

(In)Stability Properties of Limit Order Dynamics

Eyal Even-Dar *

Sham M. Kakade †

Michael Kearns ‡

Yishay Mansour §

ABSTRACT

We study the stability properties of the dynamics of the standard continuous limit-order mechanism that is used in modern equity markets. We ask whether such mechanisms are susceptible to “butterfly effects” — the infliction of large changes on common measures of market activity by only small perturbations of the order sequence. We show that the answer depends strongly on whether the market consists of “absolute” traders (who determine their prices independent of the current order book state) or “relative” traders (who determine their prices relative to the current bid and ask). We prove that while the absolute trader model enjoys provably strong stability properties, the relative trader model is vulnerable to great instability. Our theoretical results are supported by large-scale experiments using limit order data from INET, a large electronic exchange for NASDAQ stocks.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Theory

Keywords

Computational Finance, Market Microstructure

*Department of Computer and Information Science, University of Pennsylvania, Philadelphia PA. Email address: even-dar@seas.upenn.edu

†Toyota Technological Institute, Chicago IL. Email address: sham@tti-c.org

‡Department of Computer and Information Science, University of Pennsylvania, Philadelphia PA. Email address: mkearns@cis.upenn.edu

§School of Computer Science, Tel Aviv University, Tel Aviv, Israel. Email address: mansour@cs.tau.ac.il

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'06, June 11–15, 2006, Ann Arbor, Michigan.

Copyright 2006 ACM 1-59593-236-4/06/0006 ...\$5.00.

1. INTRODUCTION

In recent years there has been an explosive increase in the automation of modern equity markets. This increase has taken place both in the exchanges, which are increasingly computerized and offer sophisticated interfaces for order placement and management, and in the trading activity itself, which is ever more frequently undertaken by software. The so-called Electronic Communication Networks (or ECNs) that dominate trading in NASDAQ stocks are a common example of the automation of the exchanges. On the trading side, computer programs now are entrusted not only with the careful execution of large block trades for clients (sometimes referred to on Wall Street as program trading), but with the autonomous selection of stocks, direction (long or short) and volumes to trade for profit (commonly referred to as statistical arbitrage).

The vast majority of equity trading is done via the standard limit order market mechanism. In this mechanism, continuous trading takes place via the arrival of limit orders specifying whether the party wishes to buy or sell, the volume desired, and the price offered. Arriving limit orders that are entirely or partially executable with the best offers on the other side are executed immediately, with any volume not immediately executable being placed in a queue (or book) ordered by price on the appropriate side (buy or sell). (A detailed description of the limit order mechanism is given in Section 3.) While traders have always been able to view the prices at the top of the buy and sell books (known as the *bid* and *ask*), a relatively recent development in certain exchanges is the real-time revelation of the *entire* order book — the complete distribution of orders, prices and volumes on both sides of the exchange. With this revelation has come the opportunity — and increasingly, the need — for modeling and exploiting limit order data and dynamics. It is fair to say that market microstructure, as this area is generally known, is a topic commanding great interest both in the real markets and in the academic finance literature. The opportunities and needs span the range from the optimized execution of large trades to the creation of stand-alone “proprietary” strategies that attempt to profit from high-frequency microstructure signals.

In this paper we investigate a previously unexplored but fundamental aspect of limit order microstructure: the *stability* properties of the dynamics. Specifically, we are interested in the following natural question: To what extent are simple models of limit order markets either susceptible or immune to “butterfly effects” — that is, the infliction of large changes in important activity statistics (such as the

number of shares traded or the average price per share) by only minor perturbations of the order sequence?

To examine this question, we consider two stylized but natural models of the limit order arrival process. In the *absolute price* model, buyers and sellers arrive with limit order prices that are determined *independently* of the current state of the market (as represented by the order books), though they may depend on all manner of exogenous information or shocks, such as time, news events, announcements from the company whose shares are being traded, private signals or state of the individual traders, etc. This process models traditional “fundamentals”-based trading, in which market participants each have some inherent but possibly varying valuation for the good that in turn determines their limit price.

In contrast, in the *relative price* model, traders express their limit order prices *relative* to the best price offered in their respective book (buy or sell). Thus, a buyer would encode their limit order price as an offset Δ (which may be positive, negative, or zero) from the current bid p_b , which is then translated to the limit price $p_b + \Delta$. Again, in addition to now depending on the state of the order books, prices may also depend on all manner of exogenous information. The relative price model can be viewed as modeling traders who, in addition to perhaps incorporating fundamental external information on the stock, may also position their orders strategically relative to the other orders on their side of the book. A common example of such strategic behavior is known as “penny-jumping” on Wall Street, in which a trader who has an interest in buying shares quickly, but still at a discount to placing a market order, will deliberately position their order just above the current bid. More generally, the entire area of modern execution optimization [9, 10, 8] has come to rely heavily on the careful positioning of limit orders relative to the current order book state. Note that such positioning may depend on more complex features of the order books than just the current bid and ask, but the relative model is a natural and simplified starting point. We remark that an alternate view of the two models is that *all* traders behave in a relative manner, but with “absolute” traders able to act only on a considerably slower time scale than the faster “relative” traders.

How do these two models differ? Clearly, given any fixed sequence of arriving limit order prices, we can choose to express these prices either as their original (absolute) values, or we can run the order book dynamical process and transform each order into a relative difference with the top of its book, and obtain identical results. The differences arise when we consider the stability question introduced above. Intuitively, in the absolute model a small perturbation in the arriving limit price sequence *should* have limited (but still some) effects on the subsequent evolution of the order books, since prices are determined independently. For the relative model this intuition is less clear. It seems possible that a small perturbation could (for example) slightly modify the current bid, which in turn could slightly modify the price of the next arriving order, which could then slightly modify the price of the subsequent order, and so on, leading to an amplifying sequence of events.

Our main results demonstrate that these two models do indeed have dramatically different stability properties. We first show that for *any* fixed sequence of prices in the absolute model, the modification of a single order has a bounded

and extremely limited impact on the subsequent evolution of the books. In particular, we define a natural notion of distance between order books and show that small modifications can result in only constant distance to the original books for all subsequent time steps. We then show that this implies that for almost any standard statistic of market activity — the executed volume, the average price execution price, and many others — the statistic can be influenced only infinitesimally by small perturbations.

In contrast, we show that the relative model enjoys no such stability properties. After giving specific (worst-case) relative price sequences in which small perturbations generate large changes in basic statistics (for example, altering the number of shares traded by a factor of two), we proceed to demonstrate that the difference in stability properties of the two models is more than merely theoretical. Using extensive INET (a major ECN for NASDAQ stocks) limit order data and order book reconstruction code, we investigate the empirical stability properties when the data is interpreted as containing either absolute prices, relative prices, or mixtures of the two. The theoretical predictions of stability and instability are strongly borne out by the subsequent experiments.

In addition to stability being of fundamental interest in any important dynamical system, we believe that the results described here provide food for thought on the topics of market impact and the “backtesting” of quantitative trading strategies (the attempt to determine hypothetical past performance using historical data). They suggest that one’s confidence that trading “quietly” and in small volumes will have minimal market impact is linked to an implicit belief in an absolute price model. Our results and the fact that in the real markets there is a large and increasing amount of relative behavior such as penny-jumping would seem to cast doubts on such beliefs. Similarly, in a purely or largely relative-price world, backtesting even low-frequency, low-volume strategies could result in historical estimates of performance that are not only unrelated to future performance (the usual concern), but are not even accurate measures of a hypothetical past.

The outline of the paper follows. In Section 2 we briefly review the large literature on market microstructure. In Section 3 we describe the limit order mechanism and our formal models. Section 4 presents our most important theoretical results, the 1-Modification Theorem for the absolute price model. This theorem is applied in Section 5 to derive a number of strong stability properties in the absolute model. Section 6 presents specific examples establishing the worst-case instability of the relative model. Section 7 contains the simulation studies that largely confirm our theoretical findings on INET market data.

2. RELATED WORK

As was mentioned in the Introduction, market microstructure is an important and timely topic both in academic finance and on Wall Street, and consequently has a large and varied recent literature. Here we have space only to summarize the main themes of this literature and to provide pointers to further readings. To our knowledge the stability properties of detailed limit order microstructure dynamics have not been previously considered. (However, see Farmer and Joshi [6] for an example and survey of other price dynamic stability studies.)

On the more theoretical side, there is a rich line of work examining what might be considered the game-theoretic properties of limit order markets. These works model traders and market-makers (who provide liquidity by offering both buy and sell quotes, and profit on the difference) by utility functions incorporating tolerance for risks of price movement, large positions and other factors, and examine the resulting equilibrium prices and behaviors. Common findings predict negative price impacts for large trades, and price effects for large inventory holdings by market-makers. An excellent and comprehensive survey of results in this area can be found in [2].

There is a similarly large body of empirical work on microstructure. Major themes include the measurement of price impacts, statistical properties of limit order books, and attempts to establish the informational value of order books [4]. A good overview of the empirical work can be found in [7]. Of particular note for our interests is [3], which empirically studies the distribution of arriving limit order prices in several prominent markets. This work takes a view of arriving prices analogous to our relative model, and establishes a power-law form for the resulting distributions.

There is also a small but growing number of works examining market microstructure topics from a computer science perspective, including some focused on the use of microstructure in algorithms for optimized trade execution. Kakade et al. [9] introduced limit order dynamics in competitive analysis for one-way and volume-weighted average price (VWAP) trading. Some recent papers have applied reinforcement learning methods to trade execution using order book properties as state variables [1, 5, 10].

3. MICROSTRUCTURE PRELIMINARIES

The following expository background material is adapted from [9]. The market mechanism we examine in this paper is driven by the simple and standard concept of a *limit order*. Suppose we wish to purchase 1000 shares of Microsoft (MSFT) stock. In a limit order, we specify not only the desired volume (1000 shares), but also the desired price. Suppose that MSFT is currently trading at roughly \$24.07 a share (see Figure 1, which shows an actual snapshot of an MSFT order book on INET), but we are only willing to buy the 1000 shares at \$24.04 a share or lower. We can choose to submit a limit order with this specification, and our order will be placed in a queue called the *buy order book*, which is ordered by price, with the highest offered unexecuted buy price at the top (often referred to as the *bid*). If there are multiple limit orders at the same price, they are ordered by time of arrival (with older orders higher in the book). In the example provided by Figure 1, our order would be placed immediately after the extant order for 5,503 shares at \$24.04; though we offer the same price, this order has arrived before ours. Similarly, a sell order book for sell limit orders is maintained, this time with the lowest sell price offered (often referred to as the *ask*) at its top.

Thus, the order books are sorted from the most competitive limit orders at the top (high buy prices and low sell prices) down to less competitive limit orders. The bid and ask prices together are sometimes referred to as the *inside market*, and the difference between them as the *spread*. By definition, the order books always consist exclusively of *unexecuted* orders — they are queues of orders hopefully waiting for the price to move in their direction.



Figure 1: Sample INET order books for MSFT.

How then do orders get (partially) executed? If a buy (sell, respectively) limit order comes in above the ask (below the bid, respectively) price, then the order is matched with orders on the opposing books until either the incoming order’s volume is filled, or no further matching is possible, in which case the remaining incoming volume is placed in the books.

For instance, suppose in the example of Figure 1 a buy order for 2000 shares arrived with a limit price of \$24.08. This order would be partially filled by the two 500-share sell orders at \$24.069 in the sell books, the 500-share sell order at \$24.07, and the 200-share sell order at \$24.08, for a total of 1700 shares executed. The remaining 300 shares of the incoming buy order would become the new bid of the buy book at \$24.08. It is important to note that *the prices of executions are the prices specified in the limit orders already in the books*, not the prices of the incoming order that is immediately executed. Thus in this example, the 1700 executed shares would be at different prices. Note that this also means that in a pure limit order exchange such as INET, market orders can be “simulated” by limit orders with extreme price values. In exchanges such as INET, any order can be withdrawn or canceled by the party that placed it any time prior to execution.

Every limit order arrives atomically and instantaneously — there is a strict temporal sequence in which orders arrive, and two orders can never arrive simultaneously. This gives rise to the definition of the *last price* of the exchange, which is simply the last price at which the exchange executed an order. It is this quantity that is usually meant when people casually refer to the (ticker) price of a stock.

3.1 Formal Definitions

We now provide a formal model for the limit order pro-

cess described above. In this model, limit orders arrive in a temporal sequence, with each order specifying its limit price and an indication of its type (buy or sell). Like the actual exchanges, we also allow cancellation of a standing (unexecuted) order in the books any time prior to its execution. Without loss of generality we limit attention to a model in which every order is for a single share; large order volumes can be represented by 1-share sequences.

DEFINITION 3.1. Let $\Sigma = \langle \sigma_1, \dots, \sigma_n \rangle$ be a sequence of limit orders, where each σ_i has the form $\langle n_i, t_i, v_i \rangle$. Here n_i is an order identifier, t_i is the order type (buy, sell, or cancel), and v_i is the limit order value. In the case that t_i is a cancel, n_i matches a previously placed order and v_i is ignored.

We have deliberately called v_i in the definition above the limit order value rather than price, since our two models will differ in their interpretation of v_i (as being absolute or relative). In the absolute model, we do indeed interpret v_i as simply being the price of the limit order. In the relative model, if the current order book configuration is (A, B) (where A is the sell and B the buy book), the price of the order is $\text{ask}(A) + v_i$ if t_i is sell, and $\text{bid}(B) + v_i$ if t_i is buy, where by $\text{ask}(X)$ and $\text{bid}(X)$ we denote the price of the order at the top of the book X . (Note v_i can be negative.)

Our main interest in this paper is the effects that the modification of a small number of limit orders can have on the resulting dynamics. For simplicity we consider only modifications to the limit order values, but our results generalize to any modification.

DEFINITION 3.2. A k -modification of Σ is a sequence Σ' such that for exactly k indices i_1, \dots, i_k $v_{i_j} \neq v'_{i_j}$, $t_{i_j} = t'_{i_j}$, and $n_{i_j} = n'_{i_j}$. For every $\ell \neq i_j$, $j \in \{1, \dots, k\}$ $\sigma_\ell = \sigma'_\ell$.

We now define the various quantities whose stability properties we examine in the absolute and relative models. All of these are standard quantities of common interest in financial markets.

- $\text{volume}(\Sigma)$: Number of shares executed (traded) in the sequence Σ .
- $\text{average}(\Sigma)$: Average execution price.
- $\text{close}(\Sigma)$: Price of the last (closing) execution.
- $\text{lastbid}(\Sigma)$: Bid at the end of the sequence.
- $\text{lastask}(\Sigma)$: Ask at end of the sequence.

4. THE 1-MODIFICATION THEOREM

In this section we provide our most important technical result. It shows that in the absolute model, the effects that the modification of a single order has on the resulting evolution of the order books is extremely limited. We then apply this result to derive strong stability results for all of the aforementioned quantities in the absolute model.

Throughout this section, we consider an arbitrary order sequence Σ in the absolute model, and any 1-modification Σ' of Σ . At any point (index) i in the two sequences we shall use (A_1, B_1) to denote the sell and buy books (respectively) in Σ , and (A_2, B_2) to denote the sell and buy books in Σ' ; for notational convenience we omit explicitly superscripting by the current index i . We will shortly establish that at all times i , (A_1, B_1) and (A_2, B_2) are very “close”.

Although the order books are sorted by price, we will use (for example) $A_1 \cup \{a_2\} = A_2$ to indicate that A_2 contains

an order at some price a_2 that is not present in A_1 , but that otherwise A_1 and A_2 are identical; thus deleting the order at a_2 in A_2 would render the books the same. Similarly, $B_1 \cup \{b_2\} = B_2 \cup \{b_1\}$ means B_1 contains an order at price b_1 not present in B_2 , B_2 contains an order at price b_2 not present in B_1 , and that otherwise B_1 and B_2 are identical. Using this notation, we now define a set of *stable* system states, where each state is composed from the order books of the original and the modified sequences. Shortly we show that if we change only one order’s value (price), we remain in this set for any sequence of limit orders.

DEFINITION 4.1. Let ab be the set of all states (A_1, B_1) and (A_2, B_2) such that $A_1 = A_2$ and $B_1 = B_2$. Let $\bar{a}b$ be the set of states such that $A_1 \cup \{a_2\} = A_2 \cup \{a_1\}$, where $a_1 \neq a_2$, and $B_1 = B_2$. Let $a\bar{b}$ be the set of states such that $B_1 \cup \{b_2\} = B_2 \cup \{b_1\}$, where $b_1 \neq b_2$, and $A_1 = A_2$. Let $\bar{a}\bar{b}$ be the set of states in which $A_1 = A_2 \cup \{a_1\}$ and $B_1 = B_2 \cup \{b_1\}$, or in which $A_2 = A_1 \cup \{a_2\}$ and $B_2 = B_1 \cup \{b_2\}$. Finally we define $S = ab \cup \bar{a}b \cup a\bar{b} \cup \bar{a}\bar{b}$ as the set of stable states.

THEOREM 4.1. (1-Modification Theorem) Consider any sequence of orders Σ and any 1-modification Σ' of Σ . Then the order books (A_1, B_1) and (A_2, B_2) determined by Σ and Σ' lie in the set S of stable states at all times.




Figure 2: Diagram representing the set S of stable states and the possible movements transitions in it after the change.

The idea of the proof of this theorem is contained in Figure 2, which shows a state transition diagram labeled by the categories of stable states. This diagram describes all transitions that can take place *after* the arrival of the order on which Σ and Σ' differ. The following establishes that immediately after the arrival of this differing order, the state lies in S .

LEMMA 4.2. If at any time the current books (A_1, B_1) and (A_2, B_2) are in the set ab (and thus identical), then modifying the price of the next order keeps the state in S .

PROOF. Suppose the arriving order is a sell order and we change it from a_1 to a_2 ; assume without loss of generality that $a_1 > a_2$. If neither order is executed immediately, then we move to state $\bar{a}b$; if both of them are executed then we stay in state ab ; and if only a_2 is executed then we move to state $\bar{a}\bar{b}$. The analysis of an arriving buy order is similar. \square

Following the arrival of their only differing order, Σ and Σ' are identical. We now give a sequence of lemmas showing




Figure 3: The state diagram when starting at state $\bar{a}b$. This diagram provides the intuition of Lemma 4.3

that following the initial difference covered by Lemma 4.2, the state remains in S forever on the remaining (identical) sequence. We first show that from state $\bar{a}b$ we remain in S regardless the next order. The intuition of this lemma is demonstrated in Figure 3.

LEMMA 4.3. *If the current state is in the set $\bar{a}b$, then for any order the state will remain in S .*

PROOF. We first provide the analysis for the case of an arriving sell order. Note that in $\bar{a}b$ the buy books are identical ($B_1 = B_2$). Thus either the arriving sell order is executed with the same buy order in both buy books, or it is not executed in both buy books. For the first case, the buy books remain identical (the bid is executed in both) and the sell books remain unchanged. For the second case, the buy books remain unchanged and identical, and the sell books have the new sell order added to both of them (and thus still differ by one order).

Next we provide an analysis of the more subtle case where the arriving item is a buy order. For this case we need to take care of several different scenarios. The first is when the top of both sell books (the ask) is identical. Then regardless of whether the new buy order is executed or not, the state remains in $\bar{a}b$ (the analysis is similar to an arriving sell order).

We are left to deal with case where $ask(A_1)$ and $ask(A_2)$ are different. Here we discuss two subcases: (a) $ask(A_1) = a_1$ and $ask(A_2) = a_2$, and (b) $ask(A_1) = a_1$ and $ask(A_2) = a'$. Here a_1 and a_2 are as in the definition of $\bar{a}b$ in Definition 4.1, and a' is some other price. For subcase (a), by our assumption $a_1 < a_2$, then either (1) both asks get executed, the sell books become identical, and we move to state ab ; (2) neither ask is executed and we remain in state $\bar{a}b$; or (3) only $ask(A_1) = a_1$ is executed, in which case we move to state $\bar{a}\bar{b}$ with $A_2 = A_1 \cup \{a_2\}$ and $B_2 = B_1 \cup \{b_2\}$, where b_2 is the arriving buy order price. For subcase (b), either (1) buy order is executed in neither sell book we remain in state $\bar{a}b$; or (2) the buy order is executed in both sell books and stay in state $\bar{a}b$ with $A_1 \cup \{a'\} = A_2 \cup \{a_2\}$; or (3) only $ask(A_1) = a_1$ is executed and we move to state $\bar{a}\bar{b}$. \square

LEMMA 4.4. *If the current state is in the set $\bar{a}\bar{b}$, then for any order the state will remain in S .*

LEMMA 4.5. *If the current configuration is in the set $\bar{a}\bar{b}$, then for any order the state will remain in S .*

The proofs of these two lemmas are omitted, but are similar in spirit to that of Lemma 4.3. The next and final lemma deals with cancellations.

LEMMA 4.6. *If the current order book state lies in S , then following the arrival of a cancellation it remains in S .*

PROOF. When a cancellation order arrives, one of the following possibilities holds: (1) the order is still in both sets of books, (2) it is not in either of them and (3) it is only in one of them. For the first two cases it is easy to see that the cancellation effect is identical on both sets of books, and thus the state remains unchanged. For the case when the order appears only in one set of books, without loss of generality we assume that the cancellation cancels a buy order at b_1 . Rather than removing b_1 from the book we can change it to have price 0, meaning this buy order will never be executed and is effectively canceled. Now regardless the state that we were in, b_1 is still only in one buy book (but with a different price), and thus we remain in the same state in S . \square

The proof of Theorem 4.1 follows from the above lemmas.

5. ABSOLUTE MODEL STABILITY

In this section we apply the 1-Modification Theorem to show strong stability properties for the absolute model. We begin with an examination of the executed volume.

LEMMA 5.1. *Let Σ be any sequence and Σ' be any 1-modification of Σ . Then the set of the executed orders (ID numbers) generated by the two sequences differs by at most 2.*

PROOF. By Theorem 4.1 we know that at each stage the books differ by at most two orders. Now since the union of the IDs of the executed orders and the order books is always identical for both sequences, this implies that the executed orders can differ by at most two. \square

COROLLARY 5.2. *Let Σ be any sequence and Σ' be any k -modification of Σ . Then the set of the executed orders (ID numbers) generated by the two sequences differs by at most $2k$.*

An order sequence Σ' is a k -extension of Σ if Σ can be obtained by deleting any k orders in Σ' .

LEMMA 5.3. *Let Σ be any sequence and let Σ' be any k -extension of Σ . Then the set of the executed orders generated by Σ and Σ' differ by at most $2k$.*

This lemma is the key to obtain our main absolute model volume result below. We use $edit(\Sigma, \Sigma')$ to denote the standard edit distance between the sequences Σ and Σ' — the minimal number of substitutions, insertions and deletions or orders needed to change Σ to Σ' .

THEOREM 5.4. *Let Σ and Σ' be any absolute model order sequences. Then if $edit(\Sigma, \Sigma') \leq k$, the set of the executed orders generated by Σ and Σ' differ by at most $4k$. In particular, $|volume(\Sigma) - volume(\Sigma')| \leq 4k$.*

PROOF. We first define the sequence $\tilde{\Sigma}$ which is the intersection of Σ and Σ' . Since Σ and Σ' are at most k apart, we have that by k insertions we change $\tilde{\Sigma}$ to either Σ or Σ' , and by Lemma 5.3 its set of executed orders is at most $2k$ from each. Thus the set of executed orders in Σ and Σ' is at most $4k$ apart. \square

5.1 Spread Bounds

Theorem 5.4 establishes strong stability for executed volume in the absolute model. We now turn to the quantities that involve execution prices as opposed to volume alone — namely, $average(\Sigma)$, $close(\Sigma)$, $lastbid(\Sigma)$ and $lastask(\Sigma)$. For these results, unlike executed volume, a condition must hold on Σ in order for stability to occur. This condition is expressed in terms of a natural measure of the *spread* of the market, or the gap between the buyers and sellers. We motivate this condition by first showing that without it, by changing one order, we can change $average(\Sigma)$ by any positive value x .

LEMMA 5.5. *There exists Σ such that for any $x \geq 0$, there is a 1-modification Σ' of Σ such that $average(\Sigma') = average(\Sigma) + x$.*

PROOF. Let Σ be a sequence of alternating sell and buy orders in which each seller offers p and each buyer $p + x$, and the first order is a sell. Then all executions take place at the ask, which is always p , and thus $average(\Sigma) = p$. Now suppose we modify only the first sell order to be at price $p + 1 + x$. This initial sell order will never be executed, and now all executions take place at the bid, which is always $p + x$. \square

Similar instability results can be shown to hold for the other price-based quantities. This motivates the introduction of a quantity we call the *second spread* of the order books, which is defined as the difference between the prices of the second order in the sell book and the second order in the buy book (as opposed to the bid-ask difference, which is commonly called the spread). We note that in a liquid stock, such as those we examine experimentally in Section 7, the second spread will typically be quite small and in fact almost always equal to the spread.

In this subsection we consider changes in the sequence only after an initialization period, and sequences such that the second spread is *always* defined after the time we make a change. We define $s_2(\Sigma)$ to be the maximum second spread in the sequence Σ following the change.

THEOREM 5.6. *Let Σ be a sequence and let Σ' be any 1-modification of Σ . Then*

1. $|lastbid(\Sigma) - lastbid(\Sigma')| \leq s_2(\Sigma)$
2. $|lastask(\Sigma) - lastask(\Sigma')| \leq s_2(\Sigma)$

where $s_2(\Sigma)$ is the maximum over the second spread in Σ following the 1-modification.

PROOF. We provide the proof for the last bid; the proof for the last ask is similar. The proof relies on Theorem 4.1 and considers states in the stable set S . For states ab and $\bar{a}b$, we have that the bid is identical. Let $bid(X)$, $sb(X)$, $ask(X)$, be the bid, the second highest buy order, and the ask of a sequence X . Now recall that in state $\bar{a}b$ we have that the sell books are identical, and that the two buy books are identical except one different order. Thus

$$bid(\Sigma) + s_2(\Sigma) \geq sb(\Sigma) + s_2(\Sigma) \geq ask(\Sigma) = ask(\Sigma') \geq bid(\Sigma').$$

Now it remains to bound $bid(\Sigma)$. Here we use the fact that the bid of the modified sequence is at least the second highest buy order in the original sequence, due to the fact that the books are different only in one order. Since

$$bid(\Sigma') \geq sb(\Sigma) \geq ask(\Sigma) - s_2(\Sigma) \geq bid(\Sigma) - s_2(\Sigma)$$

we have that $|bid(\Sigma) - bid(\Sigma')| \leq s_2(\Sigma)$ as desired.

In state $\bar{a}b$ we have that for one sequence the books contain an additional buy order and an additional sell order. First suppose that the books containing the additional orders are the original sequence Σ . Now if the bid is not the additional order we are done, otherwise we have the following:

$$bid(\Sigma) \leq ask(\Sigma) \leq sb(\Sigma) + s_2(\Sigma) = bid(\Sigma') + s_2(\Sigma),$$

where $sb(\Sigma) \leq bid(\Sigma')$ since the original buy book has only one additional order.

Now assume that the books with the additional orders are for the modified sequence Σ' . We have

$$bid(\Sigma) + s_2(\Sigma) \geq ask(\Sigma) \geq ask(\Sigma') \geq bid(\Sigma'),$$

where we used the fact that $ask(\Sigma) \geq ask(\Sigma')$ since the modified sequence has an additional order. Similarly we have that $bid(\Sigma) \leq bid(\Sigma')$ since the modified buy book contains an additional order. \square

We note that the proof of Theorem 5.6 actually establishes that the bid and ask of the original and modified sequences are within $s_2(\Sigma)$ at *all* times.

Next we provide a technical lemma which relates the (first) spread of the modified sequence to the second spread of the original sequence.

LEMMA 5.7. *Let Σ be a sequence and let Σ' be any 1-modification of Σ . Then the spread of Σ' is bounded by $s_2(\Sigma)$.*

PROOF. By the 1-Modification Theorem, we know that the books of the modified sequence and the original sequence can differ by at most one order in each book (buy and sell). Therefore, the second-highest buy order in the original sequence is always at most the bid in the modified sequence, and the second-lowest sell order in the original sequence is always at least the ask of the modified sequence. \square

We are now ready to state a stability result for the average execution price in the absolute model. It establishes that in highly liquid markets, where the executed volume is large and the spread small, the average price is highly stable.

THEOREM 5.8. *Let Σ be a sequence and let Σ' be any 1-modification of Σ . Then*

$$|average(\Sigma) - average(\Sigma')| \leq \frac{2(p_{max} + s_2(\Sigma))}{volume(\Sigma)} + s_2(\Sigma)$$

where p_{max} is the highest execution price in Σ .

PROOF. The proof will show that every execution in Σ besides the execution of the modified order and the last execution has a matching execution in Σ' with a price different by at most $s_2(\Sigma)$, and will use the fact that $p_{max} + s_2(\Sigma)$ is a bound on the price in Σ' .

Referring to the proof of the 1-Modification Theorem, suppose we are in state $\bar{a}b$, where we have in one sequence (which can be either Σ or Σ') an additional buy order b and an additional sell order a . Without loss of generality we assume that the sequence with the additional orders is Σ . If the next execution does not involve a or b then clearly we have the same execution in both Σ and Σ' . Suppose that it involves a ; there are two possibilities. Either a is the modified order, in which case we change the average price

difference by $(p_{max} + s_2(\Sigma))/volume(\Sigma)$, and this can happen only once; or a was executed before in Σ' and the executions both involve an order whose limit price is a . By Lemma 5.7 the spread of both sequences is bounded by $s_2(\Sigma)$, which implies that the price of the execution in Σ' was at most $a + s_2(\Sigma)$, while execution in Σ is at price a , and thus the prices are different by at most $s_2(\Sigma)$.

In states $\bar{a}b, a\bar{b}$ as long as we have concurrent executions in the two sequences, we know that the prices can differ by at most $s_2(\Sigma)$. If we have an execution only in one sequence, we either match it in state $\bar{a}b$, or charge it by $(p_{max} + s_2(\Sigma))/volume(\Sigma)$ if we end at state $\bar{a}b$.

If we end in state $ab, \bar{a}b$ or $a\bar{b}$, then every execution in states $\bar{a}b$ or $a\bar{b}$ were matched to an execution in state $\bar{a}b$. If we end up in state $\bar{a}b$, we have the one execution that is not matched and thus we charge it $(p_{max} + s_2(\Sigma))/volume(\Sigma)$. \square

We next give a stability result for the closing price. We first provide a technical lemma regarding the prices of consecutive executions.

LEMMA 5.9. *Let Σ be any sequence. Then the prices of two consecutive executions in Σ differ by at most $s_2(\Sigma)$.*

PROOF. Suppose the first execution is taken at time t ; its price is bounded below by the current bid and above by the current ask. Now after this execution the bid is at least the second highest buy order at time t , if the former bid was executed and no higher buy orders arrived, and higher otherwise. Similarly, the ask is at most the second lowest sell order at time t . Therefore, the next execution price is at least the second bid at time t and at most the second ask at time t , which is at most $s_2(\Sigma)$ away from the bid/ask at time t . \square

LEMMA 5.10. *Let Σ be any sequence and let Σ' be a 1-modification of Σ . If the volume(Σ) ≥ 2 , then*

$$|close(\Sigma) - close(\Sigma')| \leq s_2(\Sigma)$$

PROOF. We first deal with case where the last execution occurs in both sequences simultaneously. By Theorem 5.6, both the ask and the bid of Σ and Σ' are at most $s_2(\Sigma)$ apart at every time t . Since the price of the last execution is their asks (bids) at time t we are done.

Next we deal with the case where the last execution among the two sequences occurs only in Σ . In this case we know that either the previous execution happened simultaneously in both sequences at time t , and thus all three executions are within the second spread of Σ at time t (the first execution in Σ by definition, the execution at Σ' from identical arguments as in the former case, and the third by Lemma 5.9). Otherwise the previous execution happened only in Σ' at time t , in which case the two executions are within the the spread of Σ at time t (the execution of Σ' from the same arguments as before, and the execution in Σ must be inside its spread in time t).

If the last execution happens only in Σ' we know that the next execution of Σ will be at most $s_2(\Sigma)$ away from its previous execution by Lemma 5.9. Together with the fact that if an execution happens only in one sequence it implies that the order is in the spread of the second sequence as long as the sequences are 1-modification, the proof is completed. \square

5.2 Spread Bounds for k-Modifications

As in the case of executed volume, we would like to extend the absolute model stability results for price-based quantities to the case where multiple orders are modified. Here our results are weaker and depend on the k -spread, the distance between the k th highest buy order and the k th lowest sell order, instead of the second spread. (Looking ahead to Section 7, we note that in actual market data for liquid stocks, this quantity is often very small as well.) We use $s_k(\Sigma)$ to denote the k -spread. As before, we assume that the k -spread is always defined after an initialization period.

We first state the following generalization of Lemma 5.7.

LEMMA 5.11. *Let Σ be a sequence and let Σ' be any 1-modification of Σ . For $\ell \geq 1$, if $s_{\ell+1}(\Sigma)$ is always defined after the change, then $s_\ell(\Sigma') \leq s_{\ell+1}(\Sigma)$.*

The proof is similar to the proof of Lemma 5.7 and omitted. A simple application of this lemma is the following: Let Σ_ℓ be any sequence which is an ℓ -modification of Σ . Then we have $s_2(\Sigma_\ell) \leq s_{\ell+2}(\Sigma)$. Now using the above lemma and by simple induction we can obtain the following theorem.

THEOREM 5.12. *Let Σ be a sequence and let Σ' be any k -modification of Σ . Then*

1. $|lastbid(\Sigma) - lastbid(\Sigma')| \leq \sum_{\ell=1}^k s_{\ell+1}(\Sigma) \leq ks_{k+1}(\Sigma)$
2. $|lastask(\Sigma) - lastask(\Sigma')| \leq \sum_{\ell=1}^k s_{\ell+1}(\Sigma) \leq ks_{k+1}(\Sigma)$
3. $|close(\Sigma) - close(\Sigma')| \leq \sum_{\ell=1}^k s_{\ell+1}(\Sigma) \leq ks_{k+1}(\Sigma)$
4. $|average(\Sigma) - average(\Sigma')| \leq \sum_{\ell=1}^k \left(\frac{2(p_{max} + s_{\ell+1}(\Sigma))}{volume(\Sigma)} + s_{\ell+1}(\Sigma) \right)$

where $s_\ell(\Sigma)$ is the maximum over the ℓ -spread in Σ following the first modification.

We note that while these bounds depend on deeper measures of spread for more modifications, we are working in a 1-share order model. Thus in an actual market, where single orders contain hundreds or thousands of shares, the k -spread even for large k might be quite small and close to the standard 1-spread in liquid stocks.

6. RELATIVE MODEL INSTABILITY

In the relative model the underlying assumption is that traders try to exploit their knowledge of the books to strategically place their orders. Thus if a trader wants her buy order to be executed quickly, she may position it above the current bid and be the first in the queue; if the trader is patient and believes that the price trend is going to be downward she will place orders deeper in the buy book, and so on.

While in the previous sections we showed stability results for the absolute model, here we provide simple examples which show instability in the relative model for the executed volume, last bid, last ask, average execution price and the last execution price. In Section 7 we provide many simulations on actual market data that demonstrate that this instability is inherent to the relative model, and not due to artificial constructions. In the relative model we assume that for every sequence the ask and bid are always defined, so the books have a non-empty initial configuration.

We begin by showing that in the relative model, even a single modification can double the number of shares executed.

THEOREM 6.1. *There is a sequence Σ and a 1-modification Σ' of Σ such that $\text{volume}(\Sigma') \geq 2\text{volume}(\Sigma)$.*

PROOF. For concreteness we assume that at the beginning the ask is 10 and the bid is 8. The sequence Σ is composed from n buy orders with $\Delta = 0$, followed by n sell orders with $\Delta = 0$, and finally an alternating sequence of buy orders with $\Delta = +1$ and sell orders with $\Delta = -1$ of length $2n$. Since the books before the alternating sequence contain $n+1$ sell orders at 10 and $n+1$ buy orders at 8, we have that each pair of buy sell order in the alternating part is matched and executed, but none of the initial $2n$ orders is executed, and thus $\text{volume}(\Sigma) = n$. Now we change the first buy order to have $\Delta = +1$. After the first $2n$ orders there are still no executions; however, the books are different. Now there are $n+1$ sell orders at 10, n buy orders at 9 and one buy order at 8. Now each order in the alternating sequence is executed with one of the former orders and we have $\text{volume}(\Sigma') = 2n$. \square

The next theorem shows that the spread-based stability results of Section 5.1 do not also hold in the relative model. Before providing the proof, we give its intuition. At the beginning the sell book contains only two prices which are far apart and both contain only two orders, now several buy orders arrive, at the original sequence they are not being executed, while in the modified sequence they will be executed and leave the sell book with only the orders at the high price. Now many sell orders followed by many buy orders will arrive, such that in the original sequence they will be executed only at the low price and in the modified sequence they will be executed at the high price.

THEOREM 6.2. *For any positive numbers s and x , there is sequence Σ such that $s_2(\Sigma) = s$ and a 1-modification Σ' of Σ such that*

- $|\text{close}(\Sigma) - \text{close}(\Sigma')| \geq x$
- $|\text{average}(\Sigma) - \text{average}(\Sigma')| \geq x$
- $|\text{lastbid}(\Sigma) - \text{lastbid}(\Sigma')| \geq x$
- $|\text{lastask}(\Sigma) - \text{lastask}(\Sigma')| \geq x$

PROOF. Without loss of generality let us consider sequences in which all prices are integer-valued, in which case the smallest possible value for the second spread is 1; we provide the proof for the case $s_2(\Sigma) = 2$, but the $s_2(\Sigma) = 1$ case is similar.

We consider a sequence Σ such that after an initialization period there have been no executions, the buy book has 2 orders at price 10, and the sell book has two orders at price 12 and 2 orders with value $12+y$, where y is a positive integer that will be determined by the analysis. The original sequence Σ is a buy order with $\Delta = 0$, followed by two buy orders with $\Delta = +1$, then $2y$ sell orders with $\Delta = 0$, and then $2y$ buy orders with $\Delta = +1$. We first note that $s_2(\Sigma) = 2$, there are $2y$ executions, all at price 12, the last bid is 11 and the last ask is 12. Next we analyze a modified sequence. We change the first buy order from $\Delta = 0$ to $\Delta = +1$. Therefore, the next two buy orders with $\Delta = +1$

are executed, and afterwards we have that the bid is 11 and the ask is $12+y$. Now the $2y$ sell orders are accumulated at $12+y$, and after the next y buy orders the bid is at $12+y-1$. Therefore, at the end we have that $\text{lastbid}(\Sigma') = 12+y-1$, $\text{lastask}(\Sigma') = 12+y$, $\text{close}(\Sigma') = 12+y$, and $\text{average}(\Sigma') = \frac{y}{y+2}(12+y) + \frac{2}{y+2}(12)$. Setting $y = x+2$, we obtain the lemma for every property. \square

We note that while this proof was based on the fact that there are two consecutive orders in the books which are far (y) apart, we can provide a slightly more complicated example in which all orders are close (at most 2 apart), yet still one change results in large differences.

7. SIMULATION STUDIES

The results presented so far paint a striking contrast between the absolute and relative price models: while the absolute model enjoys provably strong stability over any fixed event sequence, there exist at least specific sequences demonstrating great instability in the relative model. The worst-case nature of these results raises the question of the extent to which such differences could actually occur in real markets. In this section we provide indirect evidence on this question by presenting simulation results exploiting a rich source of real-market historical limit order sequence data. By interpreting arriving limit order prices as either absolute values, or by transforming them into differences with the current bid and ask (relative model), we can perform small modifications on the sequences and examine how different various outcomes (volume traded, average price, etc.) would be from what actually occurred in the market. These simulations provide an empirical counterpart to the theory we have developed. We emphasize that *all such simulations interpret the actual historical data as falling into either the absolute or relative model*, and are meaningful only within the confines of such an interpretation. Nevertheless, we feel they provide valuable empirical insight into the potential (in)stability properties of modern equity limit order markets, and demonstrate that one's belief or hope in stability largely relies on an absolute model interpretation. We also investigate the empirical behavior of mixtures of absolute and relative prices.

7.1 Data

The historical data used in our simulations is commercially available limit order data from INET, the previously mentioned electronic exchange for NASDAQ stocks. Broadly speaking, this data consists of practically every single event on INET regarding the trading of an individual stock — every arriving limit order (price, volume, and sequence ID number), every execution, and every cancellation of a standing order — all timestamped in milliseconds. It is data sufficient to recreate the precise INET order book in a given stock on a given day and time.

We will report stability properties for three stocks: Amazon, Nvidia, and Qualcomm (identified in the sequel by their tickers, AMZN, NVDA and QCOM). These three provide some range of liquidities (with QCOM having the greatest and NVDA the least liquidity on INET) and other trading properties. We note that the qualitative results of our simulations were similar for several other stocks we examined.

7.2 Methodology

For our simulations we employed order-book reconstruction code operating on the underlying raw data. The basic format of each experiment was the following:

1. Run the order book reconstruction code on the original INET data and compute the quantity of interest (volume traded, average price, etc.)
2. Make a small modification to a single order, and recompute the resulting value of the quantity of interest.

In the absolute model case, Step 2 is as simple as modifying the order in the original data and re-running the order book reconstruction. For the relative model, we must first pre-process the raw data and convert its prices to relative values, then make the modification and re-run the order book reconstruction on the relative values.

The type of modification we examined was extremely small compared to the volume of orders placed in these stocks: namely, the deletion of a single randomly chosen order from the sequence. Although a deletion is not 1-modification, its edit distance is 1 and we can apply Theorem 5.4. For each trading day examined, this single deleted order was selected among those arriving between 10 AM and 3 PM, and the quantities of interest were measured and compared at 3 PM. These times were chosen to include the busiest part of the trading day but avoid the half hour around the opening and closing of the official NASDAQ market (9:30 AM and 3:30 PM respectively), which are known to have different dynamics than the central portion of the day.

We run the absolute and relative model simulations on both the raw INET data and on a “cleaned” version of this data. In the “cleaned” we remove all limit orders that were canceled in the actual market prior to their execution (along with the cancellations themselves). The reason is that such cancellations may often be the first step in the “repositioning” of orders — that is, cancellations of the order that are followed by the submission of a replacement order at a different price. Not removing canceled orders allows the possibility of modified simulations in which the “same” order ¹ is executed twice, which may magnify instability effects. Again, it is clear that neither the raw nor the “cleaned” data can perfectly reflect “what would have happened” under the deleted orders in the actual market. However, the results both from the raw data and the clean data are qualitatively similar. The results mainly differ, as expected, in the executed volume, where the instability results for the relative model are much more dramatic in the raw data.

7.3 Results

We begin with summary statistics capturing our overall stability findings. Each row of the tables below contains a ticker (e.g. AMZN) followed by either -R (for the uncleaned or raw data) or -C (for the data with canceled orders removed). For each of the approximately 250 trading days in 2003, 1000 trials were run in which a randomly selected order was deleted from the INET event sequence. For each quantity of interest (volume executed, average price, closing price and last bid), we show for the both the absolute and

¹Here “same” is in quotes since the two orders will actually have different sequence ID numbers, which is what makes such repositioning activity impossible to reliably detect in the data.

relative model the average percentage change in the quantity induced by the deletion.

The results confirm rather strikingly the qualitative conclusions of the theory we have developed. In virtually every case (stock, raw or cleaned data, and quantity) the percentage change induced by a single deletion in the relative model is many orders of magnitude greater than in the absolute model, and shows that indeed “butterfly effects” may occur in a relative model market. As just one specific representative example, notice that for QCOM on the cleaned data, the relative model effect of just a single deletion on the closing price is in excess of a full percentage point. This is a variety of market impact entirely separate from the more traditional and expected kind generated by trading a large volume of shares.

| Stock | Date | <i>volume</i> | | <i>average</i> | |
|--------|------|---------------|--------|----------------|---------|
| | | Rel | Abs | Rel | Abs |
| AMZN-R | 2003 | 15.1% | 0.04% | 0.3% | 0.0002% |
| AMZN-C | 2003 | 0.69% | 0.087% | 0.36% | 0.0007% |
| NVDA-R | 2003 | 9.09% | 0.05 % | 0.17% | 0.0003% |
| NVDA-C | 2003 | 0.73% | 0.09 % | 0.35% | 0.001% |
| QCOM-R | 2003 | 16.94% | 0.035% | 0.21% | 0.0002% |
| QCOM-C | 2003 | 0.58% | 0.06% | 0.35% | 0.0005% |

| Stock | Date | <i>close</i> | | <i>lastbid</i> | |
|--------|------|--------------|---------|----------------|---------|
| | | Rel | Abs | Rel | Abs |
| AMZN-R | 2003 | 0.78% | 0.0001% | 0.78% | 0.0007% |
| AMZN-C | 2003 | 1.10% | 0.077% | 1.11% | 0.001% |
| NVDA-R | 2003 | 1.17% | 0.002 % | 1.18 % | 0.08% |
| NVDA-C | 2003 | 0.45% | 0.0003% | 0.45% | 0.0006% |
| QCOM-R | 2003 | 0.58% | 0.0001% | 0.58% | 0.0004% |
| QCOM-C | 2003 | 1.05% | 0.0006% | 1.05% | 0.06% |

In Figure 4 we examine how the change to one the quantities, the average execution price, grows with the introduction of greater perturbations of the event sequence in the two models. Rather than deleting only a single order between 10 AM and 3 PM, in these experiments a growing number of randomly chosen deletions was performed, and the percentage change to the average price measured. As suggested by the theory we have developed, for the absolute model the change to the average price grows linearly with the number of deletions and remains very small (note the vastly different scales of the y-axis in the panels for the absolute and relative models in the figure). For the relative model, it is interesting to note that while small numbers of changes have large effects (often causing average execution price changes well in excess of 0.1 percent), the effects of large numbers of changes levels off quite rapidly and consistently.

We conclude with an examination of experiments with a mixture model. Even if one accepts a world in which traders behave in either an absolute or relative manner, one would be likely to claim that the market contains a mixture of both. We thus ran simulations in which each arriving order in the INET event streams was treated as an absolute price with probability α , and as a relative price with probability $1 - \alpha$. Representative results for the average execution price in this mixture model are shown in Figure 5 for AMZN and NVDA. Perhaps as expected, we see a monotonic decrease in the percentage change (instability) as the fraction of absolute traders increases, with most of the reduction already being realized by the introduction of just a small population of absolute traders. Thus even in a largely relative-price world, a

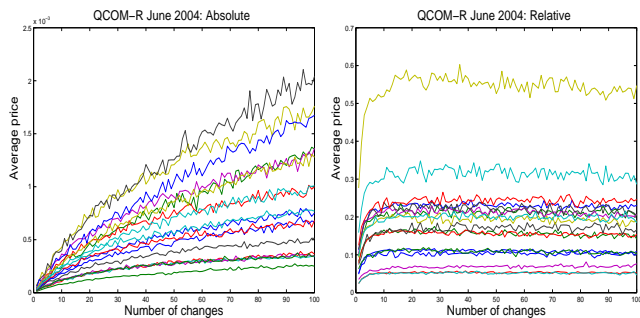


Figure 4: Percentage change to the average execution price (y-axis) as a function of the number of deletions to the sequence (x-axis). The left panel is for the absolute model, the right panel for the relative model, and each curve corresponds to a single day of QCOM trading in June 2004. Curves represent averages over 1000 trials.

small minority of absolute traders can have a greatly stabilizing effect. Similar behavior is found for closing price and last bid.




Figure 5: Percentage change to the average execution price (y-axis) vs. probability of treating arriving INET orders as absolute prices (x-axis). Each curve corresponds to a single day of trading during a month of 2004. Curves represent averages over 1000 trials.

For the executed volume in the mixture model, however, the findings are more curious. In Figure 6, we show how the percentage change to the executed volume varies with the absolute trader fraction α , for NVDA data that is both raw and cleaned of cancellations. We first see that for this quantity, unlike the others, the difference induced by the cleaned and uncleaned data is indeed dramatic, as already suggested by the summary statistics table above. But most intriguing is the fact that the stability is *not* monotonically increasing with α for either the cleaned or uncleaned data — the market with maximum instability is not a pure relative price market, but occurs at some nonzero value for α . It was in fact not obvious to us that sequences with this property could even be artificially constructed, much less that they would occur as actual market data. We have yet to find a satisfying explanation for this phenomenon and leave it to future research.

8. ACKNOWLEDGMENTS

We are grateful to Yuriy Nevmyvaka of Lehman Brothers in New York for the use of his INET order book reconstruction code, and for valuable comments on the work presented

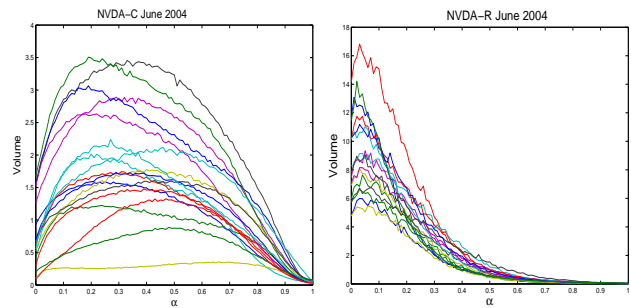


Figure 6: Percentage change to the executed volume (y-axis) vs. probability of treating arriving INET orders as absolute prices (x-axis). The left panel is for NVDA using the raw data that includes cancellations, while the right panel is on the cleaned data. Each curve corresponds to a single day of trading during June 2004. Curves represent averages over 1000 trials.

here. Yishay Mansour was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778, by a grant from the Israel Science Foundation and an IBM faculty award.

9. REFERENCES

- [1] D. Bertsimas and A. Lo. Optimal control of execution costs. *Journal of Financial Markets*, 1:1–50, 1998.
- [2] B. Biais, L. Glosten, and C. Spatt. Market microstructure: a survey of microfoundations, empirical results and policy implications. *Journal of Financial Markets*, 8:217–264, 2005.
- [3] J.-P. Bouchaud, M. Mezard, and M. Potters. Statistical properties of stock order books: empirical results and models. *Quantitative Finance*, 2:251–256, 2002.
- [4] C. Cao, O.Hansch, and X. Wang. The informational content of an open limit order book, 2004. AFA 2005 Philadelphia Meetings, EFA Maastricht Meetings Paper No. 4311.
- [5] R. Coggins, A. Blazejewski, and M. Aitken. Optimal trade execution of equities in a limit order market. In *International Conference on Computational Intelligence for Financial Engineering*, pages 371–378, March 2003.
- [6] D. Farmer and S. Joshi. The price dynamics of common trading strategies. *Journal of Economic Behavior and Organization*, 29:149–171, 2002.
- [7] J. Hasbrouck. Empirical market microstructure: Economic and statistical perspectives on the dynamics of trade in securities markets, 2004. Course notes, Stern School of Business, New York University.
- [8] R. Kissell and M. Glantz. *Optimal Trading Strategies*. Amacom, 2003.
- [9] S.Kakade, M. Kearns, Y. Mansour, and L. Ortiz. Competitive algorithms for VWAP and limit order trading. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 189–198, 2004.
- [10] Y.Nevmyvaka, Y. Feng, and M. Kearns. Reinforcement learning for optimized trade execution, 2006. Preprint.