these variables in terms of tick size Δ instead of dollars.

We also make following simplifying assumptions: submitted orders cannot be cancelled or modified; submitted limit orders must be price improving – i.e. they must narrow the spread at least by one tick; buyers and sellers arrive alternatively with certainty. Let's denote the price paid by a buyer as p_{buyer}, price paid by a seller as p_{seller}.

Buyer can either buy at the ask or submit a limit order, which will be entered on the buy side of the book creating a narrower spread. The argument for a seller is symmetrical.

$$p_{buyer} = a - j, \quad p_{seller} = b + j, \text{ where } j \in [0, \dots s-1].$$

We call an order that creates a spread of size j a *j-limit order*. j = 0 is a market order. T(j) is the expected time-to-execution for a j-limit order. Since the cost of waiting is linear in time, we assume it's $\delta T(j)$. The expected profit of a trader i ($i \in 1,2$), as a function of a limit order he submits is the following (in dollars):

$$\Pi_I(j) = \begin{cases} V_{buyer} - p_{buyer}\Delta - \delta_i T(j) = (V_{buyer} - a\Delta) + j\Delta - \delta_i T(j), \text{ if i is a buyer} \\ \\ p_{seller}\Delta - V_{seller} - \delta_i T(j) = (b\Delta - V_{seller}) + j\Delta - \delta_i T(j), \text{ if i is a seller} \end{cases}$$

where V_{buyer} and V_{seller} are the valuations of buyers and sellers. The part of the profits in parenthesis is determined by valuations and current inside quotes – factors exogenous to the trader's order submission problem. Therefore, both buyers and sellers have to solve an identical optimization problem for a spread of size s:

$$\max \pi_i(j) \equiv j\Delta - \delta_i T(j), \text{ for all } j \in [0,\dots, s\text{-}1]. \qquad (15)$$

The optimal placement strategy thus assigns a j-limit order $j \in [0,\dots, s\text{-}1]$ for every possible spread size s. We denote the order placement strategy of a trader of type i as $o_i(.)$.

We now determine the function for the expecting waiting time T(j). Suppose that a trader places a j-limit order. Then let's denote by $\alpha_k(j)$ a probability that the next trader responds with a k-limit order, where $k \in [0, \dots, j\text{-}1]$. It can be shown that

$$T(j) = \begin{cases} 1/\lambda, \text{ if } j = 1, \\ 1/\alpha_0(j)[1/\lambda + \sum_{k=1}^{j-1} \alpha_k(j)T(k)], \text{ if } \alpha_0(j) > 0 \text{ and } j > 1, \\ \infty, \text{ if } \alpha_0(j) = 0 \text{ and } j > 1. \end{cases}$$

Remember that a buyer's arrival will be always followed by an arrival of a seller and vice versa; therefore, when the first trader establishes a spread of one tick (j=1), the following trader has no choice but to submit a limit order, which makes time-to-execution equal to the expected time of the next trader's arrival, which is $1/\lambda$.

From (15) the profit from submitting a market order is zero; therefore, traders will submit limit orders only when the price improvement $j\Delta$ exceeds the cost of waiting

$\delta_i T(j)$. When a trader submits a limit order, he has to wait at least one period ($1/\lambda$), which makes the smallest cost for a trader of type i equal to δ_i/λ. This implies that the smallest spread trader i can establish is the smallest integer j_i^R, such that $\pi_i\ j_i^R = j_i^R \Delta - \delta_i/\lambda \geq 0$. We refer to j_i^R as *reservation spread* and write it as:

$$j_i^R = \lceil \delta_i/\lambda\Delta \rceil,\ i \in [1,2].$$

When we consider patient and impatient traders, by definition $j_1^R \leq j_2^R$. In the case when $j_1^R = j_2^R$ we call the two types of traders *indistinguishable*. Let's consider this homogenous case. By definition of the reservation spread, all traders would prefer to submit a market order when the spread is less then j_i^R. This implies that the waiting time for a j_i^R-limit order is one time period. Thus

$$\pi_i\ j_i^R \geq 0 \text{ for } i \in [1,2].$$

But this also means that when a trader sees a spread larger then j_i^R, he will ever submit a market order, but will go with a j_i^R-limit order instead. Therefore, the probability of execution for any limit order with $j > j_i^R$ is zero. The optimal order placement strategy in this homogenous case is the following: for any spread s, if $s > j_i^R$, traders will submit j_i^R-limit orders, and when $s < j_i^R$, they will submit market orders. In this rather simplistic market, the spread will alternate between very large (A-B) and small (j_i^R), and transactions will take place only when the spread is small. Perhaps surprisingly, this is one of the patterns that have been identified by [Bias et al., 1995] in the study of transactions on the Paris Bourse (presently Euronext).

Let's now turn to the case of heterogeneous traders, where $j_1^R < j_2^R$. It can be shown that in the simplest case, the state variable s can be partitioned into three regions:

(1) $s < j_1^R$ – all traders submit market orders,

(2) $j_1^R < s < s_c$ – patient traders submit limit orders, impatient – market orders,

(3) $s > s_c$ – all traders submit limit orders,

where s_c is some "cutoff" spread. Furthermore, we can partition the space of possible spreads into q spreads ($2 \leq q \leq K$), $n_1 < \ldots < n_q$, such that when a patient trader sees a spread equal to n_{h+1}, he will submit such a limit order that will shrink the to n_h. In such equilibrium, the expected waiting time can be expressed as follows:

$$T(n_1) = 1/\lambda;\ T(n_h) = 1/\lambda(1 + 2\sum_{k=1}^{h-1} r^k), \text{ for } h = 2, \ldots, q-1, \tag{16}$$

and

$$T(j) = T(n_h), \text{ when } n_{h-1} \leq j < n_h,$$

Where $r = \theta/(1-\theta)$ is the proportion of patient traders to impatient traders. Limit order does not execute until all limit orders at lower prices execute. Therefore, choosing the size of the new spread is tantamount to choosing the order's placement in the priority queue.

Suppose a trader arrives to the market and sees a spread of n_{h+1}. In the equilibrium, he is supposed to submit an order that will reduce the spread to n_h. But this means that it must be more profitable for the trader to submit an n_h-limit order rather than an order that results in a tighter spread and faster execution. The following condition must be satisfied:

$$n_h\Delta - T(n_h)\delta_1 \geq n_{h-1}\Delta - T(n_{h-1})\delta_1,$$

or

$$\Psi_h = n_h - n_{h-1} \geq [T(n_h) - T(n_{h-1})]\delta_1/\Delta. \tag{17}$$

Now, when the next trader arrives and observes a spread of n_h, it must be optimal for him to submit an n_{h-1}-limit order:

$$n_{h-1}\Delta - T(n_{h-1})\delta_1 \geq (n_h - 1)\Delta - T(n_h - 1)\delta_1,$$

thus

$$\Psi_h < [T(n_h) - T(n_{h-1})]\delta_1/\Delta - 1. \tag{18}$$

Combining (17) and (18) and substituting (16), we can define Ψ_h as:

$$\Psi_h = \lceil (T(n_h) - T(n_{h-1}))\delta_1/\Delta \rceil = \lceil 2r^{h-1}\delta_1/(\lambda\Delta) \rceil. \tag{19}$$

We call Ψ_h "*spread improvement*" when the spread is equal to n_h. It determines the aggressiveness of the incoming trader's limit order when the observed spread is n_h. Equation (19) has a simple econometric explanation: it equates price concession Ψ_h with the improvement in the waiting cost $\lceil (T(n_h) - T(n_{h-1}))\delta_1/\Delta \rceil$. This equation also tells us that the set of equilibrium spreads is given by:

$$n_1 = j_1^R; \qquad n_q = A - B; \quad n_h = n_1 + \sum_{k=2}^{h} \Psi_h \text{ for } 1 < h < q. \qquad (20)$$

Equation (20) specifies for each spread size by how much will it be improved with an arrival of a new patient trader. Ψ_h increases as the proportion of patient traders, θ, increases, waiting cost δ_1 increases, and order arrival rate λ decreases.

This model can also tell us about *market resiliency* – how fast does a market reverts to a "competitive" j_1^R spread. We can define resiliency as a probability that the spread will return to j_1^R in the next time step. For example, in the market with homogeneous agents, R = 1, whereas in the case of heterogeneous traders it may take multiple steps to reach a competitive spread. We can use our model to show that markets dominated by patient traders ($r > 1$) are more resilient and have tighter spreads than those dominated by impatient traders ($r < 1$). For an empirical study of market resiliency see [Degryse et al., 2001], where the authors run several regressions to quantify the impact of large ("aggressive") orders on the limit order market. Our current model also gives us a number of other insights regarding the size of the spread, time-to-execution, intraday patterns, etc. – see [Foucault et al., 2003] for more details. Although we have made several assumptions in our derivations, they should not limit the pertinence of our conclusions – it can be shown that allowing for cancellation or submission at less competitive prices will not alter the end result. Several studies, such as [Bouchaud et al., 2002], suggest that the relationship between the distance from the inside market and the probability of execution are non-linear; this can be incorporated in the current model, but will make it more algebraically complex.

What are the implications of our model for an automated trader? First of all, it confirms the need for an intermediary in a limit-order market: too many impatient traders lead to large spreads, high trading costs; presence of an electronic dealer in such market will mitigate this tendency by providing additional committed liquidity. From the technical point of view, we can now analytically quantify the trade-off between the price improvement/concession and execution time. This model uses proportions of patient and impatient traders in the market to calculate order placement strategies, but these values are not observable. What we can do instead is estimate the trade-off between time-to-execution and order aggressiveness directly from the transaction records. It is very important to remember that the current approach concentrates on a single aspect of the limit order trading: price improvement vs. time-to-execution, while ignoring a number of other realities, such as inventory and information effects. This means that results presented here are to be used in conjunction with other models, such as the ones from earlier sections. Creating a comprehensive model of limit order trading is very difficult because of the interplay of a great number of factors. For example, [Harris, 1997] explicitly models informed traders, but his overall model ends up being highly abstract and not very helpful to a market making agent who's mostly interested in "recipe"-type models.

As previously mentioned, estimation of the time-to-execution is central to many limit order trading models, and [Lo et al., 2000] suggest another approach to calculating this parameter. The method is based on the concept of a "first-passage time" of the stock price through the limit price. The idea is very simple: if we treat the stock price (transaction price, for example) as a stochastic process, then the time-to-execution is the time when the price process first hits the limit price level.

It is a generally acceptable practice to model a stock price as a geometric Brownian motion with drift:

$$dP(t) = \alpha P(t)dt + \sigma P(t)dW(t),$$

where α and σ are constants. A limit buy submitted at time t_0 at a price P_l will execute over the period $[t_0, t_0 + t]$ only if the smallest observed price over that period P_{min} is less than or equal to P_l. A well-known formula can be modified as follows:

$$\Pr(P_{min} \le P_l \mid P(t_0) = P_0) = 1 - \Phi\,[(\log(P_0/P_l) + \mu t)/\sigma\sqrt{t}\,) + (P_l/P_0)^{2\mu/(\sigma\sigma)}\Phi\,[(\log(P_0/P_l) + \mu t)/\sigma\sqrt{t}\,),\ P_l < P_0, \qquad (21)$$

where $\mu = \alpha - 0.5\sigma^2$, $\Phi(.)$ is the standard normal cumulative distribution function (CDF). A very similar expression can be written for sell orders. If we define T as the time-to-execution for a limit order, then (21) yields the following CDF, F(t) for a given T:

$$F(t) = \Pr(T \le t \mid P(t_0) = P_0, P_l, \mu, \sigma, X) = \Pr(P_{min} \le P_l) \qquad \text{(buys)}$$

$$F(t) = \Pr(T \le t \mid P(t_0) = P_0, P_l, \mu, \sigma, X) = \Pr(P_{max} \ge P_l) \qquad \text{(sells)},$$

where X is a vector of (optional) explanatory variables, such as previous trade, daily volume, closing price, etc. We can compare this theoretical distribution with the real-world one to calibrate the parameters we include into X. A number of statistical techniques can be used to assess the relative importance of various factors. Parameters like μ and σ can also be estimated from historical data using maximum likelihood learning:

$$\mu = 1/(N\tau) \sum_{j=1}^{N} r_j; \qquad \sigma^2 = 1/N \sum_{i=1}^{N} [(r_j - \mu\tau)^2/\tau],$$

where N is the number of observations, $r_j = \log P_j - \log P_{j-1}$ (compounded return), and τ is the time interval. This is a way for an automated market maker to estimate the probability of a limit order execution, given the current state of the book. This goal is very similar in spirit to the one of the first model. While this approach may not be as intuitively compelling as the earlier one, it is more general, and it uses observable variables. One setback of this approach is that we can pack an arbitrary number of features into vector X (i.e. he second best orders, number of orders in the book, depth at various levels, etc.) to achieve a near perfect fit. We have to pay extra attention to the feature selection process to make sure that the output model remains robust.

c. More Comprehensive Model

While a complete model of limit order trading, which includes information and inventory effects, monitoring costs, order cancellations and resubmissions, will be prohibitively complex, here we present an approach that adds one more level of complexity to the kind of model described earlier. In addition to the price improvement vs. time-to-execution trade-off, we need to address the information aspect inherent to any type of trading. This goes back to the point made in the beginning of this chapter about how a dealer's quote can be regarded as a free option to the rest of the market to trade at a pre-determined price. This option will be exercised when the value of the underlying security moves above or below this price, so that profitable trading opportunities arise. The same holds true for limit orders: when new information arrives to the market (be it public or private), outstanding limit orders can be "picked-off" by opportunistic traders, if they are not cancelled in time. One way to combat this problem is to constantly monitor all outstanding quotes (for market makers) or limit orders (all traders), which can be expensive and has to be included into the overall cost of trading – see [Foucault et al. 2003b] for a structured analysis of monitoring costs. Furthermore, it is not always possible to cancel and re-submit an order even with constant monitoring. It's the same argument as in simple information models where you can never tell with certainty if your counterparty is informed or not. When a limit order transacts, it can be either someone demanding liquidity (and you taking advantage of the spread), or someone considers your price advantageous (and you lose money since the underlying value has moved). This reality must be taken into account by an electronic market maker in the limit order market. Below we present a model developed in [Hollifield et al., 2003b] and expanded in [Hollifield et al., 2003a].

At time t, one risk neutral trader arrives to the market and has an opportunity to submit market or limit order. He wishes to trade q_t shares, and his valuation of the stock is

$$v_t = y_t + u_t,$$

where y and u are public (common) and private components respectively. The common portion of the valuation changes with new information:

$$y_{t+1} = y_t + \delta_{t+1}, \qquad \text{where } E[\delta_{t+1}] = 0.$$

The innovation of the common value is conditioned on a finite set of prior innovations (i.e. we may observe ARIMA type of relationship between current and past values). Traders can have different private values for portfolio and other exogenous reasons. All private values are drawn from the same continuous distribution:

$$\Pr_t(u_t \le u) = G_t(u), \tag{22}$$

where time subscript t indicates that values are conditioned on the information available before period t and previous historical values of private information.

At a random time $t + \tau_{cancel}$, the trader who arrived at t is forced to cancel all his unfilled limit orders. The cancellation time is not revealed to the trader in advance, but is drawn instead from a distribution conditional on same statistics as private an public information. If Υ is the maximum possible lifetime of the order, then

$$\Pr_t(\tau_{cancel} \leq \Upsilon < \infty) = 1.$$

The trader's quantity q_t is independent of v_t and is again drawn from a distribution conditioned on prior events.

In choosing between a market and a limit order, the trader observes the state of the limit book, current common value, and a history of value innovations. The trader pays cost c per share to submit an order of either type. The decision indicators $d_{t,s}^{sell} \in \{0,1\}$ for $s = 0, 1, \ldots, S$, and $d_{t,s}^{buy} \in \{0,1\}$ for $b = 0, 1, \ldots, B$. The trader can choose from a bounded set of order prices $S < \infty$, $B < \infty$. If he decides to submit a market buy, then the trade price is the current best (lowest) ask, and $d_{t,0}^{buy} = 1$. If the trader wishes to submit a buy order s ticks below the best ask, then $d_{t,s}^{buy} = 1$. It works in the same way for sell orders. If the trader doesn't want to transact at time t, then $d_{t,s}^{sell} = d_{t,s}^{buy} = 0$ for all s.

Let's say a trader with valuation v_t submits a buy order for q_t shares at a price $p_{t,b}$, which is b ticks below the best ask – $d_{t,b}^{buy} = 1$. Define $dQ_{t,t+\tau}$ as a number of shares that transact before the cancellation. This is an unknown quantity, since τ is a random variable; we do know, however that $dQ_{t,t+\tau} = 0$ for all $t > \tau$. In any event, the total quantity bought/sold cannot exceed the total quantity submitted through the limit order:

$$\sum_{r=1}^{Y} dQ_{t,t+\tau} \leq q_t.$$

The payoff to the trader at time $t + \tau$ from transacting $dQ_{t,t+\tau}$ at a price $p_{t,b}$ is

$$dQ_{t,t+\tau}(y_{t+\tau} + u_t - p_{t,b}) = dQ_{t,t+\tau}(v_t - p_{t,b}) + dQ_{t,t+\tau}(y_{t+\tau} - y_t).$$

$dQ_{t,t+\tau}(v_t - p_{t,b})$ represents the payoff that would result from immediate execution, $dQ_{t,t+\tau}(y_{t+\tau} - y_t)$ is the portion attributed to the change in value overtime. When we sum over all possible execution times and encorporate transaction costs, the aggregate realized payoff is

$$U_{t,t+\Upsilon} = \sum_{\tau=0}^{Y} dQ_{t,t+\tau}(v_t - p_{t,b}) + \sum_{\tau=0}^{Y} dQ_{t,t+\tau}(y_{t+\tau} - y_t) - q_t c. \qquad (23)$$

Execution probability for a buy order can be defined as follows:

$$\psi_t^{buy}(b, q_t) \equiv E_t[\sum_{\tau=0}^{Y} dQ_{t,t+\tau}/q_t \mid d_{t,b}^{buy} = 1, q_t], \qquad (24)$$

and the picking off risk as

$$\xi_t^{buy}(b, q_t) \equiv E_t[\sum_{\tau=0}^{\Upsilon} (dQ_{t,t+\tau}/q_t)(y_{t+\tau} - y_t) \mid d_{t,b}^{buy} = 1, q_t]. \quad (25)$$

Note that conditional probabilities (24) and (25) do not depend on the trader's private value; probability of execution of the market order is one, risk of being picked off is zero.

The trader's *expected payoff* is the expected value of (24), conditional on the trader's information, the state of the order book, the common value, and past values:

$$E_t[U_{t,t+\Upsilon} \mid d_{t,b}^{buy} = 1, u_t, q_t] = q_t\psi_t^{buy}(b, q_t)(v_t - p_{t,b}) + q_t\xi_t^{buy}(b, q_t) - q_t c.$$

The trader submits the order that maximizes his expected payoff, conditional on his information, private value, and order quantity:

$$\max \sum_{s=0}^{S} d_{t,s}^{sell} E_t[U_{t,t+\Upsilon} \mid d_{t,s}^{sell} = 1, u_t, q_t] + \sum_{b=0}^{B} d_{t,b}^{buy} E_t[U_{t,t+\Upsilon} \mid d_{t,b}^{buy} = 1, u_t, q_t]$$

subject to

$$d_{t,s}^{sell} \in \{0,1\} \text{ for } s = 0, 1, \dots, S, \; d_{t,s}^{buy} \in \{0,1\} \text{ for } b = 0, 1, \dots, B,$$

and

$$\sum_{s=0}^{S} d_{t,s}^{sell} + \sum_{b=0}^{B} d_{t,b}^{buy} \leq 1. \quad (26)$$

The last condition ensures that only one order is submitted.

It can be shown that if one trader has private value of u, and optimally submits a limit buy b ticks below the best ask, and another trader with private value u\` > u optimally submits an order b\` ticks below the best ask, then $\psi_t^{buy}(b`, q_t) \geq \psi_t^{buy}(b, q_t)$ and b\` ≤ b. Similar results hold on the sell side as well. This means that optimal order submission strategy depends on the trader's private value. We can therefore partition the set of valuations into intervals, such that all traders wishing to trade the same quantity with valuations in the same interval will submit the same order. Let's define the threshold valuation $\theta_t^{buy}(b, b`, q_t)$ as a valuation of a trader who is indifferent between submitting a buy order at price $p_{t,b}$ and a buy at price $p_{t,b`}$:

$$\theta_t^{buy}(b, b`, q) = p_{t,b} + [(p_{t,b} - p_{t,b`})\psi_t^{buy}(b`, q) + (\xi_t^{buy}(b`, q) - \xi_t^{buy}(b, q))]/ (\psi_t^{buy}(b, q) - \psi_t^{buy}(b`, q)).$$

The threshold valuation for a buy order at price $p_{t,b}$ and not submitting any order at all:

$$\theta_t^{buy}(b, NO, q) = p_{t,b} + (-\ \xi_t^{buy}(b, q) + c)/\ \psi_t^{buy}(b, q).$$

Similar thresholds can be defined for sell orders:

$$\theta_t^{sell}(s, s`, q) = p_{t,s} + [(p_{t,s`} - p_{t,s})\psi_t^{sell}(s`, q) + (\xi_t^{sell}(s, q) - \xi_t^{sell}(s`, q))]/ (\psi_t^{sell}(s, q) - \psi_t^{sell}(s`, q)),$$

and

$$\theta_t^{sell}(s, NO, q) = p_{t,s} + (-\ \xi_t^{sell}(s, q) + c)/\ \psi_t^{sell}(s, q).$$

Threshold valuation between a buy and a sell order is

$$\theta_t(s, b, q) = p_{t,s} + [(p_{t,b} - p_{t,s})\psi_t^{buy}(b, q) + (\xi_t^{sell}(s, q) - \xi_t^{buy}(b, q))]/ (\psi_t^{sell}(s, q) - \psi_t^{buy}(b, q)).$$

Let $B_t^*(q)$ index the set of buy order prices that are optimal for some trader who wishes to trade q shares at time t,

$$B_t^*(q) = \{b \mid d_t^{buy*}(b, u, q) = 1 \text{ for some } u\},$$

with elements $b_{i,t}^*(q)$ for i = 1, …, I, ordered by execution probabilities. $S_t^*(q)$ is the same index for sell orders.

It can be further shown that

$$\theta_t^{buy}(b_{1,t}^*(q), b_{2,t}^*(q), q) > \theta_t^{buy}(b_{2,t}^*(q), b_{3,t}^*(q), q) > \ldots > \theta_t^{buy}(b_{I-1,t}^*(q), b_{I,t}^*(q), q),$$

$$\theta_t^{sell}(s_{J-1,t}^*(q), s_{J,t}^*(q), q) > \theta_t^{sell}(s_{J-2,t}^*(q), s_{J-1,t}^*(q), q) > \ldots > \theta_t^{sell}(s_{1,t}^*(q), s_{2,t}^*(q), q),$$

$$\theta_t^{buy}(b_{I-1,t}^*(q), b_{I,t}^*(q), q) > \theta_t(s_{J,t}^*(q), b_{I,t}^*(q), q) > \theta_t^{sell}(s_{J-1,t}^*(q), s_{J,t}^*(q), q).$$

To describe the optimal decision rule, we define the marginal thresholds (trading vs. not trading) for sellers and buyers as

$$\theta_t^{buy}(\text{Marginal}_t(q), q) = \max(\theta_t(s_{J,t}^*(q), b_{I,t}^*(q), q), \theta_t^{buy}(b_{I,t}^*(q), NO, q)),$$

$$\theta_t^{sell}(\text{Marginal}_t(q), q) = \min(\theta_t(s_{J,t}^*(q), b_{I,t}^*(q), q), \theta_t^{sell}(s_{J,t}^*(q), NO, q)).$$

If marginal thresholds for buyer and seller are equal, all traders find it optimal to submit an order. Otherwise, some traders may refrain from trading.

Finally, we arrive to the following optimal submission strategy:

$$d_t^{buy*}(b, u, q) = 1, \text{ if } \begin{cases} b = b_{1,t}^{*}(q) \text{ and} \\ \theta_t^{buy}(b_{1,t}^{*}(q), b_{2,t}^{*}(q), q) \le y_t + u < \infty, \\ \text{or} \\ b = b_{i,t}^{*}(q) \text{ for } i = 2, \ldots, I-1 \text{ and} \\ \theta_t^{buy}(b_{i,t}^{*}(q), b_{i+1,t}^{*}(q), q) \le y_t + u < \theta_t^{buy}(b_{i-1,t}^{*}(q), b_{i,t}^{*}(q), q), \\ \text{or} \\ b = b_{I,t}^{*}(q) \text{ and} \\ \theta_t^{buy}(\text{Marginal}_t(q), q) \le y_t + u < \theta_t^{buy}(b_{I-1,t}^{*}(q), b_{I,t}^{*}(q), q). \end{cases}$$

$$d_t^{sell*}(b, u, q) = 1, \text{ if } \begin{cases} s = s_{1,t}^{*}(q) \text{ and} \\ -\infty \le y_t + u < \theta_t^{sell}(s_{1,t}^{*}(q), s_{2,t}^{*}(q), q), \\ \text{or} \\ s = s_{j,t}^{*}(q) \text{ for } j = 2, \ldots, J-1 \text{ and} \\ \theta_t^{sell}(s_{j-1,t}^{*}(q), s_{j,t}^{*}(q), q) \le y_t + u < \theta_t^{sell}(s_{j,t}^{*}(q), s_{j+1,t}^{*}(q), q), \\ \text{or} \\ s = s_{I,t}^{*}(q) \text{ and} \\ \theta_t^{sell}(s_{J-1,t}^{*}(q), s_{J,t}^{*}(q), q) \le y_t + u < \theta_t^{sell}(\text{Marginal}_t(q), q). \end{cases}$$

Otherwise,

$$d_t^{buy*}(b, u, q) = d_t^{sell*}(b, u, q) = 0.$$

Let $V_t(y_t + u, q)$ be the indirect utility function for a trader at time t with valuation $y_t + u$ and quantity q. This function can be computed by substituting the optimal order submission strategy into the trader's objective function (26). V_t has the following properties:

- it's a positive convex function of $y_t + u$,
- if $d_t^{buy*}(b, u, q) = 1$, then for u` > u, $V_t(y_t + u', q) > V_t(y_t + u, q)$,

- if $d_t^{sell*}(b, u, q) = 1$, then for $u` < u$, $V_t(y_t + u', q) > V_t(y_t + u, q)$,

Here is an example of such function:

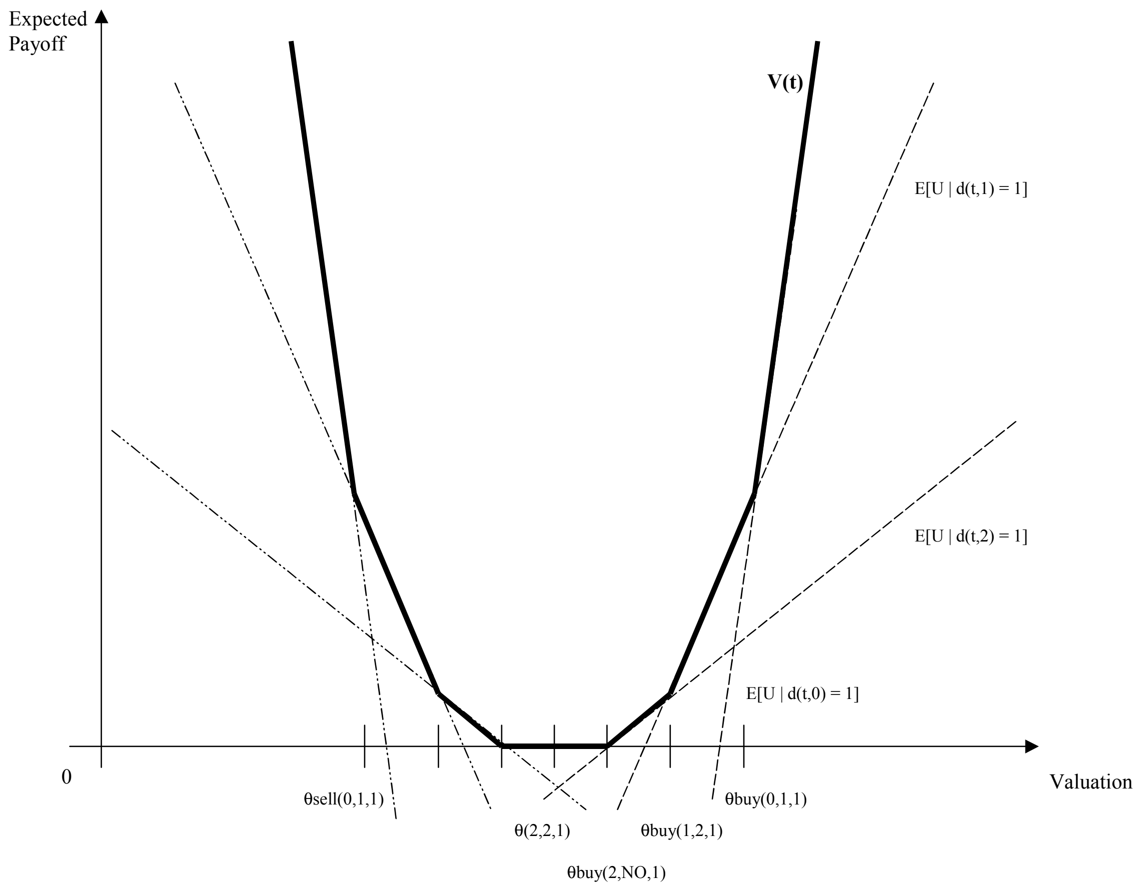

Figure 10.

This graph essentially depicts, for a given valuation, the trade-off between how far is the submitted order from the opposing best quote and the expected profit for submitting such an order. Buy orders are on the right side of the graph, which corresponds to higher valuations, with sell orders being on the left, with low valuations. Note that the closer is the order to the opposite side of the market, the steeper is the payoff function for that order. In our model, payoff is a linear function of the trader's valuation, with the execution probability as a slope. The indirect utility function V_t is plotted in a thick line and represents the optimal order submission strategy. The increase in cost c will result in a parallel downward shift in payoff functions for all orders. A change in the picking-off risk or execution probability for a specific order, shifts its payoff function down, with the rest of schedules remaining unchanged. Intersections of these payoff functions are the thresholds $\theta_t(.)$ from the discussion above. It is obvious from Figure 10 how the trader's optimal order changes when his valuation passes through the threshold. The horizontal section of V_t represents the valuation window where no orders are submitted.

Now that we have derived the optimal order submission strategy, we can define the probability of seeing an order of a given size a certain distance away from the opposite side of the market:

$$\Pr_t(d_t^{sell*}(s_{1,t}^*(q), u_t, q) = 1 \mid q) = \Pr(y_t + u_t < \theta_t^{sell}(s_{1,t}^*(q), s_{2,t}^*(q), q) \mid q)$$

$$= G_t(\theta_t^{sell}(s_{1,t}^*(q), s_{2,t}^*(q), q) - y_t), \tag{27}$$

where G comes from equation (29). For j = 2, …, J – 1:

$$\Pr_t(d_t^{sell*}(s_{j,t}^*(q), u_t, q) = 1 \mid q) = G_t(\theta_t^{sell}(s_{j,t}^*(q), s_{j+1,t}^*(q), q) - y_t) - G_t(\theta_t^{sell}(s_{j-1,t}^*(q), s_{j,t}^*(q), q) - y_t), \tag{28}$$

and for j = J:

$$\Pr_t(d_t^{sell*}(s_{J,t}^*(q), u_t, q) = 1 \mid q) = G_t(\theta_t^{sell}(\text{Marginal}_t(q), q) - y_t) - G_t(\theta_t^{sell}(s_{J-1,t}^*(q), s_{J,t}^*(q), q) - y_t). \tag{29}$$

Same expressions are derived for buy orders. Theoretical probabilities (27) through (29) allow us to explain certain empirical dependencies present in the limit order markets. For example [Bias et al., 1995] found that a larger ask depth increases the probability of a market sell and decreases the probability of a limit sell. This can be explained by the fact that the increased depth decreases the execution probability of a one tick limit order, which increases the threshold valuation for a market sell. This, in turn, increases the probability of seeing a market sell, as (27) – (29) tell us. We can similarly explain the observation that wider spreads induce more limit orders and autocorrelations in the order flow.

The very nice feature of this model is that although it is built around the trader's individual preferences, we don't need to use those when we apply this model to real-world data. All we need are the execution probabilities, which tie in with the rest of the model through (27) – (29) and are easy to estimate through transaction records.

For an electronic trader, this model helps to understand the mechanics of the order submission behavior of other market agents, and, second, to ensure that his own quotes lie on the optimal frontier V_t. One specific application: an opportunistic liquidity provider can scan a limit order book and "fill it in" with additional liquidity on all the efficient levels. The one shortcoming of this model, from the practical point of view, is its exogenous treatment of the quantity submitted. This issue has been addressed in other studies.

[Sandas, 2001] internalizes the order quantity by recognizing the fact that every limit order submitted by a trader should be expected to at least break even. What we can do is, for each possible price, calculate the maximum quantity for which expected profit – like in equations (15) and (26) – is positive. This will be the optimal quantity to submit for a competitive market maker attempting to replenish the book. More specifically, [Sandas, 2001] incorporates information effects of trading, like those in section B, into their model. There, the expected profit is

$$0.5(p_t - c - X_t - \alpha(q + \lambda))e^{-q/\lambda},$$

where c is a fixed per-trade cost, q is the quantity submitted, X_t is the common value of a security, e is the price innovation, α and λ are coefficients. The optimal quantity to be submitted at the best sell price level p_1 in this case is

$$Q_1 = (p_t - c - X_t)/\alpha - \lambda.$$

The same number can be computed for all other price levels. To get rid of the unobservable common value X_t, we can differentiate transaction prices and use the momentum to replace fluctuations in the intrinsic value.

Another explicit solution for the depth of the market maker's quotes is offered in [Kavajecz, 1998]. It is again based on the information effects of trading, but the optimal bid and ask prices and quantities attached to those quotes are expressed as functions of the number of informed and uninformed traders in the market place. Although such variables are not easily observable in the real world, these ratios can be estimated from past transactions.

[Ahn et al., 2001] adopts a different approach to modeling the quoted depth by tying it to the short-term ("transitory") volatility of the stock. Such dependence can be established by a simple regression:

$$q_t = \alpha + \beta\, RISK_{t-1} + \theta\, NTRADE_{t-1} + \gamma\, TIME_t + \rho\, q_{t-1} + \varepsilon_t,$$

where q is depth, RISK is volatility, NTRADE is net order flow, TIME is a dummy variable to account for the time of the day, ε is the idiosyncratic component, and α, β, θ, γ, and ρ are coefficients. This approach represents a step away from the pure market microstructure theory, since it employs proxy variables.

The final point to bring up in connection with the limit order trading, is its possible link with technical analysis, as pointed out in [Kavajecz and Odders-White, 2002]. While academic literature has historically kept technical analysis in low regard, and in practice it still remains more of an art than science, market microstructure research may offer some theoretical explanations for the "trading rules" employed by technical analysts. One illustrative example is a strong statistical correlation between the "support levels" and "resistance levels" predicted by technical analysis and the presence of "depth peaks" in limit order books. Indeed when there is extra liquidity present on one side of the market, it is difficult to trade through it; therefore, if we only look at transaction prices, it may appear as if they reached some invisible barrier, while in fact we have just stumbled upon a pocket of liquidity. This finding has several implications: first, technical analysis is not necessarily at odds with the academic concept of efficient markets; and, second, perhaps we can use technical analysis to find liquidity instead of predict prices, which maybe very beneficial for an electronic dealer.

D. Summary.

In this chapter, we established the theoretical foundation of Market Microstructure and suggested formal models to describe the process of price formation and behavior of different market participants.

We take the concept of perfectly efficient markets as our starting point and describe how prices and investors should behave with no market frictions. We mention a large number of empirical studies that show the limitations of the Efficient Markets Hypothesis and prices as random walks. We explain these inconsistencies by market frictions, such as spreads, market impact, institutional trading, and so on. We examine the effects of various frictions on transaction prices – how they disrupt the random walk behavior.

We offer a number of formal models that allow us to explain and quantify the short-term behavior of prices, spreads, volumes and other microstructure variables. We started in the dealership market environment, where all trading happens through a designated market maker. In this simplified setting, we presented the two fundamental problems that a dealer faces: inventory management and information misbalance. We developed analytical models for quote updates that address each of these problems separately and then combined these approaches by decomposing the bid-ask spread into inventory, information, and order processing components. We have also examined liquidity provision in limit order markets – a much more realistic environment of great practical interest. We pointed out the inherent trade-offs in limit order trading and presented two theoretical models: the first one mainly showed how to quantify the trade-off between the price improvement and execution probability, while the second one also added the cost of waiting, order cancellation, and the "picking-off" risk. We have suggested a way to endogenize the depth of the dealer's quote in this setting.

This evolution of models represents a theoretical roadmap to creating an electronic trader: we have detailed all the risks a trader is facing, showed how to quantify and manage them, and proposed several ways to combine different solutions into one coherent approach.

Empirical Analysis: Reinforcement Learning Approach

So far in this thesis, we gave a brief overview of the mechanics of modern financial markets, surveyed previous work in the field of market microstructure, and described our experimental setup. We then laid out the theoretical foundation for our work: we showed how real world markets can deviate from perfectly efficient behavior, and suggested a number of theoretical models capable of capturing trader's behavior on a microstructure level. In this chapter, we present our empirical methods and result, which we consider our main contribution to the fields of Automation and Market Microstructure.

First, we explain the transition from theoretical models of agent's behavior to empirical studies of prices, spreads, quotes, volumes, and other observable variables. We then present our initial results that were obtained by following conventional empirical techniques. In analyzing these results, we point out inherent shortcomings in what we call "explanatory" studies: most currently-used methods aim to simply find relationships among microstructure variables. We claim that such attitude is ineffective from the practical point of view and that we need to go one step further – to not just establish that A influences B, but also to determine how such relationship can be exploited to make marketplace interactions more optimal. Our hope is that this normative approach to market microstructure research will not only lead to higher well-being of market participants who optimize their actions, but will also create more efficient and stable financial markets. Our optimization technique uses reinforcement learning to derive efficient trading strategies. We describe the algorithm, present our own explanatory model of market microstructure, and report empirical results – the performance of our algorithm on different stock, under various constraints and market conditions. Finally we analyze the strength and weaknesses of our normative method.

A. Motivation for Normative Approach

1. Current Empirical Methods and Sample Models

In the previous chapter, we have presented a number of theoretical models. While they do provide an insight into the inner-working of financial markets, their practical application is questionable. In other words, it is not clear how well these models fit the real-world markets. The problem with theoretical models is that they are built around quantities that are difficult or impossible to observe and record: proportion of informed and uninformed traders, trader's risk preferences, order arrival processes, "preferred" portfolios, inventory levels, and so on and so forth. Therein lies the problem with modeling market participants – a model becomes difficult to reconcile with available empirical data.

Financial data sources that are easily accessible to both researchers and practitioners (i.e. exchanges, ATS, news services, etc), provide information about prices, quotes, spreads, volumes, limit order book composition, etc, and it is the goal of empirical models to make inferences about market microstructure effects from this observable information. While theoretical models build out complex structures of interacting economic agents, an alternative approach to looking at market events is to just treat prices, quotes, spreads, and volumes as streams of numbers generated by some abstract processes without explicitly worrying about the real-world aspects that influence these variables. Therefore, when we examine time series of prices and quotes, we regard them as a system characterized by auto- and cross-correlations of a very general nature. This takes us back to the discussion of how market frictions affect prices: if we observe consistent negative autocorrelations in transactions prices, we can infer the presence of liquidity-providing agents in the marketplace, but if we see positive autocorrelations in our data, we assume that momentum trader have more influence, and so on.

The goal is to determine if there is some dependency in the time-indexed series of observations of a single variable or a relationship among a number of different variables. When a statistical dependency is established, we can go back to the market microstructure theory and attempt to explain it using the same factors as in theoretical models, thereby bringing our finding of purely statistical nature back to the world of traders, investors, and exchanges. Since we do not make any fundamental ex-ante assumptions about the forces that influence the behavior of microstructure variables, we often refer to such empirical models as *relaxed* models (as opposed to *structured* or theoretical models). Relaxed models can be used on past transaction records for explanatory purposes (i.e. spreads revert to mean), but they can also be employed for financial forecasting. If we can describe behavior of some variables as a function (of either its own past values or values of other variables), then we can extend this function into the future and therefore predict the evolution of such variable.

A quick example to illustrate the connection between structured and relaxed models. Suppose we examine two time series – transaction prices and dealer's quotes – for a particular stock, and find out that price covariance is $-0.25s^2$ and that quote covariance is zero (s is the spread). Armed with this statistical information, we can turn to the microstructure theory, which unambiguously tells us that we are dealing with the kind of market where the dealer only faces fixed order-processing costs. Here, just by looking at a series of numbers, we are able to make an inference about participating agents' behavior.

a. Estimating Information Component of Trades

Having explained the general idea behind empirical approach to market microstructure, we can now turn to specific models like the ones developed in [Hasbrouck, 1991a] and expanded in [Hasbrouck, 1991b]. Similarly to the theoretical information model from the Market Microstructure chapter, the goal here is to determine the extent of information misbalance in the market. In the current case, however, we have to do it without relying on dealers' inventory, proportion of informed and uninformed traders, etc, but use only readily observable variables – trade prices and best bids and offers quoted in the market. The key insight of the new model is that inventory effects and price fluctuations from market impact are inherently transient, whereas information effects communicate the real value of a security and thus should persist over long periods of time. Another important observation is that if there is some private information to be inferred from a trade, it should be inferred not from the *total* trade, but only from the part of the trade that is *unanticipated* given previous history of prices, quotes, and trades. To summarize, within this model, the information impact of the trade is defined as a persistent impact on the stock price, which results from an unexpected portion of a trade. The statistical technique employed to detect these impacts in time series of prices and quotes is called vector autoregression or VAR. For more information on VAR and other time series statistical tools, see [Yafee and McGee, 2000].

The primary price variable in this model is the mid-point of the bid-ask spread. Its evolution can be defined as:

$$r_t = (P_t^B + P_t^A)/2 - (P_{t-1}^B + P_{t-1}^A)/2.$$

Transaction costs are fixed and symmetrical, and therefore they are reflected in the size of the spread and don't affect r in any way. This variable is, however, greatly affected by the arrival of both public and private information, which makes the assessment of the impact of any particular trade quite problematic. It is still possible to do this in some average sense, if we assume that quote revision is a stable function of trade. If we assume that such dependence is linear, for example $r_t = bx_t + \varepsilon_t$, where x_t is the trade and ε_t is public information, then the price impact b can be estimated via a regression. Since the impact of a trade needs not be instantaneous, we have to introduce lagged values into the model. Our view of r_t should look more like this:

$$r_t = a_1 r_{t-1} + a_2 r_{t-2} + \ldots + b_0 x_t + b_1 x_{t-1} + \ldots + \varepsilon_{1,t}, \tag{1}$$

where a's and b's are coefficients, x's are trades (actually signed trade indicators in this particular model), and $\varepsilon_{1,t}$ is the residual. This implies that the current quote revision is a function of previous revisions and trades that these revisions have induced. In theory, this sequence can be infinite, but in practice we truncate it. The trades are described in almost the same fashion:

$$x_t = c_1 r_{t-1} + c_2 r_{t-2} + \ldots + d_1 x_{t-1} + d_2 x_{t-2} + \ldots + \varepsilon_{2,t}. \tag{2}$$

It is also a function of past trades and quotes, and the term $\varepsilon_{2,t}$ captures the unanticipated component of a trade, since all the other components are already known. $\varepsilon_{2,t}$ should not be interpreted, however, solely as the private information of the informed traders because it also contains the "noise" from uninformed trading. Together equations (1) and (2) specify a bivariate vector autoregression model. To use this statistical technique, several transformations may have to be applied to the transactions data, like first differencing or taking logs. We additionally define $\alpha_m(\varepsilon_{2,0})$ as the sum of the predicted quote revisions between t = 0 and t = m:

$$\alpha_m(\varepsilon_{2,0}) = \sum_{t=0}^{m} E[r_t \mid \varepsilon_{2,0}].$$

This can be interpreted as the information revealed by the trade and is effectively the sum of coefficients of r's. This is the main construct of this particular model. VAR modeling strategy applied to trades and quotes allows us to distinguish between private information (the trade innovation) and public information (quote revision innovation). It also lets us to assess both instantaneous and lagged impact of a transaction.

Let's say that we have estimated the following model:

$$r_t = -0.118r_{t-1} - 0.011r_{t-2} - 0.013r_{t-3} + 0.014x_t + 0.007x_{t-1} + \varepsilon_{1,t},$$

$$x_t = -3.647r_{t-1} - 1.689r_{t-2} - 0.505r_{t-3} + 0.167x_{t-1} + 0.12x_{t-2} + 0.078x_{t-2} + \varepsilon_{2,t}.$$

What can we conclude from this form? The most significant coefficients are those of x_t, …, x_{t-n} in the r_t equation. For example, they imply that on average the midpoint of the dealer's quote is raised by \$0.014 immediately following a buy order. Similarly, coefficients at longer lags are positive but decreasing. Another finding is a strong positive autocorrelation in trades, reflected by x_t, …, x_{t-n} coefficients in the x_t equation. This means that sells tend to follows sells, and buys follow buys. This can be explained by traders adjusting to new information or by large trades being broken into smaller ones.

b. Spread Estimation through Proxy Variables.

In the above discussion, we replaced a rigid structural model with statistical relationships in a time series. There we have used fairly specialized high-frequency data – prices, quotes, and trade records – which can be noisy, difficult to obtain, and datasets are usually extremely large. We can go one step further in the application of the time series techniques by using proxy variables that can capture similar information, but are easier to observe or collect. In this example, we go back to examining the sources of the bid-ask spread.

Regardless of the microstructure sources of the bid-ask spread, empirical research has clearly established a strong relation between the spread in a security and trading characteristics of that security – volatility, volume, prices, etc. Market microstructure theory focuses on inventory and information effects, and uses variables such as order flow, dealer's holdings, supply and demand elasticities, etc., but in practice all the important effects can be captured by a number of proxy variables that are easily observable. Most popular variables in this category include activity variables such as trading volume, risk indicators such as volatility, company-specific variables like size and industry, and some others, such as trading pressure, price discreteness, etc.

Here's one example of cross section relation that accurately estimates the spread (from [Stoll, 2000]):

$$s/P = a_0 + a_1 \log V + a_2\sigma^2\ a_3 \log MV + a_4 \log P + a_5 \log M + a_6 Avg|I| + e, \quad (3)$$

where s is a quoted half-spread, P – closing price, V – daily volume, σ^2 – daily return variance from the prior year, MV – market value of the company, N – number of trades per day, I – daily imbalance between volume at the ask and the bid, e – error term. This explains over 79% of the cross section variation in the NYSE traded stocks. Results obtained from the coefficients are not surprising: spreads are lower for stocks with higher volume, lower volatility, higher price, and lower trading imbalances.

While these findings may not be directly useful for an automated trader, they show how hard-to-observe microstructure variables can be replaced by some more readily available proxies. A very similar approach can also be applied to transaction prices prediction, as demonstrated in [Huang and Stoll, 1994]. A regression very similar to (3) uses futures prices, trade size, overall volume, trade history, etc. to predict securities prices over a very short period of time. This model does not explicitly address transaction costs, so arbitrage opportunities my not arise; but his short-term predictability can still be used intelligently to optimize trading, as we will demonstrate later in this chapter.

While the power of these relatively simple statistical methods aimed at extracting the microstructure relationships from empirical data is apparent, certain caution must be exercised in their application. While the time series techniques can be regarded as powerful and easy-to-use "black boxes" that can seamlessly produce price predictions and quote-revision rules, one always must remember a very significant assumption that these techniques are based on. They all assume that relationships and patterns discovered in the time series at hand are stable – i.e. they will persist over extended periods of time. If transaction prices had a certain autocorrelation trend in the past, this trend will persist in the future with same parameters (coefficients). This needs not be true for any real-world process, and especially not in the securities markets. This is the reason why it is advisable to compare empirical results with predictions of structured models. Theoretical models try to capture the underlying causes for price and quotes behavior, and they will adjust when certain factors change, whereas statistical analysis just looks for past patterns and projects them into the future.

2. Initial Experimental Work

In this sub-section, we will briefly present the results of some very straight-forward applications of empirical methods to the data available to us (see the Experimental Setup chapter for details). Our aim here is to highlight the inherent limitations of standard empirical approaches, which we subsequently address in our Reinforcement Learning method.

a. ARMA models

The purpose of this analysis is to investigate whether straightforward application of statistical methods to the data available to us will reveal the microstructure behavior predicted by theoretical models. There are essentially two approaches that can be used: first, we can take a time series of a single variable and try to find some underlying structure there (i.e. the series tends to revert to its historical mean with a certain lag); second, we can take several variables and look for dependencies among them (i.e. mid-spread leads transaction prices by a certain number of time steps). If any useful information is discovered through this type of statistical studies, it can then be incorporated into an automated trading strategy. Below we outline the necessary steps – pre-processing, filtering, and model fitting, which can later be extended to more complex examples.

We tested the following hypotheses: (1) can past transaction prices help predict future transaction prices, (2) can spread size (together with past transaction prices) help predict future transaction prices, and (3) can the mid-spread help predict transaction prices. To run these experiments, we used frequently sampled MSFT (Microsoft Corp. stock) transaction data and order book evolution history during different time periods of one day. Essentially, we looked at three time series: transaction price, size of bid-ask spread, and midpoint of the bid-ask spread. We first fit the univariate ARMA (autoregressive moving average) model into each series searching for some underlying structure, and then used the spread size and the mid-spread to see if they can help model the transaction price.

All this analysis has been performed within SAS statistics software package – use [Delwiche and Slaughter, 2001] as a reference. For the detailed description of time series models see [Yafee and McGee, 2000]. Very briefly we introduce two basic stochastic processes. First, we make an assumption that adjacent entries in a time series are related to one another via some sort of a process, which can be described mathematically. There are many ways this can be done, but we are mostly interested in two types of time series models: moving average MA(q) and autoregressive process AR(p).

Under the one-step moving average process MA(1), the current output Y_t is influenced by a random innovation e_t plus the innovation from the previous time step:

$$Y_t = e_t - \theta_1 e_{t-1}.$$

The lag between t and t-1 needs not be one step, but can be any lag q or multiple lags. Another process can be such that the current output is determined by previous value plus some innovation:

$$Y_t = \phi_1 Y_{t-1} + e_t.$$

We call this an autoregressive process, which again can have an arbitrary lag p. (θ and ϕ are parameters for MA and AR processes respectively). These two processes put together form our main tool – process ARMA(p, q), which is simply a sum of the autoregressive and moving average components. Once again, notice that no microstructure variables enter this notation – in relaxed models, we deal only with numbers, regardless of where they came from.

We evaluate goodness-of-fit using three standard criteria: loglikelihood, AIC, and SBC. The first one is essentially a logarithm of the mean square error. AIC (Akaike Information Criterion) and SBC (Schwartz Bayesian Criterion) penalize mean square error with the number of features in the model. AIC = $\exp(2k/T)$MSE and SBC = T^kMSE, where k is the number of features, T is the number of observations, and MSE is the mean square error.

Our experimental findings mostly confirm accepted principles from the Market Microstructure chapter:

(1) markets do appear efficient (at least in a very liquid stock such as MSFT) showing little or no structure beyond white noise;
(2) size of bid-ask spread exhibits a fairly prominent AR(1) behavior in most cases;
(3) spread size does not help in transaction price forecasting;
(4) mid-spread is, in fact, useful for transaction price modeling, but only over extremely short time periods (3-15 seconds if that).

Although none of the above is revolutionary, these experiments highlight the power of multivariate ARMA models in market microstructure analysis. The exact same approach can help investigate more complex relationships: does volume misbalance signal upcoming price movement, does higher volatility lead to larger spreads, etc.

For this particular experiment, we have used the dataset includes MSFT transaction prices over a course of a single day. This main dataset (MSFT.DAY) serves as a base for the smaller time-of-the-day dependent datasets (see Table 1 below). In order to concentrate on the short-term behavior of the transaction price, we selected three one-hour time periods during the day: beginning (10:30 am – 11:30 am), middle (12:30 pm – 1:30 pm), and end (2:30 pm – 3:30 pm). Note that we avoided using the opening and closing hour because, presumably, price behavior during these periods will be significantly different from "normal" rest-of-the-day behavior. The first 3 datasets created (MSFT.MORNING, MSFT.NOON, and MSFT.EVENING) are just the subsets of the master dataset in the indicated time periods. Since the price is sampled every 3 seconds, each of them contains 1,200 observations. In case if such sampling is too frequent, we also created 3 more datasets for the same time periods, but with the price sampled every 15 seconds. These datasets are called MSFT.MSHORT, MSFT.NSHORT, and MSFT.ESHORT and contain 240 observations each.

The second collection of data that we examined was a list of top bids and asks from the order book sampled at the same time as the transaction price. We used this data to create two more time series: the size of the bid ask spread, calculated as (Ask-Bid), and the mid-point of the spread: (Ask+Bid)/2. The later is often used in market microstructure theory as a proxy for the "true price" of a security. Then we went through the same steps as for transaction prices and ended up with 12 smaller time series: for both the size and the mid-spread we had MORNING, NOON, and EVENING periods sampled at 3 and 15 seconds each.

Tables 1 and 2 below summarize basic statistics for transaction prices and bid-ask spread respectively. The mid-spread dataset's statistics are essentially the same as those for transaction prices (Table 1), and thus are not reproduced here.

Name	**Time**	**N**	**Mean**	**Min**	**Max**
DAY	9:30-16:00	7653	25.7281	25.328	25.940
MORNING	10:30-11:30	1200	25.6798	25.551	25.770
MSHORT	10:30-11:30	240	25.6797	25.560	25.770
NOON	12:30-13:30	1200	25.8167	25.740	25.853
NSHORT	12:30-13:30	240	25.8169	25.740	25.853
EVENING	14:30-15:30	1200	25.8488	25.761	25.940
ESHORT	14:30-15:30	240	25.8489	25.761	25.936

Table 1.

We can clearly observe from Table 2 the "U-shaped pattern" of the bid-ask spread, which is mentioned on many occasions: average spread is the largest in the middle of the day and is tighter in the morning and afternoon. The same holds for the maximum spread as well.

Name	**Rate (sec)**	**N**	**Mean**	**STDev**	**Min**	**Max**
DAY	3	7653	0.013639	0.0072	0.01	0.054
MORNING	3	1200	0.013039	0.00605	0.01	0.033
MSHORT	15	240	0.012846	0.0059	0.01	0.033
NOON	3	1200	0.01401	0.007384	0.01	0.054
EVENING	3	1200	0.013276	0.006776	0.01	0.043

Table 2.

In order for the ARMA model to be applicable, the time series have to be stationary – in simpler terms, we had to remove the trend and render the volatility homoskedastic (roughly constant). All the series that involve prices (transaction or mid-spread) have unit root in them and must be first differenced. Dickey-Fuller tests in SAS ARIMA procedure prove that this is sufficient. It is much less clear, however, if taking logs of prices is in order to stabilize the series volatility. Results for several significance tests – log likelihood, AIC, and SBC – are presented in Table 3 for both regular prices and their logs.

Series	**Log Likelihood**	**AIC**	**SBC**
DAY	30621.56	-61231.12	-61189.46
Log	30608.10	-61204.19	-61162.54
MORNING	4660.74	-9309.48	-9278.94
Log	4660.77	-9309.54	-9279.01
MSHORT	752.908	-1493.82	-1472.96
Log	752.843	-1493.69	-1472.83
NOON	5288.68	-10565.36	-10534.82
Log	5288.57	-10565.14	-10534.60
NSHORT	870.565	-1729.13	-1708.27
Log	870.532	-1729.06	-1708.20
EVENING	4841.19	-9670.39	-9639.85
Log	4841.06	-9670.11	-9639.58
ESHORT	794.220	-1576.44	-1555.58
Log	794.137	-1576.27	-1555.42

Table 3.

It appears that taking logs is not necessary for prices (transactions and mid-spread), but the difference is marginal. The reason for this: over short time periods, it is safe to assume that the drift is zero.

Another issue is whether it is appropriate to work with actual prices, or should returns be used instead. The later approach is customary in financial literature, but may not matter for the kind of data we are using. To test which method is more appropriate, we initially fit all the ARMA models (see below) to transactions data using actual prices, and then replaced prices with log-returns, but left all the models parameters unchanged. Both approaches yielded the same results, so we chose to work with actual prices for other experiments. We also determined that we needed to take logs (but not first difference) of the spread size for all time series. Spread size models had a significant intersection term, while prices did not. This can be attributed to first differencing of prices.

We found it very challenging to fit an ARMA model to a time series of transaction prices, since they look very much as white noise. Surprisingly, however, when we extended the number of time periods to be examined by our model from 20 to 250 for 3 second series and to 50 for 15 seconds series, we found significant autoregressive terms that lag from 6 to 13 minutes:

Model	**P**	**Q**
MORNING	1,3,4	1,3,4
MSHORT	23,37	0
NOON	23,257,258	0
NSHORT	1,42,43	0
EVENING	6,117,118,230	0
ESHORT	1,46	0

Table 4.

While we are very much inclined to discard these results as nonsensical from the market microstructure point of view (a price at the next period depends on a price 10 minutes ago, but on nothing in between – seems unlikely), but these results have strong statistical support. Every single one of the parameters above is statistically significant (t-value is greater than 2), and both SBC and AIC are lower for the above models than for the base p=0, q=0 model; and, finally, the residuals are generally improving in most cases compared to the white noise model. Overall, if there is any underlying structure for transaction prices, it is almost certainly an autoregressive (as opposed to a moving average) relationship.

Unlike transaction prices, spread size showed much more structure in correlograms: most of them look very similar to AR(1) model. AR(1) turns out first- or second-best model in AIC/SBC scoring, but some low-level (1 or 2) MA process seems to be present as well. Here are the parameters that we estimated:

Model	**P**	**Q**
MORNING	1	0
MSHORT	0	1
NOON	1	2
NSHORT	1	0

EVENING	1	1
ESHORT	0	1

Table 5.

This autoregressive behavior has a coherent explanation from the market microstructure point of view: as the spread narrows, it becomes cheaper for traders to “step over” the spread and transact immediately with outstanding limit orders; by definition, this will remove orders from the book and thus widen the spread. As the spread gets wider, submitting market orders becomes more expensive, and traders resorts to posting limit orders inside the wide spread, which, in turn, shrinks the spread.

We next attempted to use the spread size and mid-spread as exogenous variables that help predict the transaction price. Whereas we did manage to find lags that make the spread size significant for the transaction price estimation, the new models’ SBC and AIC were always higher than the ones from the univariate model. Therefore, we reject our hypothesis that the spread size can be helpful for transaction price forecasting.

The mid-spread turns out to be a much more helpful variable especially when sampled every 3 seconds, which certainly is not surprising. We had to fit an ARMA model to the mid-spread series as well, again resulting in mostly AR models. After adding the mid-spread to the transaction price forecasting, we can conclude that in general, knowing the mid-spread at time t is useful for forecasting the transaction price at time t+1. SBCs and AICs are lower than without the exogenous variable, and lags are significant, but the residuals still leave a lot to desire in both cases. Does this finding have any practical significance? Not very likely, since one variable is leading the other one by an extremely short time period (plus lots of structure remains unexplained).

In all our experiments we obtained a vast amount of information describing significance of various coefficients, goodness of fit, behavior that still remains unexplained, some predictions, etc.; most of this data can also be plotted. But we are not reproducing all this numbers here because of the shear volume, and also because our primary goal is to demonstrate the process of finding out if there is some relationship between various microstructure variables and not to explicitly forecast stock prices or other variables.

We have described the basic idea behind the relaxed models and shown a simple application of time series techniques to market microstructure modeling, using the data available to us. The main contribution of these experiments is the proof of applicability of multivariate ARMA models to the market microstructure research where we are dealing with discretely sampled data. We also found some structure in the spread size, which can mean that this variable is actually forecastable – a fact that can be used in creating automated trading agents. And finally, while our efforts have confirmed that prices are hard (read impossible) to forecast, the same needs not be true for other microstructure variables.

b. Naïve Market Making

Our next experiment is to apply the insights from theoretical models and our time series experience to construct a simple trading strategy. The goal of this exercise is to demonstrate the complexity of this domain ant to highlight the fact that microstructure relationships unearthed by relaxed models are not easily exploited in practice. One microstructure effect predicted by theoretical models and confirmed by empirical methods is that over very short horizon, transaction prices tend to bounce between the bid and the ask. Here we attempt to construct a naïve market making strategy that aims to profit from this bid-ask bounce.

We propose a simple dealership model by decomposing the problem that the electronic market maker is facing into two components: establishing the bid-ask spread and updating it. We further suggest a coarse subdivision of the update methods. The first step to creating an electronic market maker is the understanding of the responsibilities of a securities dealer. As previously discussed, the primary objective of a market maker is to continuously update the bid-ask spread. Doing this correctly is the key to making profits: the spread has to be positioned in such a way that trades occur at the bid as often as at the ask, thus allowing the dealer to "buy low and sell high". The bid and the ask quotes are supposed to straddle the "true price" of a security [Ho and Stoll, 1981], and the difference between the two is the dealer's revenue. However, the "true price" is difficult to determine or model, plus it's not even clear if such quantity exists in the first place.

Therefore, the first potential problem for a market maker (either human or artificial) is to decide where to establish the initial spread.

There are, essentially, two ways to approach this dilemma. The first, hard way is to perform the actual valuation of a security that is being traded: if it's a stock, try to determine the value of the company using Corporate Finance methods (cash flows, ratios, etc.); if it's a bond, determine the present value of the promised payments, and so on and so forth. An entire new set of issues not discussed here arises if the market maker's valuation differs from the consensus among the rest of market participants. If there is no established market, or the market is very illiquid, then doing the valuation may be the only approach. Fortunately, the majority of the modern securities markets employ limit orders in some capacity. The buy and sell limit order books serve as a fairly accurate representation of the current supply (sell queue) and demand (buy queue) for the traded security. Presented with such supply-demand schedule, the market maker can determine the consensual value of a security. In the simplest case, the market maker can observe the top of each book – the best (highest) buy and the best (lowest) sell – also known as the "inside market". He can safely assume that the market's consensus of the true value of the security lies somewhere between these two numbers. If the best bid is \$25.21, and the best ask is \$25.30 then the consensual price of the stock is somewhere in this interval. Now, the market maker can use the top of each book as a reference point for positioning his initial quotes – i.e. establish his own spread at \$25.23-28 – and then start updating the bid-ask spread as the books evolve with new order arrivals, transactions and cancellations.

Updating the spread with a goal of maintaining profitability is the essence of the market making. For the purpose of our experiment here, we can broadly classify market-making strategies into "predictive" and "non-predictive". The former try to foresee the upcoming market movements (from order book misbalances, or from many other possible patterns) over a very short horizon, and adjust the spread according to these expectations, while the latter do not attempt to look forward, but are based solely on the information about the current inside market (top of each book). The non-predictive strategies are inherently simpler, but are worth considering nonetheless. order to make a case for the non-predictive strategies being even worth considering.

Non-predictive approach to market making is based on an assumption that an electronic market maker can revise his quotes quickly enough to take advantage of the bid-ask bounce without having to predict the next price move. To give an example, if the price of a stock is going up consistently for an hour, it doesn't mean that everyone is buying (or that all arriving orders are buy orders); selling is going on as well, and the price (along with the inside market in the order books) moves down as well as up. Figure 1 illustrates this scenario: while there is a general upward trend (the dotted line), we can see the simultaneous evolution of the order book, and transactions happening at the top of buy queue (market sale, dealer's purchase), then the sell queue (market purchase, dealer's sale), then buy, then sell again. While price is generally going up, the stock is being sold as well as bought.

How does the market maker fit into this scenario? By maintaining his quotes (the bid and the ask) on both sides of the market, at or close to the top of each order book, the market maker expects to get "hit" or transact at his bid roughly as often as at his ask because of these fluctuations. This way, after buying at the bid (low) and selling at the ask (high), the dealer receives the profit equal to the bid-ask spread for the two trades, or half-the-spread per trade, which is the fundamental source of the market maker's revenue. In the context of Figure 1, suppose that the top order in each queue is the dealer's; in this case, the dealer buys at $25.10, then sells at $25.18 (8 cents per share profit), then buys at $25.16 and sells for $25.26 (10 cents per share profit). If each transaction involves 1,000 shares, and all this happens over several seconds, it becomes clear that market making can be profitable.

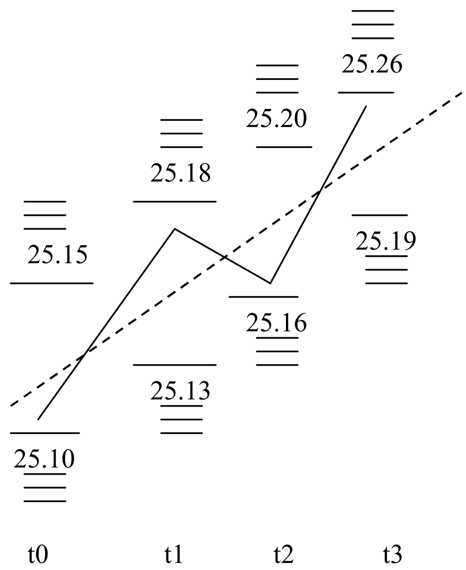

Figure 1.

Now, having understood the nature of the dealer's income, we can re-formulate our task: adjust the bid-ask spread in such a way that the orders generated by other market participants will transact with the dealer's bid quote and the dealer's ask quote with the same frequency. In our example, we are looking for an algorithm to maintain the dealer's quotes on top of each queue to capture all incoming transactions. To facilitate thinking about the spread update, we can say that at any given point in time, the dealer has two possible actions: move the spread up or down relatively to its midpoint, and resize the spread – make it wider or narrower – again relatively to its midpoint. He may also want to change the "depth" of his quote – a number of shares he is committed to buy or sell. Let's put resizing and depth aside for the time being and assume that the size of the spread (and the inside market) is constant, and that the only thing the market maker does is moving the spread up and down as the price changes (state of the order book evolves). In our example, the stock price is steadily going up overall, while actually fluctuating around this general climb. If the market maker wants to capture the "buy low, sell high" opportunity, then his spread should also continuously move up straddling the stock price.

But how can the dealer tell at any given time looking forward that it's time to move the spread up and by how much? The non-predictive family of electronic trading strategies would argue that he cannot and needs not do so. Non-predictive strategies postulate that while there are some patterns (streaks where the stock is either rising or falling) globally, the local evolution of the stock price is a random walk. If this random walk hits the bid roughly as often as it hits the ask, then the market maker makes a profit. If one subscribes to the theory that the short-term evolution of supply and demand for a security is random, then it is understood that an uptick in the stock price is as likely to be followed by a downtick as by another uptick. This implies the futility of trying to incorporate expectations of the future supply/demand shifts into the model governing the bid-ask spread updates. If the above assumption holds, and if the market maker is actually able to operate quickly enough, then the trading strategy can be very simple. All the market maker needs to do is to maintain his bid and ask quotes symmetrically distant from the top of each book. As the orders arrive, transact, or get cancelled, the market maker has to revise his quotes as fast as possible, reacting to changes in such a way that his profitability is maintained.

In principle, the dealer should be market neutral – i.e. he doesn't care what direction the market is headed – he is only interested in booking the spread. On the other hand, the dealer is interested in knowing how the inside market will change over the next iteration in order to update his quotes correctly. The way the non-predictive strategies address this is by assuming that the inside market after one time step will remain roughly at the same level as the current inside market (that's the best guess we can make, in other words). Therefore, being "one step behind" the market is good enough if one is able to react quickly to the changes. Such is the theory behind this class of strategies, but in practice this turns out to be more complicated.

Here is a general outline of an algorithm that implements a generic non-predictive strategy; at each iteration:

(1) Retrieve the updated order book;
(2) Locate an inside market;
(3) Submit new quotes (buy and sell limit orders), positioned relatively to the inside market;
(4) Cancel previous quote.

Having defined this very simplistic model of market making, we now explore the ad-hoc parameterization of the actual implementation. There are three main factors, or parameters, that determine a non-predictive strategy: position of the quote relative to the inside market, depth of the quote (number of shares in the limit order that represents the quote), and the time between quote updates.

Timing is, perhaps, the simplest out of the three parameters to address. In the spirit of the theoretical non-predictive model presented in the previous section, the market maker wants to respond to changes in the market as soon as possible, and therefore, the time between the updates should be as close to zero as the system allows. Despite this property, it is still useful to think of the update timing as a parameter that should be minimized: i.e. the computational cycle should be performed as fast as possible, and communication delays between the dealer and the market (how long does it take for an updated quote to show up in the limit order book) should also be minimized. In our experiment, the computational cycle is extremely short – under 1 second – because of the inherent simplicity of the algorithm, but we set the communication delay to be non-trivial to demonstrate the real world complexities. It takes about 3 to 5 seconds for the order to get inserted into the book, and about the same amount of time for the order to get cancelled (if not transacted) after it appears in the book. This is one of the market frictions, which should not be overlooked. The dealer wants to access the market as quickly as possible, but such delays can prevent him from operating on a scale small enough to capture the small fluctuations. Therefore, real systems where these delays can be decreased can potentially be more effective and produce better results than our simulated setup.

Positioning the quote relative to the rest of the order book is the most important aspect. We use a simple distance metric – number of cents by which the dealer's quote differs from the top [non-dealer] order in the appropriate book. We decided to start our strategy implementation from a fairly well-know, albeit somewhat controversial practice of "penny jumping".

…
25.56 – 300
25.55 – 1000
25.35 – 200
25.30 – 150 ← Top sell
25.29 – 500 ← Dealer's ask

Bid-Ask Spread

25.22 – 500 ← Dealer's bid
25.21 – 200 ← Top buy
25.19 – 300
25.15 – 785
25.10 – 170
…

Figure 2.

In general, penny jumping occurs when a dealer, after entering his customer's order into the order book, submits his own order, which improves the customer's limit price by a very small amount. The dealer effectively "steps in front" of his customer: the customer's potential counterparty will now transact with the dealer instead; thus the dealer, arguably, profits from the customer's information, and, in some sense, trades ahead of the customer, although at a price improvement over the customer's limit order. Such practice is not exactly illegal (because the client's potential counterparty does get a better price by transacting with the dealer instead of the customer), but is considered unethical, and became the center of the recent NYSE investigation/review [Ip and Craig, 2003].

In our case, we are simply undercutting the current inside market (or the "de facto" bid-ask spread) by one cent on both sides. The dealer's bid improves the current best bid by a penny, and the dealer's ask does the same on the sell side – Figure 2 shows that if the inside market is 25.21 – 30, our electronic market maker's orders will make it 20.22 – 29 (the size of the bid-ask spread goes from 9 to 7 cents). This way, the market maker is guaranteed to participate in any incoming transaction up to the size specified in the depth of his quote. Since ECN market makers have no designated customers and we are working in a simulated environment, our strategy is much less controversial then its real world counterpart.

We expect the following behavior from this strategy: the revenue (P&L) should rise slowly over time (because profit per share is tiny) as in Figure 3a, but our result look more like Figure 3b: while the inventory does fluctuate around zero, the strategy gradually loses money over the course of a trading day.

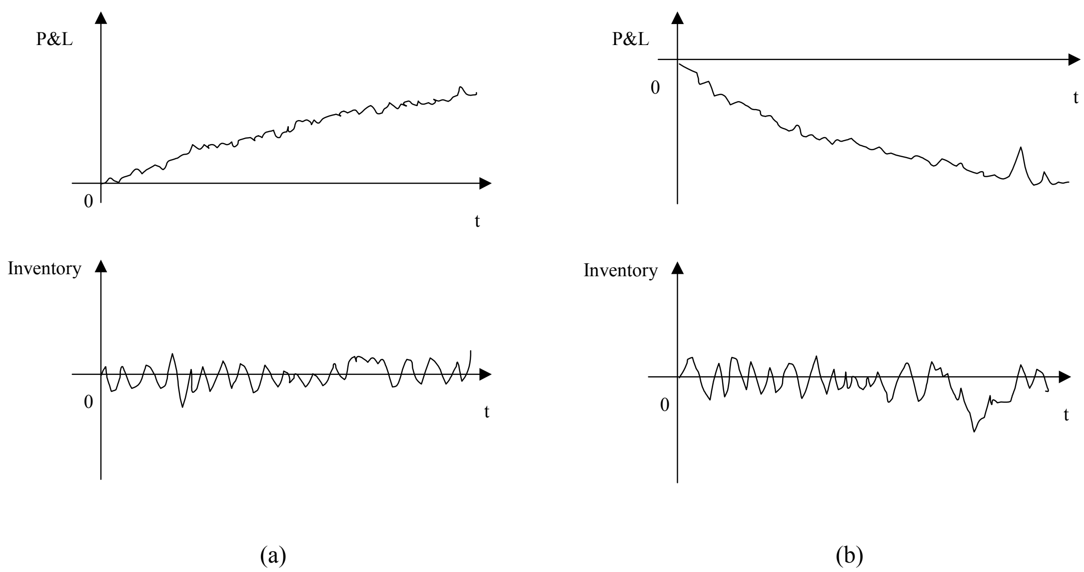

Figure 3.

The most fundamental problem is presented in Figure 4: while we base our decision on the state of the book at time t_0, the outcome of our decision gets placed in a book at time t_1, which may or may not be different from the original t_0 book. We have already touched upon this problem in the update time discussion above. Essentially, the non-predictive market-making strategy places an implicit bet that the book at t_1 will be fairly close to the book at t_0, or at least close enough to preserve the profitable property of dealing. What actually happens in our experience with penny jumping, is that the inside market is tight already, plus the book changes somewhat over in 3 seconds, and so, oftentimes, both the bid and the ask quotes (limit orders) issued at t_0 end up on the same side of the market at t_1 (Figure 4). Then one of the orders transacts, and the other ends up buried deep in the order book.

…	…
25.56 – 300	25.56 – 300
25.55 – 1000	25.55 – 1000
25.35 – 200	25.35 – 200
25.30 – 150	25.30 – 150
25.29 – 500	**25.29 – 500**
	25.25 – 100
25.27 – 500	25.24 – 300
25.26 – 200	
25.22– 1000	25.22 – 750
25.19 – 300	25.19 – 300
25.15 – 785	25.15 – 785
25.10 – 170	25.10 – 170
…	…
t0	t1

Figure 4.

If we find ourselves in this situation on a more or less regular basis throughout the day, we end up paying the spread instead of profiting from it. This explains why our actual P&L pattern mirrors the expected pattern. We discern three main reasons for the penny jumping fiasco: (1) making decisions in one book, acting in another; this is further aggravated by (2) the "frictions" of the simulator; and, finally, (3) spreads are extremely tight leaving little or no profit margin.

Tight spreads (inside markets) deserve further notice. The way we have defined our penny jumping strategy implies that the spread has to be at least 3 cents, but since the stock used this experiment is a very liquid Microsoft (MSFT), oftentimes during the day the spread becomes smaller than 3 cents. This forces our strategy to "sit out" for extended periods of time, which doesn't improve its profitability. The size of the spread is closely related to the "decimalization" of the US equity markets, which was implemented in 2001 and is still under scrutiny. Now stocks trade in increments of 1 cent, as opposed to 1/16 of a dollar, on both NYSE and NASDAQ. From the perspective of a non-predictive market-making strategy, this can also have positive effects: when stepping in front of someone else's order, one wants to be as close as possible to the original price. Decimalization actually helps here, since undercutting by 1/16 of a dollar is much riskier than undercutting by 1/100 of a dollar. But it also makes the spread a lot tighter, cutting into the dealers' profits [Barclay et al., 1999].

…	…
25.56 – 300	25.56 – 300
25.55 – 1000	25.55 – 1000
25.35 – 200	25.35 – 200
25.33 – 500	**25.33 – 500**
25.30 – 150	25.30 – 150
	25.25 – 100
	25.24 – 300
25.26 – 200	
25.23 – 500	**25.23 – 250**
25.22– 1000	25.22 – 1000
25.19 – 300	25.19 – 300
25.15 – 785	25.15 – 785
25.10 – 170	25.10 – 170
…	…
t0	t1

Figure 5.

Does this mean that the non-predictive strategies are inherently money-losing? Not at all – one small change can bring the profitability back. We don't really have to undercut the inside market, instead we can put our quotes at the inside market or even deeper in their respective books. This makes the dealer's spread larger (more profits per trade), but can reduce fairly drastically the overall volume flowing through the dealer. Essentially, one has to find a balance between the potential profitability and volume. In practice, by putting the quotes 1 – 3 cents away from the inside market works well, or, at least, alleviates the concerns that make penny jumping unprofitable. The dealer's spread is much wider now, so even when the quotes get put into a different book with a significant delay, more often than not they still manage to straddle the inside market and therefore preserve the "buy low, sell high" property. Figure 5 shows the exact same scenario as Figure 4, but with wider dealer quotes. While this is certainly good news, there are still several issues that expose the vulnerability of the non-predictive strategies to certain market conditions. For example, inventory management becomes an important consideration.

In theory, the market maker should buy roughly as frequently as he sells, which implies that his stock inventory should fluctuate around zero. The dealer makes money on going back and forth from long to short position. Since he gets a fraction of a penny for each stock traded, the dealer naturally wants to "move" as many stocks as possible to compensate in volume for thin margins. Therefore, the dealer would prefer to set the depth of his quote – the third fundamental parameter – as high as possible. Potentially, all trading in a stock could flow through the market maker. In practice, however, this doesn't always work out. If a stock price is going up consistently for some period of time, what ends up happening is that the dealer's ask gets hit more often than his bid. The dealer winds up with a [potentially large] short position in a rising stock – he is taking a loss. Again, the same issues that were discussed earlier in this section are in play. Plus, at times the main assumption behind the non-predictive strategies just doesn't hold: for example, when a stock "crashes" there are actually no buyers in the marketplace, and the entire market-making model is simply not valid any more. Exchanges halt trading in the stock when this happens, but the dealer will have probably taken a considerable loss by then. Also, if a dealer accumulates a large position in a stock, he becomes vulnerable to abrupt changes in supply and demand (price fluctuations) – i.e. if a market maker has a significant long position, and the stock price suddenly falls, then he's taking a loss. And finally, there are some real-world operational issues, like certain pre-determined limits on exposure. Securities firms, for example, can prohibit their traders from holding an inventory of more than 100,000 shares long or short. The bottom line is: there is a trade-off for the market maker. On one hand, he wants to post deep quotes and have a large inventory to move back and forth from one side of the market to the other, but then he doesn't want to become exposed by having a large position that cannot be easily liquidated or reversed.

To reconcile these conflicting goals, we have implemented some ad-hoc rules to manage the dealer's inventory. We have tested a number of such approaches. The most straightforward one is to impose some global limit – i.e. no position in excess of 20,000 shares long or short. When the market maker reaches the said limit, he stops posting a quote on the side of the market that will make him go over the limit. The problem with this approach is that when the limit is reached, the market-making revenue model no longer holds. The dealer is only active on one side of the market, and is exposed to market movements by holding a large inventory, which makes this approach not very practical. One can also manage inventory by simply varying the depth of the dealer's quote: if the depth is 300 shares, the market maker is less likely to accumulate excess inventory than if the depth were 5,000 shares. This can certainly be effective – by setting the depth low enough, the dealer doesn't have to worry about the inventory side effects. However, as shown earlier, shallow quote translates into less volume and less revenue. Therefore, by getting rid of the inventory risk, the market maker gives up the necessary revenue to continue its operation. The compromise can be reached by establishing some "schedule" of quote depth as a function of inventory. Example:

Inventory (Absolute)	**Depth of Quote**
0 to 20,000 shares	5,000 shares
20,000 to 50,000 shares	1,000 shares
More than 50,000 shares	500 shares

Table 6.

We have also tried an actual functional dependence between the dealer's inventory and the depth of his quote. For example: Depth = 5,000 – max(0, (Inventory – 20,000)/Inventory*1,000) means that the quote starts at 5,000 shares and decreases gradually when the inventory exceeds 20,000 shares. One may also be tempted to decrease the quote on the side of the market where the excess inventory is being accumulated, while leaving the other side unchanged, but this will go against the definition of market making. This will induce the reduction of the position, but the lack of symmetry will cut into profits on future trades. Theoretically, this general approach of balancing the inventory through the depth of quote should work; in practice, however, it is very difficult to calibrate. The schedule and the formula given above are entirely ad-hoc; while they generally "make sense", how can we tell that various levels, decreases, coefficients, etc. are the optimal numbers for this case? The usual statistical/ML optimization techniques are not very effective here, since there is no straightforward relationship between the profitability (the outcome) and the depth of the quote because many other factors, such as the size of the spread, are in play. Therefore, while this approach is sound, it's difficult to implement effectively.

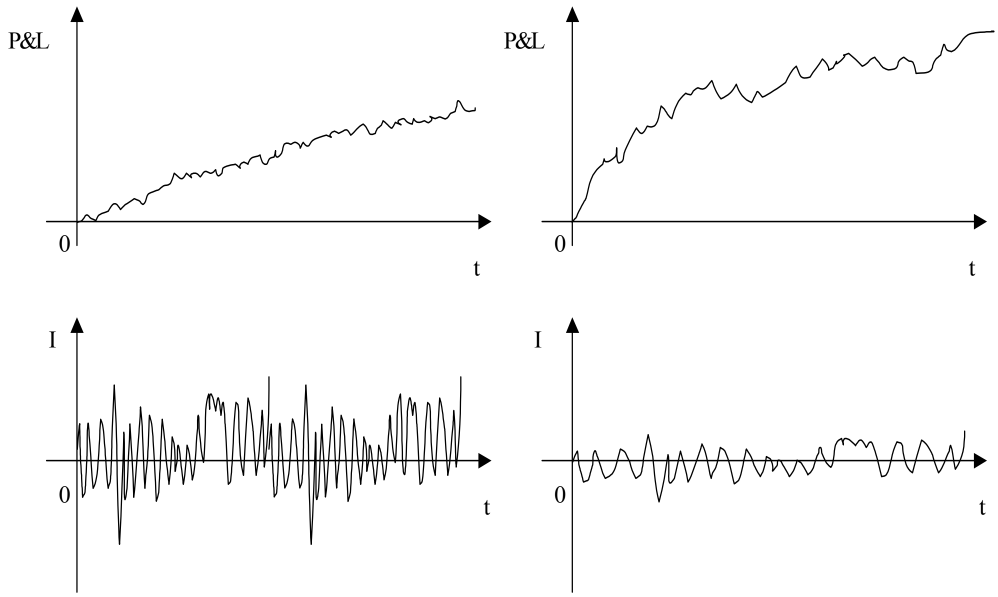

Figure 6.

The one method that we believe to be practical is mitigating the inventory effects through resizing the spread. If there is too much buying (the dealer's ask is being hit too often, and he accumulates a short position), then moving the ask deeper into the sell book should compensate for this. Also, if the stock is going in one direction continuously, this approach will force the spread to also be continuously revised upwards, using the inventory misbalance as a signal. Just like it was the case with the depth of the quote, there is a number of ways to implement this method. We can statically set the spread wider, which will increase the profit margin, decrease the risk of inventory accumulation, but will also decrease the overall volume, and thus generate a P&L pattern not much different from the earlier one (see Figure 6).

We found it more effective to establish a functional dependence between how deep inside the book should the quote be and the stock inventory. Similar ideas were proposed before – [Ho and Stoll, 1981] suggest an analytical solution to a similar spread update problem, which involved solving a differential equation. We use a formula similar to the one in the depth of quote discussion: Distance from the inside market = MinimumDistance + alpha* max(0, Inventory–InitialLimit) /Inventory * MinimumDistance. The two main parameters to determine here are alpha and

InitialLimit. (Minimum Distance is fixed separately, guided by the volume vs. profit margins trade-off). When the position is within the InitialLimit, the quote is always set MinimumDistance away from the market, but if the inventory gets outside the limit, we start putting pressure on it to move in the opposite direction (see Figure 7).

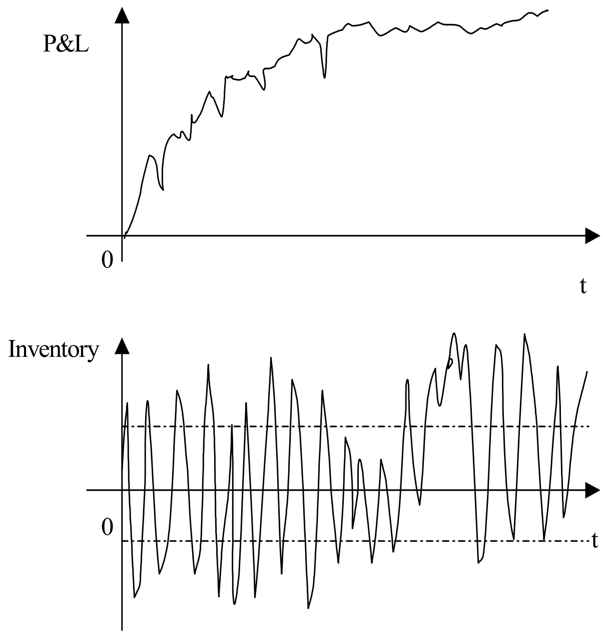

Figure 7.

The "rubber band" analogy is appropriate here: when the inventory gets too large, it's being pushed back by a rubber band – the band can be expanded more, but it becomes harder to do the further you go. Parameter alpha regulates the "stiffness" of the band – the higher you set it, the less expandable the band becomes. Experimentally, we've determined that it is beneficial to make both the InitialLimit and alpha relatively low. Figure 8 shows a fairly typical performance of this strategy: inventory swing often from a large positive to a large negative position, generating solid profits.

Day	**P&L**
April 28	-3,781
April 29	-4,334
April 30	4,841
May 1	-15,141
May 2	-3,036
May 5	6,405
May 6	33,387
May 7	24,021

May 8	1,380
May 9	-7,252
Total	36,490

Table 7.

Implementing and testing the non-predictive market-making strategies, we arrived at a number of conclusions: faster updates allow to follow the market more closely and increase profitability; to combat narrow spread and time delays, we can put the quote deeper into the book, although at the expense of the trading volume; trading volume can be increased with deeper quotes; inventory can be managed effectively by resizing the spread. We have also found out, however that the non-predictive strategies do not solve the market-making problem completely. The performance of a market-making strategy with complete functionality over 10 trading days from April 28th to May 9th 2003 is summarized in Table 7.

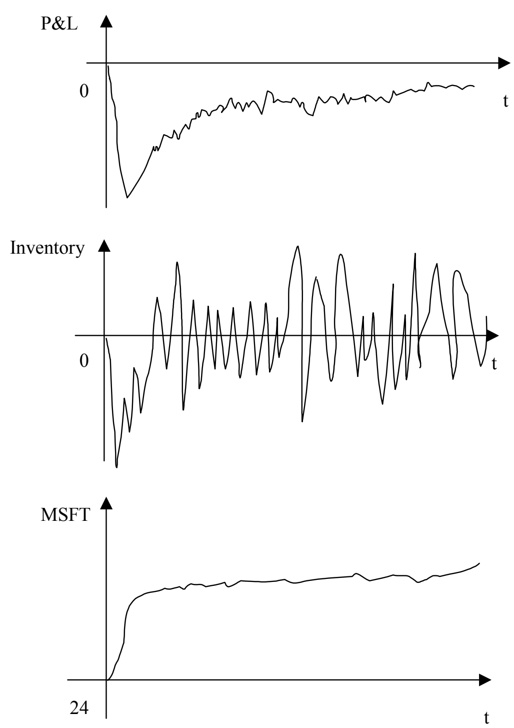

Figure 8.

As you can see, the outcome is not exactly stellar. So, what is the problem, if all the issues from Section 6 are taken into account? What happened early in the morning on

May 9th (see Figure 8) exemplifies the general shortcoming of non-predictive strategies: the price keeps going up, the market maker cannot get his quotes out of the way fast enough, accumulates a large short position, and loses a lot of money. All this happens in 10 minutes. The same scenarios can be observed on other money-losing days. This has been mentioned before, but we are back to the two fundamental problems. First, even an electronic market maker fails to operate on a small enough time scale to take advantage of the short-term fluctuations in supply and demand; second, there are times when such fluctuations just don't exist, and the entire premise behind the non-predictive market making just no longer holds. These are the realities that have to be accounted for. We will have to use some predictive instruments – order book misbalances, past patterns, or both – in order to solve these problems.

3. Normative Proposal

Having surveyed existing empirical methods and having conducted sample empirical studies of our own, we conclude that time series analysis and other statistical techniques can be powerful tools in determining (a) patterns in microstructure variables evolution, and (b) statistical relationships among variables. Unfortunately, such approach to market microstructure analysis also introduces significant problems:

(1) since the time series analysis is disconnected from fundamental microstructure realities, it can be challenging to decide which lags to look at, what are the meaningful combinations of explanatory variables, etc.

(2) high-frequency microstructure-level data is necessarily "messy" and has to be pre-processed because of the requirements imposed by statistical models (first differencing, unit root removal, outliers, and so on). But this pre-processing and smoothing may eliminate the very "frictions" we aim to study.

(3) microstructure is a complex environment, which is also very noisy. Therefore, researchers are bound to find spurious correlations in the data – just like the 10-min lagged positive autocorrelation we have observed

(4) perhaps most importantly, statistical models are inherently *explanatory* – while they may tell us with high level of confidence that high price volatility leads to larger spread, they do not make recommendations about how such "market imperfection" (since we consider any deviation from the random walk behavior an imperfection) can be exploited or eliminated

(5) we also saw from our naïve market making effort how difficult it can be to translate the insights gained from market microstructure research (either theoretical or empirical) into a profitable trading strategy

(6) also from our experience with market making, we found out that an automated strategy must be forward looking an anticipate impending changes: even with microstructure-induced frictions, markets are efficient enough to ensure that no free money is left just laying around

In light of the above issues, we propose a different approach to market microstructure research, which can address the current inadequacies and be generally geared towards practical applications. We call it a *normative* method because of its prescriptive (as opposed to descriptive) nature.

Normative approach includes the following features:

(1) predictive: forward looking model based on historical data – optimal actions should adjust to market conditions according to patterns learned from past records

(2) explanatory: learns patterns in the evolution of a single microstructure variable as well as dependencies among multiple variables; unlike time series analysis, our approach does not require an apriori specification of which variable should influence which, but the process of variable selection is still exogenous to the model (i.e. we have to determine ourselves which factors should be considered). Sample microstructure variables:

- transaction prices
- quotes
- volumes
- volatility
- order book structure

- many others

(3) prescriptive: based on learned relationships, it determines optimal actions that a trading agent should take in order to profit from uncovered predictability in market variables. This requires defining some *scoring system*, which encodes the goal of a trading strategy and allows us to distinguish between "good" and "bad" policies and actions.

(4) conditional: these learned actions can be conditioned on the state of the world – i.e. they are functions of both market movements and the agent's internal variables and constraints. For example, actions can be conditioned on any of the following

- time constraints
- holdings
- measure of risk
- "distance" to the strategy's goal
- prices
- volumes
- and so on

(5) flexible: not tied to any particular problem, but rather serving as a general framework that can detect microstructure effects, optimize agent's actions to take advantage of possible patterns, all while working towards some goal. Possible applications:

- execution
- market making
- technical trading
- short-term investing
- etc

(6) multi-agent: allows interaction of multiple strategies on the same data

We now present the full implementation of the above method in the domain of efficient trade execution.

B. Trade Optimization Through Reinforcement Learning

1. Assessing Execution and Measuring Trading Costs

Here we explain the practical importance of efficient execution, show how to quantify execution quality, point out inherent trade offs in this domain, and provide motivation for using RL to derive optimal trading strategies. The purpose of this subsection is to impress upon the reader how important and omnipresent the problem of trade optimization is, and why it necessitates a principled and normative approach. Our treatment here is similar to [Glantz and Kissell, 2003].

The importance of optimal – or at least efficient – trade execution cannot be underestimated, since it affects every level of investment activity. Investment strategies are developed in "laboratory" back-testing environments, and oftentimes use closing or other bullet prices to compute their performance. While such strategies can show stable "paper" profits, markets are competitive, and thus expected profit margins are already thin. So when we introduce real-world trading frictions into the equation, we risk seeing these profits diminish further or disappear altogether. The two main limitations are that "paper" strategies would like to transact at a single price each day (open or close, for example) and to operate at large volumes (to leverage thin profit margins). These two preferences are unfortunately unattainable in the real world – we have to take into account such factors as the bid-ask spread, volatility, price impact, and many others. In general, it is the intraday price uncertainty and payment for liquidity that comprise strategy implementation costs.

How important is it to trade efficiently? Imagine an investment manager who expects one particular stock currently trading at $20 to go up 10 percent over the next month to $22. To take advantage of this appreciation, the manager has to incur trading costs both now, acquiring a long position, and in the future, when converting his holdings back into cash. Suppose he is able to buy at $20.50; then the price goes up to $22 as predicted, but the manager can cash in only at $21.50. When these trading costs are taken into account, the investment return turns out to be about 4.9 percent – less than half of what was expected. While this example maybe somewhat exaggerated, it highlights an important reality: investment returns are significantly and adversely affected by trading costs. It is our high-level goal to get these costs as close to zero as possible.

Current most widely used tools for trade execution are market orders, VWAP engines (volume-weighted average price – a quantitative tool that promises transaction price that is close the average transaction price in the entire market), or human brokers. Market orders can result in significant price impacts, brokers charge commissions, and VWAP is not well-suited for many situations – i.e. in rapidly rising or falling markets, or when traders have a short- to medium-term view (being either contrarians or a trend followers). A more general execution system, which is better attuned to the high-level strategy requirements, can provide the most cost-effective execution for all trading strategies.

Figure 9 shows a schematic representation of an execution system as a middle agent between higher-level strategies and actual markets. Here we can see that the first use of an execution system is to serve as a cost-estimator in a "feed-back loop" of a strategy development process. Let's say either a portfolio manager wants to buy certain quantity of stock over a given time period. Before going ahead with this particular transaction, we can use the execution system to give us a pre-trade estimation of transaction costs. These estimations can then be passed back up and included into the total return optimization process, and this procedure then gets re-iterated until a suitable compromise is found. Example: a high-level strategy wishes to buy a large number of shares in an illiquid stock; the execution system informs it that to acquire this position in a specified interval will require 1 per cent price concession. Having this information, high-level strategy can decrease the order size and/or increase the time window, which will mitigate transaction costs, while still pursuing the profit opportunity in the stock.

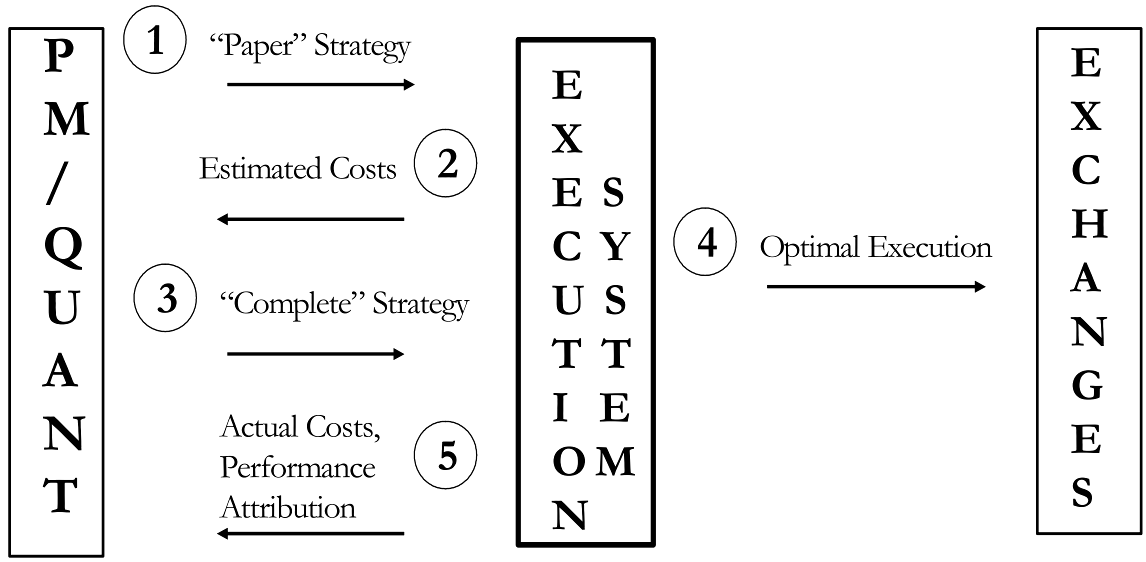

Figure 9.

The second "leg" of the execution solution provided by such system is the actual automated trading. In this phase, having come up with a priori (pre-trade) ideal size and time window, the execution system makes order-submission decisions: whether to be passive or aggressive, what types of orders to use, take or provide liquidity, and many others. In this stage, we have to combine our historical beliefs with current market conditions to trade optimally. Unlike a human trader in this situation, a computer can handle many securities at the same time, pay attention to a multitude of changing variables, and strictly follow all the constraints imposed on it. Unfortunately – again unlike a human trader – a machine cannot adjust to market events that were not foreseen by its creators; therefore this stage should be monitored on a regular basis.

Finally, this same system can assist in the post-trade analysis, performance attribution, and future execution optimization. When the transaction is complete, we can look back and measure the actual incurred costs, and then try to discern which ones were a result of poor judgment and faulty assumptions or just came from the realized uncertainty (volatility, chance, bad luck, etc.). Furthermore, over a longer horizon, such analysis can help determine which portion of the P&L can be attributed to the high-level investment idea and which to the (un)lucky implementation. The insights gained from this information can be later used to refine both the execution mechanism and the high-level strategy.

What are the main factors that must be accounted for during the execution? It is probably fair to say that liquidity is the most important influence on execution costs. The central question is: how many shares can one buy/sell without moving the price adversely. Higher liquidity means having more available volume on each price level, which in turn means that trading can be done faster with shorter exposure to adverse price movements. To take this line of reasoning to its logical extreme – with infinite liquidity, one can transact arbitrary size instantaneously at no cost.

A trader can chose to either take or provide liquidity. The former can be accomplished with market orders, which involves paying the spread, incurring price impact, but guaranteeing execution. The other option is to use limit orders, which can result in price improvement, but carry execution uncertainty with them. And these considerations are only a tip of the liquidity iceberg.

As it is evident from the discussion above, price volatility is also a major concern. The more volatile is the stock, the higher is the likelihood that it will move in an unfavorable direction resulting in loses. Therefore, a general rule of thumb is to transact faster in more volatile securities to minimize the price risk. Somewhat related to the volatility consideration, we may also want to explicitly consider if a trend is present in the market. This will affect the aggressiveness of our execution: for example, if we are acquiring a long position in a rapidly rising market, a trader is probably better off using market orders, paying the spread and making some impact than sitting back with limit orders and watching the market move away from him.

Finally, all these issues should be considered within the context of an entire portfolio of securities to be transacted – the goal is to minimize the *residual* liquidity and price risk. For example, the two securities with the highest volatility in the portfolio appear very risky taken separately, but may serve as natural hedge together, and thus their execution could be slowed down in favor of other securities.

Now that we have described the general investing-trading workflow and enumerated the factors that influence execution, we now classify different dimensions of trading costs and explain how they fit into our framework.

Transaction costs can be decomposed into visible and invisible. Many investors concentrate just on visible costs, while they account only for a fraction of total costs. In our framework, visible component is mostly exchange fees and commissions. It is trivial to incorporate these factors into the total return analysis by simply adding or subtracting them from observed prices.

We can also identify the following invisible costs: delay cost, trend cost, market impact cost, and opportunity costs. Delay cost represents the unfavorable change in price from the time when the investment decision has been made and the time when execution commenced.

$$\text{Delay Cost} = X * (P_d - P_0),$$

where X is the trade size, P_d is the price when the decision was made, and P_0 is the price at the start of the execution. For automated strategies, this cost can be trivialized, since time from decision to execution can be made arbitrarily small. Human traders, however, should be aware of this cost and coordinate their research and trading activities.

The price movement cost should be incorporated into the execution of those strategies that explicitly account for price trends:

$$\text{Price Movement Cost} = \Sigma\, (x_j * P_j) - X * P_j,$$

where x_j is a number of shares transacted at time j, and P_j is the corresponding transaction price. As noted earlier, transaction speed should be modified depending on favorable or unfavorable price trends.

Market impact is the most important and the most complicated aspect of the transaction costs. It can be further subdivided into two components:

$$\text{Market Impact} = \text{Temporary Impact} + \text{Permanent Impact}.$$

Temporary impact is caused by the demand for immediate liquidity, for which the opposite side of the market (liquidity providers) has to be reimbursed through price concessions. The permanent impact is caused by the information component of the trade. Each transaction can be interpreted as a probabilistic signal to the rest of the market about the intrinsic value of the underlying security. Therefore, a large market buy will first of all pay a significant “slippage” cost for taking liquidity from the sell side and also drive the price higher in a longer term.

Both of these effects are difficult to quantify, model, or even discern. This is especially true in the case of information impact. In most of our models we will implicitly set the long-term impact to zero, since we assume that our trading does not affect other market participants' supply and demand for a security. Most of our work is concerned with minimizing the immediate (temporary) impact.

Another cost to consider is the opportunity cost or the cost of non-execution:

$$\text{Opportunity Cost} = (X - x_j) * (P_n - P_d),$$

where P_d is a decision price and P_n is the price of the last execution. This cost component comes from the market not being able to absorb the entire order size. This concern can be mitigated by incurring a much higher market impact (can be very expensive and not always possible) or going through the pre-trade "feed-back loop" as described earlier and picking the appropriate size (however, market conditions can change rapidly, so this pre-trade estimation is not a guarantee).

The main idea behind the efficient execution is to aggregate all these costs and figure out a way to minimize them: act rapidly, take trends into account, transact at the "right" size, balance passive and aggressive orders, etc. But while concentrating on costs, we must not forget about the corresponding risks. Probably the most important source of uncertainty is the price volatility – i.e. a probability that the price of unexecuted shares will move against the trader. Stock (or portfolio) variance and standard deviation are the appropriate quantitative measures of these risks. If we want to assign a dollar value to our price risk, we can use Value-at-Risk (VAR) methodology over the appropriate time interval.

The other source of execution risk is what we like to call "volume volatility". This refers to a possibility that the actual volume available for transaction will differ from its historical levels, which in turn will result in a different execution price. There is little to no research devoted to quantifying volume volatility, but we believe that this is a very important aspect of the execution, and that we can use similar approaches as for the price volatility. For example, we can calculate an average available volume for a given distance away from the inside spread, and then compute the average mean squared error to get the standard deviation.

The challenge is to balance the trade-off between increased returns (or decreased costs) and increased risks. For example, using a limit order instead of a market order may promise price improvement, but it also comes with a possibility of non-execution and being forced to transact later at unfavorable prices. Thus, we may have higher expected average returns, but a wider distribution of these returns (i.e. higher risk). There are several ways we can connect risks and returns. First, we could calculate Sharpe ratio for a given strategy – essentially return divided by its standard deviation. The problem with this approach is that when the standard deviation is close to zero – i.e. in the case of market orders – Sharpe ratio will approach infinity, which will in turn bias our optimization. Our second option is to use some trade-off coefficient λ, and to optimize the following function $f = \text{return} + \lambda * \text{risk}$. One of the suggested values for λ is 0.3, based on related research. The shortcoming of this method is that we are introducing an extra exogenous variable, which we will have to fit our data or tweak it manually, thus rendering our results suspect almost by definition. Probably the most sound way to go about analyzing the risk-return trade-off is by building an "efficient execution frontier" (similar Markowitz's efficient frontier from the portfolio analysis field), where we plot all possible risk-return combinations in the x-y coordinates. We can allow a trader to pick a desired risk tolerance level, and then find a strategy with the highest expected return, which corresponds to the preferred risk level. Or, the trader can specify the desired return, and we can locate a strategy, which achieves this return on average with minimal possible risk.

The main goal of developing such system is to present the "total solution" to trade execution. We would like to build an intelligent mechanism, which quantifies embedded costs and risk and optimizes trading accordingly. We see this trading middle layer as a support tool, which should permeate every level of trading activity – from the inception of a given strategy to the actual trade execution.

2. Singe-Period Model

In our implementation, we purposefully take a myopic view of trading in order to concentrate on the market microstructure effects. We separate problems of investing and trading (i.e. execution of investment objective). Decision to either buy or sell a particular stock is exogenous to our model, and our goal is to achieve the most favorable execution price, given a high-level investment directive. For example, if the investment objective is to buy a stock – presumably because it is expected to go up in price – we try to achieve the lowest possible execution price, when the objective is to sell, we would like to sell at the highest price possible. Such ideological separation of trading and investment is not arbitrary, but reflects the real-world practices. The main reason for this de-coupling is the difference in time horizons. Investment managers tend to analyze stock returns over longer time periods – at least days, sometimes years – whereas the process of acting on this analysis (actual buying and selling of stocks) must happen as quickly as possible. While extremely large trades may take multiple days to execute, average trade execution happens over several minutes. Therefore, investment managers have to concentrate on company fundamentals that are likely to affect the stock's long-term prospects, while traders deal with short-term fluctuations in market supply and demand, which is the domain of market microstructure and our main interest here.

We have adopted a straightforward methodology for quantify this cost. We compare the average execution price p^* with a mid-spread price at the beginning of trading $p_0 = (a_0 + b_0)/2$. For the sake of cross-sectional comparison, this price difference is further normalized and is expressed in basis points. We refer to it as execution strategy's *score* $s = (p^* – p_0)/p_0$. Example: the market at the beginning of the time period is \$24.12 (ask) – \$24.18 (bid); if we execute a buy order at the ask, our price differential is (\$24.18 - \$24.15)/\$24.15 = 0.001242, or 12 basis points. We strive to minimize this score: when s = 0, we have transacted at the mid-spread price, and according to our framework, such transaction is costless.

As discussed earlier, trading costs have multiple dimensions: fees, commissions, spreads, delay costs, information impact, etc. Since execution problem is separate from the investment analysis in our setup, we can disregard the information effects of trading, and accounting for the fixed visible costs such as broker fees is trivial. Our scoring system allows us concentrate on only those costs that arise from the market microstructure mechanics: the bid-ask spread (the difference between the lowest sell and the highest buy in the limit order book) and the immediate market impact – having to execute against multiple standing orders with progressively more disadvantageous prices. Technically, price trends – the expected price movement during the execution period – should be a part of trading costs analysis, but we assume that over multiple iterations these price movements cancel out. In other words, the expected change in the mid-spread price is zero.

We now explain the microstructure sources of execution costs and propose a general idea how to minimize them. The simplest execution option for the trader is to transact with limit orders that are already present in the market. In order to do so, the trader submits a market order; it is said that he *demands liquidity*. If the trader wants to trade up to the volume available at the best quote, then the spread is the only cost he has to pay for liquidity, but if he demands more volume than quoted at the best bid or offer (BBO), he must accept further price concessions – i.e. transact at inferior prices to compensate other traders for providing additional volume.

A	B
…	…
25.56 – 300	25.56 – 300
25.55 – 1000	**25.55 – 1000**
25.35 – 200	**25.35 – 200**
25.30 – 150 ← Best sell	**25.30 – 150** ← Best sell
↕	↕
25.25 - mid-spread	25.25 - mid-spread
↕	↕
25.20 – 200 ← Best buy	25.20 – 200 ← Best buy
25.19 – 300	25.19 – 300
25.15 – 785	25.15 – 785
25.10 – 170	25.10 – 170
…	…

A **B**

.**Figure 10. Bid-ask spread and market impact.**

To illustrate this mechanism, in Figure 10 we reproduce a snapshot of a hypothetical limit order book. In case A, the trader wants to buy 100 shares. He can do so at the ask, by paying \$25.30 per share. According to our scoring system, the impatient trader pays 20 bp for liquidity: (\$25.30-\$25.25)/\$25.25 = 20 bp. In case B, we see the same book, but now the trader needs to buy 1000 shares. The best offer is good for only 150 shares, so the trader has to "*walk the book*" by paying increasingly higher prices for additional volume: \$25.30 for 150 shares, \$25.35 for 200 shares, and \$25.55 for the remaining 650 shares. The average transaction price now is \$25.4725 per share or 88 basis points. This situation illustrates the general microstructure reality: the more shares trader wants to transact, the higher his trading costs are. In addition to paying the bid-ask spread, large market orders also have *market impact* – post-trade prices move in the direction of the trade.

The only way to avoid paying for liquidity and still achieve the execution goal is to *supply liquidity* to other traders in the market. This is accomplished by submitting limit orders. In the above example, the trader can potentially achieve costless execution by submitting an order priced at \$25.25. The key word is "potentially", since execution of a limit order is not guaranteed. There is always a possibility that the market will move away from the limit order (prices will go up in this case), and then the trader will have to "chase the market" and transact at even more unfavorable terms. Limit orders provide price improvement compared to market orders, but they come with an inherent risk of non-execution, which can ultimately lead to higher trading costs. It is our goal to quantify this trade-off and to derive an order-submission strategy that minimizes trading costs.

To formalize the optimal execution problem: given an order to acquire V (positive or negative) shares of stock X over at most M minutes, we have to come up with an execution policy $P^*(X, V, M)$ with the lowest expected cost: $E[s(P^*)] < E[s(P)],\ \forall\ P \in \mathbf{\Pi}$, where $\mathbf{\Pi}$ is the set of all possible policies. We now present the essence of our approach to execution: first, we describe how to determine which limit order price results in the most advantageous execution price; second, we introduce risk into our analysis; and finally, we will combine the two to derive the "efficient pricing frontier".

a. Expected Execution Price

Let us revisit the basic setting for our execution problem: a high-level investment strategy issues a directive to acquire V shares of some stock, and this position must be entered within a time horizon of H seconds.

This task can be executed using the following trading strategies:

(1) Submit a market order for the entire amount immediately. This guarantees both the execution and the amount of cash paid (respectively, received), but has to pay for liquidity demanded

(2) Wait until the end of the time period, hoping for a favorable price move, and then go to the market with the entire amount. This may achieve price improvement over (1), but has exposure to price volatility, and still has to pay transaction costs;

(3) Submit a limit order at the beginning of the time period. This order may execute completely, partially, or not at all; then we must submit a market order for the remainder of the shares (if any) at the end of the interval.

All these strategies end up with the same position after H seconds, but will have spent different amounts of cash. Therefore, if we plot the measure of transaction costs that we have defined earlier, then we can find which one is the most efficient and if there is some general relationship between various strategies' performance.

If during the execution the trader steps over the spread and demands liquidity from the other side of the market, returns are negative thus representing transaction costs. If the trader's limit order is priced below the initial mid-spread and is later executed, then returns are positive ("transaction savings").

We submit orders from high-priced to low-priced and record average returns for each strategy. As expected, these returns tend to peak around a certain price level, which consequently represents the optimal pricing level to achieve the most advantageous execution price. A representative order price-return curve is shown in Figure 11:

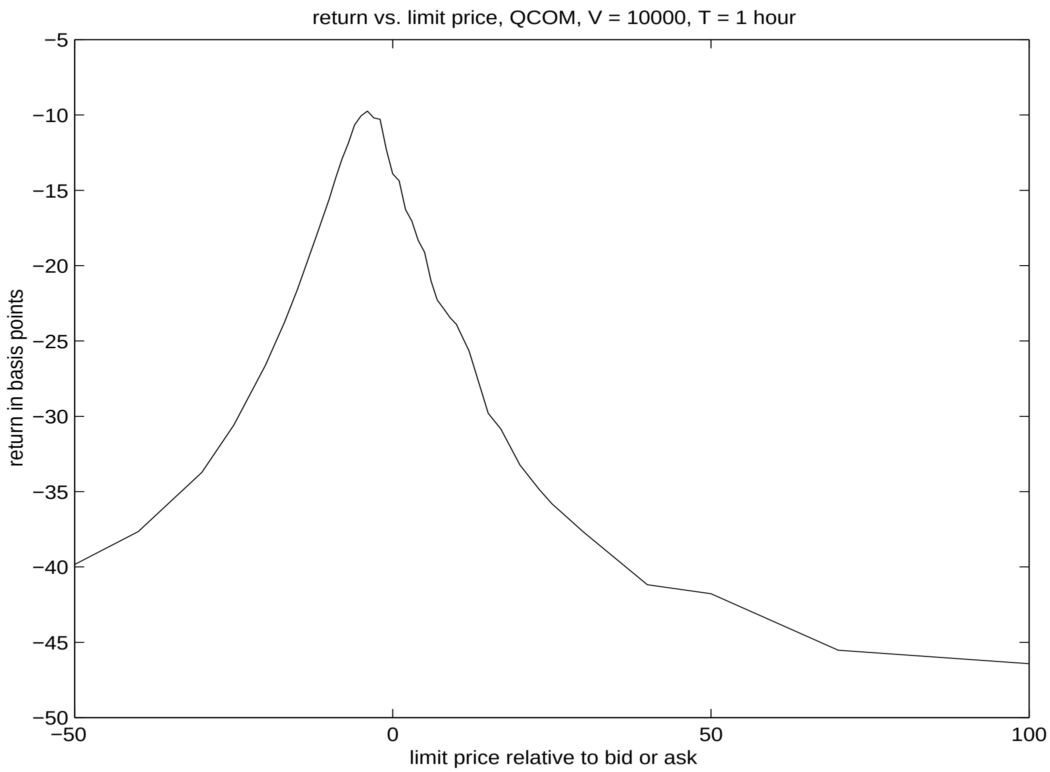

Figure 11. The peak in the curve represents pricing strategy that produces the most favorable expected execution price.

This is a realistic curve for many stocks and situations, but for concreteness, the above graph is a plot of a scenario where the task is to acquire 10,000 shares of Qualcomm stock (QCOM) within one hour.

On the x-axis, we plot where the limit order stands relatively to its own side of the market. For example, x = -70 means that the order is submitted at a limit price of best bid minus 70 cents (in other words, it is 70 cents deep within the buy book); x = 10 means that the limit price is best bid plus 10 cents – i.e. the order is submitted either within the bid-ask spread, or it is a "marketable" order, which transacts with the sell book until the specified number of shares is bought or the limit price is reached. In the later case, the shares that are not executed immediately get placed at the top of the buy book. On the y-axis, we plot the aforementioned "returns" from our execution strategies, expressed in basis points. They represent transaction costs of trading a given block and thus are negative.

By this construction, the parts of the curve on the extreme right and left of the graph correspond to Strategy 1 and Strategy 2 respectively: on the right, we have orders with high prices that get executed immediately and completely; on the left, orders are priced so far away from the inside market that they never get executed, thus forcing a market order at the end of the trial. The peak in the middle of the graph supports our main thesis – superior execution price can be achieved by using limit orders. The peak represents the lowest possible transaction cost that can be achieved, and the corresponding x-value is the optimal limit order price. Therefore, the main message of Figure 11 should be interpreted as follows: if you want to acquire 10,000 shares of QCOM over 60 minutes and seek the most attractive *expected* execution price, you should submit a limit order at the best bid minus 5 cents. There are some fundamental reasons for the curve to look like it does – we present an explanatory theoretical model in the Appendix.

Notice that the entirety of our returns curve is below zero. This means that transaction costs are always present – i.e. limit orders help to "lose less money" as opposed to generating profit opportunities on their own. While this may be somewhat counterintuitive, one has to remember that our results are averaged over many trials; therefore, in many cases limit orders end up not being executed, and the trader is forced to incur all the regular costs of a market order at the end of the time period.

b. Risk

This brings us to the second major point – returns alone do not tell the entire story. While it may be tempting to just adopt the previous conclusion that the optimal order should be submitted exactly at the price where returns peak and to end the discussion at that, we have to remember that higher returns come with higher risks. In our case, we are mostly concerned with the risk of non-execution and being forced to transact at a later time at an inferior price. And while the risk of non-execution, the mid-spread volatility, and the volume volatility are all slightly different concepts, we study them jointly by defining risk as the standard deviation of returns.

We are looking to understand the nature of this variable. A risk profile that corresponds to the returns curve from Figure 11 is shown in Figure 12. We simply plotted the standard deviation of returns (y-axis) – which are averaged in Figure 11 – for every limit order price (x-axis).

A couple of observations about the shape of the curve. First, it generally slopes upwards from left to right, which means that the deeper you hide your order in the book, the less likely it is to execute before the end of the allotted time interval, and the higher is the uncertainty around the final price.

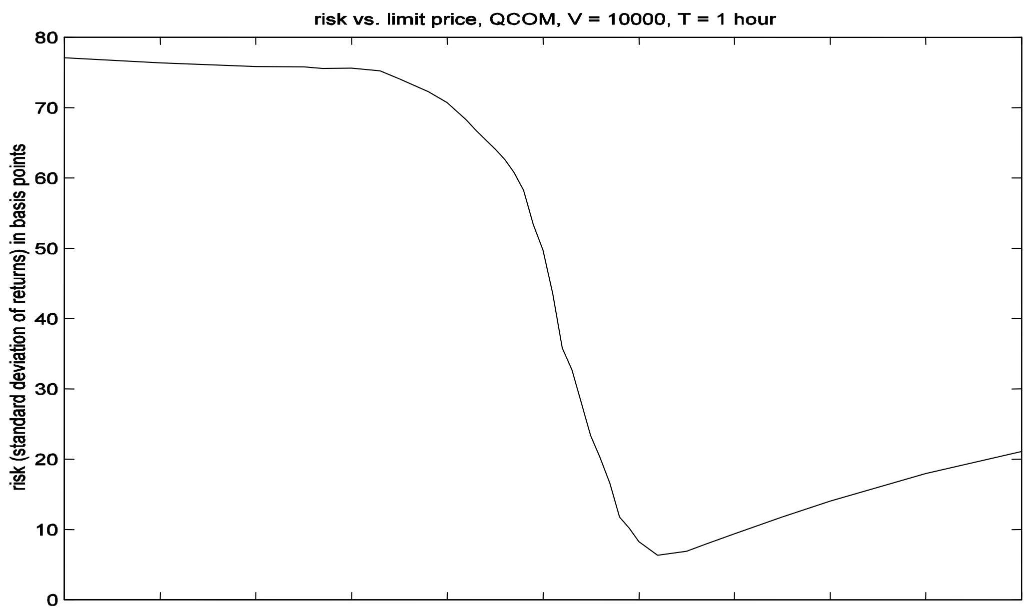

Figure 12. Returns become more uncertain the further we move from the inside market.

This shape is partly a result of our choice for the scoring system. Since we "mark" our position to the beginning of the time period, the further in time we get from the starting point (time to execution is proportional to the distance from the inside market for limit orders), the wider becomes the distribution of our "returns". Thus the general upward trend in our risk profile.

The more curious aspect of the graph is the "dip" in risk between the pure market orders and the non-marketable limit orders. This is partly an artifact of our simulation setup, but it is also grounded in reality. Market orders (limit orders with prices of bid plus several dollars) sweep the sell book for the entire size at once; therefore they trade through multiple price levels with volume getting smaller and more volatile as we move away from the inside market. The upside of this strategy is that execution is guaranteed. On the other hand, "marketable" limit orders (not as highly priced, with limit prices of bid plus a few cents) transact with the top of the sell book where volume is the highest and then leave the residual shares sitting on top of the buy book. The execution is still almost guaranteed, but the transaction price is now capped at some more reasonable level.

More formally, we can think of two sources of price risk in this setup: volume volatility, which causes highly-priced orders to trade through more price levels thus making average price less attractive, and the risk of non-execution for orders that are too passive, which translates into inferior and uncertain prices in the future. Therefore, the decreasing portion of the risk curve shows that, for certain price levels, you can decrease your volume volatility risk without increasing your non-execution risk – i.e. moderately-priced orders still virtually guarantee the execution while capping transaction costs. If this theory is correct, we have an immediate practical implication for these experimental results – it is optimal to use moderately-priced limit orders instead of full-out market orders when demanding liquidity.

At the same time, we have to be aware that we have arrived at this conclusion in an artificial testing environment and not in the real world. In other words, we are making an implicit assumption that when we submit an aggressive limit order, which leaves a large “hole” in the book behind it, incoming orders on the opposite side of the market will transact with this attractively priced order. Another possibility is that the “liquidity hole” behind the highly-priced order will fill up with the same-side orders and then the inside market will move away from our order, thereby increasing the risk of non-execution. The bottom line: in the real world, we should still see this kind of risk “dip”, but probably to a lesser extent. See [Weber and Rosenow, 2003] for an empirical study of how order books re-fill after large trades.

c. Efficient Pricing Frontier

Now that we have described both returns and risk profiles, we need a method for combining the two measures so that we can optimize them together to derive the actual optimal limit order price. A standard performance measure from the investment literature

– Sharpe ratio [Sharpe, 1964] – does not work very well for us. Sharpe ratio in its simplest form is just returns divided by the standard deviation. Since the standard deviation for Strategy 1 can be made very close to zero (transaction costs are known immediately and with certainty), this will bias the ratio heavily in favor of these strategies.

In order to perform a meaningful comparison among alternative strategies, we borrow another popular tool from the classic Finance Theory – Markowitz' efficient frontier [Markowitz, 1952]. This methodology was developed to show the trade-off between the risk and return in an investment portfolio: in order to achieve higher returns, investor has to assume more risk. The same holds true for our domain – to get price improvement the trader has to employ a riskier strategy.

To plot a risk-return profile, we place every possible execution strategy on a two-dimensional graph, where x-axis represents standard deviation, and y-axis – returns. By connecting all the strategies together, we get a plot presented in Figure 13. This profile is a combination of results from Figures 11 and 12. The semi-circle shape of the graph can be explained by the "dip" in the risk function described above. This shape has one important implication: many trading strategies from our setup are sub-optimal. Only the top part of the risk-return profile where the increase in risk results in higher expected returns should be considered in the strategy selection process. This is what we call an "efficient pricing frontier" (a similar concept is used in [Kissell and Glantz, 2003]). In this example, this is the upward-sloping section of the curve, which connects the point of minimum risk (8, -18) to the point of maximum returns (29, -9). For any other point along the curve, we can always find either less risk for the same expected return, or higher expected return for the same level of risk, or both. In terms of actual trading strategies, we conclude that it only makes sense to price limit orders in the interval [best bid -5, best bid +11].

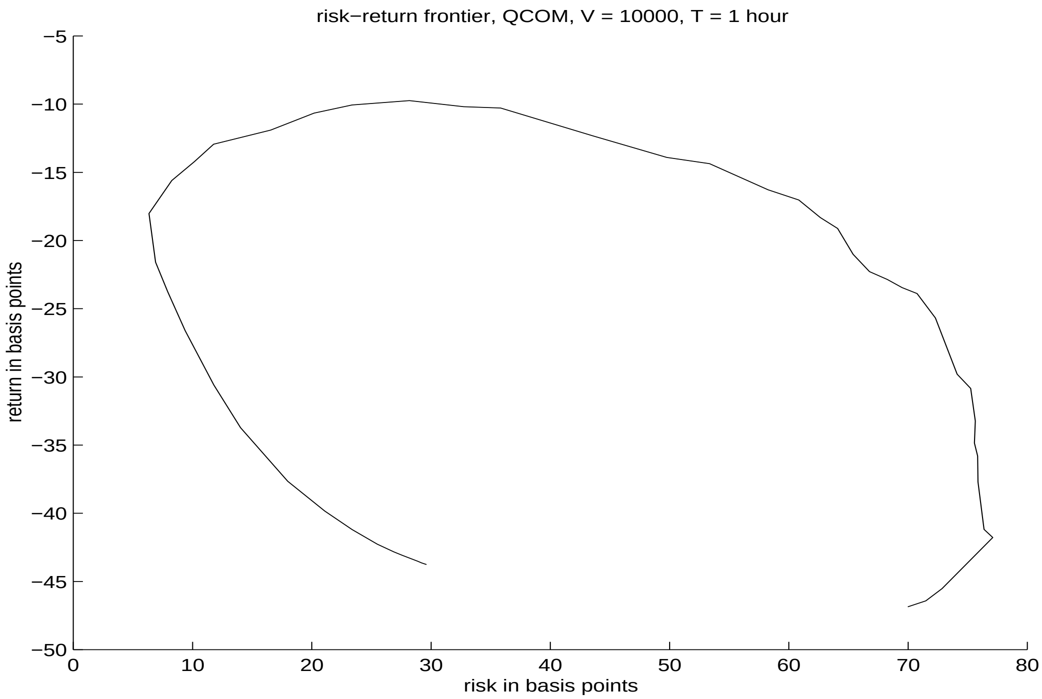

Figure 13. Trade-off between risk and return. Only the top portion where higher risk results in higher returns should be used for execution.

This is the pivotal part of our analysis: we run historical simulations to construct return and risk curves, combine the two into the risk-return profile, and ultimately extract the efficient pricing frontier. Using this frontier, the trader can do one of two things – pick a target level of returns (price improvement) and find a strategy that will deliver these returns in expectations with minimal risk; or he can select a level of risk he is comfortable with and get the strategy, which will deliver the highest expected return for that level. Obviously, after picking a point on the frontier, we need to refer back to the returns and risk graphs to determine the corresponding limit price.

d. Results

In this section we present a summary empirical results for the single-step model. Our goal here is two-fold: to show practical applications of the suggested model and to point out various microstructure variables that must be taken into account during the analysis. We first examine the effects of modifying the inputs of the execution strategy – order size, execution window, and time of the day, and then explore the real strength of our approach: conditioning the execution on the state of the market – trading volume and book depth in this case. For every variable we examine we provide plots of returns, risk, and pricing frontiers for each of the four stocks (MSFT, QCOM, NVDA, and MERQ).

Perhaps the most straightforward parameter of the execution strategy is the order size. Everything else kept constant, it is more expensive to trade larger orders. Figure A1 illustrates this point. All trading is done within one-hour period; each solid line represents 1,000 shares, dashed line 5,000 shares, and dash-dot line 10,000 shares. Across all stocks, trading a smaller volume is clearly less expensive than a larger quantity. Returns and frontier curves are therefore stacked horizontally without intersecting. Furthermore, trading larger orders is riskier, as we can see from the second column in Figure A1 – larger orders put a larger dollar value at the risk of an adverse price movement during the execution period, thus again making large-order risk curves dominate those of smaller orders.

From the position of the peaks in return curves and from the shape of the pricing frontier (they are shifting to the left with increasing size), we can conclude that the trader has to price his orders more aggressively for larger quantities. Across different stocks, the importance of liquidity also becomes apparent. Trading a large block of very liquid MSFT is more expensive than a smaller block, but only by a few basis points, whereas transacting large quantities of illiquid MERQ becomes prohibitively expensive – we see a large fallout in the case of 10,000 shares.

While we do conclude that trading large volumes is costly, it is difficult to propose a clear remedy – most of the time acquiring or selling of a significant block of securities is a necessity. One way to address this issue is to split a large order into several pieces and transact them sequentially.

If we divide a large order into multiple small orders, we must reduce the execution window for each smaller piece. This effect is explored in Figure A2. Trading here is performed over 60 minutes (solid line), 10 minutes (dashed line), and 1 minute (dash-dot line); and every transaction is for 1,000 shares. Not surprisingly, return graphs show that it is more expensive to transact over shorter time intervals – limit orders remain in the book only briefly thus making it less likely that transaction price will reach the limit level, forcing the trader to submit market orders and incur price impact. The case of the time window is not as clear-cut, however, as that of the order size. While transacting on longer time scale is less expensive, it is also more risky – the solid line dominates the others in both return and risk plots for all stocks. This means that none of the three strategies is strictly superior, and therefore the choice of the time window depends on the trader's attitude towards risk and return.

The efficient frontier plot for MSFT illustrates this point: if the trader picks the target risk level of 15, then he should buy 1,000 shares over 10 minutes (the dashed line corresponds to highest expected returns for that level of risk), whereas if the he can tolerate the risk level of 20, then he should transact over 60 minutes and expect higher returns (smaller transaction costs). We observe the same pattern for other stocks as well. Final observation: similar to the order size, shorter execution time necessitates more aggressive order pricing.

One other variable that the trader can potentially control is the time of the day when the execution is performed. Temporal liquidity patterns are well document: there is more volume right after the open and before the close than in the middle of the day. We are trying to answer a slightly different question: should we consider time of the day as a separate input variable, which can influence the outcome of our analysis? In other words, do curves differ significantly in the morning and in the afternoon? Figure A3 seems to suggest that time of the day indeed makes a difference. There we transact 1,000 shares over 60 minute intervals starting at 11 am (solid line), 12 pm (dashed), and 2 pm (dash-dot). We avoided open and close on purpose, since it is widely believed that the price formation process and liquidity dynamics are different during those times. While Figure A3 does show that curves can differ significantly from one time period to another, it is impossible to make meaningful generalizations. It does not appear that trading at a particular time of the day is more profitable than at some other time.

We contrast the above conclusion with the results from Figure A4, where we traded the same quantity of stock, starting at the same time, but during a 10-minute window. In this case, all curves are very close to one another; time of the day effect is much less pronounced. (Since the situation is the same for every stock, we show only QCOM and NVDA here.) Therefore, if the trader is planning on transacting over a long time period, he should take the time of trade into account, otherwise, this variable can be disregarded.

The real strength of our approach is presented in Figures A5 and A6, where we condition our optimization on the state of the market. In all the previous experiments, the optimal pricing frontiers that we have derived are “unconditional” – i.e. we use all the data available to us in order to construct the curves. It may be more informative and practical to condition our results on a specific state of the market by using only those parts of data that conform to desired conditions. In Figure A5, we create two sets of curves – for days with high and low transaction volume (solid and dashed line correspondingly). It appears that it is cheaper to trade on high-volume days, but it is also more risky. High volume means more liquidity and thus smaller market impact, but surges in volume are also correlated with higher volatility thus making adverse price movements more likely. Just as in the case of different execution windows, efficient pricing frontiers intersect in a non-trivial way, and thus the choice of an optimal pricing strategy depends on the trader’s risk tolerance. Also, orders should be priced more aggressively on low-volume days. This is consistent with our previous findings.

In Figure A6, we want to see how our results change when we submit our orders into a “thick” (solid line) and “thin” (dashed line) books. (We define the depth of a book by the total volume within 20 cents from the inside market). Results are similar to those in the transaction volume conditioning experiments, but somewhat less clear. While it does appear that it’s better to transact in a thick book, this relationship does not hold uniformly across all stocks. This leads us to believe that the depth of the book may be not as significant of a variable as trading volume, when it comes to limit order pricing. In any event, our goal here is to demonstrate how our optimization technique can be applied to specific market conditions. There are many other conditioning schemes that can be informative: low vs. high price volatility, low vs. high volume volatility, directional market, and so on. We explain how our model can be extended even further in the next section.

3. Multi-Period Model: Reinforcement Learning Approach

a. Formal Definition

While the single-period model of trade execution provides us with multiple valuable insights into microstructure trading (lower costs can be achieved through limit order trading, lower costs come with higher uncertainty, only a small subset of order prices is optimal, etc), it is not very practical. Traders do not submit limit orders and allow them sit in the book for an hour without reacting to market developments. Therefore, the first requirement for a more complete framework is an ability to update limit orders during the allotted execution interval. Second, if we are to allow order revisions, we must deal with partial execution – i.e. we need to know how to optimally price an order for some fraction of V shares. Finally, we might want to trade differently under different market conditions: for example, we should submit more aggressive orders when the bid-ask spread is large. In other words, limit price should be conditioned on the current state of the market.

In the previous section, optimizing trade execution was equivalent to pricing a single limit order. Now, finding an optimal policy P* amounts to determining a score-maximizing limit order price for every state of the world that can be encountered during execution. In the real world, "states" are incredibly complex – they consist of inventory holdings, full supply and demand schedules (limit order books), their historical evolution, full record of past transaction prices, time of the day, shares remaining to trade, external events (index and interest rates movements, for example), etc. It is certainly not feasible to determine the optimal action for every possible reality. We address this problem by reducing the complexity and dimensionality of our state space by attempting to identify and retain only those features (state variables) that influence our decision process: time left to execute, shares remaining, size of the spread, and so on. Each policy P is then represented by an action for every state of this new limited space. Our goal is to search over all possible strategies to find the one with the highest score. If we have identified R relevant features, we then have to perform a search over an R-dimensional space, for which we can use many available optimization techniques. Such flexibility is an attractive feature of this model.

In practical terms, we are looking for simple "rules" (patterns) induced by microstructure effects: submit a limit order when the spread is large, market order when it's small; make a limit order more aggressive when remaining inventory is large; submit more aggressive orders as allotted time runs down, and so on. This general framework of defining a tractable state space and performing optimization over it needs not be limited to just the efficient execution problem – we can apply it to market making, trend following, and many other strategies. The problem we aim to solve is two-fold: while searching over a multi-dimensional state space is challenging on its own, the real art in this setting is to determine the appropriate state representation – i.e. which features matter the most in a given problem.

We now give a more formal definition for the multi-period optimization problem. If s(P) is the score of some policy P, then our target policy $P^* = \arg\max_P s(P)$. In turn, any given policy P is a mapping from every possible state $[x_0, \ldots, x_M]$ of our state space X to actions $[a_0, \ldots, a_L]$; we can write $P(x_m) = a_l$. Possible actions correspond to different limit order prices. To determine an optimal price for a given state, we choose $P(x_m) = \arg\max_a s(P \mid P(x_m) = a)$. Changing an action in a state x_m from a_1 to a_2 results in a new policy P_2, such that $P_1(x_m) = a_1$, $P_2(x_m) = a_2$, and $P_1(x) = P_2(x)$, for all other states $x \neq x_m$. Every state $x \in X$ is a collection of attributes (state variables) that uniquely describe this state. Such attributes can be time, inventory, spread, volume, price, and many others. Time t and inventory v are determined during the simulation, while other attributes come from the empirical limit order books available to us; therefore we write $x_m = \langle t, v, o_1, \ldots, o_R \rangle$, where o_r is a state variable derived the current market state as described by the order book. The policy's score s is determined as before: initial price $p_0 = (ask_0 + bid_0)/2$; execution price is the volume-weighted price received/paid for shares sold/bought: $p = (p_1*v_1 + \ldots + p_N*v_N)/V$. The price difference is then normalized by the initial price in order to be comparable across stocks and time periods: $s = (p-p_0)/p_0$. Also as in the single-step model, we have to submit a market order at the end of the trading period for any shares that remain unexecuted to fulfill our constraints. If we define x_T as any state where $t = T$, then $P(x_T) = a_L$, where a_L corresponds to the most aggressive limit price.

b. Policy Iteration

The most straightforward and general approach to finding an optimal policy is to search over the state space of all possible policies. The following algorithm, known as Policy Iteration enables us to do just that:

Initialize policy P: $P = P_{init}$

Repeat:

- Select a state x_m using some *heuristic*
- Find an *optimal action* a^*
- $P \leftarrow P' \mid P'(x_m) = a^*$

Until *stopping criteria*

Output $P^* \leftarrow P$

While the algorithm itself may look trivial, its performance hinges on how we initialize it, how we select states to update, how we decide on an optimal action for a given state, and when do we stop.

Initialization is the easiest part – in theory, we should be able to start with an arbitrary policy and arrive to the optimal one. A good search algorithm should find a policy with the maximum score regardless of its starting point. In practice, we use a number of simple strategies for both initialization and to serve as a baseline to judge the performance of subsequent policies:

(1) submit a market order at t_0 for all states

(2) submit a market order at $t = T$ for all states; and

(3) use the optimized policy from the one-step model.

The last one involves submitting a [non-marketable] limit order at some price at t_0 and submitting a market order at $t = T$ for those states where v (shares left) is not zero.

Once we have picked a state for consideration, we must decide what action to assign to this state (at which limit price to submit an order). When we submit an order, it is always for the full remaining volume v_m, with one exception explained below. We make the following assumption: there is a single action a_{max} for which $S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max})$ is the greatest; furthermore, $S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max}) > S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max-1}) > S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max-2}) > \ldots$ and $S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max}) > S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max+1}) > \ldots S(\mathbf{P} \mid \mathbf{P}(x_m) = a_{max+2}) > \ldots$ In other words, we assume that there is a single action in state x_m which maximizes the score S of policy **P**, keeping all other states fixed, and S decreases monotonically as we move away from a_{max} in both directions. For example, if optimal price is 10, prices of 9, 8, and 7 will result in increasingly unfavorable scores as will prices 11, 12, and 13. In order to find such score-maximizing value, we can perform a simple hill-climbing search: try one price, try a higher price and a lower price, and then continue in the direction of increasing score until observing the first decrease; the second-to-last value is the optimal one we are looking for. It is important to remember that this action (limit price) is optimal only in relation to the current policy – i.e. we keep actions in all other states fixed and re-run our policy over the entire data set every time we try a new price in the current state.

Empirically, the above assumption proves very plausible with a few exceptions: sometimes a market order results in a higher score than an aggressive limit order, and an empty action ("do nothing") performs better than a passive limit order. We try these two actions after performing the hill-climbing search.

For the Policy Iteration (PI) machine learning approach this is perhaps the most important aspect: how do we select states to optimize and when do we stop iterating? There is no guaranteed upper limit for the number of iterations that are needed for PI to converge to the optimal policy. So in the worst case, we may have to try every combination of every possible action in every possible state, which is not realistic. We use the number of visits to decide which state to optimize: actions in often-visited states are more likely to influence the policy's score than those in infrequently visited states. We start with the most visited state, determine the best action for that state, go to the second-most visited state, and so on. To prevent overfitting, we stop the process after reaching states with fewer visits than some threshold. In order to speed the process up, we may also wish to sample from the visitation distribution drawing higher-visited states with a higher probability. Even when we have optimized every state in our state space, we probably have not found the truly optimal policy yet, since we've been optimizing each state separately, holding the other states fixed. This means that we have to repeat the entire process over again, perhaps multiple times. We should stop when we go through the entire state space without seeing an improvement in the policy's score from the previous iteration. The bad news is that there is no telling of how many iterations may be necessary to find an optimal policy.

In order to combat the above difficulty, we can use the structure of the problem to formulate an algorithm which can arrive to an optimal policy after a bounded number of iterations. The key observation to make is the following: in determining an optimal action for a state, it doesn't matter how we have arrived to that state, but what we will do in this and subsequent states. In other words, we are in a Markovian environment, where history doesn't matter. Note that this is only an assumption, which hinges on the ability of our state variables to capture all the information necessary for decision making.

If this Markovian assumption holds, once we have determined an optimal action for a state under the current policy, AND all states at later times are known to have optimal actions assigned to them, then we have a globally optimal action for the current state and need not re-visit that state in the future. For example, in states with $t = T$ we always have to submit a market order, so we can consider those states "optimal by construction", since no improvement is possible there. Now, we can optimize states with $t = T-1$, knowing that subsequent states are optimal; once we are done with $t = T-1$, we know that we will never have to re-optimize those states again. Knowing that all actions are globally optimal at $t = T$ and $t = T - 1$, we can optimize states at $t = T - 2$ and then all the way to $t = 0$. Once we have optimized $t = 0$, knowing that all later actions are optimal, we have found an optimal strategy.

Things become a little bit more difficult in the implementation stage. Since we are using the real-world data, not all states are accessible on all levels – for example, if we start with a policy that submits a market order at t_0, we will only encounter states with $v = 0$ on subsequent time levels during our first path through the state space. Once we start tweaking actions in various states, we will encounter new never-before-visited states as a consequence of updated actions. In the above example, once we try being less aggressive at t_0, we will start seeing states with non-zero inventory levels at later times. If we are optimizing a time level n under our algorithm, and new states appear at levels greater than n, we can no longer assume that our actions at $t = n$ are optimal, since now there are un-optimized states at subsequent levels. We have to go back and re-optimize all levels that could have been affected. The recursive algorithm is below:

```
Recursive_Opt (time level n)
{
        Optimize every state at level n
        Find the highest level t with new states
        For each level from t  to n
                Recursive_Opt (level)
}

Main()
}
        For time levels from t = N to t = 0
                Recursive_Opt(current level)
}
```

The above pseudo-code is just a skeleton, there are many subtleties to actual implementation.

Without reporting the complete evaluation of Policy Iteration, we show the evolution of the learning process on the training set and point out inherent difficulties with this "head on" approach. We run different learning algorithms on a training set, and observe how we discover policies with increasing scores as we are optimizing actions in individual states. Below is a comparison between two learning algorithms: first relies solely on visitation statistics when deciding which state to optimize next (we refer to it as v-stat); second optimizes states by time cluster, starting with $t = T$ and moving backwards to $t = 0$ (we call it t-cluster). T-cluster is a faster approximation of the Markovian-optimal algorithm outlined in previous section.

Both algorithms are initialized with the same baseline strategy – a limit order is submitted at $t = 0$, left in the book until $t = T$, and then converted into a market order if necessary. We have described how to obtain an optimal price for such one-period limit order in previous section.

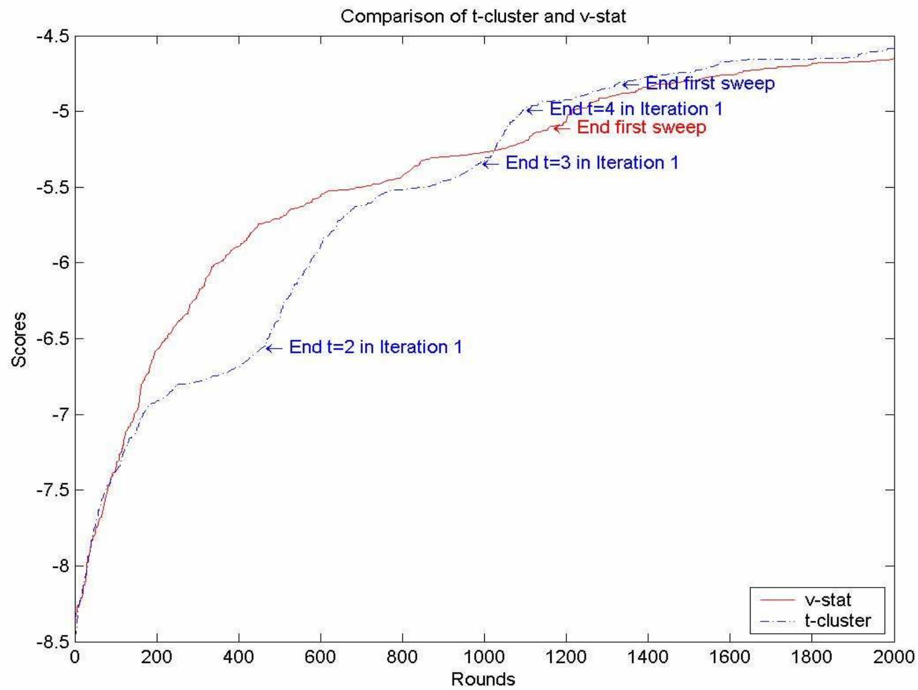

Figure 14. T-cluster vs. v-stat.

It is clear from Figure 14 that our baseline strategy, when applied to the training set (AMZN 01/03 – 12/03), produces a score of -8.5. Both algorithms start from the same point, and v-stat learns faster initially by picking the most promising states, but it levels out eventually. Output from t-cluster consist of distinct segments - we can see it moving from one time cluster to another (in our state space t measures time *remaining* and thus t = 2 is the last decision point, and t = 4 is the starting point). We can see that t-cluster outperforms v-stat in the long run. Both algorithms improve to about -4.5 bp, which is significant.

As mentioned earlier, both algorithms are just approximations of a truly optimal algorithm. After both t-cluster and v-stat reach their stopping points, further improvement is still possible. Figure 15 shows that additional optimization can be achieved if the same two algorithms were re-initialized with their respective outputs from the previous run.

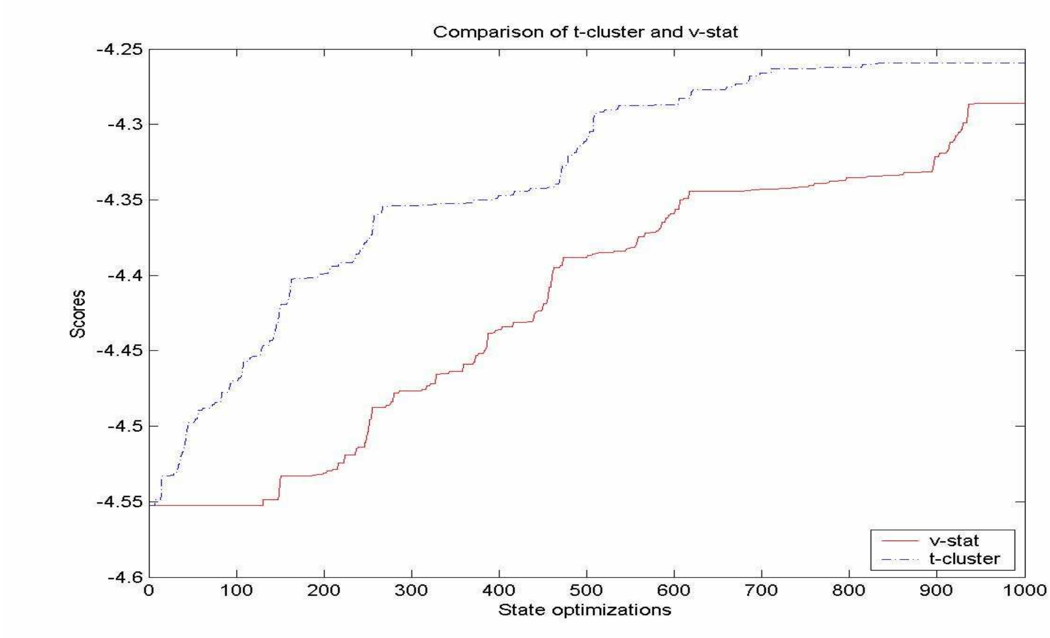

Figure 15. Further Optimization Is Possible.

From the slope of the curves in the above graphs, we can already conclude that visitation statistic is an effective heuristic function for choosing a state to be optimized: curves are steep early on and then they level out. We, however, provide a more structured analysis of our heuristic function. Below is the histogram, which shows how a number of visits to a state relate to an average expected improvement from such state – see Figure 16. The general trend is that optimizing often-visited states is more likely to lead to a superior strategy. This graph is based on statistics generated by the v-stat algorithm after a single iteration over the entire state space.

So far we have limited our discussion to the performance of our algorithms on the training set. We saw that we are able to achieve a significant improvement compared to our “seed” policy. But the real reason we do machine learning is to achieve superior execution on the future, never-before-seen data – i.e. the test set. One of our concerns is that on the training set, we are optimizing states that have only been visited once or twice. Improvements that we see from such states are likely incidental rather than systematic, and therefore such states should be avoided. Such behavior is also justified by our analysis of the heuristic function – states with few visits are unlikely to produce tangible improvements in overall performance.

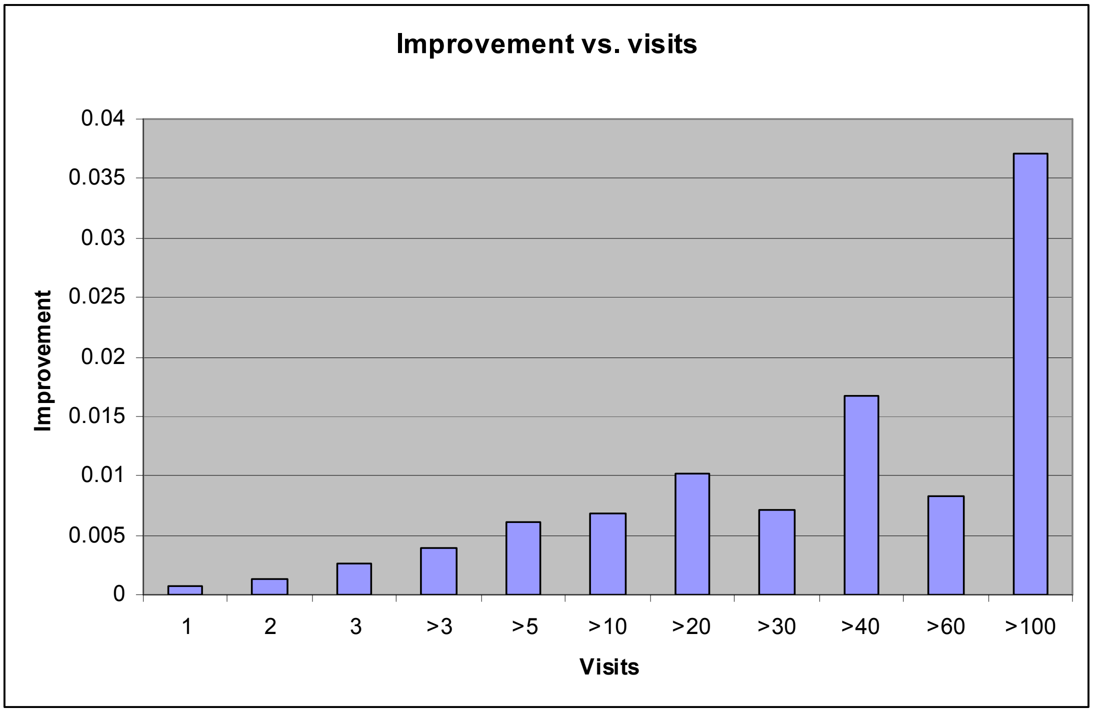

Figure 16. Evaluating Heuristic Function.

Figure 17 compares the training set performance of the regular t-cluster algorithm and its version that avoids states that have been visited less then 3 times (t-cluster-v2):

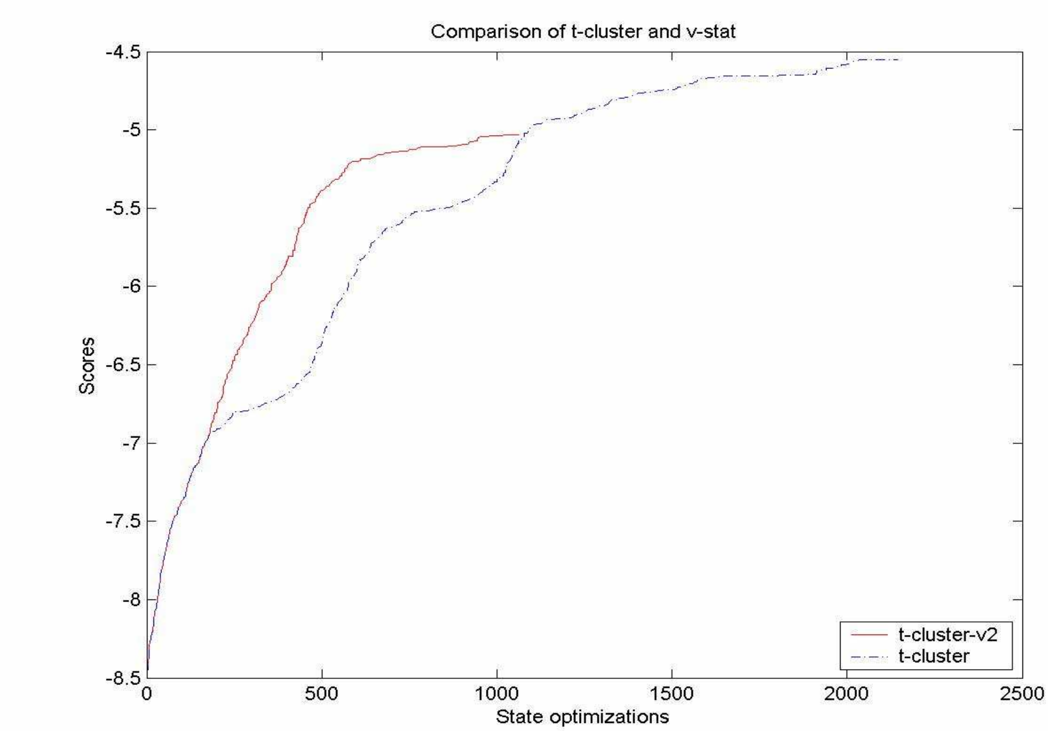

Figure 17. All States vs. States with 3+ Visits

It shows that while the new algorithm does not achieve the same performance as the one which looks at every possible state, it takes fewer iterations to finish.

One other property that we are looking for in a search/optimization method, is that, seeded with two different policies, it should arrive to the same optimal policy (or, in any event, come very close). Unfortunately, in our experiments this is not always the case. The two graphs here show the distribution of states with full inventory (Figure 18) and no inventory (Figure 19) when we apply two final policies to the available data.

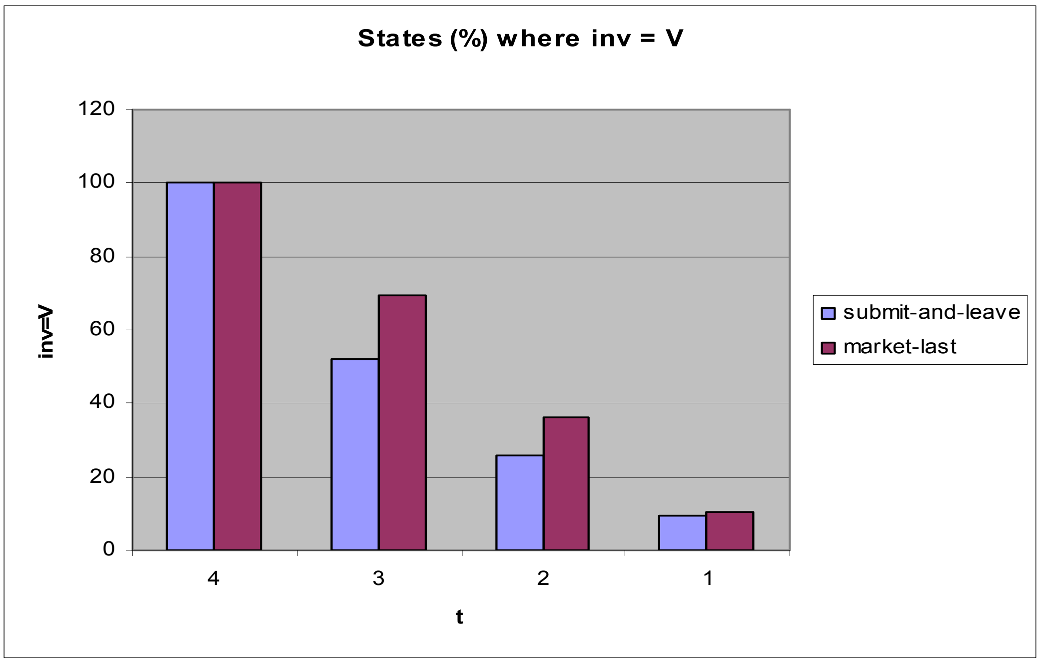

Figure 18. Percentage of States with Full Inventory.

One policy has evolved from a "submit-and-leave" baseline (blue), while the other one was seeded with the "market order at the last step" strategy. Both final policies start with full inventory at t = 4 (state space notation) and end up with no inventory at the final step – all this by construction. But they differ in their behavior during intermediate steps, where they are still biased towards their seed strategies. The blue policy is more aggressive early in the experiment, which is reflected by it having relatively more states with no inventory and fewer states with full inventory. The other policy is biased in other direction towards being aggressive at the end, which can be explained by its evolution from the "market order at the last step" strategy. Both "optimal" policies failed to merge into a single policy.

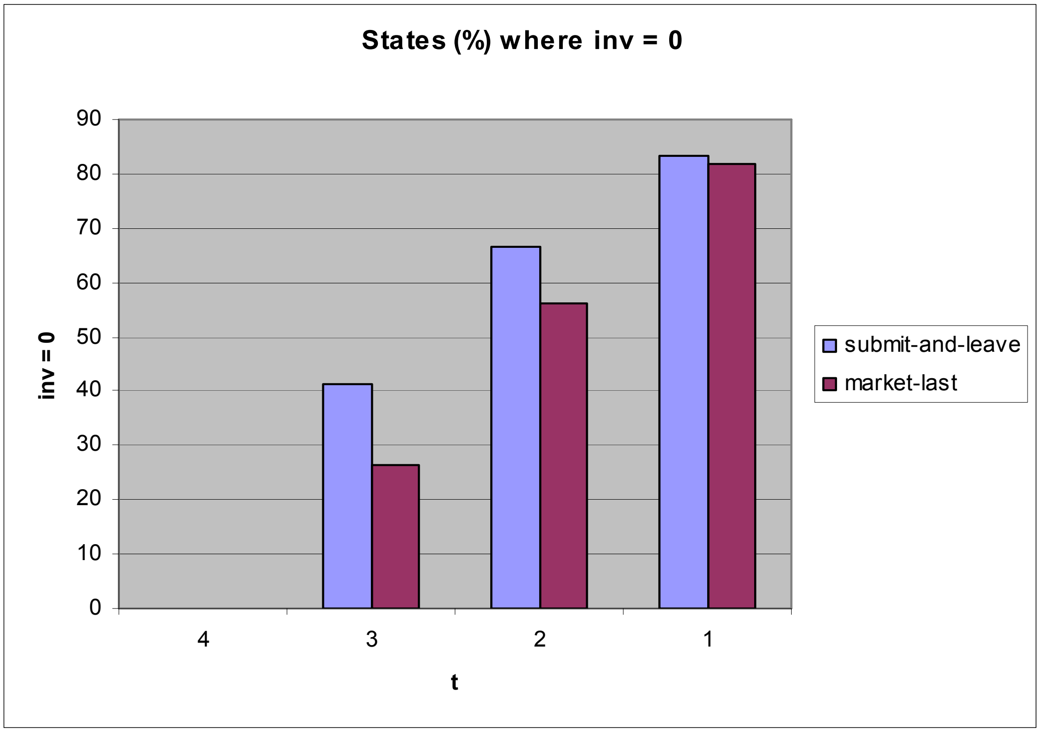

Figure 19. Percentage of States with No Inventory.

Applying policy iteration to execution optimization certainly has its merits: it is simple, general, allows to modify individual heuristics without changing the main algorithm, previous actions of the trading agent affect future states. But the shortcomings are numerous: extensive use of heuristics, multiple visits to the same state, even after reaching the stopping criteria, further improvement is still possible, failure to converge to the same policy when seeded with different initial policies. Perhaps the biggest problem is the running time: we have R state variables, each of them can take H values, and in each state we can take L actions, in the worst case we will have to go through our dataset L^(H^R) times. Plus datasets we work with are very large because of the high-frequency nature of the data. Costs of computing an optimal or even a near-optimal strategy will be prohibitively high unless we can drastically reduce the search space.

c. Dynamic Programming Approach

To address the limitations of Policy Iteration, we propose a new reinforcement learning algorithm, which takes advantage of the structure of the execution problem, and which is both fast and uses data efficiently. The feature that allows us to learn more efficiently is the Markovian nature of trade execution: if our state space is properly defined, our optimal action at any given point in time is independent of what actions we have taken previously. In other words, the history of state visits during the execution is irrelevant – no matter how we have arrived into a certain state our action there depends on that state only. Under this assumption, we need to optimize every state only once: if all states reachable from the current state are know to have optimal actions assigned to them, then once we determine an optimal action in this state, such action will remain globally optimal no matter what changes are made to other states (the ones that are not reachable from the current state).

How do we determine for any given state, which states are independent of the current state? We exploit the fact that with every action that we take, we advance once step closer to the end of the allotted execution period. If we have T decision points (chances to update our order), then no sequence of actions can be longer than T+1 (T update points, plus the mandatory market order at the end). By Markovian property, optimal action in a state with $t = \tau$ is independent of actions in all states with $t \leq \tau$. Extending this logic, optimal actions in states with $t = T$ are completely independent from all other actions. Indeed, when time runs out, we are forced to submit a market order for all unexecuted shares to bring our inventory to the target level V (no mater what else we do between $t = 0$ and $t = T-1$). Such time dependence allows us to solve our optimization problem dynamically: having established optimal actions for all states with $t = T$ (market order in this case), we now have all the information we need to determine optimal action for all states with $t = T-1$; having done that, we move one time step back to $t = T-2$, and so on until $t = 0$. Having arrived at $t = 0$, we have a globally optimal policy, under the assumption that our state space is a Markovian environment, of course. Another significant assumption that our conclusion relies upon is that our own actions do not affect the behavior of other market participants (they do not disturb the order flow).

According to our formal notation, our actions do not affect state variables $o_1 \ldots o_R$, but they do influence t (with every action that we take t increases by 1) and v (as we gradually execute our order v decreases monotonically); we refer to t and v as "private" variables, and we call $o_1 \ldots o_R$ "market" variables. We exploit this separation property to use available data more efficiently (and thus reduce overfitting) and to make sure that we optimize every feasible state within our dataset.

Say we are optimizing execution over a two-minute period, with T = 4, which means that we update our limit order every 30 seconds. Such learning parameters partition our data into two-minute long episodes, with 4 30-second sub-periods within each episode. According to our dynamic programming approach, we need to start optimization with t = T = 4, which implies that during our first passage through our dataset we should skip to the end of each episode and determine the cost of a market order there. But there is no good reason to use this little of available data – if we estimate the cost of a limit order every 30 seconds, it does not make our estimate biased in any way, but improves its accuracy, since we now use 4 times as many episodes. The same applies to t = 3: we do not skip the first 3 sub-periods, but use every 30-second interval to simulate what happens to a limit order; after 30 seconds, we need to know the cost of submitting a market order for the residual volume, but we have already estimated it in the previous iteration. So again, we do not bias our results and use data more efficiently.

The second issue is the remaining inventory. When we are estimating the cost of trading in a particular state, we have all state variables specified except for v: we determine $o_1 \ldots o_R$ by converting the limit book information, and we set t according to our dynamic programming algorithm as just described, but v can take any value, since it is dependent on actions at earlier time steps. This means that we have to try every possible value of $v \in [0, V]$ in every state we are optimizing. This ensures that every possible state that can be generated from our dataset gets visited.

Finally, we explain our methodology for determining an optimal action in a given state. Our approach is similar to Q-learning (see [Mitchell, 1999]): in every state we encounter, we try all possible actions and update the expected cost associated with taking each action and following the optimal strategy afterwards. Taking an action results in an immediate payout if any number of shares get executed, and it transfers us into a new state one time step later. Since our learning moves backwards in time, this new state has already been optimized and we know the expected cost of following the optimal strategy from that state. More formally, our cost update rule is the following:

$$c(x, a) = n/(n+1) * c(x, a) + 1/(n+1) * [c_{im}(x, a) + \arg\max_{p} c(y, p)],$$

where x is the initial state, a is the action taken, y is the new state, n is the number of times we have tried a in x, and c is the cost.

In order to learn an optimal strategy through our dynamic programming approach, we need to go through the dataset T*V*L times (L is the number of actions in each state). This is a much smaller number than the worst-case scenario under Policy Iteration, and it has several interesting properties. First, the running time of our algorithm is independent of the number of market variables, and, second, we can increase T and V arbitrarily without risking overfitting because of the efficient use of our data. Figure 20 demonstrates this point.

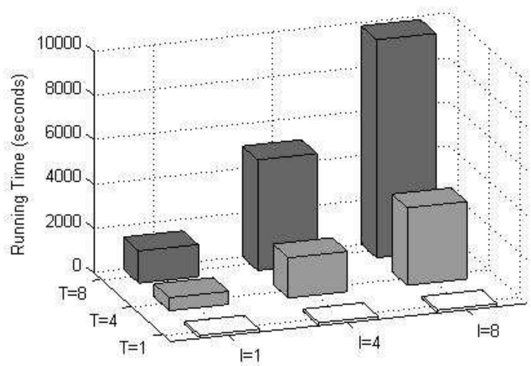

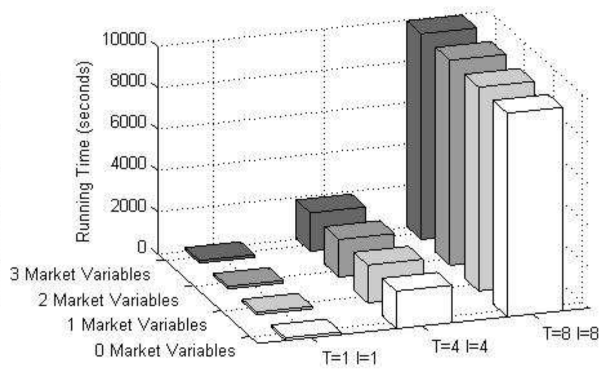

Figure 20.

Here is our algorithm in pseudo-code:

```
Optimal_strategy (V, T, L){
    For t = T to 0 {
        While (not end of data){
            Generate o1 … oR from the current limit order book
            For v = 0 to V {
                For a = 0 to L {
                    Calculate immediate cost
                    Determine the new state
                    Update c(<t,v, o1 … oR>, a)
                }
            }
```

}

}

Back-out optimal strategy by selecting the lowest-cost action in every state

}

4. Exploring the State Space

In this section, we describe our experimental results. We analyze performance of trade execution models of progressively increasing complexity. As previously explained, we partition our variables into exogenous, private, and market variables. Exogenous variables – stock, trade size, and time-to-execution – are not a part of our execution model, but rather a description of the task at hand. Private variables regulate how often we update the limit order and how crudely or finely we reason about unexecuted inventory. We use market variable to see whether we should act differently under different market conditions.

To keep our empirical results consistent, we fix the exogenous variables for all experiments. We investigate three stocks: AMZN (Amazon Corp), NVDA (NVIDIA Corp), and QCOM (Qualcomm); two order sizes: 5,000 shares and 10,000 shares; and two execution horizons: 4 minutes and 8 minutes. When we make changes in the composition of our state space by modifying private and market state variables, we test each new setting on all 12 combinations of exogenous variables:

AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA	QCOM	QCOM	QCOM	QCOM
2 min	2 min	8 min	8 min	2 min	2 min	8 min	8 min	2 min	2 min	8 min	8 min
5,000	10,000	5,000	10,000	5,000	10,000	5,000	10,000	5,000	10,000	5,000	10,000

Table 8.

In the text, we present our results selectively here due to space constraints, but we provide full results in the appendix. Tables A1 and A2 show comprehensive RL output for all combinations in Table 8. Tables A3 and A4 report results for AMZN and NVDA only in more "exotic" settings. In the Appendix we report training time, expected training score (internal RL algorithm score), actual training score (optimal policy applied to the training set, different from expected score), standard deviation of the actual score, number of episodes in the training set, test score (optimal policy applied to the test set), its standard deviation, and number of episodes in the test set. We train on several sets of different length in order to control for overfitting.

We observe the following dependencies among exogenous variables in all our experiments. All other parameters kept equal, (1) NVDA is the least liquid stock, and thus is the most expensive to trade, QCOM is the most liquid and the cheapest to trade; (2) trading larger orders is always more costly than trading smaller ones; (3) having less time to execute a trade results in higher costs.

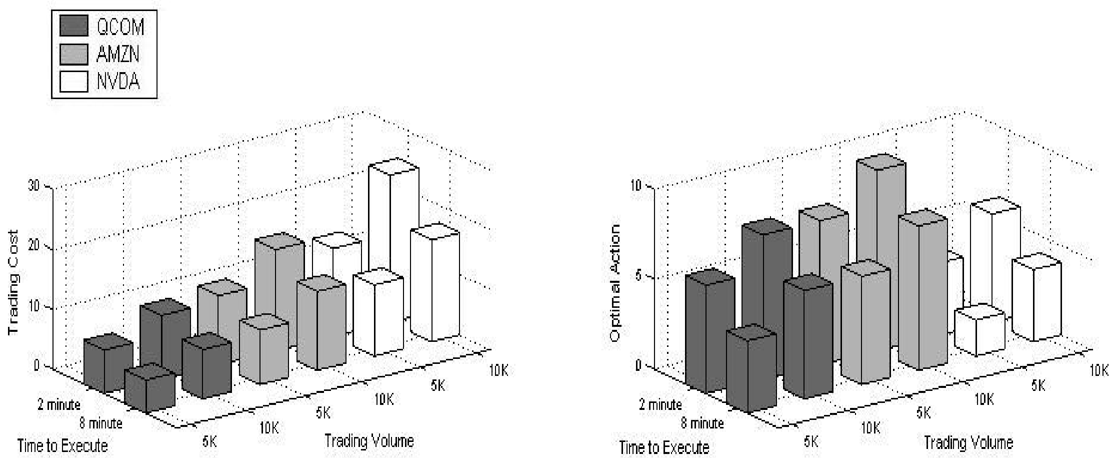

Figure 21. Cost and optimal actions across exogenous variables.

Figure 21 summarizes all of these dependencies. Trading costs in the left graph are based on the one-period model, and optimal limit order prices plotted on the right correspond to the cost-minimizing limit order prices as derived in Figure 11, but created separately for each stock-volume-time setting. We have also determined that for larger volumes and longer execution periods microstructure effects are more pronounced – i.e. we get more consistent improvements when using more complex models. We can see this in Figure 22, where we achieve better relative performance in experiments with 10,000 shares. Table A1 provides more details, including risk associated with optimal actions.

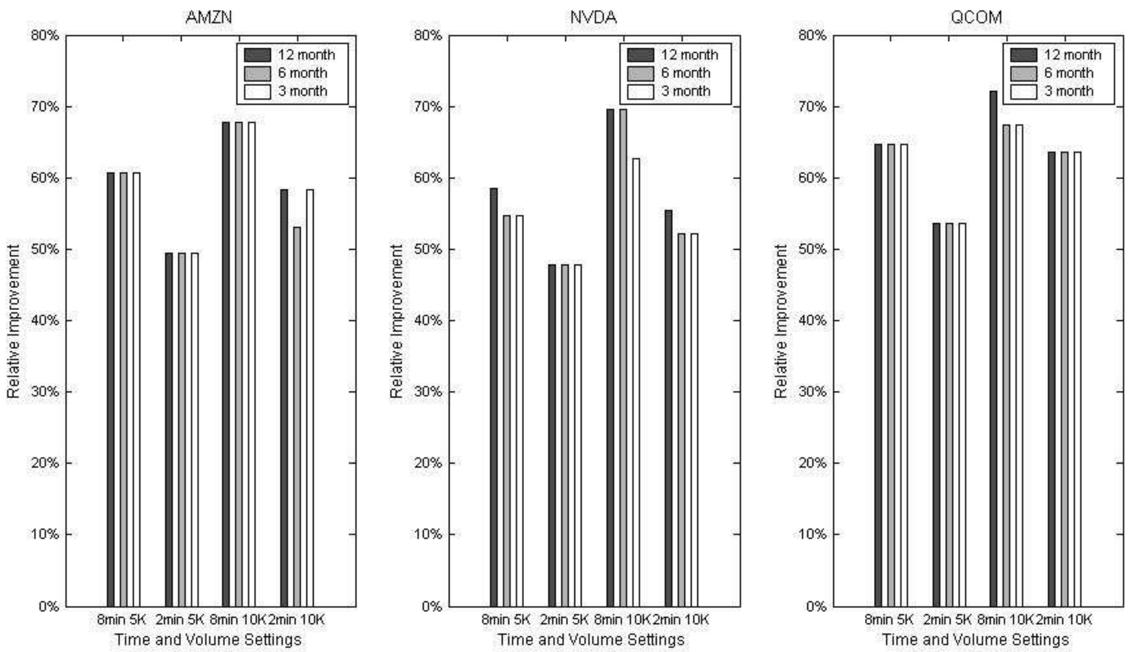

Figure 22. One-period model vs. market order (more volume = more improvement).

Figure 22 also illustrates our training set – test set methodology: in order to keep track of whether we have sufficient amount of data to learn a given model without overfitting, for every configuration of the state space we train our RL on 12, 6, and 3 months worth of data. In a simple setting such as one-step model here, 3 months of data is largely sufficient, but all things being equal the more data we use to train the RL, the more robust is our performance on the test set. Another effect at work here – how recent our training data is. Even in more complex settings that we describe later in the section, small training set continue to be sufficient (if the results are positive for 12 months, they will also be positive for 3 months). Since all training sets directly precede the test set, 3 months data is relatively more "fresh" than 12 months data, which explains this better-than-expected performance on shorter training sets.

We use the above results from the one-step model as a benchmark for more complex settings, where we are allowed to update the order during execution and to act differently with various levels of inventory. Figure 23 shows a comparison between our benchmark (left panel) and scenarios where we can update the limit order 4 times and can differentiate among 4 levels of inventory (T = 4, I = 4 – middle panel), and perform 8 updates and distinguish 8 inventory levels (T = 8, I = 8 – right panel).

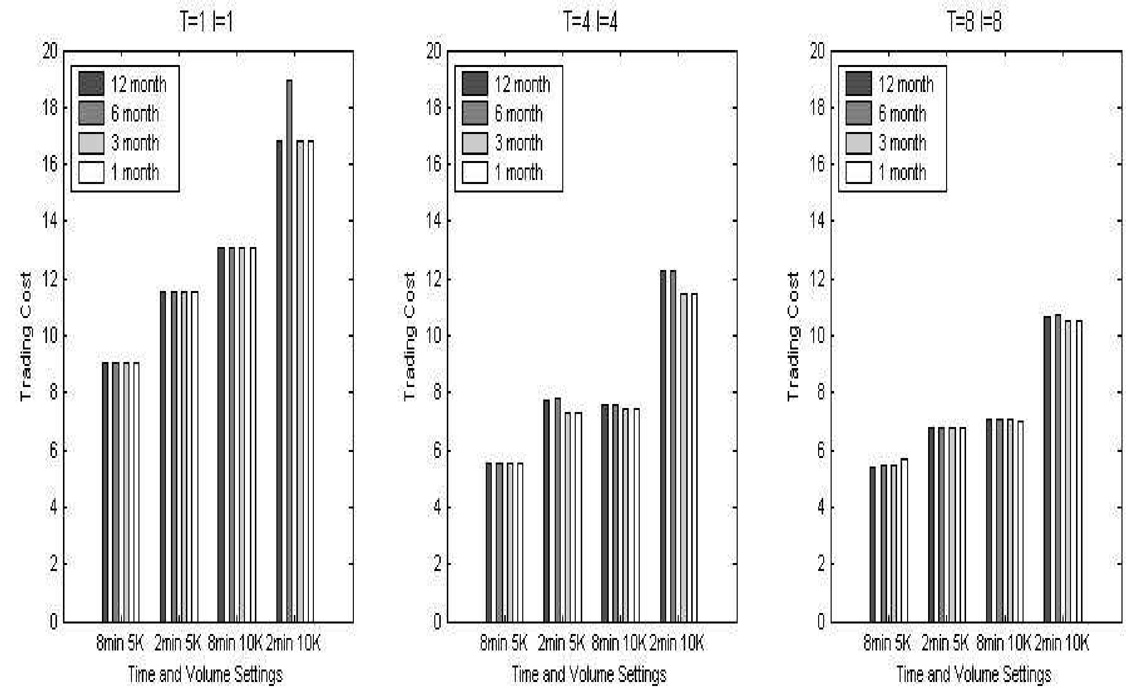

Figure 23. Increasing T and I leads to lower costs.

Modifying the state space by increasing either T or I leads to lower costs.

T=4 I=1	27.16%	T=8 I=1	31.15%
T=4 I=4	30.99%	T=8 I=4	34.90%
T=4 I=8	31.59%	T=8 I=8	35.50%

Table 9. Relative improvement of T&I over S&L.

We report relative improvements for all T&I experiments in Table 9. Another conclusion that we draw from this data is that increasing the number of updates T produces relatively larger improvement in performance than increasing the inventory resolution I. For more details, refer to tables A2 through A4.

We have also determined that increasing T gives us relatively more improvement:

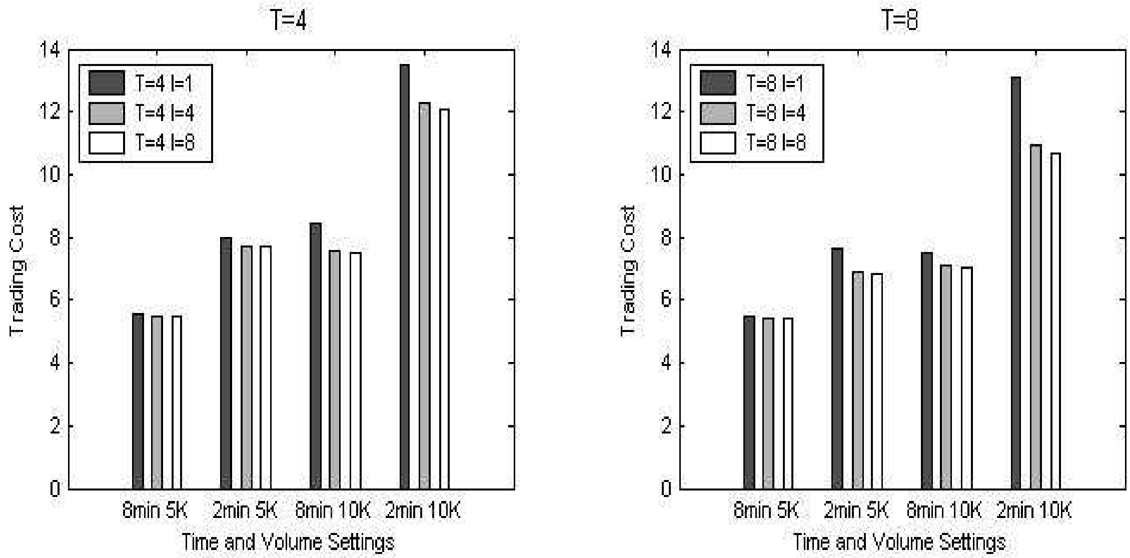

Figure 24. Increasing T provides marginally more improvement than I.

We ran additional experiments where we hold T constant and change I, or hold I constant and change T – see Figure 24. Increasing I (moving between bars in the same panel) helps reduce trading costs, but by a small amount, whereas the increase in T (moving between panels) is more pronounced. When we combine this insight with the running time information from Figure 20 that increasing either T or I affects running time roughly equally, we can conclude that if our training is time constrained, we should expand T instead of I.

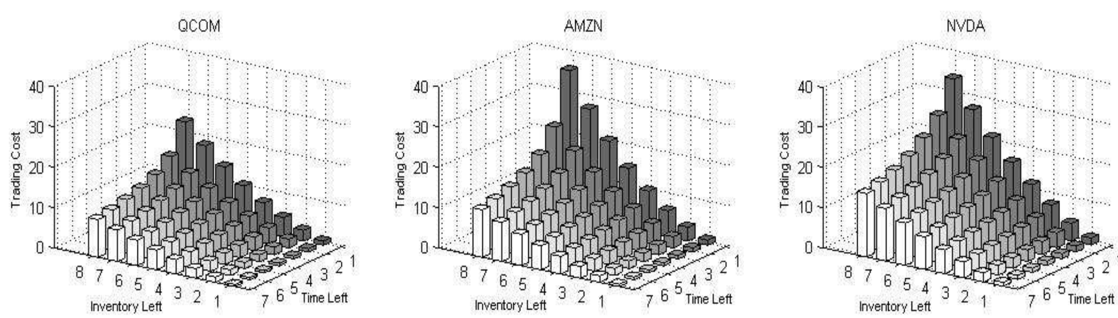

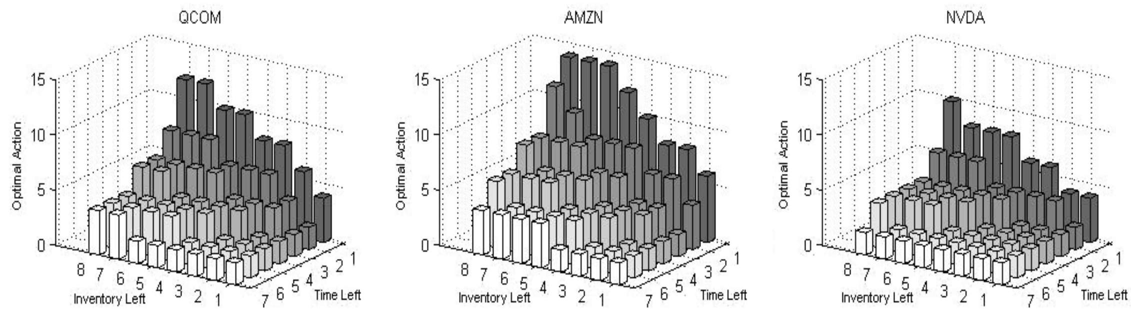

Figure 25. Cost and optimal actions for T=8, I=8 (10K, 2 min).

In Figure 25 we plot both expected trading costs (top) and corresponding optimal actions (bottom) for a setting where we have 8 order updates and 8 inventory levels. We show all three stocks, but we fix trading volume at 10,000 shares and execution time at 2 minutes. Essentially, each panel in figure 10 corresponds to a single bar in Figure 21. We can see that in this setting we will take different actions depending on which state we end up during execution. The more inventory we still have to transact, the more cost we expect to incur and the more aggressively we have to price our orders. The same relationship holds for remaining execution time – as time runs out, trading becomes more expensive, and we have to act more aggressively in order to avoid incurring the cost of submitting a market order at the end. Expected transaction cost and order aggressiveness are directly related in all settings. The curious case is NVDA: while it is the stock with highest expected transaction costs, it is also a stock with the most passive actions.

To explain where this relationship between cost and aggressiveness come from, in Figure 26 we show how we derive optimal actions from the q-values, which are updated during learning. We plot cost as a function of actions, and pick the action which corresponds to the nadir of the curve, except that now we have multiple curves and they "bottom out" in different locations.

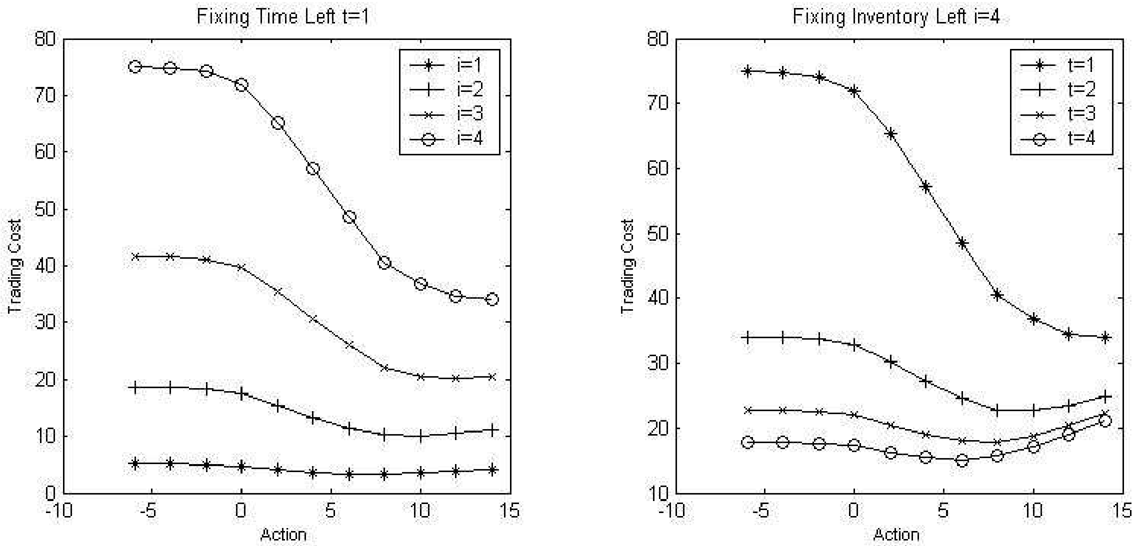

Figure 26. Q-values: curves change with inventory and time (AMZN, T=4, I=4)

Here again, we show a single representative snapshot from our results: in a setting with T = 4 and I = 4, we hold time constant (left panel) and show how q-values evolve for different levels of inventory, and then we fix inventory (on the right) and plot how q-values depend on remaining time. Shapes of q-value functions explain the relationship between state variables, costs and actions in Figure 25: for large inventories and little time remaining, the entire q-value function shifts upwards, which reflects higher expected cost for all actions; but we also observe that the nadir of the cost curve shifts more to the right, which indicates that optimal action must be more aggressive in this situation. Observe that these horizontal shifts in optimal action are more pronounced when we change t – this explains the relationship in figure 24 and support our conclusion that T is a relatively more "valuable" variable than I.

Having demonstrated the power of RL applied to execution optimization, we now add market variables to our state space to investigate whether optimal actions should be contingent on market conditions. We summarize our findings in Table 10, which shows improvement over using just the T and I variables (averaged over all stocks, sizes, and execution periods) for each market variable we tested. All these variables are derived from the empirical order books that RL observes during learning.

(1) bid volume – total volume at the highest buy price

(2) ask volume – total volume at the lowest sell price

(3) bid-ask volume misbalance – bid volume minus ask volume

(4) bid ask spread – bid-ask (in cents)

(5) price level – percentage price increase/decrease compared to the beginning of the episode

(6) immediate market order cost - how much would it cost to submit a market order for the balance of inventory immediately instead of waiting until the end of the episode (a measure of liquidity beyond the bid-ask spread)

(7) signed transaction volume – total volume of market buy orders minus total volume of market sell orders within a given time interval (a measure of supply/demand misbalance)

(8) price volatility – standard deviation of the mid-spread price

(9) spread volatility – standard deviation of the size of the bid-ask spread

(10) signed incoming volume – total volume added to the buy book minus total volume added to the sell book within a given time period

(11) fast/slow market – number of actions (order submissions, executions, and cancellations) that happened within a given period

In the majority of our experiments, "given period" is the 15 seconds preceding the decision point. While each of the variables above can take many values, we constrain them to only 3 – 5 alternatives. "High" (2), "low" (0), and "average" (1) is the most frequently-used encoding. This conversion is accomplished by comparing a realized value to some historical average, and placing the result into one of three bins depending on its sign and magnitude. We also explored combinations of 2 and 3 market variables. Specific scores from introducing each market variable into the RL framework can be found in Tables A5 and A6.

Bid Volume	-0.06%	Ask Volume	-0.28%
Bid-Ask Volume Misbalance	0.13%	Bid-Ask Spread	**7.97%**
Price Level	0.26%	Immediate Market Order Cost	**4.26%**
Signed Transaction Volume	2.81%	Price Volatility	-0.55%
Spread Volatility	1.89%	Signed Incoming Volume	0.59%
Fast/Slow Market	1.5%	Spread+ImmCost+Signed Vol	**12.85%**
Spread + Immediate Cost	**8.69%**		

Table 10. Additional cost reduction when introducing market variables

The first insight from this data is that finding profitable opportunities based on market patterns is very difficult, which is consistent with the Neo-Classical Finance Theory and the Efficient Market Hypothesis. We do, however, succeed in finding several variables (highlighted) that allow us to reduce trading costs even further if we condition our actions on them in addition to time and inventory. Figure 27 plots the relative improvement of strategies that consider the bid-ask spread, the cost of submitting a market order immediately, and both market variables together.

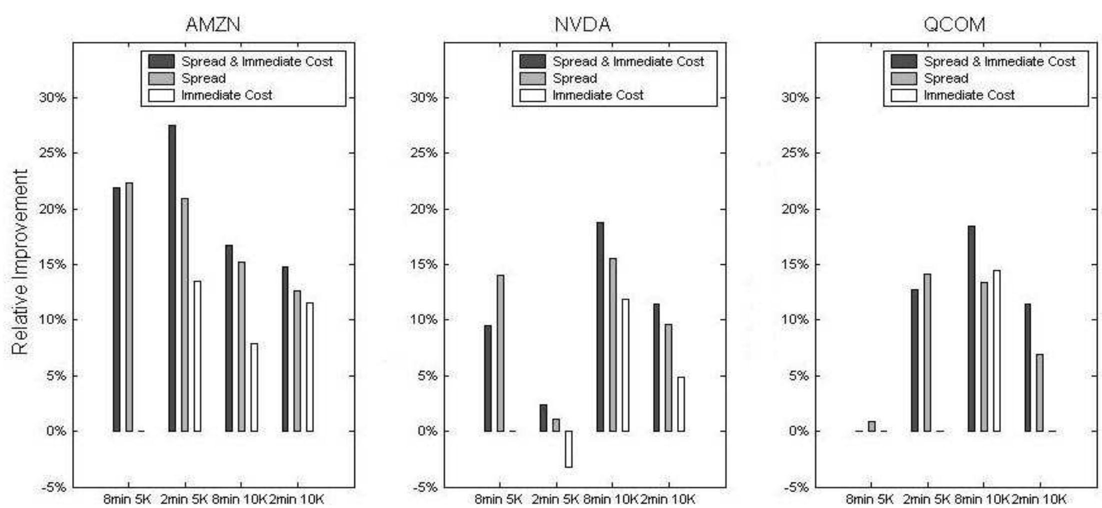

Figure 27. Improvement brought by market variables.

In order to provide this improvement in performance, the RL algorithm must learn different actions for different values of market variables. Figure 28 is similar in spirit to Figure 25: it shows how optimal actions depend on the bid-ask spread and the cost of market order submission, while keeping other variables (V, H, T, and I) fixed. Larger spreads and lower costs of submitting a market order require more aggressive actions. Other "useful" market variables induce similar dependencies.

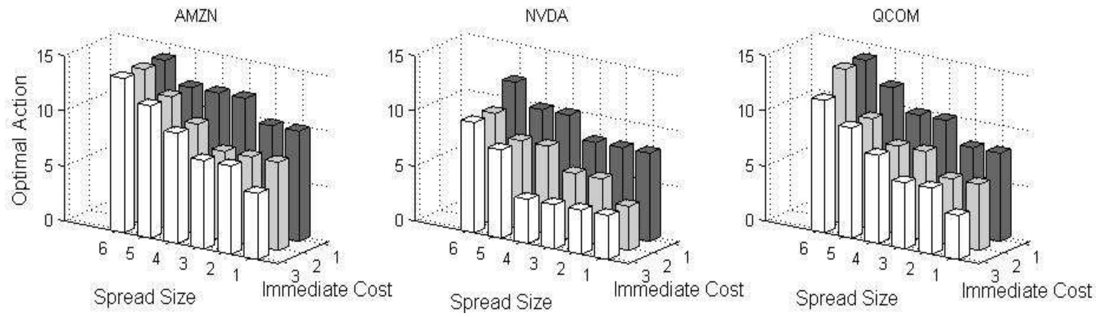

Figure 28. Large spreads and small market order costs induce aggressive actions.

As it was the case with T and I, it is the shape of Q-value functions that is responsible for difference of optimal actions across states. In the left panel of Figure 29, we see how the location of the optimal action on the cost curve shifts to the right as the spread size increases. On the right, we show why some market variables do not improve the performance of our algorithm: while volume misbalance is a predictor of future trading costs (as indicated by difference in Q-value functions), we cannot take advantage of this predictability, since cost-minimizing actions are identical in all three cases.

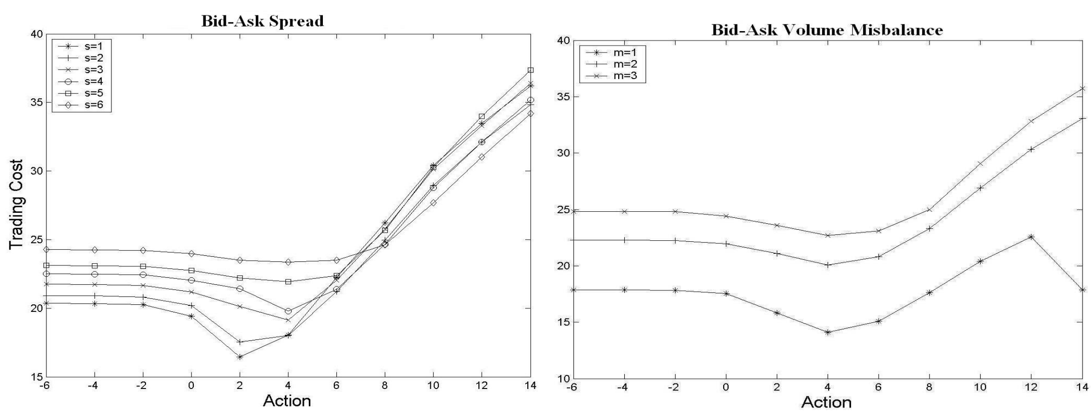

Figure 29. Q-values: different action vs. cost indicator

5. Model Extensions

a. Executing at the Open or Close

By comparing execution in the morning (9:30 – 10:30 am), evening (3:00 – 4:00 pm), and during the rest of the day (10:30 am – 3:00 pm), we are looking to answer the following two questions:

(1) When execution is the cheapest?

(2) Should execution strategies change during different parts of the trading day?

Quick answers: (1) the close is the cheapest, then the middle of the day, then the open; (2) yes, optimal policies differ among different portions of the trading day.

The detailed answer to the first question is in Table A7 where we compare optimized execution costs at the open and close relative to the middle of the day. For example, if we want to buy or sell 5,000 shares of AMZN over 2 minutes by using the submit-and-leave strategy, our expected trading cost is about 11 bp during the day, 13 bp at the open (12% increase in cost), and 9 bp at the close (21 % reduction in cost). Across the board, the open is the most expensive, and the close is the least expensive.

This relationship is plotted in the following graph (for AMZN only):

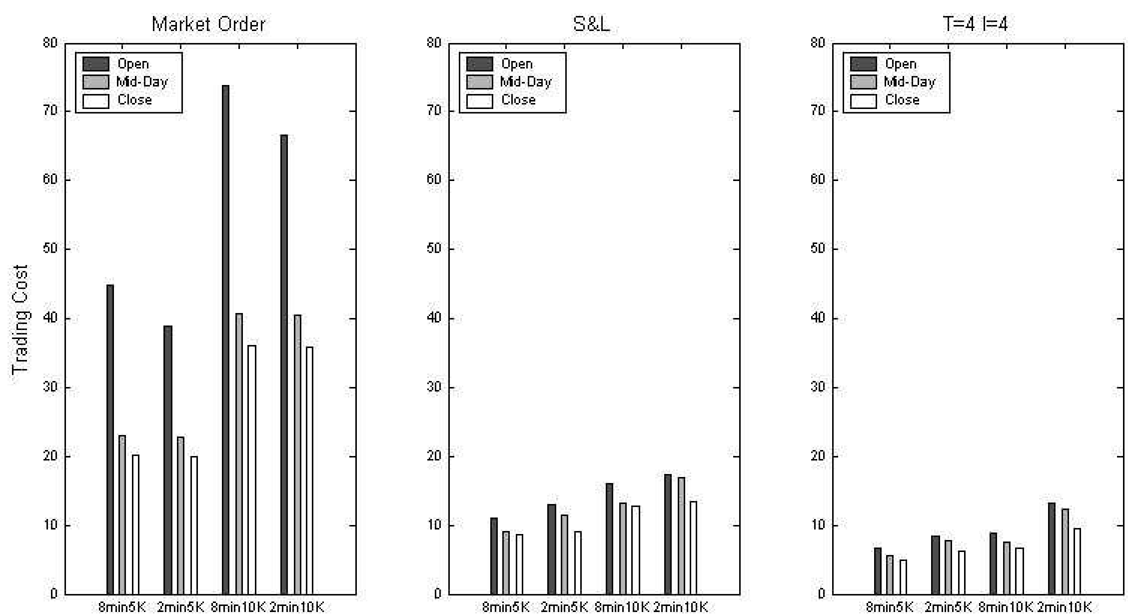

Figure 30. Cost Comparison across Strategies

While the cost ranking close-middle-open persists as we increase the number of update points from zero to one and then to four, the magnitude of difference between periods shrinks. There is an explanation for such pattern. Optimized execution strives to balance out a price impact from a market order and an uncertain price improvement from a limit order. Early in the day, there is very little (in relative terms) pre-committed liquidity on the limit order books, thus the outsized cost of submitting a market order. But at the same time, the open is associated with the highest intra-day volatility, which increases the probability of limit order being executed, and makes limit order submission more attractive. Therefore, the intra-day differences in execution costs are smaller for those strategies that employ both limit and market orders.

This brings us to the second question: should the prices of limit orders be contingent on the trading day period? Yes, they should be, as shown in Figure 31. Here we have optimized trade execution for the exact same setting (AMZN, 5K shares, 8 min, 4 order updates), but for different times of the trading day. In the middle of the day we submit limit orders inside the bid-ask spread (only positive numbers on the z-axis), while at the open we can be more patient and submit orders inside the book on our own side of the market and then count on high volatility to hit them.

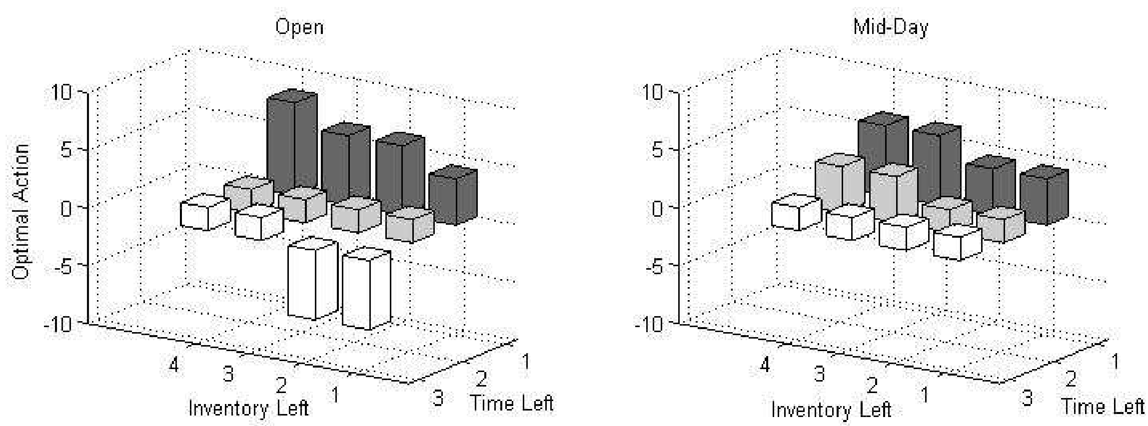

Figure 31. Different Strategies (AMZN)

Although, to be completely fair, we do not see this difference in policies across the border: when we perform the exact same experiment as above, but for NVDA instead of AMZN, optimized policies come out very similar:

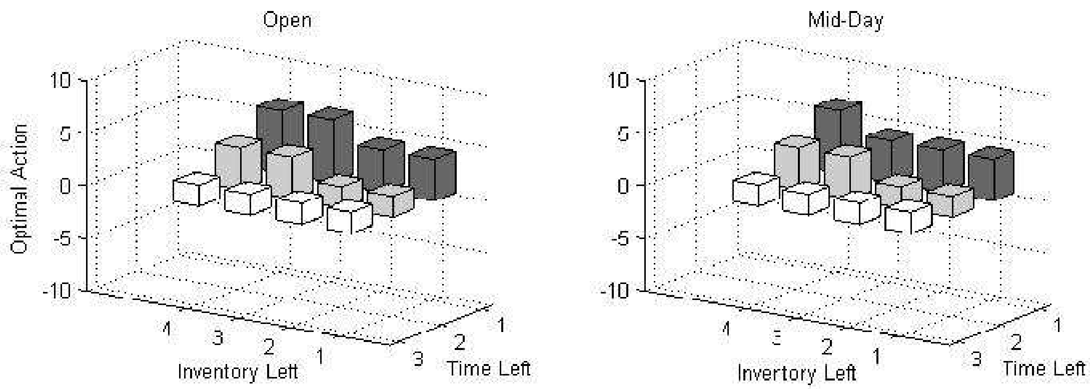

Figure 32. Similar Strategies (NVDA)

Our general recommendation is the following: if one has the data and time to optimize separately for the open, close, and mid-day, then he should do so, otherwise it is probably okay to use the same strategy during the entire day – results will not be catastrophic.

b. Optimized Execution vs. Simple VWAP

The second type of experiments we have conducted is a systematic comparison between our "learned" optimized execution strategies and a very simple version of VWAP algorithm. In the RL-based execution, we always submit limit orders for the entire volume that must be executed, and then update them during the execution period. In the VWAP setting, the order is split into equal slices (the number of slices in VWAP is equal to the number of order updates in RL), each slice sits at the top of own-side book for some pre-determined period, then it gets converted into a market order, then the same process happens to the next slice, and so on.

Table A8 shows that in general RL-based execution outperforms simple VWAP by an appreciable margin, but this difference disappears as the number of updates/slices increases. This table should be interpreted just like Table A7. If we want to buy 5,000 shares of AMZN over 2 min by following the learned strategy (submit a limit order for 5,000 shares and revise it up to 4 times depending on time period and remaining inventory), then our expected trading cost is 7.7 bp; if we want to achieve the same goal through VWAP (break the order into four 1,250 share slices, let each slice sit at the bid for 30 seconds and then convert it into the market order) then our expected trading score is 8.05 bp, which is 4% worse than RL.

Comparison between RL and VWAP is also plotted in Figures 33 and 34:

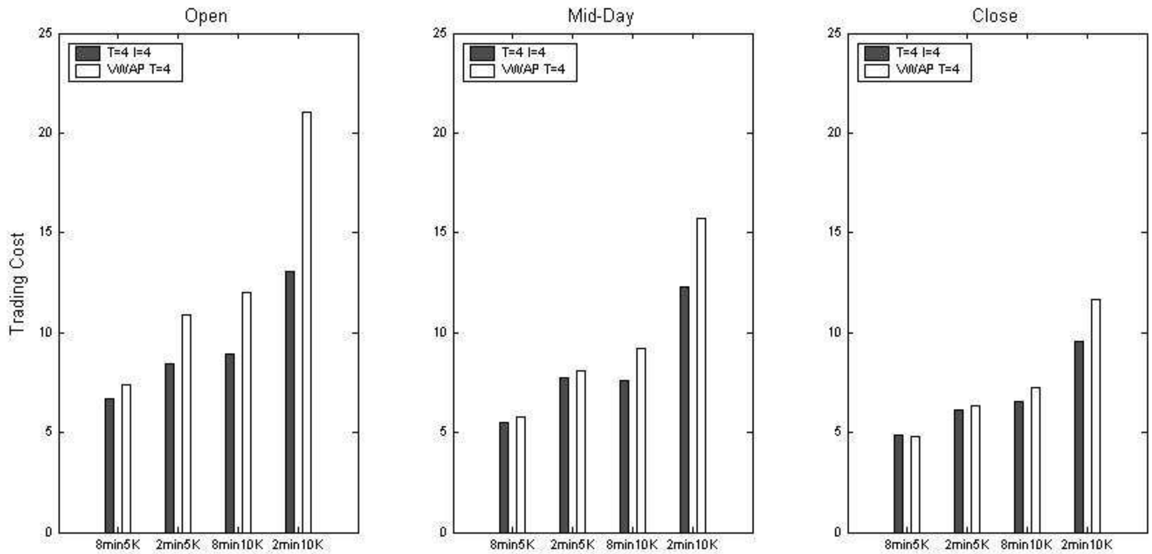

Figure 33: RL Is Better with 4 Updates/Slices (AMZN, different periods)

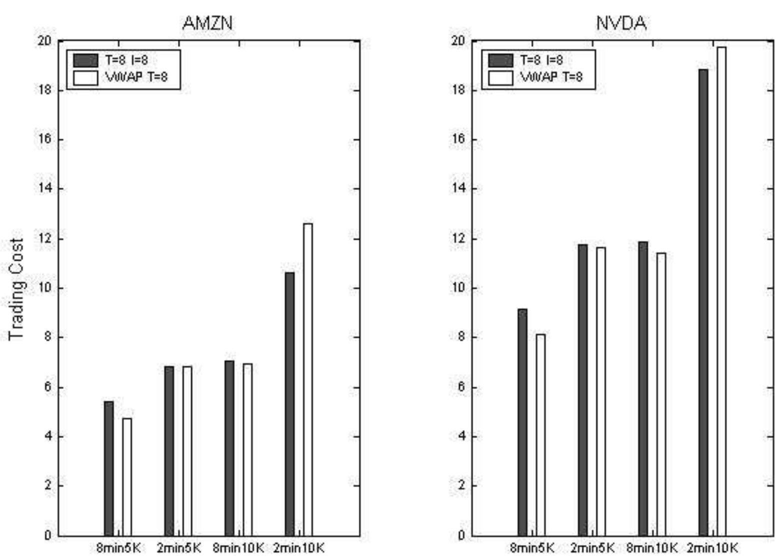

Figure 34: VWAP Can Outperform RL with 8 Updates (AMZN, NVDA, mid-day)

As stated previously, RL outperforms VWAP with 1 or 4 updates, but can fall behind with 8 (and presumably more) updates. We have to note, however, that such head-to-head comparison is not appropriate, since VWAP explicitly breaks an order into slices, while RL reasons about an order as a whole. Here is an argument why RL is a superior approach: if we tell the RL algorithm to break the order into slices, this will make it equivalent to a series of back-to-back submit-and-leave executions. And we have already shown that RL is superior to the simple VWAP in a single update setting. In VWAP's defense, it is a much simpler ad-hoc strategy, which takes zero time to train and still delivers decent results.

c. Maximizing Sharpe Ratios

As previously noted, our state-based RL optimization method needs not focus exclusively on execution cost minimization. Lower costs always come with higher risk, and so traders may be interested in either risk minimization, or some combination of cost and risk minimization. Generally speaking, if the trader's utility function is easy to specify, then our goal is to optimize the following: min(cost + R_{coef} *risk), where R_{coef} is a coefficient of risk aversion.

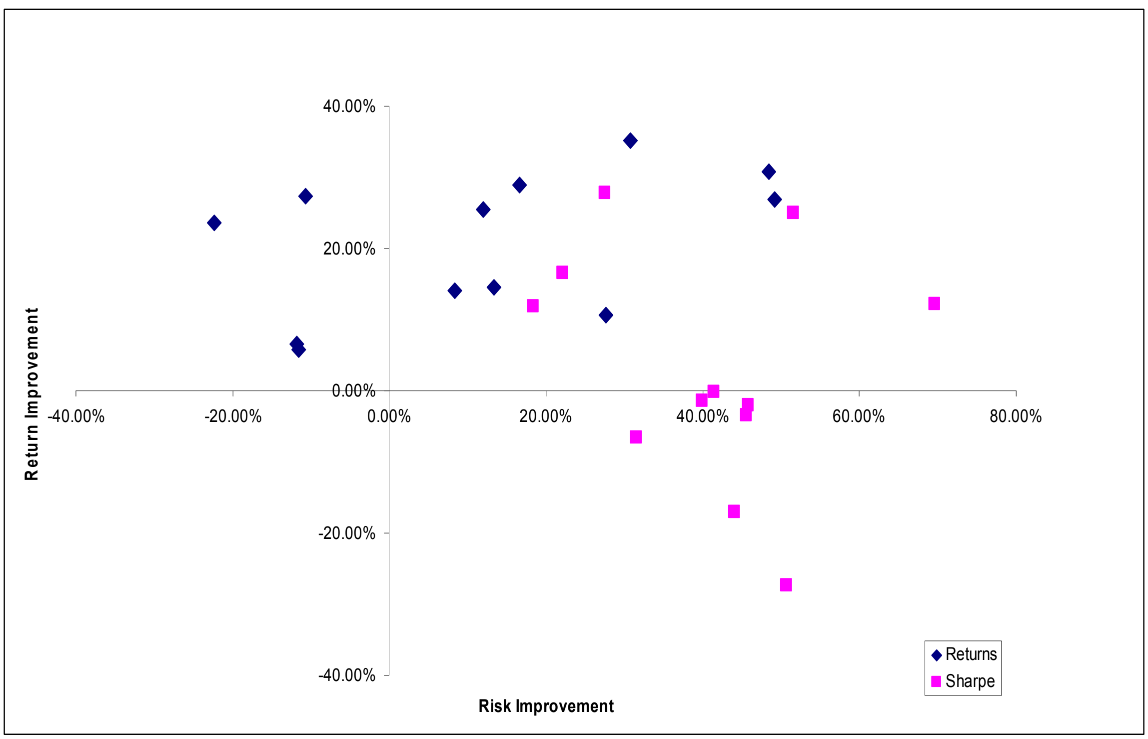

Figure 35. Sharpe Ratio Maximization: Risk-Return Trade-off

Since coefficients of risk aversion are highly-dependent on individual traders and market scenarios, we optimize instead a widely accepted measure or risk-return trade-off known as "Sharpe ratio". Sharpe ratio = Expected Return/Standard Deviation – maximizing this measure, we strive to achieve the highest return per unit of risk. Defining Sharpe ratios in our context is non-trivial: first, we convert costs into returns (so they can be maximized), and then we make all returns positive (since negative Sharpe ratios carry no economic meaning) by picking the lowest return, setting it to zero, and transforming other returns accordingly.

Our findings are summarized in Table A9. There we compare performance of a baseline strategy (submit-and-leave) and two optimized strategies with 4 decision points and 4 inventory levels. The first optimized strategy maximizes returns (minimizes costs), while the second maximizes Sharpe ratio (most return per unit of risk). For each strategy, we report its expected return and the standard deviation of this return. We compare returns and standard deviations of the two optimized strategies to those of the baseline strategy. The graphical representation of our results can be found in Figure 35. The main message here is that risk reduction comes at the expense of returns – we can see that those points that correspond to Sharpe ratio-maximizing strategies are situated lower and to the right of the return-maximizing data points.

The important question is: which parameter should be optimized in the case of efficient execution? In our opinion, it should still be returns (cost minimization). Yes, higher returns do come with higher risk, but in most cases trade execution is an activity that is performed very frequently; therefore, over long term significant negative results will be offset by significant positive results. So, when optimized execution is used with many trades over the course of the day (week, month, or year) volatility effects will become unnoticeable, while overall returns will improve. Another way of saying this is that we can assume that in the case of optimized execution traders are risk neutral and thus care about returns only.

d. Other Strategies and Multi-Strategy Learning

While it is largely out of the scope for this thesis, in this last section we will show how to expand the same analytical framework that we have used for efficient execution to other microstructure strategies, such as market making. We will then explain how to use multiple rounds of joint learning in order to find multi-strategy market equilibria.

Market makers – a.k.a. dealers – simultaneously post bid and ask quotes (submit buy and sell orders) to the market, thus providing liquidity to other market participants and hoping to benefit from the bid-ask spread. This strategy can be re-cast into the same state-based format as the efficient execution. We can set a time horizon H, within which a market maker can update (cancel and re-post) his quotes T times every H/T seconds; this will allow his inventory to fluctuate from –T to T, as measured in the number of shares Q being quoted. At the end of the episode, the market maker gets rid of any accumulated inventory through either a market order or some optimized execution mechanism.

Similarly to the efficient execution scenario, the goal of the reinforcement learning here is to determine the position of each quote within the limit order book. Unlike the execution algorithm, the market maker is responsible for two quotes, but we make life easier for ourselves by assuming that the dealer spread size is fixed. This way the RL still needs to determine only one optimal value – how to position the dealer's spread relative to the inside market's mid spread. So from the machine learning standpoint, we are largely facing the same problem as before. A sample market-making strategy that can be learned is shown in Figure 36.

Time Left

Inventory Leve	8	7	6	5	4	3	2	1
-7								8
-6							8	8
-5						8	8	8
-4					6	6	8	8
-3				6	6	6	6	8
-2			6	6	6	6	6	6
-1		6	6	6	6	6	6	6
0	-2	-2	-2	-2	-2	-2	0	0
1		-2	-2	-2	-2	-2	-2	-2
2			-2	-2	-2	-4	-4	-8
3				-4	-4	-4	-8	-8
4					-8	-8	-8	-8
5						-8	-8	-8
6							-8	-8
7								-8

Figure 36. Sample MM policy

Positive numbers in the policy mean that the dealer's mid-spread is higher than the inside market's mid-spread, negative – below. Here we can see that the market maker is posting such quotes that will induce his holdings to revert to zero: if the dealer has a long position, he will shift his quotes downwards to make his asking price relatively more attractive and thus induce public buying; when the dealer is short, he'll do the opposite and raise his bid.

How were we able to learn such intuitively compelling policy? From the computational standpoint, at the very last step (position liquidation) the dealer will be punished for carrying excessive inventory, since his high market impact will result in a negative score. This dependency (unbalanced inventory = negative score) then gets propagated to earlier time steps, thus resulting in this quote-shifting policy.

The reason why the policy from Figure 36 is not symmetric is because it was trained on a short time period (1 month) and thus is sensitive to the price trend. We also have to mention that the policy is unrealistic for other reasons – it is acting over a very short time horizon for a dealer (5 min), has very few decision points, reasons crudely about its inventory holdings, and uses no market variables. It is therefore not exactly shocking that this strategy doesn't make money. But it is still optimal in the sense that other strategies in the same class would lose even more money. In any event, we do not explicitly concern ourselves with profitability here, since this market making presentation is a proof of concept and not a full optimization work.

Inventory Leve \ *Time Left*	8	7	6	5	4	3	2	1
-7								8
-6							8	8
-5						8	8	8
-4					8	8	8	8
-3				8	8	8	8	8
-2			8	8	8	8	8	8
-1		8	6	6	6	8	8	8
0	0	0	0	0	0	0	2	0
1		0	0	0	0	0	2	0
2			-2	-2	0	-2	-4	-8
3				-2	-2	-2	-6	-8
4					-6	-6	-6	-8
5						-6	-6	-8
6							-8	-8
7								-8

Figure 37: MM policy with buying pressure

This brings us to the second point – multi-agent learning. Imagine that we have trained a market making agent to maintain balanced quotes as described above, and then we introduce another agent, which represents a large buyer, into the market. The buyer agent can follow some fixed policy (i.e. simple VWAP), or it can be trained using our RL framework. In any event, such agent will introduce buying pressure into the market, and will render the policy previously learned by the dealer less profitable (or more unprofitable, as in our case). Now, we can give the dealer an opportunity to re-learn his optimal policy on the previously-available data, but with the other agent's orders imbedded in it. Now the dealer revises his quotes upwards to account for the buying pressure, as shown in Figure 37.

We can now re-train the execution agent on the data that includes the dealer's revised quotes, then re-train the dealer once again with the new execution policy, and so on. In our extremely simplified setting, the dealer learned to hide his quotes deep enough into the book so that he doesn't have to interact with the execution agent. We would like to emphasize once again that this is simply a "toy" example, which shows how to apply our method, but a complete exploration of multi-strategy equilibria is well outside the scope of this thesis.

Research Contributions and Future Directions

Here we reflect upon the most important contributions of this thesis. We will highlight the importance of efficient execution and suggest specific applications for our method, we will discuss the insights that we have gained into the economics of financial markets and into applications of RL. We will conclude with a discussion of future directions for this research.

A. Practical Significance and Applications

Who should be interested in implementing and applying our reinforcement learning framework? Admittedly, our work is geared towards institutional investors – entities that have direct access to the market, pay little or no commissions, and trade at considerable volumes. The key insight into this problem is that we are aiming to save fractions of pennies on execution. There are essentially 2 ways to capitalize on this:

(1) if you are trading billions, this will save you hundred thousands of dollars (even if you are paying other fees)

(2) if you have direct market access, these savings are "free money" that you are giving away to other market participants (even if you are trading on a smaller scale)

Large mutual funds, pension funds, and some hedge funds fall in the first category. They should incorporate our analysis of trading costs into their high level investment/portfolio optimization cycle to have a better estimate of how turnover and execution speed affect their overall performance. This method will also help these institutions to issue more precise execution instructions to their traders or brokers – i.e. execute over 2 days, avoid the open, be more aggressive in stock A, and so on.

Investment banks, trading firms, and execution boutiques fall into the second category. Their mandate is to obtain the best price in the market place for their customers of for a proprietary portfolio. And this is what this research is all about – they can apply our RL framework to past data to come up with specific trading strategies that deliver the best execution.

All that being said, implications of our research extend well beyond the institutional investor community. Almost everyone has some funds in 401(k) plans, other retirement plans, mutual funds, etc. Therefore, individual investors should be aware of these market microstructure dynamics and make sure that whoever manages their money does everything possible to reduce these hidden costs. Otherwise, if investment manager does not have proper incentives to address this problem, these costs will be just passed on to the investors in a form of lower returns.

The second important implication is the overall market efficiency. Aggregate market behavior is difficult to predict, so we can only speculate here. If more and more agents will try to act "optimally" in the market place, this will alter the price formation process, which will require different "optimal" strategies, which will influence the marketplace again, and so on. This is clearly a dynamic on-going process of market efficiency evolution. Our guess is that if everyone were to embark on a pursuit of optimal market behavior, this would drive spreads closer to zero, make prices even closer to a random walk, will make information dissemination even faster… In other words, our financial markets will approach the frictionless ideal of the perfectly efficient markets.

Moving from market participants to actual systems and strategies: what specific areas can benefit from this RL analysis of execution costs? Once again, we believe that trading cost accounting must be pervasive: it should be included in portfolio selection, securities pricing, risk management, strategy development, and trade execution proper. In fact, it is difficult to think of a single important area of Finance where market microstructure can be completely disregarded. It all stems from the fact that market mechanisms and overall market quality determine actual transaction prices and thus influence ultimate levels of wealth. In portfolio selection, correlations and other dependencies are determined from past transaction prices, which are themselves dependent upon trade sizes, times and other execution parameters. Therefore, considerations such as spreads and liquidity must be taken into account. Similarly, derivatives pricing is most frequently performed by estimating the costs of a hedging strategy used to replicate the payoffs of the derivative. This hedging cost must be derived from prices at which trades can take place, so once again, issues like market impact must be taken into account. We have already suggested how RL can be used to improve the VWAP execution strategy, and how it can be used for market making.

B. Academic Contributions and Innovations

The most far-reaching implication of this work is the automation of microstructure-based trading activities. We have conveniently placed investment decision making outside of our model, thus allowing us to concentrate fully on trade implementation. This is an environment where robots (electronic agents) are much better positioned to both make decisions and take actions than human traders, because of the overwhelming amount of information and the extremely high speed with which it gets updated. This research direction contributes to further automation of market activities and general market efficiency.

From the economic point of view, we offer a way to formalize some intuitions about how markets work. For example, it is a generally accepted idea that when spreads are large, one should submit a limit order inside the spread. In our framework, we can create a corresponding market variable, confirm or disprove the conventional wisdom based on historical data, and even tell exactly where inside the spread should the order be placed. We emphasize the systematic nature of our approach – we strive to eliminate ad-hoc decisions from all levels of market activities. This means that our optimized strategies deliver the best expected performance that can be learned from historical data.

Another contribution to the Market Microstructure research community is our empirical study of various microstructure factors. We showed which variables can help improve strategies' performance, and which cannot, and also explained why. We broke the mold by taking our investigation one step further – traditional econometric studies simply verify if one variable influence some other variable or variables. We have also explored if we can take advantage of such relationship – i.e. should our optimal action change as a function of a given variable. Our "surprise find" was the fact that predictability does not necessarily equates with a change of strategy based on that variable: while volume misbalance is a good predictor of future prices, we were not able to take advantage of this information.

Finally, we have suggested how our method can be applied to many other scenarios, which has implications for individual market participants and financial markets as a whole.

As for the AI contribution, we specify a quantitative framework that allows algorithms to reason about very complex domains by reducing the problem's dimensionality. We show how to convert intractably complicated real world information into a lower-dimensional state space. We then define policies over such space and search for the optimal one.

We have developed a novel reinforcement learning algorithm specifically tailored to market microstructure optimization. It combines the best features of dynamic programming and Q-learning. It derives an optimal policy (given a set of assumptions) while visiting the minimal number of states and efficiently re-using available data. We have applied this algorithm to very large data sets of high-frequency data. It is important to remember that successful large-scale applications of machine learning are relatively rare, and therefore we interpret the fact that our approach yields immediately-applicable results as a major accomplishment.

Among less significant contributions is our simulator, which allows reconstructing full limit order books from past data and then integrating artificial orders into the market in the same way it is done in the real world. And we have also carefully constructed a theoretical foundation for the automation of microstructure activities.

C. Future Directions

As we have mentioned throughout this document, there are many directions in which this research can progress. The most immediate need, in our opinion, is the automated feature selection and internalization of other parameters. Right now, specification of trade size, time period, number of updates, state variables and their resolutions all come from outside the model. A superior execution system should be able to determine many of these inputs by itself. Specifically, we can envision a neural network, which takes an entire order book (or several order books) and then determines which features are important in a given context.

Similarly, right now we submit limit orders for all the remaining shares and adjust expected execution volume through price (more aggressive order = more volume). This is very likely suboptimal and unrealistic, especially for large orders. Therefore, order size can be another decision variable for RL alongside with the order price.

This execution framework should also be extended to multiple stocks and general portfolios – by acknowledging that a direct price impact in one stock is likely to influence the price of other stocks, we can trade multiple stocks together with a greater efficiency than if we were to treat them separately.

Pure execution aside, we have also demonstrated how to apply this flexible state-based framework to other trading problems and to multi-strategy equilibrium research. Since our foray in this direction was very cursory, a lot still remains to be done. Market making can be improved from both the theoretical standpoint (source of profits, significance of different parameters) and the implementation standpoint (proper bellman's equation, all input parameters, market variables, etc). We can suggest other strategies that can be amenable to the same analysis: one-sided trading, statistical arbitrage, automated hedging, just to name a few.

We also believe that strategies equilibrium research has a great potential: in a market with many heterogeneous actors, we can perform multiple rounds of learning for each strategy, and see if equilibrium can be achieved. Studying combinations of strategies also permits us to evaluate relative strengths and weaknesses of individual strategies and strategy classes. Ultimately, this type of research maybe able to resolve one of the open questions mentioned earlier: what would happen to financial markets if all market participants had embarked on a pursuit of acting optimally. In general, multi-agent approach can give us a better insight into the process of price formation in securities markets.

To conclude, we believe that in addition to immediate contribution, our research has great potential for future advances in both empirical microstructure studies and in market equilibrium research.

Bibliography

1. INET ECN. www.inetats.com.
2. New York Stock Exchange. www.nyse.com.
3. Penn-Lehman Automated Trading Project (PLAT). www.cis.upenn.edu/~mkearns/projects/plat.html.
4. Penn Exchange Simulator (PXS). www.cis.upenn.edu/~mkearns/projects/plat-api.html.
5. *Best Execution: Executing Transactions in Securities Markets on Behalf of Investors*, 2002, Collection of essays, European Asset Management Association.
6. Ahn, H., Bae, K., and Chan, K., 2001, *Limit Orders, Depth, and Volatility: Evidence from the Stock Exchange of Hong Kong*, Journal of Finance 55, 767 – 788.
7. Almgren, R. and Chriss, N., 1999, *Value Under Liquidation*, Risk.
8. Almgren, R. and Chriss, N., 2000, *Optimal Execution of Portfolio Transactions*, Journal of Risk.
9. Amihud, Y. and Mendelson, H., 1980, *Dealership Market: Market Making with Inventory*, Journal of Finance 42, 533-553.
10. Back, K. and Baruch, S., 2004, *Limit-Order Markets and Uniform-Price Auctions: An Irrelevance Proposition*, working paper, University of Utah
11. Barclay M., Christie W., Harris, J., Kandel, E., and Schultz, P., *The Effects of Market Reform on the Trading Costs and Depth of NASDAQ Stocks*, Journal of Finance 54, 1999.
12. Barclay, M., Hendershott, T., and McCormick, T., 2002, *Information and trading on Electronic Communication Networks*, working paper, University of Rochester.
13. Becker, B., Lopez, E., Berberi-Doumer, V., Cohn, R. and Adkins, A., 1992, *Automated Securities Trading*, Journal of Financial Services Research, 6, 327-341.
14. Bertsimas, D. and Lo, A., 1998, *Optimal Control of Execution Costs*, Journal of Financial Markets 1, 1-50.

15. Bertsimas, D., Hummel, P. and A. Lo, 2000, *Optimal Control of Execution Costs for Portfolios*, Computing in Science & Engineering 1, 40-53.
16. Bias, B., Hilton, P., and Spatt, C., 1995, *An Empirical Analysis of the limit Order book and the Order Flow in the Paris Bourse*, Journal of Finance, 50, 1655-1689.
17. Bias, B., Glosten, L., and Spatt, C., 2004, *Market Microstructure: a Survey of Microfoundations, Empirical Results and Policy Implications*, Journal of Financial Markets.
18. Black, F., 1971, *Toward a Fully Automated Exchange*, Financial Analysts Journal, Nov.-Dec. 1971.
19. Bloomfield, R., O'Hara, M., and Saar, G., 2003, *The "Make or Take" Decision in an Electronic Market: Evidence on the Evolution of Liquidity*, working paper, Stern School of Business, New York University.
20. Boehmer, E., G. Saar, and L. Yu, 2004, *Lifting the Veil: An Analysis of Pre-Trade Transparency at the NYSE*, Journal of Finance, forthcoming.
21. Bouchaud, J.-P., Mezard, M., and Potters, M., 2002, *Statistical Properties of Stock Order Books: Empirical Results and Models*, Quantitative Finance 2, 251 – 256.
22. Campbel, J. Y., Lo, A. W., and MacKinlay, A. C., 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
23. Chan, N. and Shelton, C., 2001, *An Electronic Market-Maker*, AI Memo 2001, MIT.
24. Chau, M., 2002, *Dynamic Trading and Market-Making with Inventory Costs and Private Information*, working paper, ESSEC.
25. Choi, J.Y., Salandro, D., and Shastri, K., 1988, *On the Estimation of Bid-Ask Spreads: Theory and Evidence*, Journal of Financial and Quantitative Analysis 23, 219-230.
26. Chordia, T. and Subrahmanyam, 1995, *Market Making, the Tick Size, and Payment-for-Order Flow: Theory and Evidence*, Journal of Business 68, 543-575.
27. Chordia, T., Roll, R., and Subrahmanyam, A., 2002, *Order Imbalance, Liquidity, and Market Returns*, Journal of Financial Economics.

28. Christie, W. and Schultz, P., 1994, *Why do NASDAQ market makers avoid odd-eighth quotes?* Journal of Finance 49, 1813-1840.
29. Chung, K., Van Ness, B., Van Ness, R., 1999, *Limit Orders and the Bid-Ask Spread*, Journal of Financial Economics 53, 255-287.
30. Cohen, K., Maier, S., Schwartz, D., Whitcomb, D., 1979, *Market Makers and the Market Spread: A Review of Recent Literature*, Journal of Financial and Quantitative Analysis 14, 813-835.
31. Cohen, K., Maier, S., Schwartz, D., Whitcomb, D., 1981, *Transaction Costs, Order Placement Strategy, and Existence of the Bid-Ask Spread*, Journal of Political Economy 89, 287-305.
32. Copeland, T. and Galai, D., 1983, *Information Effects and the Bid-Ask Spread*, Journal of Finance 38, 1457-1469.
33. Coppejans, M. and Domowitz, I., 1999, *Screen Information, Trader Activity, and Bid-Ask Spreads in A Limit Order Market*, working paper, Duke University.
34. Dann, L., Mayers, D., and Raab, R., 1977, *Trading Rules, Large Blocks and the Speed of Price Adjustments*, Journal of Financial Economics, January 1977.
35. Das, S., 2003, *Intelligent Market-Making in Artificial Financial Markets*, AI Technical Report 2003-005, MIT.
36. Degryse, H., de Jong, F., van Ravenswaaij, M., Wuyts, G., 2001, *Aggressive Orders and the Resiliency of a Limit Order Market*, Discussion paper, Tilburg University.
37. Delwiche, L. and Slaughter, S., 2003, *The Little SAS Book: A Primer*, SAS Publishing.
38. Dennert, J., 1993, *Price Competition between Market Makers*, Review of Economic Studies 60, 735-751.
39. Domowitz, I., 2001, *Liquidity, Transaction Costs, and Reintermediation in Electronic Markets*, Journal of Financial Services Research 22.
40. Dubil, R., 2002, *Optimal Liquidation of Large Security Holdings in Thin Markets*, Research in International Entrepreneurial Finance and Business Ventures at the Turn of the Third Millennium, Academy of Entrepreneurial Finance, special proceedings issue.

41. Dutta, P. and Madhavan, A., 1997, *Competition and Collision in Dealer Markets*, Journal of Finance 52, 245-276.
42. Easley, D. and O'Hara, M., 1987, *Price, Trade Size, and Information in Securities Markets*, Journal of Financial Economics 19, 69-90.
43. Easley, D. and O'Hara, M., 1992, *Time and the Process of Security Price Adjustment*, Journal of Finance 47, 577-605.
44. Ellul, A., C. Holden, P. Jain, and R. Jennings, 2003, *Determinants of Order Choice on the New York Stock Exchange*, working paper, Indiana University.
45. Fama, E., 1970, *Efficient Capital Markets: A Review of Theory and Empirical Work*, Journal of Finance, 1970.
46. Fama, E., Fisher, L., Jensen, M., and Roll, R., 1969, *The Adjustment of Stock Prices to New Information*, International Economic Review, February 1969.
47. Fama, E., and French, K., 1992, *The Cross-Section of Expected Stock Returns*, Journal of Finance, June 1992.
48. Foucault, T., Kadan, O., and Kandel, E., 2003a, *Limit Order Book as a Market for Liquidity*, CERP working paper.
49. Foucault, T., Roell, A., and Sandas, P., 2003b, *Market Making with Costly Monitoring: An Analysis of the SOES Controversy*, Review of Financial Studies 16, 345-384.
50. Glosten, L., 1987, *Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices*, Journal of Finance 42, 1293-1307.
51. Glosten, L., 1994, *Is the Electronic Open Limit Order Book Inevitable?*, Journal of Finance 49, 1127 – 1167.
52. Glosten, L. and Milgrom, P., 1985, *Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders*, Journal of Financial Economics 13, 71 – 100.
53. Goettler, R., Parlour, C., Rajan, U., 2003, *Equilibrium in a Dynamic Limit Order Market*, working paper, GSIA, Carnegie Mellon University.
54. Goettler, R., Parlour, C., Rajan, U., 2004, *Information Acquisition in a Limit Order Market*, GSIA Working Papers, Carnegie Mellon University, Tepper School of Business.

55. Golstein, M. and Kavajecz, K., 2000, *Eighths, sixteenths and market depth: Changes in tick size and liquidity provision on the NYSE*, Journal of Financial Economics, 56, 125-149.
56. Griffiths, M. D., Smith, B. F., Turnbull, D. A. S., White, R. W., 2000. *The Costs and Determinants of Order Aggressiveness*, Journal of Financial Economics 56, 65-88.
57. Hakansson, N., Beja, A., Kale, J., 1985, *On the Feasibility of Automated Market Making by a Programmed Specialist*, Journal of Finance 40, 1-20.
58. Handa, P. and Schwartz, R., 1996, *Limit Order Trading*, Journal of Finance 51, 1835-1861.
59. Handa, P., Schwartz, R., and Tiwari, A., 2002, *Quote Setting and Price Formation in an Order Driven Market*, Journal of Financial Markets.
60. Harris, L., 1996, *The Economics of Best Execution*, working paper, 1996.
61. Harris, L., 1997, *Optimal order Submission Strategies In Some Stylized Trading Problems*, Financial Markets, Institutions, and Instruments 7.
62. Harris, L., 2003, *Trading & Exchanges: Market Microstructure for Practitioners*, Oxford University Press.
63. Hasbrouck, J., 1988, *Trades, Quotes, Inventories, and Information*, Journal of Financial Economics 22, 229-252.
64. Hasbrouck, J., 1991a, *Measuring the Information Content of Stock Trades*, The Journal of Finance 46, 179-207.
65. Hasbrouck, J., 1991b, *The Summary Informativeness of Stock Trades: An Econometric Analysis*, Review of Financial Studies 4, 571-595.
66. Hasbrouck, J., 2004, *Empirical Market Microstructure*, Teaching notes, New York University.
67. Hasbrouck, J., and Harris, L., 1996, *Market vs. Limit Orders: The SuperDot Evidence on Order Submission Strategy*, Journal of Financial and Quantitative Analysis, 31.
68. Hasbrouck, J. and Saar, G., 2002, *Limit Orders and Volatility in a Hybrid Market: The Island ECN*, working paper, Stern School, NYU.

69. Hasbrouck, J. and Sofianos, G., 1993, *The Trades of Market Makers: An Empirical Analysis of NYSE Specialists*, Journal of Finance 48, 1565-1593.
70. Ho, T. and Stoll, H., 1981, *Optimal Dealer Pricing Under Transactions and Return Uncertainty*, Journal of Financial Economics 9, 47-73.
71. Ho, T. and Stoll, H., 1983, *The Dynamics of Dealer Markets Under Competition*, Journal of Finance 38, 1053-1074.
72. Hollifield, B., Miller, R., Sandas, P., 2003a, *Empirical Analysis of Limit Order Markets*, GSIA working paper, Carnegie Mellon University.
73. Hollifield, B., Miller, R., Sandas, P., Slive, J., 2003b, *Liquidity Supply and Demand in Limit Order Markets*, CEPR discussion paper.
74. Huang, R. and Stoll, H., 1994, *Market Microstructure and Stock Return Predictions*, Review of Financial Studies 7, 179-213.
75. Huang, R. and Stoll, H., 1996, *Dealer Versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE*, Journal of Financial Economics 41, 313-357.
76. Huang, R. and Stoll, H., 1997, *The Components of the Bid-Ask Spread: A General Approach*, Review of Financial Studies 10, 995-1034.
77. Huang, Pu, Scheller-Wolf, Alan, and Sycara, Katia, 2002, *Design of a Multi-unit Double Auction E-market*, Computational Intelligence 18.
78. Huberman, G. and Stanzl, W., 2001, *Optimal Liquidity Trading*, Yale School of Management Working Paper.
79. Ip, Greg and Craig, Suzanne, *NYSE's 'Specialist' Probe Puts Precious Asset at Risk: Trust*, The Wall Street Journal, April 18, 2003.
80. Irvine, P., Benston, G., Kandel, E., 2000, *Liquidity Beyond the Inside Spread: Measuring and Using Information in the Limit Order Book*, working paper, September 2000.
81. Kakade, S., Kearns, M., Mansour, Y., Ortiz, L., 2003, *Competitive Algorithms for VWAP and Limit Order Trading*, working paper, University of Pennsylvania.
82. Kaniel, R., and Liu, H., 2004, *So What Orders Do Informed Traders Use?*, working paper. University of Texas at Austin.

83. Kavajecz, K., 1998, *A Specialist's Quoted Depth as a Strategic Choice Variable*, working paper, University of Pennsylvania.
84. Kavajecz, K. and Odders-White, E., 2002, *Technical Analysis and Liquidity Provision*, Review of Financial Studies (forthcoming).
85. Kearns, M., and Ortiz, L., 2003, *The Penn-Lehman Automated Trading Project*, IEEE Intelligent Systems, 2003.
86. Kim, A. J. and Shelton, C. R., 2002, *Modeling Stock Order Flows and Learning Market-Making from Data*, AI Memo 2002-009, MIT.
87. Kissell, R., and Glantz, 2003, M., *Optimal Trading Strategies: Quantitative Approaches For Managing Market Impact and Trading Risk*, AMACOM.
88. Klusch, M. and Sycara, K., 2001, *Brokering and Matchmaking for Coordination of Agent Societies: A Survey*. In Coordination of Internet Agents, A. Omicini et al. (eds.), Springer.
89. Kraus, A. and Stoll, H, 1972, *Price Impact of Block Trading on the New York Stock Exchange*, Journal of Finance 27, 569-588.
90. Krokhmal, P. and Urysaev, S., 2003, *A Sample-Path Approach to Optimal Position Liquidation*, Research Report #2003-5, University of Florida, Gainsville.
91. Lintner, J., 1965, *Security Prices and Maximal Gains from Diversification*, Journal of Finance, December 1965.
92. Lo, A., MacKinlay, C., 2001, *A Non-Random Walk Down Wall-Street*, Princeton University Press.
93. Lo, A., MacKinlay, C., Zhang, J., 2000, *Econometric Models of Limit-Order Executions*, National Bureau of Economic Research working paper.
94. Lowenstein, R., 2000, *When Genius Failed: The Rise and Fall of Long-Term Capital Management*, Random House.
95. Madhavan, A., 1992, *Trading Mechanisms in Securities Markets*, Journal of Finance 47, 607-641.
96. Madhavan, A., 2000, *Market Microstructure: A Survey*, Journal of Financial Markets, 205-258.
97. Madhavan, A., 2002, *Market Microstructure: A Practitioner's Guide*, Financial Analysts Journal, vol. 58, no. 5, 28-42.

98. Madhavan, A., 2002b, *VWAP Strategies*, Transaction Performance, Spring '02.

99. Madhavan, A. and Smidt, S., 1991, *A Bayesian Model of Intraday Specialist Pricing*, Journal of Financial Economics 30, 99-134.

100. Madhavan, A. and Smidt, S., 1993, *An Analysis of Changes in Specialist Inventories and Quotations*, Journal of Finance 48, 1595-1628.

101. Malkiel, B., 1999, *A Random Walk Down Wall Street*, seventh edition, W. W. Norton and Company.

102. Markowitz, H., 1952, *Portfolio Selection*, Journal of Finance, 7 (1), 77-91.

103. Merton, R. C., 1990, *Continuous-Time Finance*, Blackwell Publishing.

104. Mitchell, T., 1997, *Machine Learning*, WCB/McGraw-Hill.

105. Modigliani, F. and Cohn, R., 1979, *Inflation, Rational Valuation and the Market*, Financial Analysts Journal 35, 24-44.

106. Mossin, J., 1966, *Equilibrium in a Capital Asset Market*, Econometrics, October 1966.

107. Nevmyvaka, Y., Kearns, M., Papandreou, A., and Sycara, K., 2004, *Optimized Execution Via Order Book Simulation*, working paper, CMU.

108. Nevmyvaka, Y., Kearns, M., Papandreou, A., and Sycara, K., 2005, *Electronic Trading in Order-Driven Markets: Efficient Execution*, in the Proceedings of IEEE International Conference on E-Commerce Technology.

109. Nevmyvaka, Y. and Sycara, K., 2005, *Market Microstructure: Time Series Analysis*, in the Proceedings of the Fifth International Conference on Computational Intelligence in Economics and Finance.

110. Nevmyvaka, Y., Sycara, K., and Seppi, D., 2003, *Electronic Market Making: Initial Investigation*, In the Proceedings of Third International Workshop on Computational Intelligence in Economics and Finance, 2003.

111. O'Hara, M., 1995, *Market Microstructure Theory*, Blackwell Publishers.

112. O'Hara, M. and Oldfield, G., 1986, *The Microeconomics of Market Making*, Journal of Financial and Quantitative Analysis 21, 361-376.

113. Ozenbas, D., Schwartz, R., and Wood, R., *Volatility in U.S. and European Equity Markets: An Assessment of Market Quality*, in International Finance 5, no. 3, 437-461.

114. Papageorgiou, C., 1997, *High Frequency Time Series Analysis and Prediction Using Markov Models*, in CIFEr '97.

115. Parlour, C., 1998, *Price Dynamics in Limit Order Markets*, Review of Financial Studies 11, 789-816.

116. Parlour, C. and Rajan, U., 2002, *Payment for Order Flow*, Journal of Financial Economics 68, 379-411.

117. Parlour, C. and Seppi, D. J., 2003, *Liquidity-Based Competition for Order Flow*, Review of Financial Studies 16, 301-343.

118. Peterson, M., and E. Sirri, 2002, *Order Submission Strategy and the Curious Case of Marketable Limit Orders*, Journal of Financial and Quantitative Analysis.

119. Poggio, T., Lo, A., LeBaron, B., and Chan, N., 1999, *Agent-Based Models of Financial Markets: A Comparison with Experimental Markets*, MIT Artificial Markets Projects, Paper No. 124.

120. Roll, R., 1984, *A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market*, Journal of Finance 39, 1127-1139.

121. Ranaldo, A., 2004, *Order Aggressiveness in Limit Order Book Markets*, Journal of Financial Markets 7, 53-74.

122. Rosu, I., 2004, *A Dynamic Model of the Limit Order Book*, Working paper, University of Chicago.

123. Rust, J., Miller, J., and Palmer, R., 1994, *Characterizing effective trading strategies: Insights from a computerized double auction tournament*. Journal of Economic Dynamics and Control, 18:61-96.

124. Sandas, P., 2001, *Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market*, Review of Financial Studies 14, 705 – 734.

125. Schiler, R., 1981, *Do Stock Prices Move Too Much to Be Justified by Subsequent Changes in Dividends?*, in American Economic Review 71, June 1981, 421-436.

126. Schwartz, R., Byrne, J., and Colaninno, A., 2004, *A trading Desk's View of Market Quality*, Kluwer Academic Publishers.

127. Schwartz, R., Byrne, J., Aidan, J., and Colaninno, A., 2005, *Coping with Institutional Order Flow*, Springer.

128. Schwartz, R., Francioni, R., 2004, *Equity markets in Action: The Fundamentals of Liquidity, Market Structure and Trading*, John Wiley and Sons.

129. Seppi, D. J., 1997, *Liquidity Provision with Limit Orders and Strategic Specialist*, Review of Financial Studies 10, 103-150.

130. Sharpe, W., 1964, *Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk*, Journal of Finance, September.

131. Smith, Adam, 1776, *The Wealth of Nations*, Collier, New York.

132. Smith, J., 2001, *Market vs. Limit Order Submission Behavior at a NASDAQ Market Maker*, NASDAQ Economic Research.

133. Steward and Russell, 1998, *Artificial Intelligence: The Modern Approach*, Weily.

134. Stoll, H., 1989, *Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests,* Journal of Finance 44, 115-134.

135. Stoll, H., 1999, *Alternative Views of Market Making*, in Market Making and Changing Structure of Securities Industries, by Amihud, J., Ho, T., and Schwartz, R.

136. Stoll, H., 2000, *Friction*, Journal of Finance 55, 1479-1514.

137. Stoll, H., 2001, *Market Microstructure*, Vanderbilt University working paper, for Handbook of the Economics of Finance.

138. Stoll, H. and Whaley, R., 1990, *Stock Market Structure and Volatility*, Review of Financial Studies 3, 37-71.

139. Thomas, J. D., 2003, *News and Trading Rules*, Ph. D. Thesis, SCS, Carnegie Mellon University.

140. Vayanos, D., 2001, *Strategic Trading in a Dynamic Noisy Market*, Journal of Finance 56, 131-171.

141. Wang, X., 2002, *On an Open Limit order Book – How Can We Use It or … Can We?*, Working Paper, Penn State University.

142. Weber, P. and Rosenow, B., 2003, Order book approach to price impact, working paper.

143. Wellman, M., Greenwald, A., Stone, P., and Wurman, P., 2003, *The 2001 Trading Agent Competition*, Electronic Markets 13.

144. Yafee, R., and McGee, M., 2000, *Introduction to Time Series Analysis and Forecasting*, Academic Press.

145. Zlot, R.M., Stentz, A., Dias, M.B., and Thayer, S., 2002, *Multi-Robot Exploration Controlled by a Market Economy*, IEEE International Conference on Robotics and Automation.

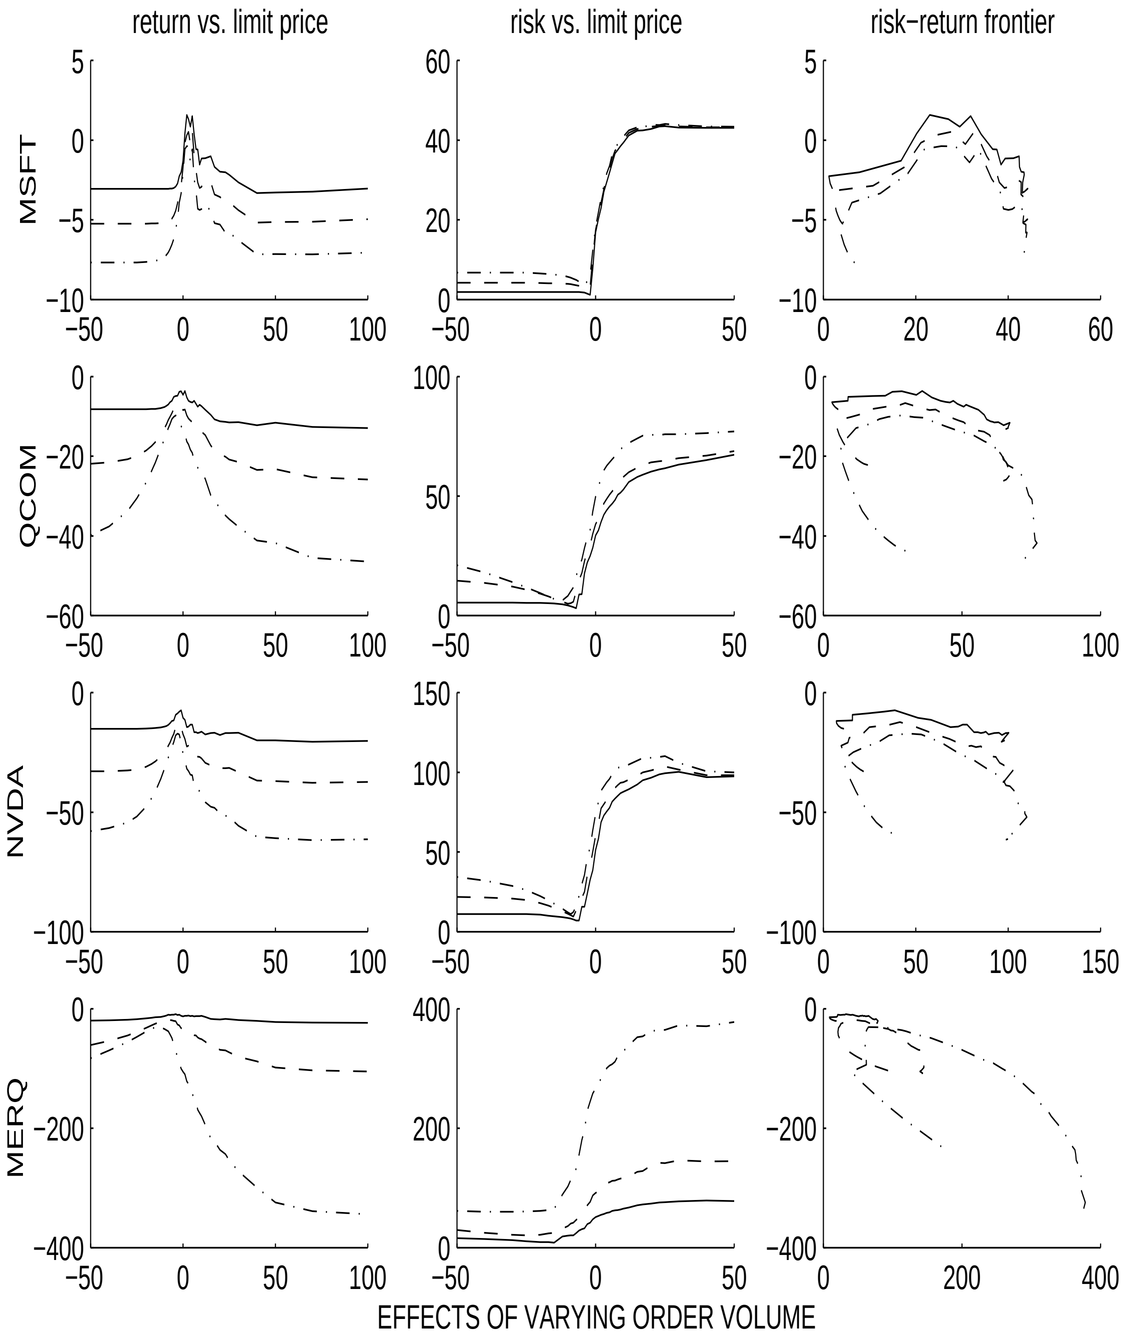

Figure A1. Transacting larger orders is both more expensive and riskier.

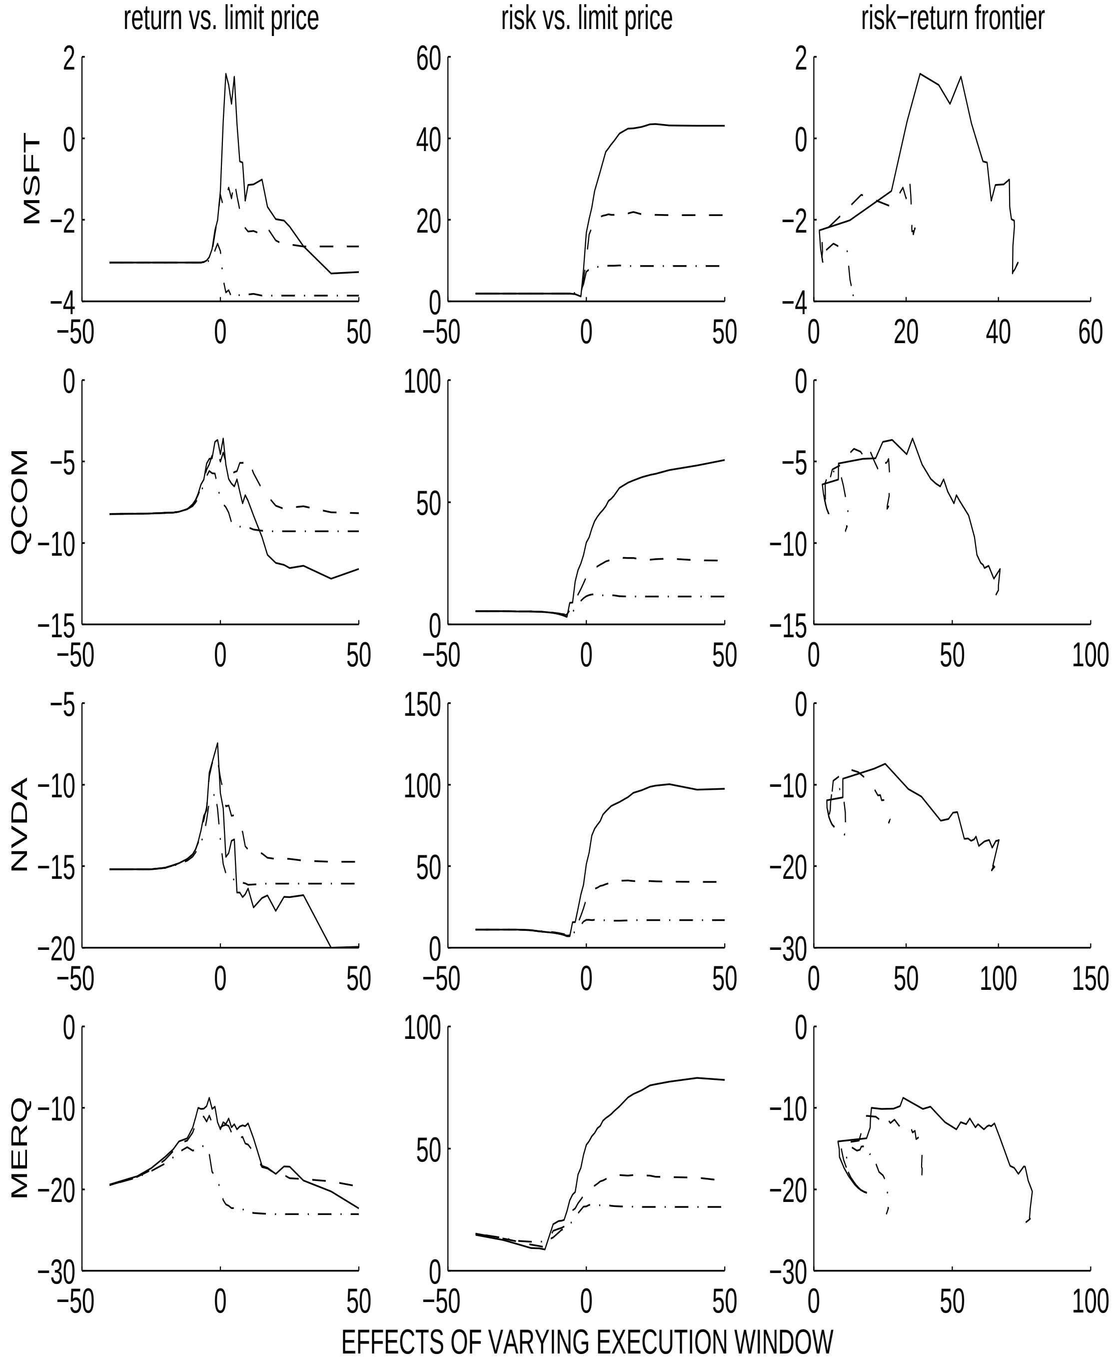

Figure A2. Execution over shorter time periods is more expensive, but less risky. Efficient pricing frontiers intersect, making optimal strategy dependent on the trader's risk preference.

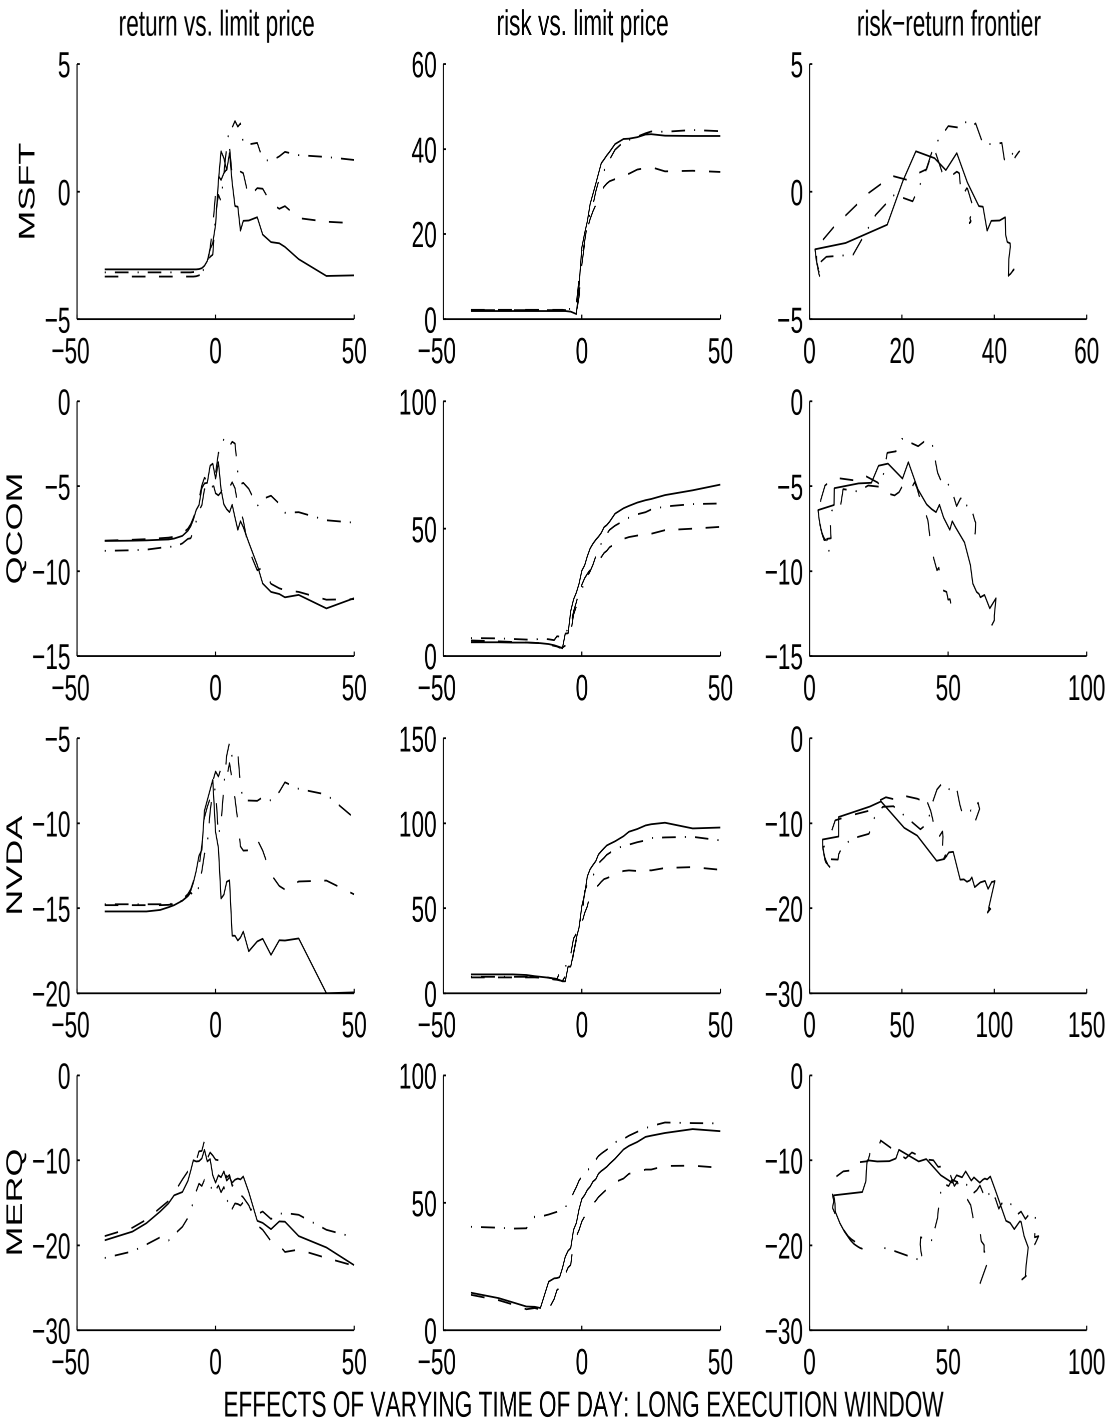

Figure A3. When executing over 60 minute period, time of the day should be used as one of the model's inputs.

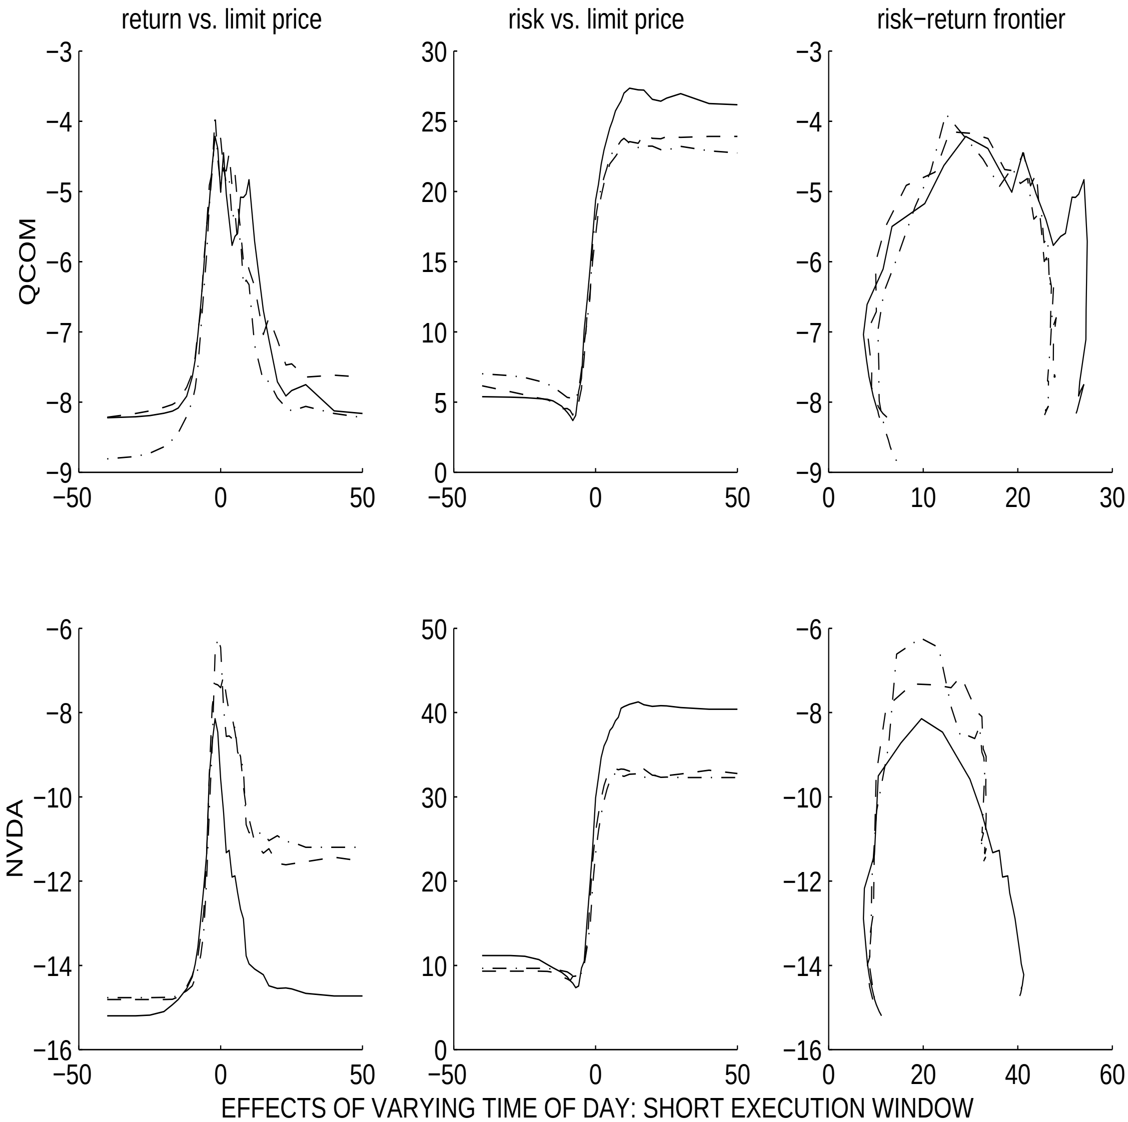

Figure A4. When executing over 10 minute period, time of the day can be disregarded.

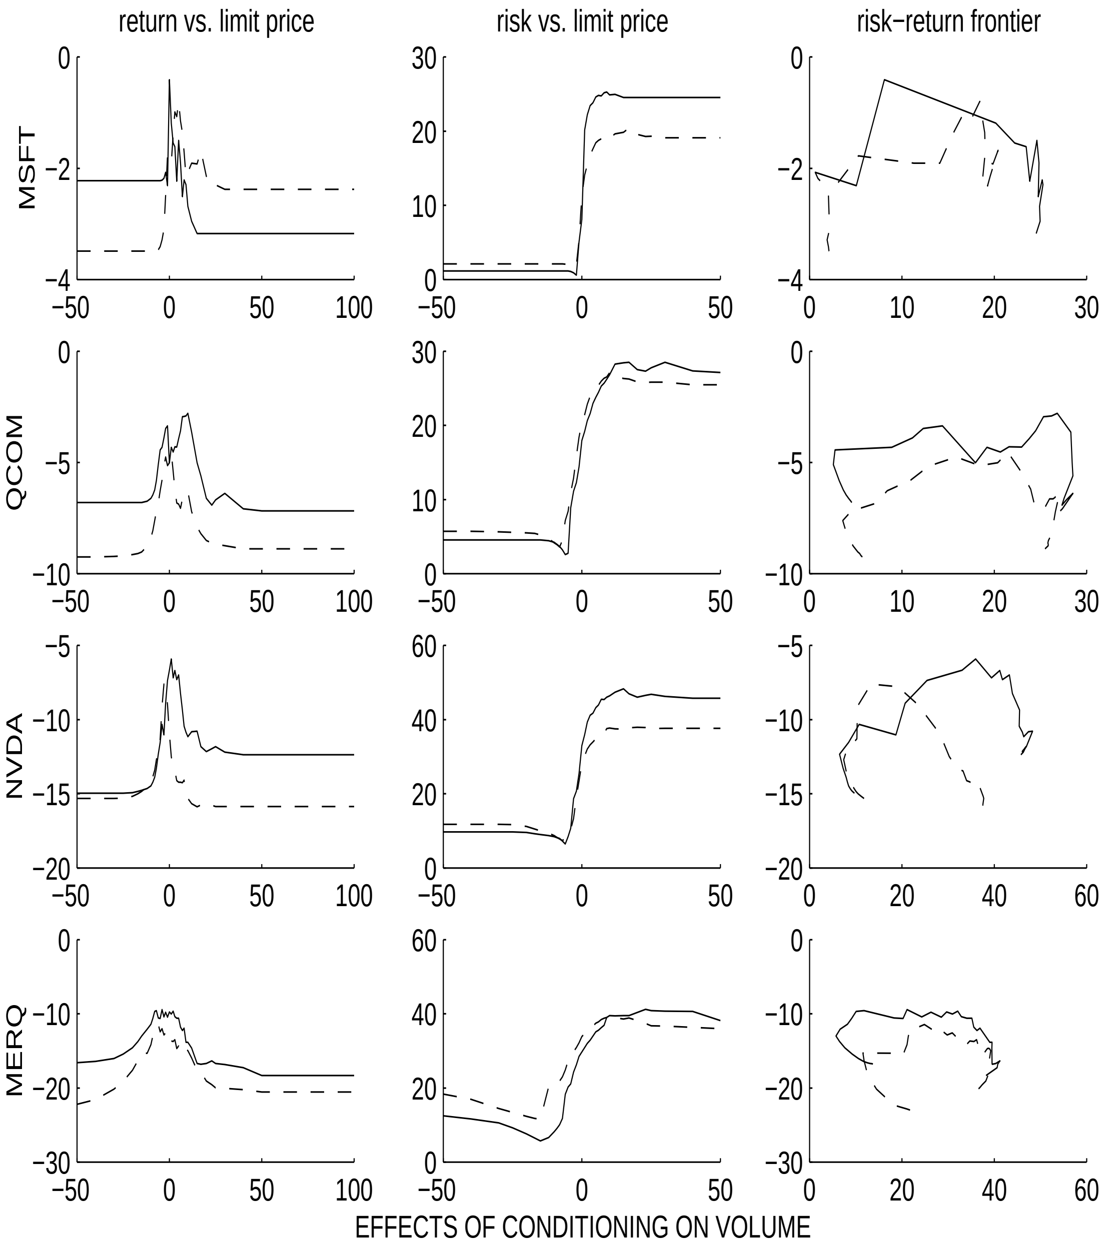

Figure A5. Transacting on a high-volume day is less expensive and less risky.

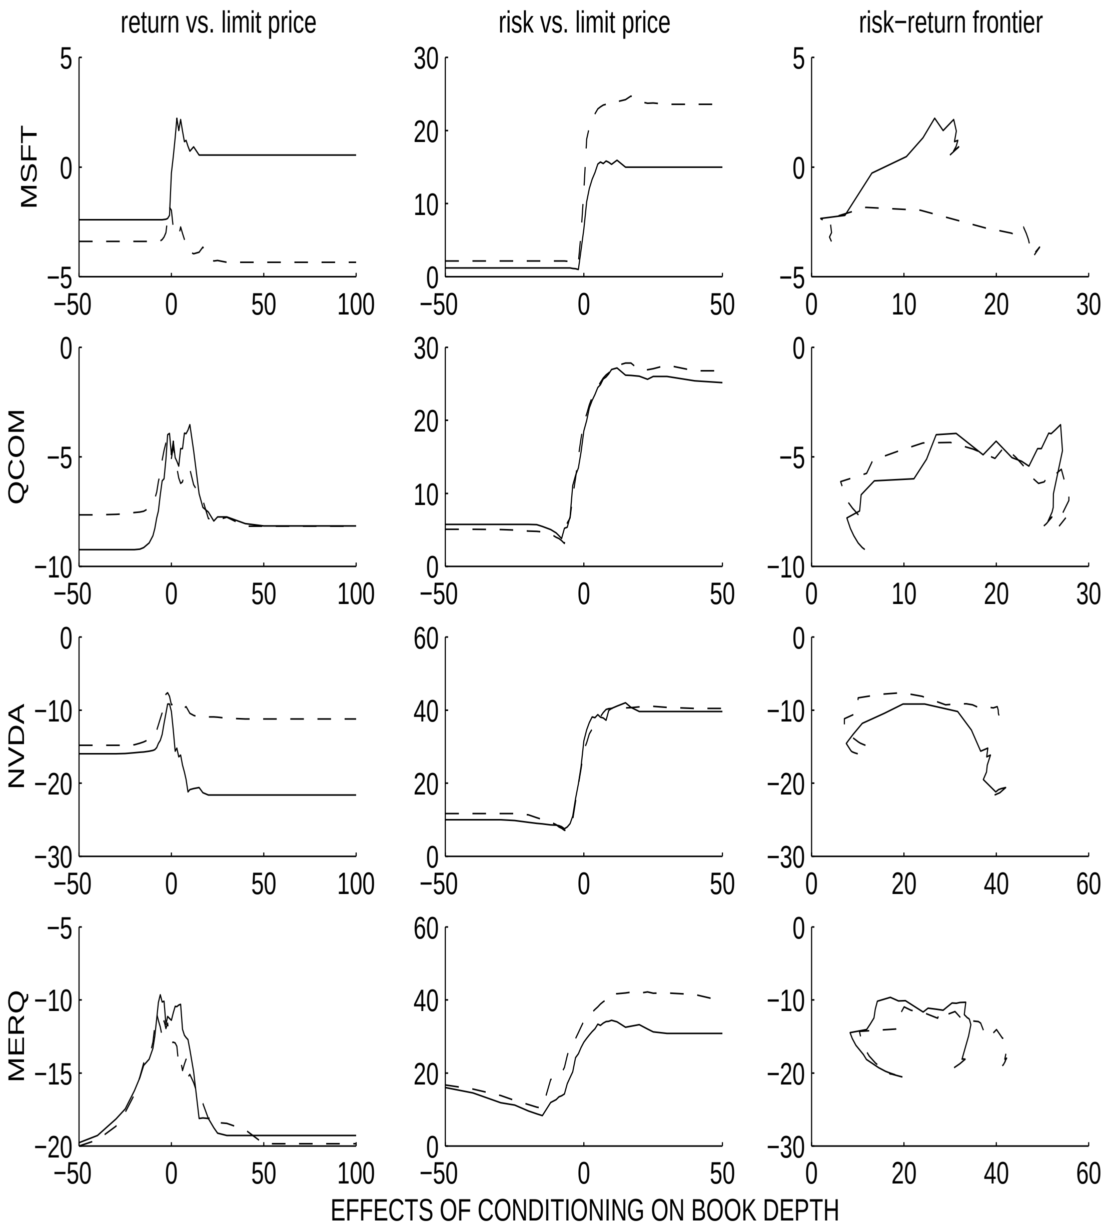

Figure A6. "Thick" book may be preferable to "thin" book.

	AMZN				NVDA				QCOM			
type	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA	QCOM	QCOM	QCOM	QCOM
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=0, no decision points, go straight to the market												
12 months												
training time	12 sec	11 sec	12 sec	11 sec	7 sec	7 sec	7 sec	7 sec	15 sec	15 sec	15 sec	14 sec
expected score	-37.75	-38.08	-75.81	-76.61	-34.40	-34.42	-57.89	-57.95	-25.52	-25.88	-49.21	-49.88
actual score	-37.75	-38.08	-75.40	-76.18	-34.40	-34.42	-57.89	-57.95	-25.26	-25.65	-48.80	-49.49
stdev	30.60	30.07	60.82	61.32	28.73	28.26	52.05	51.04	20.70	21.04	35.57	37.27
#of rounds	35837.00	8959.00	35785.00	8947.00	35837.00	8959.00	35837.00	8959.00	35806.00	8952.00	35726.00	8934.00
test score	-22.80	-22.95	-40.39	-40.57	-27.97	-28.56	-54.83	-55.78	-15.56	-15.60	-29.07	-29.31
stdev	12.80	13.13	19.29	19.54	17.97	18.71	31.46	32.13	11.28	11.28	25.11	26.91
#of rounds	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00
6 months												
training time	8 sec	8 sec	8 sec	8 sec	3 sec	4 sec	4 sec	4 sec	8 sec	9 sec	9 sec	8 sec
expected score	-35.57	-35.66	-70.46	-71.39	-35.41	-35.30	-63.01	-62.91	-26.85	-27.24	-55.04	-55.90
actual score	-35.57	-35.66	-69.63	-70.52	-35.41	-35.30	-63.01	-62.91	-26.34	-26.77	-54.27	-55.15
stdev	32.52	31.27	64.61	67.13	34.23	33.10	65.33	63.48	25.31	25.61	42.39	44.89
#of rounds	17908.00	4477.00	17856.00	4465.00	17908.00	4477.00	17908.00	4477.00	17877.00	4470.00	17797.00	4452.00
test score	-22.80	-22.95	-40.39	-40.57	-27.97	-28.56	-54.83	-55.78	-15.56	-15.60	-29.07	-29.31
stdev	12.80	13.13	19.29	19.54	17.97	18.71	31.46	32.13	11.28	11.28	25.11	26.91
#of rounds	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00
3 months												
training time	5 sec	5 sec	5 sec	6 sec	2 sec	2 sec	2 sec	2 sec	5 sec	5 sec	5 sec	5 sec
expected score	-34.55	-34.76	-67.19	-67.99	-33.91	-33.81	-64.92	-64.96	-27.59	-27.85	-59.36	-60.34
actual score	-34.55	-34.76	-67.19	-67.99	-33.91	-33.81	-64.92	-64.96	-26.60	-26.95	-57.91	-58.94
stdev	35.98	36.42	68.46	69.84	39.21	37.48	79.59	77.08	30.68	30.88	50.45	54.04
#of rounds	9324.00	2331.00	9324.00	2331.00	9324.00	2331.00	9324.00	2331.00	9293.00	2324.00	9213.00	2306.00
test score	-22.80	-22.95	-40.39	-40.57	-27.97	-28.56	-54.83	-55.78	-15.56	-15.60	-29.07	-29.31
stdev	12.80	13.13	19.29	19.54	17.97	18.71	31.46	32.13	11.28	11.28	25.11	26.91
#of rounds	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00
T=1, 1 decision point at the beginning, submit and Leave												
12 months												
training time	126 sec	118 sec	127 sec	118 sec	79 sec	69 sec	80 sec	69 sec	163 sec	152 sec	164 sec	154 sec
expected score	-17.45	-14.22	-29.35	-21.39	-21.36	-17.87	-33.12	-25.79	-12.38	-10.24	-18.74	-15.18
actual score	-18.35	-14.29	-29.71	-20.86	-22.25	-17.38	-35.04	-25.44	-12.79	-10.49	-18.25	-14.62
stdev	19.89	21.19	32.95	29.47	25.97	33.13	45.06	46.25	16.94	20.57	21.43	23.98
#of rounds	35837.00	8959.00	35776.00	8946.00	35837.00	8959.00	35837.00	8959.00	35805.00	8951.00	35721.00	8931.00
test score	-11.52	-9.05	-16.84	-13.11	-14.62	-11.84	-24.46	-16.95	-7.22	-5.50	-10.61	-8.20
stdev	4.74	6.55	6.58	6.74	11.96	21.09	17.64	19.88	5.86	8.19	7.40	8.62
#of rounds	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17462.00	4366.00
6 months												
training time	88 sec	83 sec	88 sec	83 sec	42 sec	37 sec	42 sec	37 sec	91 sec	88 sec	92 sec	87 sec
expected score	-14.03	-11.05	-23.69	-17.04	-21.46	-17.09	-32.35	-26.46	-11.88	-10.39	-18.64	-15.06
actual score	-15.32	-12.06	-25.20	-17.89	-22.65	-18.09	-36.51	-26.93	-13.02	-11.02	-18.31	-15.77
stdev	23.39	24.60	36.23	36.13	32.11	31.39	57.52	60.70	22.06	24.48	27.94	27.10
#of rounds	17908.00	4477.00	17850.00	4464.00	17908.00	4477.00	17908.00	4477.00	17876.00	4469.00	17792.00	4449.00
test score	-11.52	-9.05	-18.95	-13.11	-14.62	-12.94	-26.26	-16.95	-7.22	-5.50	-10.61	-9.56
stdev	4.74	6.55	5.18	6.74	11.96	11.84	11.44	19.88	5.86	8.19	7.40	5.50
#of rounds	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17462.00	4366.00
3 months												
training time	58 sec	56 sec	59 sec	56 sec	22 sec	18 sec	20 sec	17 sec	55 sec	53 sec	56 sec	53 sec
expected score	-11.36	-9.10	-17.55	-13.55	-21.24	-16.61	-31.31	-26.18	-10.46	-9.34	-15.78	-13.43
actual score	-13.35	-11.23	-20.23	-16.85	-22.76	-18.96	-37.48	-31.03	-12.21	-10.45	-15.80	-14.49
stdev	27.56	28.11	44.02	43.50	38.41	37.77	74.99	74.88	27.21	28.33	32.52	32.03
#of rounds	9324.00	2331.00	9322.00	2331.00	9324.00	2331.00	9324.00	2331.00	9292.00	2323.00	9208.00	2303.00
test score	-11.52	-9.05	-16.84	-13.11	-14.62	-12.94	-26.26	-20.80	-7.22	-5.50	-10.61	-9.56
stdev	4.74	6.55	6.58	6.74	11.96	11.84	11.44	13.65	5.86	8.19	7.40	5.50
#of rounds	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17464.00	4366.00	17462.00	4366.00

Table A1. Market Order and Submit-and-Leave

	AMZN				NVDA				QCOM			
type	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA	QCOM	QCOM	QCOM	QCOM
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=4, I=4												
12 months												
training time	1809 sec	1421 sec	1828 sec	1438 sec	1324 sec	878 sec	1338 sec	894 sec	2234 sec	1832 sec	2243 sec	1841 sec
expected score	-10.30519	-8.218732	-15.13786	-11.60057	-14.38133	-11.95361	-20.15326	-16.638203	-8.203276	-6.668717	-11.47585	-9.046843
actual score	-15.82637	-10.65458	-26.48004	-15.66824	-20.42954	-14.4673	-31.84608	-20.93037	-10.9371	-8.2803	-15.66835	-10.892275
stdev	21.18075	21.68009	32.54476	28.9263	25.74364	26.05181	44.30087	44.142018	17.54767	19.20311	21.08615	23.053991
#of rounds	35837	8959	35782	8947	35837	8959	35837	8959	35805	8951	35720	8930
test score	-7.715037	-5.518708	-12.2853	-7.568164	-12.18308	-9.408827	-19.66571	-12.582156	-5.500662	-4.527955	-7.865553	-5.842769
stdev	5.953001	9.113079	6.632403	10.06406	12.19725	13.44451	15.69606	16.562039	6.574295	8.088461	7.437293	9.192935
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366
6 months												
training time	1191 sec	992 sec	1197 sec	992 sec	652 sec	466 sec	661 sec	467 sec	1215 sec	1035 sec	1223 sec	1036 sec
expected score	-7.933372	-6.35652	-11.8212	-8.736483	-14.4302	-12.21857	-19.80219	-16.902595	-7.543989	-6.049731	-10.54112	-8.210931
actual score	-13.03309	-9.161941	-21.23548	-13.44372	-21.26657	-15.44521	-33.65991	-23.057553	-11.26239	-8.837457	-15.91158	-11.423007
stdev	23.89338	24.00537	37.70236	34.9788	30.50551	31.30446	58.5022	57.353402	22.25957	23.10639	27.43339	28.293098
#of rounds	17908	4477	17853	4465	17908	4477	17908	4477	17876	4469	17791	4448
test score	-7.823954	-5.526914	-12.28526	-7.613533	-12.56472	-9.442751	-20.04646	-12.689629	-5.534185	-4.536215	-7.897446	-5.860801
stdev	5.84897	9.061242	6.682547	10.08382	10.64377	13.10437	14.94494	14.606947	6.34668	7.903355	7.284513	8.649946
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366
3 months												
training time	773 sec	660 sec	779 sec	665 sec	309 sec	221 sec	311 sec	222 sec	717 sec	627 sec	720 sec	628 sec
expected score	-6.594304	-5.587466	-9.57232	-7.383089	-14.44136	-12.76947	-19.47406	-17.256738	-6.733907	-5.508023	-9.312613	-7.285368
actual score	-10.37996	-8.49229	-16.57474	-12.37195	-21.85704	-16.92053	-35.33408	-26.061595	-10.35734	-8.62748	-13.09439	-9.97196
stdev	28.34348	28.11779	44.23046	43.58103	37.49168	38.40657	75.97311	75.322118	27.52739	27.65771	32.16504	33.479203
#of rounds	9324	2331	9324	2331	9324	2331	9324	2331	9292	2323	9207	2302
test score	-7.264452	-5.52768	-11.48544	-7.458356	-12.56472	-9.442751	-20.04646	-12.974131	-5.500662	-4.536215	-7.873972	-5.84724
stdev	7.4651	9.121019	8.397466	10.15048	10.64377	13.10437	14.94494	14.371366	6.574295	7.903355	7.424618	9.188948
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366
T=8, I=8												
12 months												
training time	9816 sec	6716 sec	9910 sec	6750 sec	8107 sec	4445 sec	8161 sec	4546 sec	11767 sec	8525 sec	11855 sec	8566 sec
expected score	-7.880431	-6.549995	-10.9983	-8.877315	-10.75337	-10.22063	-14.87648	-13.268152	-6.449149	-5.80764	-8.856292	-7.488684
actual score	-15.44026	-9.989105	-26.55501	-14.75419	-19.93484	-14.21851	-31.64038	-20.036426	-10.57034	-7.905848	-15.35224	-10.295367
stdev	21.75607	20.6435	34.65007	28.75444	25.76986	28.95561	45.53104	43.099203	17.77092	19.40049	21.72407	22.198338
#of rounds	35836	8959	35781	8947	35837	8959	35837	8959	35805	8951	35720	8930
test score	-6.803329	-5.418219	-10.6396	-7.053905	-11.76981	-9.142096	-18.86171	-11.870035	-5.237678	-4.196931	-7.513414	-5.532187
stdev	7.166585	8.160265	8.168754	9.373346	12.23989	17.01021	18.65135	14.9452	6.622117	8.488317	7.735976	8.104049
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366
6 months												
training time	6157 sec	4576 sec	6208 sec	4609 sec	3868 sec	2360 sec	3914 sec	2297 sec	6264 sec	4763 sec	6264 sec	4763 sec
expected score	-6.124149	-5.12673	-8.70919	-6.788506	-11.12257	-10.29173	-14.631	-13.649491	-8.189054	-6.717761	-8.189054	-6.717761
actual score	-12.5133	-8.683162	-20.99259	-12.82052	-20.88976	-14.93578	-32.75064	-22.344856	-15.72045	-10.63594	-15.72045	-10.63594
stdev	24.28862	23.44588	38.88254	35.07966	31.1674	30.51348	58.69407	57.082417	28.14929	27.94658	28.14929	27.946577
#of rounds	17907	4477	17852	4465	17908	4477	17908	4477	17791	4448	17791	4448
test score	-6.806294	-5.430933	-10.69785	-7.088947	-11.71916	-9.161856	-18.99202	-11.978061	-7.55554	-5.543824	-7.55554	-5.543824
stdev	7.12136	8.161247	8.190146	9.521143	11.43486	12.17449	15.57266	14.355497	7.703854	8.023804	7.703854	8.023804
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366
3 months												
training time	3912 sec	3011 sec	3926 sec	3037 sec	1827 sec	1099 sec	1829 sec	1088 sec	3611 sec	2858 sec	3611 sec	2858 sec
expected score	-5.196359	-4.60484	-7.339056	-5.875644	-11.2762	-10.6381	-14.45102	-14.026249	-7.469398	-6.131463	-7.469398	-6.131463
actual score	-9.928387	-8.011618	-15.92967	-11.96324	-21.26446	-16.43734	-34.55475	-25.084563	-12.70379	-9.373915	-12.70379	-9.373915
stdev	28.24792	27.47272	44.45797	43.81266	37.8764	37.37723	76.30452	74.740804	32.54382	33.03719	32.54382	33.037186
#of rounds	9324	2331	9323	2331	9324	2331	9324	2331	9207	2302	9207	2302
test score	-6.767322	-5.443865	-10.54443	-7.081015	-11.61237	-9.164233	-19.11069	-11.999239	-7.453059	-5.554905	-7.453059	-5.554905
stdev	7.548486	8.239586	8.663896	9.607591	11.08039	12.08312	15.60875	14.346544	7.805502	8.08682	7.805502	8.08682
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366	17464	4366

Table A2. T=4, I=4 and T=8, I=8

	AMZN				NVDA			
type	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=4								
12 months								
training time	585 sec	465 sec	591 sec	469 sec	427 sec	287 sec	431 sec	293 sec
expected score	-11.50343	-8.726163	-18.44609	-12.82165	-15.52844	-12.55719	-22.75277	-18.22534
actual score	-16.19002	-10.75379	-27.94206	-15.95464	-20.8127	-14.56087	-32.71411	-21.399612
stdev	20.72643	21.2839	32.22932	27.33408	25.63885	26.02749	44.06863	44.093118
#of rounds	35837	8959	35782	8947	35837	8959	35837	8959
test score	-7.983525	-5.5602	-13.46921	-8.445114	-12.50436	-9.447862	-20.59005	-12.788462
stdev	5.870114	9.062438	6.607187	7.556182	12.12678	13.44565	15.29512	16.449447
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
6 months								
training time	384 sec	322 sec	389 sec	324 sec	212 sec	153 sec	213 sec	153 sec
expected score	-8.84719	-6.663597	-14.25078	-9.485148	-15.50381	-12.92806	-22.48607	-18.329705
actual score	-13.23776	-9.258794	-22.09807	-13.70194	-21.65213	-15.6649	-34.36128	-23.683398
stdev	23.81108	24.00308	37.21687	33.74939	30.37491	31.23723	57.99144	57.180084
#of rounds	17908	4477	17853	4465	17908	4477	17908	4477
test score	-7.983525	-5.5602	-13.46921	-8.481047	-12.99262	-9.586072	-20.744	-13.303959
stdev	5.870114	9.062438	6.607187	7.648972	10.50328	13.08254	13.98083	14.333326
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
3 months								
training time	251 sec	215 sec	254 sec	217 sec	100 sec	72 sec	101 sec	73 sec
expected score	-7.265833	-5.763128	-11.44877	-7.775082	-15.48256	-13.51272	-22.32648	-18.694023
actual score	-10.68357	-8.561961	-17.01635	-12.58173	-22.00708	-17.17005	-36.14187	-26.454916
stdev	27.64951	28.0723	43.75982	43.36288	36.84724	38.34945	75.31526	75.259622
#of rounds	9324	2331	9324	2331	9324	2331	9324	2331
test score	-7.983525	-5.5602	-12.82008	-7.782528	-14.13519	-9.586072	-21.55262	-13.303959
stdev	5.870114	9.062438	6.98125	10.07698	8.163852	13.08254	13.06082	14.333326
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
T=8								
12 months								
training time	1447 sec	997 sec	1456 sec	1006 sec	1183 sec	660 sec	1206 sec	677 sec
expected score	-9.307929	-7.194761	-14.51927	-10.51125	-12.14078	-11.11168	-17.81213	-15.294632
actual score	-16.02825	-10.14162	-27.39166	-15.40035	-20.29871	-14.60244	-32.62856	-20.726261
stdev	20.98413	20.57799	32.38541	27.88061	25.48886	26.92721	44.2485	42.995054
#of rounds	35836	8959	35782	8947	35837	8959	35837	8959
test score	-7.607713	-5.46298	-13.08229	-7.519648	-12.00243	-9.331725	-20.344	-12.273227
stdev	6.488346	8.15738	6.452692	9.167688	11.81679	14.84637	16.04985	14.821528
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
6 months								
training time	905 sec	675 sec	912 sec	691 sec	561 sec	339 sec	566 sec	349 sec
expected score	-7.267485	-5.496761	-11.60231	-7.774574	-12.48858	-11.27093	-17.50251	-15.80372
actual score	-13.00007	-8.70825	-21.75429	-13.00664	-21.15512	-15.30733	-33.74988	-23.12107
stdev	23.7241	23.26408	37.23947	34.79942	30.67535	30.37113	58.21404	56.480317
#of rounds	17907	4477	17853	4465	17908	4477	17908	4477
test score	-7.675877	-5.472449	-13.13543	-7.32745	-12.34253	-9.372069	-20.49164	-12.532301
stdev	6.412169	8.161877	6.501058	9.399335	10.88889	11.91251	14.46624	13.605684
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
3 months								
training time	577 sec	445 sec	580 sec	453 sec	265 sec	160 sec	267 sec	161 sec
expected score	-6.146993	-4.842472	-9.690859	-6.545635	-12.66591	-11.68866	-17.49601	-16.196448
actual score	-10.35882	-8.048972	-16.81271	-12.02392	-21.80773	-16.88213	-35.43936	-25.920694
stdev	27.93571	27.47691	44.00381	43.6235	37.68366	37.19689	75.82902	74.286003
#of rounds	9324	2331	9324	2331	9324	2331	9324	2331
test score	-7.338451	-5.474057	-11.96604	-7.23321	-12.34253	-9.419283	-20.49164	-12.532301
stdev	7.031198	8.241346	7.94965	9.620456	10.88889	11.85317	14.46624	13.605684
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366

Table A3. T=4, I=1 and T=8, I=1

	AMZN				NVDA			
type	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=4, I=8								
12 months								
training time	3487 sec	2751 sec	3515 sec	2760 sec	2544 sec	1695 sec	2595 sec	1737 sec
expected score	-9.973299	-8.064756	-14.34249	-11.23931	-14.00891	-11.7745	-19.41669	-16.14066
actual score	-15.77394	-10.65061	-26.47441	-15.55143	-20.35208	-14.4673	-31.70706	-20.774783
stdev	21.21607	21.68122	33.0499	29.00816	25.87413	26.05181	44.3869	44.225327
#of rounds	35837	8959	35782	8947	35837	8959	35837	8959
test score	-7.67484	-5.517723	-12.09457	-7.504779	-12.1094	-9.408827	-19.6191	-12.509427
stdev	5.964888	9.112692	6.715363	10.0916	12.31748	13.44451	15.86211	16.692325
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
6 months								
training time	2299 sec	1913 sec	2317 sec	1921 sec	1262 sec	900 sec	1274 sec	905 sec
expected score	-7.680969	-6.268573	-11.20228	-8.521609	-14.10038	-11.99765	-18.98105	-16.473462
actual score	-13.0133	-9.158558	-21.39496	-13.42462	-21.11515	-15.38788	-33.54967	-22.845673
stdev	24.19766	24.05354	38.22461	35.09412	30.59594	31.35972	58.61641	57.498283
#of rounds	17908	4477	17853	4465	17908	4477	17908	4477
test score	-7.714102	-5.530102	-12.23337	-7.53219	-12.36889	-9.419538	-19.90433	-12.519729
stdev	5.952045	9.099	6.980418	10.10849	10.77795	13.17668	15.08425	14.813017
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
3 months								
training time	1490 sec	1274 sec	1508 sec	1281 sec	596 sec	426 sec	602 sec	429 sec
expected score	-6.413638	-5.535365	-9.101316	-7.268898	-14.11975	-12.5391	-18.63708	-16.828955
actual score	-10.35918	-8.489201	-16.42045	-12.32654	-21.69641	-16.90161	-35.24847	-25.698681
stdev	28.35297	28.125	44.40846	43.61093	37.54788	38.47233	76.07587	75.402245
#of rounds	9324	2331	9324	2331	9324	2331	9324	2331
test score	-7.216742	-5.508522	-11.23946	-7.381795	-12.36889	-9.419538	-19.90433	-12.644149
stdev	7.497673	9.119884	8.475236	10.24751	10.77795	13.17668	15.08425	14.616884
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
T=8, I=4								
12 months								
training time	4977 sec	3423 sec	5009 sec	3448 sec	4117 sec	2261 sec	4170 sec	2292 sec
expected score	-8.224111	-6.698999	-11.80437	-9.25386	-11.12431	-10.45701	-15.6199	-13.77036
actual score	-15.53851	-10.02651	-26.76873	-14.76607	-20.00879	-14.23354	-31.8905	-20.244562
stdev	21.65626	20.62162	34.4096	28.57238	25.7448	28.95463	45.40814	43.049401
#of rounds	35836	8959	35782	8947	35837	8959	35837	8959
test score	-6.910763	-5.429824	-10.94695	-7.085921	-11.83352	-9.148305	-19.11646	-12.0212
stdev	7.105615	8.159493	8.042744	9.339846	12.22426	17.0112	18.42702	14.849123
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
6 months								
training time	3122 sec	2312 sec	3143 sec	2330 sec	1946 sec	1162 sec	1964 sec	1167 sec
expected score	-6.392315	-5.211701	-9.376293	-7.01897	-11.48661	-10.53489	-15.29859	-14.200671
actual score	-12.62318	-8.690075	-21.05945	-12.72334	-20.75986	-15.03009	-32.93119	-22.406148
stdev	24.23918	23.4321	38.77729	34.87573	30.80682	30.46086	58.63378	57.072902
#of rounds	17907	4477	17853	4465	17908	4477	17908	4477
test score	-6.924778	-5.436502	-10.8761	-7.092546	-11.8214	-9.216731	-19.37059	-12.026689
stdev	7.077928	8.157921	8.096796	9.342325	10.99735	12.10115	15.37407	14.351034
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366
3 months								
training time	1975 sec	1533 sec	1995 sec	1548 sec	919 sec	549 sec	926 sec	552 sec
expected score	-5.414542	-4.658196	-7.872649	-6.023304	-11.606	-10.90631	-15.14701	-14.592361
actual score	-9.917115	-8.022994	-16.02708	-11.9293	-21.36896	-16.54619	-34.6329	-25.002016
stdev	28.14084	27.47527	44.36211	43.70095	37.81553	37.34451	76.19711	74.385237
#of rounds	9324	2331	9323	2331	9324	2331	9324	2331
test score	-6.794371	-5.447899	-10.7234	-7.086596	-11.78241	-9.22612	-19.33501	-12.030589
stdev	7.391447	8.220906	8.540461	9.538637	11.0114	12.0044	15.37784	13.800352
#of rounds	17464	4366	17464	4366	17464	4366	17464	4366

Table A4. T=4, I=8 and T=8, I=4

	AMZN				NVDA			
time period	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=1								
12 months	-11.51604	-9.04527	-16.8406	-13.1108	-14.61797	-11.84308	-24.45934	-16.95354
6 months	-11.51604	-9.04527	-18.94855	-13.1108	-14.61797	-12.94026	-26.26401	-16.95354
3 months	-11.51604	-9.04527	-16.8406	-13.1108	-14.61797	-12.94026	-26.26401	-20.79588
fast/slow market								
12 months	-10.75637	-9.04527	-16.8406	-13.1108	-14.61797	-11.84308	-24.39018	-16.95354
6 months	-10.75637	-9.04527	-17.44818	-13.1108	-14.61797	-12.94026	-26.26401	-17.36543
3 months	-10.75637	-9.04527	-16.8406	-13.1108	-14.61797	-12.94026	-26.26401	-17.36543
spread & immediate cost & transacted volume misbalance								
12 months	-8.351986	-7.073632	-13.75482	-10.27902	-13.50822	-12.02513	-20.68606	-17.14561
6 months	-8.604503	-7.439351	-13.79836	-10.87664	-13.51984	-13.0902	-21.29955	-16.72578
3 months	-8.357629	-7.896165	-13.88389	-10.87552	-13.68407	-12.03825	-21.5476	-17.00991
spread & immediate cost								
12 months	-8.517772	-7.160769	-14.51095	-10.61282	-13.36405	-11.37221	-21.24779	-16.13291
6 months	-9.084246	-7.296015	-14.54926	-11.13601	-14.10398	-11.38461	-23.23286	-16.72092
3 months	-8.350108	-7.063704	-14.34806	-10.91669	-14.26722	-11.71739	-23.27045	-16.88531
spread volatility								
12 months	-9.89373	-9.04527	-16.87618	-13.1108	-14.61797	-11.52268	-24.2935	-16.95354
6 months	-11.51604	-9.04527	-17.02616	-13.1108	-14.61797	-12.94026	-26.26401	-17.0188
3 months	-9.89373	-9.04527	-16.8406	-11.10037	-14.74131	-12.94026	-26.26401	-20.79588
price volatility								
12 months	-11.51604	-9.04527	-17.33632	-13.1108	-14.61797	-12.45651	-25.51643	-16.95354
6 months	-11.51604	-9.04527	-18.45284	-13.1108	-14.61797	-12.94026	-26.26401	-19.3515
3 months	-11.51604	-9.04527	-16.8406	-13.1108	-14.61797	-12.94026	-26.26401	-19.3515
added volume misbalance								
12 months	-11.09404	-9.04527	-16.2808	-13.1108	-14.61797	-12.30529	-25.45642	-16.95354
6 months	-11.09404	-9.04527	-17.85501	-13.1108	-14.61797	-12.94026	-26.26401	-17.60168
3 months	-11.09404	-9.04527	-16.2808	-13.1108	-14.61797	-12.94026	-26.26401	-19.36746
transacted volume misbalance								
12 months	-10.32719	-7.885921	-16.8406	-11.63941	-14.93729	-12.21554	-23.26255	-16.95354
6 months	-10.32719	-9.04527	-16.8406	-13.1108	-14.61797	-12.94026	-26.26401	-17.79279
3 months	-10.32719	-9.04527	-15.31228	-13.1108	-15.3508	-12.21554	-26.26401	-17.79279
transacted volume								
12 months	-10.48666	-9.04527	-17.12243	-13.1108	-14.61797	-11.75944	-24.45019	-16.95354
6 months	-10.48666	-9.04527	-17.12243	-13.1108	-14.82218	-12.94026	-26.26401	-17.92734
3 months	-10.48666	-9.04527	-16.8406	-13.1108	-14.95392	-12.94026	-26.26401	-17.92734
immediate cost								
12 months	-9.964776	-7.520079	-16.0063	-12.07239	-14.61797	-11.84308	-22.92884	-17.52592
6 months	-9.964776	-9.04527	-16.17301	-13.1108	-15.09414	-12.94026	-24.99227	-18.32266
3 months	-9.964776	-9.04527	-14.89597	-12.07239	-15.09414	-12.94026	-24.99227	-18.32266
misbalance & spread								
12 months	-9.116307	-7.327463	-14.68573	-11.12872	-13.19556	-11.15108	-21.75336	-16.21443
6 months	-9.116478	-7.269706	-14.71299	-11.26269	-14.43173	-11.24187	-23.30388	-16.54105
3 months	-9.109174	-7.114191	-14.64217	-11.16109	-14.44078	-11.12562	-24.06807	-17.55687
spread								
12 months	-8.668304	-7.944448	-15.49829	-11.96291	-13.29799	-11.13585	-21.52916	-17.35383
6 months	-8.668304	-7.944448	-15.49829	-11.85383	-15.25871	-11.13585	-23.80418	-17.35383
3 months	-8.695923	-7.91033	-14.50686	-11.78764	-15.25871	-11.13585	-23.80418	-17.35383
BBO misbalance								
12 months	-11.51604	-8.759237	-16.8406	-13.1108	-14.66442	-11.74844	-24.59611	-16.95354
6 months	-11.51604	-9.04527	-18.94855	-13.1108	-14.61797	-12.94026	-26.26401	-17.08206
3 months	-11.23082	-9.206044	-16.8406	-13.1108	-14.61797	-12.90002	-26.26401	-20.45165
ask volume								
12 months	-11.51604	-9.04527	-16.8406	-13.1108	-14.61797	-11.69372	-24.64439	-16.95354
6 months	-11.51604	-9.37091	-18.94855	-13.1108	-14.61797	-12.94026	-26.26401	-17.57148
3 months	-11.51604	-9.37091	-16.8406	-13.1108	-14.6728	-12.94026	-26.26401	-20.5537
bid volume								
12 months	-11.48692	-8.513385	-17.33997	-13.1108	-14.61797	-12.82726	-24.63103	-16.97339
6 months	-11.48692	-9.04527	-18.90631	-13.1108	-14.61797	-12.9095	-26.26401	-16.95354
3 months	-11.48692	-9.04527	-16.8044	-13.1108	-14.61797	-12.82726	-26.26401	-20.31455

Table A5. T=1, market variables

	AMZN				NVDA			
time period	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=4, I=4								
12 months	-7.715037	-5.518708	-12.2853	-7.568164	-12.18308	-9.408827	-19.66571	-12.58216
6 months	-7.823954	-5.526914	-12.28526	-7.613533	-12.56472	-9.442751	-20.04646	-12.68963
3 months	-7.264452	-5.52768	-11.48544	-7.458356	-12.56472	-9.442751	-20.04646	-12.97413
fast/slow market								
12 months	-7.794529	-5.518833	-12.27002	-8.131643	-12.07065	-9.420108	-19.77503	-12.56856
6 months	-7.850726	-5.530625	-12.39945	-7.750529	-12.3845	-9.412997	-19.85574	-12.5168
3 months	-7.425955	-5.518222	-11.55738	-7.532124	-12.47568	-9.43419	-19.85574	-12.72256
spread & immediate cost & transacted volume misbalance								
12 months	-7.110003	-5.696713	-10.56031	-7.531782	-11.90362	-9.559237	-18.06544	-12.30922
6 months	-7.119105	-5.754774	-10.55194	-7.495183	-11.76584	-9.58841	-18.28593	-12.56633
3 months	-6.995764	-5.824249	-10.48207	-7.601711	-11.78474	-9.611809	-18.32625	-12.78887
spread & immediate cost								
12 months	-7.234735	-5.503011	-10.72856	-7.72423	-11.78035	-9.481387	-18.08977	-12.22938
6 months	-7.251	-5.511894	-10.65036	-7.416569	-11.97821	-9.460697	-18.63431	-12.61288
3 months	-7.122112	-5.597677	-10.56924	-7.28197	-11.99514	-9.628555	-18.78835	-12.66383
spread volatility								
12 months	-7.698555	-5.509966	-12.1662	-7.599746	-12.05568	-9.420082	-19.71976	-12.51667
6 months	-7.810882	-5.518938	-12.2515	-8.258405	-12.58276	-9.426598	-19.8939	-12.69073
3 months	-7.644007	-5.515441	-11.40434	-7.54458	-12.58276	-9.444056	-19.8939	-12.75896
price volatility								
12 months	-7.72221	-5.516686	-12.22983	-8.127476	-12.09275	-9.416678	-19.66783	-12.52721
6 months	-7.81095	-5.536426	-12.27255	-8.135405	-12.56251	-9.46022	-19.833	-12.68741
3 months	-7.264452	-5.515389	-11.48275	-7.470471	-12.93049	-9.438489	-19.833	-12.90628
added volume misbalance								
12 months	-7.785574	-5.527892	-12.33182	-8.101189	-12.12913	-9.408252	-19.66571	-12.58844
6 months	-7.810839	-5.532208	-12.41848	-7.901758	-12.62139	-9.499596	-19.85057	-12.67904
3 months	-7.496091	-5.52221	-11.76505	-7.4871	-12.69111	-9.449765	-19.95614	-12.89099
transacted volume misbalance								
12 months	-7.38753	-5.59896	-11.64993	-7.700754	-12.3407	-9.663079	-19.54296	-12.85417
6 months	-7.42335	-5.488739	-11.71302	-7.711735	-12.62456	-9.42348	-19.86776	-12.59594
3 months	-7.21674	-5.886178	-11.19012	-7.352326	-12.66555	-9.540784	-19.88614	-13.12103
transacted volume								
12 months	-7.77705	-5.50584	-12.21546	-7.784465	-12.09133	-9.418884	-19.7017	-12.48006
6 months	-7.802859	-5.561377	-12.32014	-7.826031	-12.7091	-9.411614	-19.73031	-12.82245
3 months	-7.443259	-5.519187	-11.46552	-7.43615	-12.73605	-9.465169	-19.73561	-12.92601
immediate cost								
12 months	-7.269388	-5.489655	-11.18426	-7.447358	-11.72823	-9.464639	-18.67413	-12.44303
6 months	-7.271553	-5.498656	-11.29359	-7.468374	-12.05078	-9.383418	-19.08557	-12.63039
3 months	-7.131967	-5.458562	-10.9467	-7.2618	-12.07591	-9.384912	-19.12974	-12.72143
price level								
12 months	-7.770091	-5.499466	-12.24697	-7.481019	-12.18026	-9.413271	-19.70351	-12.56838
6 months	-7.799929	-5.51952	-12.27531	-7.527183	-12.49982	-9.42007	-19.80301	-12.58825
3 months	-7.228571	-5.489113	-11.43699	-7.369283	-12.54663	-9.424255	-20.03575	-12.91746
misbalance & spread								
12 months	-7.23632	-5.520025	-10.79658	-7.624017	-11.81632	-9.388548	-18.30234	-12.16575
6 months	-7.242202	-5.507667	-10.7635	-7.451851	-12.02431	-9.503094	-18.81521	-12.55038
3 months	-7.186937	-5.560466	-10.6337	-7.345989	-12.03317	-9.519668	-18.84927	-12.61299
spread								
12 months	-7.074631	-5.47304	-11.14937	-7.499175	-11.80255	-9.318248	-18.35182	-12.24052
6 months	-7.085302	-5.479379	-11.18461	-7.367948	-11.94347	-9.39259	-19.10957	-12.51334
3 months	-7.018375	-5.570015	-10.97536	-7.306992	-11.91833	-9.330069	-19.06514	-12.57665
BBO misbalance								
12 months	-7.71518	-5.523091	-12.27259	-7.663097	-12.13432	-9.396072	-19.66367	-12.51006
6 months	-7.755265	-5.523676	-12.30914	-7.716151	-12.41382	-9.444072	-19.79204	-12.67754
3 months	-7.381367	-5.524827	-11.62831	-7.445895	-12.63162	-9.443554	-20.03182	-12.70701
ask volume								
12 months	-7.735143	-5.539854	-12.2873	-7.721782	-12.18744	-9.425306	-19.69603	-12.52971
6 months	-7.770252	-5.528055	-12.32186	-7.770615	-12.61675	-9.447152	-19.79306	-12.67611
3 months	-7.376622	-5.52551	-11.72865	-7.494555	-12.62958	-9.450883	-19.79275	-12.6908
bid volume								
12 months	-7.753071	-5.508916	-12.25881	-7.573006	-12.14082	-9.433896	-19.6443	-12.5658
6 months	-7.782269	-5.530594	-12.3342	-7.616809	-12.44143	-9.45149	-20.04304	-12.71152

Table A6. T=4, I=4, market variables

	AMZN				NVDA				QCOM			
time period	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA	QCOM	QCOM	QCOM	QCOM
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=0, go straight to the market												
middle	-22.80242	-22.9548	-40.39251	-40.56605	-27.96535	-28.56462	-54.82662	-55.77596	-15.56037	-15.6034	-29.06795	-29.31355
open	-39.00535	-44.83018	-66.69617	-73.72371	-55.50906	-65.11926	-102.3339	-116.4385	-27.07475	-32.23801	-50.85694	-58.5487
	-71%	-95%	-65%	-82%	-98%	-128%	-87%	-109%	-74%	-107%	-75%	-100%
close	-19.90969	-20.10766	-35.95005	-36.17048	-23.55405	-23.60947	-44.99029	-45.44143	-14.23807	-14.45448	-25.8549	-25.98394
	13%	12%	11%	11%	16%	17%	18%	19%	8%	7%	11%	11%
T=1, submit-and-leave												
middle	-11.51604	-9.04527	-16.8406	-13.1108	-14.61797	-11.84308	-24.45934	-16.95354	-7.223714	-5.499105	-10.60618	-8.198997
open	-12.93897	-10.93055	-17.34495	-16.05533	-20.22098	-18.62836	-34.62482	-30.19747	-8.918762	-9.856424	-12.99812	-13.78512
	-12%	-21%	-3%	-22%	-38%	-57%	-42%	-78%	-23%	-79%	-23%	-68%
close	-9.053867	-8.646722	-13.33607	-12.75218	-12.7831	-9.668803	-20.73457	-15.0254	-6.923678	-5.593996	-10.1082	-8.39501
	21%	4%	21%	3%	13%	18%	15%	11%	4%	-2%	5%	-2%
T=4, I=4												
middle	-7.715037	-5.518708	-12.2853	-7.568164	-12.18308	-9.408827	-19.66571	-12.58216	-5.500662	-4.527955	-7.865553	-5.842769
open	-8.46462	-6.676184	-13.08127	-8.925382	-15.8407	-11.21487	-24.89308	-15.31719	-6.03863	-6.551096	-8.624759	-8.347163
	-10%	-21%	-6%	-18%	-30%	-19%	-27%	-22%	-10%	-45%	-10%	-43%
close	-6.110403	-4.885341	-9.590198	-6.549833	-9.660365	-8.297456	-14.87411	-10.44369	-4.785233	-4.312416	-7.03259	-5.419563
	21%	11%	22%	13%	21%	12%	24%	17%	13%	5%	11%	7%

Table A7. Middle-Open-Close Cost Comparison

	AMZN				NVDA				QCOM			
time period	AMZN	AMZN	AMZN	AMZN	NVDA	NVDA	NVDA	NVDA	QCOM	QCOM	QCOM	QCOM
	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min	2 min	8 min
	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh	5000 sh	5000 sh	10000 sh	10000 sh
T=1, submit-and-leave												
RL	-11.51604	-9.04527	-16.8406	-13.1108	-14.61797	-11.84308	-24.45934	-16.95354	-7.223714	-5.499105	-10.60618	-8.198997
VWAP	-14.4572	-10.0065	-26.47621	-17.01088	-21.24543	-16.07646	-41.88323	-29.08443	-9.883935	-7.300307	-18.02211	-12.13617
	-26%	-11%	-57%	-30%	-45%	-36%	-71%	-72%	-37%	-33%	-70%	-48%
T=4, I=4												
RL	-7.715037	-5.518708	-12.2853	-7.568164	-12.18308	-9.408827	-19.66571	-12.58216	-5.500662	-4.527955	-7.865553	-5.842769
VWAP	-8.056313	-5.796691	-15.72027	-9.203008	-12.98528	-9.695826	-22.99992	-14.35998	-5.795226	-4.458399	-10.19172	-6.514435
	-4%	-5%	-28%	-22%	-7%	-3%	-17%	-14%	-5%	2%	-30%	-11%
T=8, I=8												
RL	-6.803329	-5.418219	-10.6396	-7.053905	-11.76981	-9.142096	-18.86171	-11.87004	N/A	N/A	-7.513414	-5.532187
VWAP	-6.822898	-4.7388	-12.60162	-6.953093	-11.65042	-8.153674	-19.7719	-11.44357	N/A	N/A	-8.331147	-5.213517
	0%	13%	-18%	1%	1%	11%	-5%	4%			-11%	6%

Table A8. RL vs. VWAP Cost Comparison (mid-day)

		S&L					T=4 I=4							
time perioc	type	2 min		30 sec			2 min				30 sec			
		1000 sh		5000 sh			1000 sh				5000 sh			
Maximize return		Cost	Stdev	Cost	Stdev		Cost	Stdev	Cost%	Stdev%	Cost	Stdev	Cost%	Stdev%
AMZN	time						257 sec				361 sec			
01/03-06/(	expected	-8.48946	18.48136	-28.6855	29.36044		-5.80399		31.63%		-18.8567		34.26%	
01/03-06/(	training	-8.48946	18.48136	-28.6855	29.36044		-6.3228	16.26619	25.52%	11.99%	-19.8649	15.12749	30.75%	48.48%
07/03-12/(	test	-5.83999	15.00964	-21.935	32.13714		-5.02287	13.7584	13.99%	8.34%	-14.2077	23.31326	35.23%	30.75%
NVDA	time						247 sec				397 sec			
01/03-06/(	expected	-9.59603	19.31967	-22.6302	23.78465		-6.6		31.22%		-19.4131		14.22%	
01/03-06/(	training	-9.59603	19.31967	-22.6302	23.78465		-7.32476	23.6375	23.67%	-22.35%	-20.2367	17.20669	10.58%	27.66%
07/03-12/(	test	-9.68538	20.07093	-26.2993	35.15734		-9.05715	22.45493	6.49%	-11.88%	-22.4759	30.44082	14.54%	13.42%
QCOM	time						481 sec				609 sec			
01/03-06/(	expected	-4.99806	14.30571	-15.2819	17.97546		-3.53351		29.30%		-11.5633		24.33%	
01/03-06/(	training	-4.99806	14.30571	-15.2819	17.97546		-3.63124	15.83285	27.35%	-10.68%	-11.1738	9.142793	26.88%	49.14%
07/03-12/(	test	-5.67386	13.42747	-16.7185	26.49638		-5.34592	14.98167	5.78%	-11.57%	-11.8788	22.08217	28.95%	16.66%
Maximize Sharpe ratio														
AMZN	time						259 sec				363 sec			
01/03-06/(	expected	-8.48946	18.48136	-28.6855	29.36044		-8.34157		1.74%		-19.8591		30.77%	
01/03-06/(	training	-8.48946	18.48136	-28.6855	29.36044		-8.65541	10.02461	-1.95%	45.76%	-21.5059	14.22605	25.03%	51.55%
07/03-12/(	test	-5.83999	15.00964	-21.935	32.13714		-6.04282	8.166813	-3.47%	45.59%	-15.8411	23.28136	27.78%	27.56%
NVDA	time						248 sec				399 sec			
01/03-06/(	expected	-9.59603	19.31967	-22.6302	23.78465		-9.55964		0.38%		-21.614		4.49%	
01/03-06/(	training	-9.59603	19.31967	-22.6302	23.78465		-9.73407	11.60408	-1.44%	39.94%	-22.6538	13.92921	-0.10%	41.44%
07/03-12/(	test	-9.68538	20.07093	-26.2993	35.15734		-10.3239	13.7367	-6.59%	31.56%	-23.1665	28.68024	11.91%	18.42%
QCOM	time						484 sec				625 sec			
01/03-06/(	expected	-4.99806	14.30571	-15.2819	17.97546		-6.51629		-30.38%		-13.4948		11.69%	
01/03-06/(	training	-4.99806	14.30571	-15.2819	17.97546		-6.36176	7.060297	-27.28%	50.65%	-13.4287	5.461288	12.13%	69.62%
07/03-12/(	test	-5.67386	13.42747	-16.7185	26.49638		-6.64002	7.518885	-17.03%	44.00%	-13.9406	20.62619	16.62%	22.15%

Table A9. Sharpe Ratio Maximization: Inferior Returns

Optimal Limit Order Price Explanatory Model
(buy side, sell side is symmetrical)

- Goal: minimize average execution price **or** buy a certain quantity of stock (V) while spending as little money as possible.
- In a multi-period model: $\min[\Sigma P_t * v_t]$ for all $t \in [0, T]$, so that $\Sigma v_t = V$.
- In any given time period, our mission is the same – minimize cash out-flows from that period onwards: $\min[CF_{t+}]$
- Recursive relationship: $CF_{t+} = CF_t + CF_{(t+1)+}$ (cash flows in the current period plus all the subsequent cash flows); CF_T is the terminal value
- To simplify notation, we will drop the time subscript, and our goal is: $\min[CF + CF_+]$
- Only variable that we can control in a given time period is the price of a limit order for that period: $x \in [0, N]$.
- We have to pick an action which minimizes current and future cash flows: $\min_x[CF + CF_+]$
- What are cash flows? Execution price multiplied by execution volume: $CF = P*v$; $CF_+ = P_+*v_+$.
- $\min_x[P*v + P_+*v_+]$
- Technically, $V = v_- + v + v_+$, but since v_- (inventory already executed) is known at time t, we can write $v + v_+ = V^*$. And if we set $V^* = 1$, then $v + v_+ = 1$, or $v_+ = 1 – v$.
- $\mathbf{\min_x[P*v + P_+*(1–v)]}$. Main intuition: the more volume gets transacted in this period, the less remains for subsequent periods, and vice versa. This way, we are minimizing volume-weighted execution price.
- Let's establish the relationship between P and v, or more specifically between x and V
- Order aggressiveness:
  -------0-----------------b-----a-------------------------N---------------------->
  Order priced at 0 will never execute (v = 0), order priced at N will execute for a full amount (v = 1), b is bid, a is ask.
  Side note – in our current implementation:

----[-2000]-----[2]---0-----a---[-4]----------------[-800]------------------->

- For orders priced in [0, a) only:
    - for x =0, all cash flows are future cash flows: $P_+ * (1 – 0) = P_+$
    - orders priced more aggressively will attract more volume
    - in the *simplest possible form*, increasing price by 1 increase volume by λ
    - $v(0) = 0$; $v(1) = \lambda$; $v(2) = 2\lambda$, … $v(a -1) = (a-1)* \lambda$
    - $CF = P*v = x*v(x) = x*x*\lambda$
    - $CF_+ = P_+*(1 – x*\lambda)$
    - $\mathbf{min_x[x*x*\lambda + P_+*(1 – x*\lambda)]}$, for $x \in [0, a)$
- For orders priced in [a, N] only:
    - we now have "standing" volume already on the book in addition to the volume that will be induced by the limit order
    - when $x \in [0, a)$, $P = x$, but when $x \in [a, N])$, $P <= x$
    - $CF = CF_{standing} + CF_{induced}$
    - Assume that standing order distribution is governed by another parameter γ: volume increase by γ every price step
    - $v_s(a) = \gamma$, $v_s(a+1) = 2\gamma$, and so on
    - $CF_{standing} = \Sigma_{n \in [a, x]}\, n*v(n) = \Sigma_{n \in [a, x]}(n*(n-a+1)* \gamma)$
    - $CF_{induced} = P*v = x*v(x) = x*x*\lambda$
    - $\mathbf{min_x[\Sigma_{n \in [a, x]}(n*(n-a+1)* \gamma) + x*x*\lambda + P_+*(1 – x*\lambda)]}$, for $x \in [a, N]$
- Overall the task is
    - $\mathbf{min_x[x*x*\lambda + P_+*(1 – x*\lambda), \text{ for } x \in [0, a), \Sigma_{n \in [a, x]}(n*(n-a+1)* \gamma) + x*x*\lambda + P_+*(1 – x*\lambda), \text{ for } x \in [a, N]]}$
- There are other ways to model price-volume relationship with a single parameter
    - $v(x) = \log_\lambda(x+1)$
- Terminal value for P_T: $\Sigma_{n \in [a, N]}(n*(n-a+1)* \gamma)$, recursive solutions from here to other states
- A simple numerical example:
    - One time period
    - P_+ is a parameter (an exogenous input that can be changed)

- $\gamma = 0$ (empty book)
- λ is the second parameter

lambda:	0.1		**P+:**	5	
price	v(t)	v(t+)	CF(t)	CF(t+)	Total CF
0	0	1	0	5	5
1	0.1	0.9	0.1	4.5	4.6
2	0.2	0.8	0.4	4	4.4
3	0.3	0.7	0.9	3.5	4.4
4	0.4	0.6	1.6	3	4.6
5	0.5	0.5	2.5	2.5	5
6	0.6	0.4	3.6	2	5.6
7	0.7	0.3	4.9	1.5	6.4
8	0.8	0.2	6.4	1	7.4
9	0.9	0.1	8.1	0.5	8.6
10	1	0	10	0	10