Electronic Trading in Order-Driven Markets: Efficient Execution

Yuriy Nevmyvaka
Carnegie Mellon University
yuriy@cs.cmu.edu

Michael Kearns
University of Pennsylvania
mkearns@cis.upenn.edu

Amy Papandreou
Lehman Brothers
amerritt@lehman.com

Katia Sycara
Carnegie Mellon University
katia@cs.cmu.edu

Abstract.

In this paper, we address the importance of efficient execution in electronic markets. Due to intense competition for profit opportunities, trading costs can represent a significant portion of overall return. They must be taken into account both when a specific trade is being executed, and when a general investment strategy is being designed. We empirically demonstrate that by combining market orders (which offer immediate execution regardless of price) and limit orders (which offer uncertain execution at a specified price), we are able to obtain a superior average price than by using market orders only. Our analysis highlights the trade-off between expected price improvement from limit orders and the risk of non-execution. We show how to determine the optimal limit order price in a simplified setting and suggest how this approach can be generalized to a complete solution. All of our experimental results are obtained on an extensive collection of NASDAQ limit order data.

I. Introduction.

Execution considerations permeate every aspect of investment activity – from the development of high-level trading strategies to post-trading performance attribution – because the actual prices at which trades occur are a direct consequence of the execution mechanism in place. Every market participant should be concerned with transaction costs regardless of their investment goals and style. Ideally, one would like to have a comprehensive execution system, which aims to minimize the costs of trading. Classifying and quantifying various trading costs is a complex task (see [Kissell and Glantz, 2003]), and a comprehensive trade optimization system must solve a challenging multi-dimensional problem.

The goal of this paper is not to offer the ultimate solution to efficient execution, but rather to develop a simple building block, which can be easily implemented, quantified, and then extended into a more comprehensive system. In our study, we concentrate on a single but important aspect of the overall problem: the immediate price received or paid for a transaction of a fixed size over a fixed period of time. This makes our setup stylized and yet practical. More specifically, we ask the following question: how should one buy (respectively, sell) V shares of a given stock over T seconds while spending the least (respectively, receiving the most) cash? In the framework we consider, a trader has only three options: (1) submit a market order at time 0 for V shares, (2) submit a market order at time T for V shares, or (3) submit a limit order at time 0 for V shares and a market order for the unexecuted shares (if any) at time T. In other words, we can use a market order at the beginning of the time period, a market order at the end of the time period, or a limit order combined with a market order for residual shares. In all three cases, we will always end up with the same number of shares, but will have spent different amounts of cash entering the position. By formulating the execution problem in this fashion, we can quantitatively answer a number of important questions. Which of the three options is the most effective? How should one price a limit order to expect the most favorable execution? How can we quantify the risks of non-execution and unfavorable price movements?

We suggest a precise analytical methodology to reason about efficient execution in modern electronic markets: for every possible limit order price, we compute the ratio of the execution price to the mid-spread price at the start of the trading period; resulting "return curves" allow us to pinpoint which pricing strategy produces the most favorable expected price. We similarly build out the risk curves, defining risk as the standard deviation of the outcomes of each strategy. Risk curves elucidate the relationship between limit price and execution uncertainty. We then merge the two curves into a single function to demonstrate the trade-off between expected returns and risk. We conclude that it is only optimal to operate on a specific portion of the risk-return function, which we call "efficient pricing frontier". Thus the output of our analysis allows the trader to select some tolerable risk threshold and to determine a strategy which delivers the most favorable execution price for that level of risk. Finally, we identify the relative importance of different market microstructure variables (order size, execution window, market activity, etc.) in the optimization process, and explain how to expand our basic methodology to the multi-period dynamic execution.

To be able to assign specific numbers to the expected price improvements and risk associated with each strategy, we perform "what-if" simulations within historical limit order books. We simulate what would have happened to a hypothetical order in the real-world market, and we record the execution price at the end of each trial. (Details of this experimental methodology are provided in Section III). In our approach, we rely on the pre-committed liquidity, represented by resting limit orders, which is another aspect that sets our work apart from similar studies.

This paper is structured as follows: in the next section we will discuss related work making an emphasis on the novelty of using limit order books in our analysis. We then describe the experimental setup we use and point out its strengths and shortcomings. In Section IV, we introduce our analytical framework: how to quantify returns, risk, and the combination of the two. Section V contains a summary of our empirical results – a study across various stocks and trading conditions. We generalize our findings and suggest how more complex execution systems can be built upon them in Section VI.

II. Related Work.

We believe that the main innovation in our approach is bringing together the subjects of efficient execution and limit order trading. While each area has received attention in academic literature, we are not aware of other studies that cast the execution problem within the context of an order-driven marketplace.

Both academics and professionals have long acknowledged the importance of execution optimization. An overview of a number of quantitative execution methods is presented in [Kissell and Glantz, 2003]. This book describes a comprehensive top-down approach to evaluating trading costs and optimizing execution. Other important works on the general trade execution are [Almgren and Chriss, 1999] and [Almgren and Chriss, 2003]. [Bertsimas and Lo, 1998] suggest a dynamic programming approach to solving the execution problem, which can be further extended to the portfolio setting. The majority of studies in this area concentrate on the market impact of trading, which is the most important aspect of execution. Their approach to quantifying this effect, however, is fundamentally different from ours: other studies look at the post-trade price changes, while we analyze the ex-ante pre-committed liquidity that we see on the order books.

Analyzing and modeling limit order trading is a part of the broader discipline of market microstructure, which studies interactions among market participants and the process of price formation. The extensive overview of the subject is [Harris, 2002], which includes the introduction to the order-driven market mechanisms. One of the first inquiries into the real-world limit order markets is [Bias et al., 1995], which develops a theoretical model of limit order submissions and conducts a comprehensive empirical study of the Paris Bourse. A number of papers try to explain the rationale behind using limit orders, modeling behavior of different types of traders: [Lo et al., 2000] and [Hollifield et al., 2003].

Perhaps the closest in spirit to our work are the studies that perform empirical analyses of various order submission strategies. [Hasbrouck and Harris, 1996] and [Handa and Schwartz, 1996] compare market order and limit order trading strategies in different markets and confirm that limit orders can indeed enhance returns. It is this exact approach that we are applying to efficient execution.

III. Experimental Setup.

In our study, we used historical records from the Island Electronic Communication Network (ECN). Island (recently acquired by INET) is essentially an electronic and completely automated stock exchange which accounts for a significant volume of trading in NASDAQ stocks. In the Island files we can see every event that happened to a stock through a trading day – all order submissions, cancellations, and executions. Every action is time-stamped, and since 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 go through the record of order arrivals, cancellations, and executions, updating the state of limit order books after every event; at designated times, we insert into the order flow "artificial" orders that represent various trades we wish to investigate. We then perform all the executions and maintain order priorities in the same way it is done in the real ECN. Such a setup allows us to run "what-if" simulations in historical order books. In all simulations, we use one-and-a-half years of data – from January 2003 to June 2004 –, which is available for every stock traded on Island. While we have obtained results for a broad cross-section of stocks, for the purpose of brevity, we concentrate here on a single security – Microsoft Corp. common stock (MSFT). For an extended version of our results, see [Nevmyvaka et al., 2004].

We investigate efficient execution by examining the pre-committed liquidity, which is precisely what limit order books represent. When a trader submits a limit order, he provides 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 interval. Historical limit order books allow us to quantify and compare these two dimensions of order submissions.

IV. Analytical Framework.

In this main section we 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 re-visit 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 T seconds.

This task can be executed using the following trading strategies:

(1) submit a market order for the entire amount immediately – 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 – can achieve price improvement, 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 submit a market order for the remainder of shares (if any) at the end of the interval – can avoid paying transaction costs, otherwise becomes the worst case of Strategy 2.

All these strategies end up with the same position V after T seconds, but will have spent different amounts of cash. Therefore, if we plot levels of cash that each strategy spent (respectively, generated), we can find which one is the most efficient and if there is some general relationship between various strategies' performance.

It is important to understand first what it is exactly that we are measuring. Although we are interested in comparing levels of cash, we cannot simply express our results in dollars. For example, if a stock traded at $5 in January 2003 and at $50 in June 2004, cost savings of $0.02 will have very different implications in the two cases. Therefore, the variable that we chose to optimize is the difference between the execution price (weighted by volume) and the mid-spread price at the beginning of the time period, expressed as a fraction of the initial mid-spread price. Say the market at the beginning of the time period has a bid of $24.12 and an ask of $24.18, yielding a mid-spread price of $24.15. If we execute a buy order at the ask, our price differential is ($24.15 - $24.18)/$24.15 = – 0.001242. We express these numbers in basis points (-12 bp in this example) and for the lack of better terminology refer to them as "returns"; however, note that they are not returns in the conventional sense of the term.

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 typical order price-return curve is shown below in Figure 1.

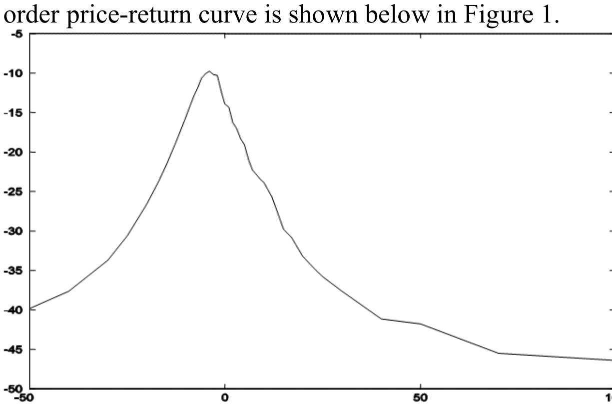

Figure 1. The peak in the curve represents pricing strategy that produces the most favorable expected execution price.

This is a realistic curve for many situations, but for concreteness, the above graph is a plot of a scenario where the task is to acquire 10,000 shares of MSFT within one hour.

On the x-axis, we plot where the limit order stands relative 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 + 10 – 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 left and right of the graph correspond to Strategy 1 and Strategy 2 respectively: on the left, we have orders with high prices that get executed immediately and completely; on the right, 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 peaked shape in the middle of the graph supports our main thesis – superior execution price can be achieved by using carefully tuned limit orders. The maximum of the curve 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 1 should be interpreted as follows: if you want to acquire 10,000 shares over 60 minutes and seek the most attractive *expected* execution price, you should submit a limit order at the best bid plus 5 cents. We discuss the shape and behavior of these curves in great detail in Section V.

Notice that the entirety of our return 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, 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 risk of non-execution, mid-spread volatility, and volume volatility are all slightly different concepts, we study them jointly by defining risk as the standard deviation of returns.

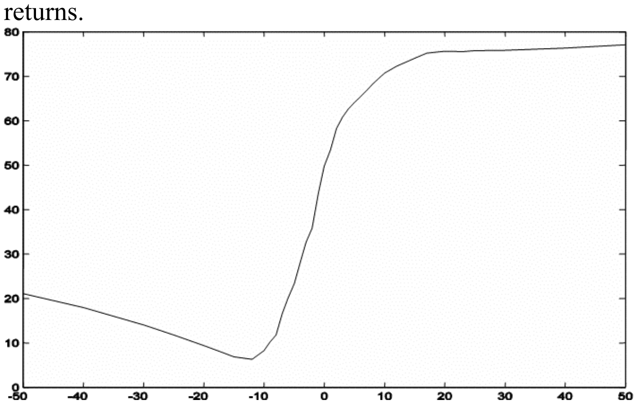

Figure 2. Returns become more uncertain the further we move from the inside market.

A risk profile that corresponds to the returns curve from Figure 1 is shown in Figure 2. We simply plotted the standard deviation of returns (y-axis) – which are averaged in Figure 1 – 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.

This shape is partly a result of our choice of 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 (e.g. 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 quasi-guaranteed, but the transaction price is now capped at some more reasonable level.

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.

In order to perform a meaningful comparison among alternative strategies, we borrowed a popular tool from the classic Finance Theory – the 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 its standard deviation, and the y-axis its returns. By connecting all the strategies together, we get the plot presented in Figure 3. This profile is a combination of results from Figures 1 and 2. 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].

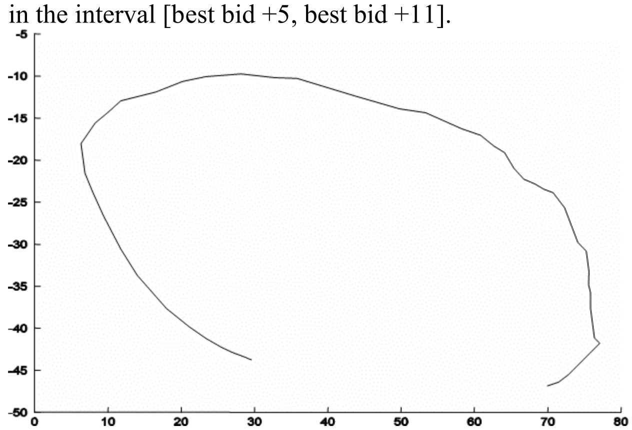

Figure 3. Trade-off between risk and return. Only the top portion where higher risk results in higher returns should be used for execution.

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 expectation 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.

V. Results.

In this section we present a summary of our results. 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.

A. Order Size.

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 4 (last page) illustrates this point. All trading is done within one-hour period; solid line represents 1,000 shares, dashed line – 5,000 shares, dash-dot line – 10,000 shares. 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 graph in Figure 4 – 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. While it is clear that trading large volumes is costly, it is difficult to propose a definite 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.

B. Time Window.

If we are to divide a large order into multiple small orders, we must reduce the execution window for each smaller piece. This effect is explored in Figure 5. Trading here is performed over 60 minutes (solid line), 10 minutes (dashed line), and 1 minute (dash-dot line); 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. 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 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). Final observation: similar to the order size, shorter execution time necessitates more aggressive order pricing.

C. Time of the Day.

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 documented: 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 6 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 6 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.

When we trade during a 10-minute window, however, all curves are much closer together, which makes time of the day effect much less pronounced. 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.

D. Market Conditions.

The real strength of our approach is presented in Figures 7 and 8, 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 7, 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 riskier. 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 8, we want to see how our results change when we submit our orders into "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. 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.

VI. Generalization and Future Directions.

While we spent the bulk of this paper laying out our basic methodology, it is important to remember that this analysis is only a building block, which can be expanded considerably. In all the experiments we have conducted, our trading strategies remained static – we submit a limit order at the beginning of the trial interval and then do not touch them until the last second. Such approach to trading is certainly unrealistic, but our findings can be extended to a more complex framework.

Say we have an hour to acquire 1000 shares of MSFT. We can look at the results of our experiments and figure out that, for example, the optimal strategy is to submit a limit order at the best bid minus one cent; which is what we do. But, as the time goes by, conditions may change: the inside market may move far away from our standing order, or our order may get partially executed, or the spread may widen, etc. But as all these events happen, at any point after the initial submission we have an option to re-evaluate our pricing strategy. We can use the shortened time horizon and updated order size to go through the same analysis as before, and then come up with a new optimal price and re-submit our order. If we perform this operation continuously, we can ensure that at every point in time we are trading at the lowest expected cost.

One way to envision this process is through a large look-up table as shown in Table 1. This matrix specifies the optimal order price for every possible combination of order size and time to execution. The table can be filled following the analysis from Section IV. The number of entries in Table 1 is likely exaggerated, as it can be built with an arbitrary resolution – i.e. sampling every 10 minutes and every 100 shares probably makes more sense.

	1000 s.	999 s.	…	1 share
60 min	Bid-10	Bid-11	…	Bid-1
59 min	Bid-9	Bid-11	…	Bid -1
…	…	…	…	…
1 min	Bid+100	Bid+99	…	Bid+1
0 min	market	market	…	market

Table 1. Look-up table for optimal price updates.

While we can adopt this data-intensive approach and run a large number of experiments to construct the look-up table exactly as described above, this is not very practical. First, it will take a significant amount of time, processing power, and raw data to create such table; then, the resulting database will be very large and difficult to maintain and update; and, finally, most information will be repetitive and thus redundant because of the similarities in trading characteristics across stocks. Therefore, the most promising development direction is to try to come up with a functional description of the optimal order placement strategy.

So far, we have been deriving our optimal price by visually determining the peak of the curve. Through our experiments in Section V, we have already gained some key insights: the limit price of the optimal order is directly related to the size of the and inversely related to the time to execution, just to mention a few dependencies. There are other factors that clearly must be considered: bid-ask spread, liquidity, volume, volume volatility, price volatility, time of the day, etc. We can run a regression using these inputs and having the optimal limit price as output. From the technical standpoint, this essentially amounts to finding coefficients for the following equation:

Limit distance = α_1*Volume + α_2*Volatility + α_3*Spread + α_4*Size + α_5*Time + …

The benefit of this approach is that we don't have to use the entire universe of stocks to come up with an accurate model; and when we know the coefficients, we can revise our estimates of the optimal price on-line. The main challenge of doing this derivation is that it will require even more data analysis to be performed on a very large dataset – we will have to quantify and keep track of spreads, volumes, price volatility, volume volatility, etc. This can be a very data-intensive and challenging process.

And, finally, to achieve even more precision and cost savings, we can introduce conditional optimality. This is essentially an extension of the conditioning experiments from Section V: in practice, we want to transact differently under different market conditions. For example, in a market with a pronounced trend we want to trade more aggressively than when a stock is essentially flat. Therefore, all our optimizations should be performed for the market conditions in which we are planning on executing. When we are building a look-up table as in Table 1 learning optimal prices for different times and sizes, we can either use all the available data, thus getting the "unconditionally optimal" prices, or we can use only those pieces of data that conform to our desired conditions. If we know beforehand that we will be transacting in a rising market, then we should use only the data from the prior rising markets, and thus learning optimal limit order prices conditioned on the uptrend in the stock. This should render our trading much more precise and efficient.

Conclusion.

In this paper, we propose a limit order book approach to efficient execution. We demonstrated how to estimate return curves, risk curves, and risk-return profiles by using historical data, and how to derive optimal pricing frontiers. Our quantitative method allows traders to optimally price their limit orders in order to minimize trading costs and control corresponding risks. Through many experiments we have studied the behavior of risk and return in this domain, highlighting the importance of a number of microstructure variables: order size, time window, liquidity, etc. And, finally, we suggested several ways how our methodology can be taken to the next level: continuous order revision, functional description of price curves, conditional optimization, and others.

We believe that our main contribution is making the first step towards a very important task of optimizing trade execution in order-driven markets. We introduce precision in the process of order submission, and more specifically, optimal limit order price determination. We emphasize, however, that our method is just a building block, which can be turned into much more significant research projects. In this regard, we remark that we are currently engaged in a large-scale application of the methods of reinforcement learning to optimized execution along the lines discussed in Section VI.

Bibliography.

1. Almgren, R. and Chriss, N., *Value Under Liquidation*, Risk, 1999.
2. Almgren, R. and Chriss, N., *Optimal Execution of Portfolio Transactions*, Journal of Risk, 2002.
3. Bertsimas, D. and Lo, A., *Optimal Control of Execution Costs*, Journal of Financial Markets 1, 1998.
4. Bias, B., Hilton, P., and Spatt, C., *An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse*, Journal of Finance, 50, 1995.
5. Handa, P. and Schwartz, R., *Limit Order Trading*, Journal of Finance 51, 1835-1861, 1996.
6. Harris, L., *Trading and Exchanges: Market Microstructure for Professionals*, Oxford University Press, October 2002.
7. Hasbrouck, J., and Harris, L., *Market vs. Limit Orders: The SuperDot Evidence on Order Submission Strategy*, Journal of Financial and Quantitative Analysis, 31, 1996.
8. Hollifield, B., Miller, R., Sandas, P., *Empirical Analysis of Limit Order Markets*, working paper, Carnegie Mellon University, 2003.
9. Nevmyvaka, Y., Kearns, M., Papandreou, A., and Sycara, K., 2004, *Optimized Execution Via Order Book Simulation*, working paper, CMU.
10. Kissell, R., and Glantz, M., *Optimal Trading Strategies: Quantitative Approaches For Managing Market Impact and Trading Risk*, AMACOM, 2003.
11. Lo, A., MacKinlay, C., Zhang, J., 2000, *Econometric Models of Limit-Order Executions*, Journal of Financial Economics, 65, 2002.
12. Markowitz, H., *Portfolio Selection*, Journal of Finance, 7 (1), 77-91, 1952.

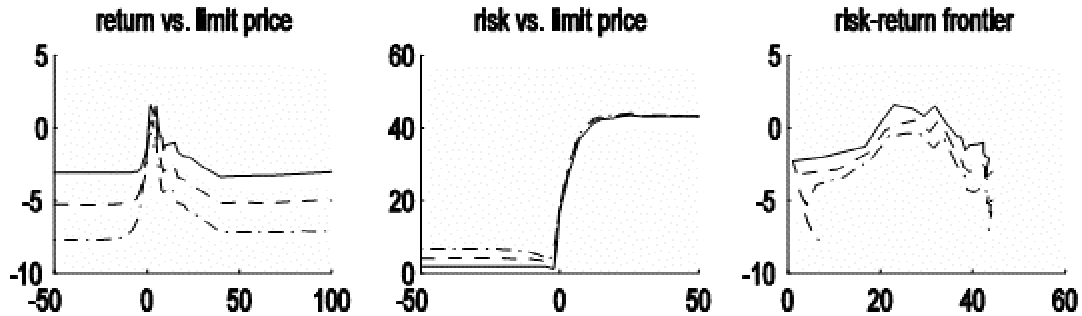

Figure 4. Transacting larger orders is both more expensive and riskier.

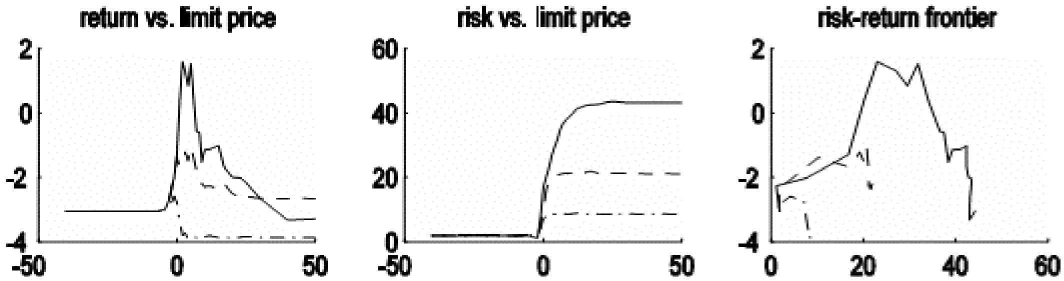

Figure 5. Execution over shorter time periods can be more expensive, but less risky.

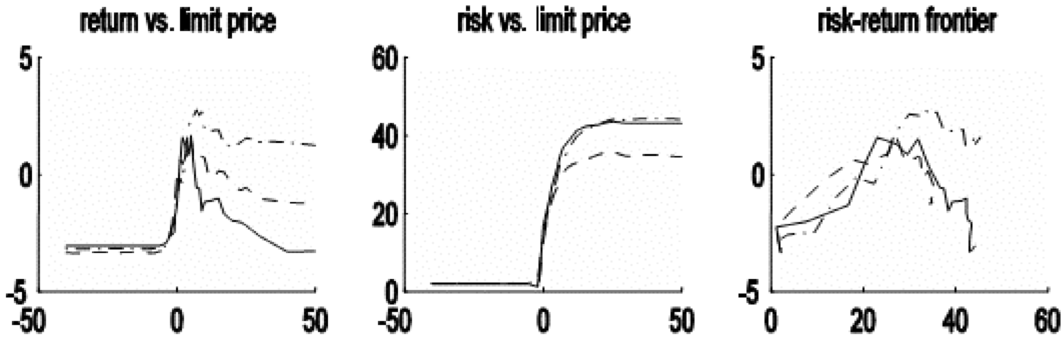

Figure 6. When executing over 60 minute period, time of the day should be used as one of the model's inputs.

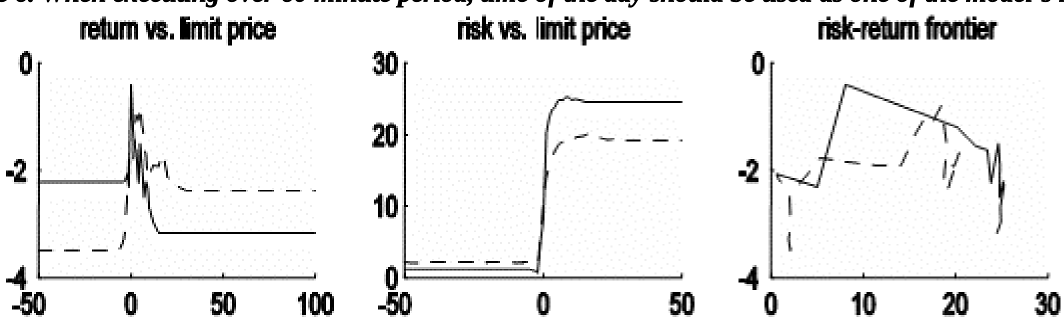

Figure 7. Transacting on a high-volume day is less expensive and less risky.

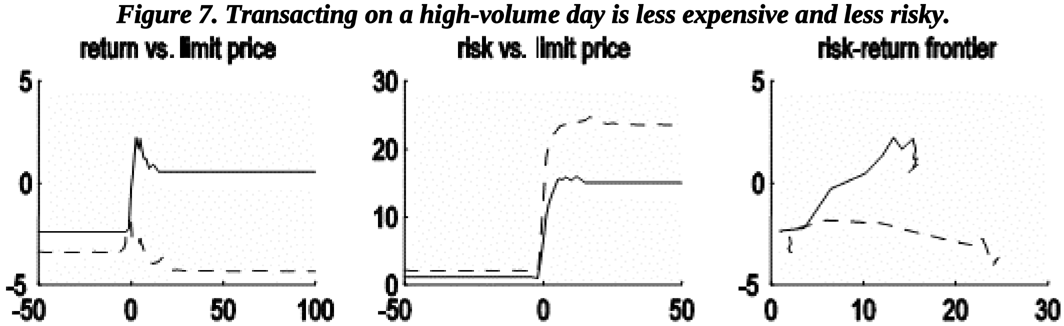

Figure 8. "Thick" book may be preferable to "thin" book.