Order Dynamcis in a High-Frequency Trading Environment*

Arne Breuer[†] Hans-Peter Burghof[‡] Julian Stitz [§]

January 2011

Preliminary – comments welcome

Abstract

We analyse order book message data in order to detect algorithmic trade activity. Previous papers usually analyse order book data with a time stamp precision of one hundredth of a second. In times of co-location, those levels of precision are not sufficient to see effects of ultra-high frequency algorithms. Our Nasdaq-supplied dataset is equipped with a time stamp precision of a billionth of a second. Thus, we 'zoom in' and analyse the sub-millisecond effects of algorithmic trading on the order book. We find evidence of algorithmic trading with the limit order lifetime, limit order revision time, and inter order placement time. In addition to that, we apply the proxies separately on exchange-traded funds and stocks to see if structured products are treated differently than common stocks.

*We thank Geir Bjønnes, Ulli Spankowski and the participants of the INFINITI Conference on International Finance 2011 for valuable comments

[†]Corresponding author. Universität Hohenheim, Lehrstuhl für Bankwirtschaft und Finanzdienstleistungen 510F, 70599 Stuttgart. Telephone: +49 711 459-22903, Fax: +49 711 459-23448. Email: arne.breuer@uni-hohenheim.de

[‡]Universität Hohenheim.

[§]Universität Hohenheim

1 Introduction

Over the last decade, high-frequency trading has become one of the most important driving factors in securities markets. Both market observers and researchers agree that a large part of the traded volume on major stock exchanges such as NASDAQ, Deutsche Börse, or the NYSE is traded by algorithms and not by humans. However, the definition of high-frequency traders is still somewhat diffuse. Hendershott et al. (2011, p. 1) define them in a very general but still accurate way as 'computer algorithms [that] automatically make certain trading decisions, submit orders, and manage those orders after submission.'

The rapid increase in high-frequency trading in the last one-and-a-half or so decades has been fuelled by the Regulation National Market System (better known as RegNMS) in the United States and the Markets in Financial Instruments Directive (MiFID) in the European Union. Although the two laws differ in details, they share the main goal to 'foster fair, competitive, efficient, and integrated equities markets and to encourage financial innovation' (Storkenmaier and Wagener, 2010, p. 3). Both RegNMS and MiFID have led to the increased market share of ECNs and multilateral trade facilities (MTFs), which have entered the competition with the established stock exchanges and have been able to gain significant parts of overall trade activity.

Many ECNs provide an open electronic limit order book. In order to provide a liquid market, they have to find a way to fill their order books. Many algorithmic strategies, especially high-frequency trading ones, generate large amounts of limit orders. Together with this liquidity generation, they generate income, making it economically very attractive as well. Consequently, ECNs often encourage high-frequency trading. They invest in fast communication and input/output technology.

This enables ECNs to provide very short times to accept, process, and respond to incoming orders, which is usually called (system) latency. For many high-frequency trading strategies, fast reaction times are a crucial factor; hence, a low latency of the stock exchanges' matching systems is a key asset. Because they find an appealing environment, many orders from algorithms are routed to ECNs. Established exchanges quickly reacted and shifted their market structure away from quote- to order-driven microstructures and also invested in IT in order to lower their latency times.

As a result of this technological arms race, latencies of well below one millisecond (10^{-3}s) are not uncommon. As we will show in the progress of this paper, high-frequency trading strategies work on the level of a few microseconds (10^{-6}s). In order to analyse their behaviour and the effects they could have on overall trading, datasets with a precision of a hundredth of a second or one millisecond are insufficient to exactly show the effects of high-frequency trading (and especially its subset HFT) on order books. Low timestamp precision causes too much information loss. We use a dataset from the NASDAQ stock exchange with a timestamp precision of one nanosecond (10^{-9}s). This enables us to accurately analyse high-frequency trading effects on the order book and thus the market environment that everybody is trading in.

It is currently being discussed, however, how large the share of high-frequency trading really is. Of course, brokerage firms, proprietary traders, and other financial institutions that implement algorithms do not have to publish their implementation of high-frequency trading strategies and are in fact very secretive about them. Thus, because usually no reliable data exists, market observers and researchers have to rely on estimations on the share of high-frequency trading. The numbers are rather diverse. For example, Mary Schapiro, SEC chairwoman, sees the share of high-frequency trading relative to market volume at '50 per cent or more' in 2010; Senator Kaufman (Democrats, Delaware) states that the market volume is 70 per cent algorithmic (Muthuswamy et al., 2011, p. 87). For Deutsche Börse's Xetra system, more reliable data exists because of its ATP, which was discontinued in 2009. Market participants who implement algorithms could sign up for it and save order fees. For detailed descriptions and analyses see, for example, Hendershott and Riordan (2009), Gsell (2009), Groth (2009), or Maurer and Schäfer (2011). There, the share of high-frequency trading floats around 50 per cent.

The analysis of high-frequency trade activity is important, because it is structurally different than human trade behaviour. With the rise of high-frequency trading, the whole market structure changes. For example, Hasbrouck and Saar (2009) state that the traditional interpretation of limit order traders as 'patient providers of liquidity' (Handa and Schwartz, 1996) has to be re-evaluated if limit order lifetimes decrease to fractions of seconds. For market participants, there are neither timely nor accurate statistics on high-frequency trading. However, traders may have reasons for the desire to know about the extent of high-frequency trading on 'their' market.

The lack of accurate data on high-frequency trading prevents us from developing an easy-to-calculate method to accurately measure high-frequency trading. Nonetheless, we try to extract information on the approximate extent of high-frequency trading from raw order book message data that does not contain any information on the source of the messages. Currently, high-frequency trade activity is not directly measurable. With this paper, we show the order activity of high-frequency algorithms and hope to help find a way to create a measure to estimate the extent of high-frequency trade activity. It can be especially useful to model an agent-driven few-type market with the integration of a stylised high-frequency trader.

Trade activity can be analysed by more than just one measure. Market participants who work intraday to trade on small price movements have to constantly revise their open positions and orders. Therefore, we analyse the structure of order strategies of high-frequency trading engines with three proxies. We choose limit order lifetime, order-revision time, and inter-order placement time to catch actively and rapidly trading computer programs. We often observe very short values for the proxies. We conclude that algorithms possibly have a built-in risk assessment for their limit orders, which we call the limit order risk function. This strictly concave function reflects the probability that a limit order does not optimally fit the current market. Over time, market factors change, which results in changed optimal limit order properties. When the limit order risk function achieves a pre-defined value, the algorithm deletes the old limit order and inserts a new one with now-optimal order properties. With this concept, we can partly explain the structure of the proxies and the difference between ETFs and common stocks.

The remainder of the paper is structured as follows: Section 2 describes the data for the empirical analysis. Section 3 provides the results of the empirical analysis. In Section 3.1, we analyse the structure of limit order lifetimes of ETFs and NASDAQ-listed common stocks, i.e., the time that passes between the insertion of a limit order and its deletion. In Section 3.2, we examine limit order revision times, i.e., the time that passes between a deletion of a limit order until the next placement for the same security. In Section 3.3, we analyse inter-order placement times, i.e., the time that passes between two placements of limit orders. For each of these proxies, we perform various analyses and use data with a timestamp precision of less than one microsecond, which enables us to very accurately research sub-millisecond observations. Section 4 concludes.

2 Dataset

The analysis of limit order book data usually employs data with a precision level of a hundredth of a second or milliseconds, i.e., 1/100 or 1/1,000 of a second, respectively. For example, to search patterns in Xetra limit order book data, Prix et al. (2007, 2008) use limit order book data with a timestamp precision of 1/100 of a second, whereas the clock of Xetra runs at 1/1,000 of a second.

We use data generously provided by the NASDAQ stock exchange. It contains all entries into the order book, so it is comparable to the one Prix et al. (2007, 2008) use. It is a protocol of the order book and every activity is stored—order insertions, deletions, partial cancellations, executions, etc. In early 2010, NASDAQ's ITCH-data's timestamp precision received an update from milliseconds to nanoseconds, or 1/1,000,000,000 of a second. This enables us to perform a search for order structure patterns at levels that, to our knowledge, have not been attempted before due to the recent increase in accuracy.

This increase in precision is more than welcome because one of the top priorities of algorithmic and especially HFT is speed. For example, because light moves at a speed of 299,792.458 km/s, a signal carrying order information from a brokerage firm in, say, Chicago, to the NASDAQ stock exchange in New York would need at least 3.87 milliseconds or 0.00387 seconds. Because the speed of light in fibre or electric signals in wires is significantly slower, these numbers are minimal values, and the actual transmission time will be much longer. In addition to this, the signal from the broker would need additional time to pass routers and other network technology, increasing the brokerage firm's latency to levels well beyond four milliseconds. If the trading strategy of the Chicago-based brokerage firm was based on ultra-high frequency models, this latency could result in significant disadvantages compared to brokerage firms located next to the exchange. Indeed, many high-frequency trading firms use the co-location service, where their engines operate from servers only a few metres away from the exchange's servers.

We analyse the limit order protocol from 22 February to 26 February 2010, i.e., five trading days. Within this time, 131,701,300 orders have been placed and 125,489,818 limit orders were deleted. That means approximately 95 per cent of the added limit orders were deleted. This figure is only approximate because some orders from the week before were deleted in this week and some limit orders that were added were not deleted until Friday's market close. We expect the error rate to be small, though. The number of deletions of last week's limit orders and insertion of limit orders that stay alive over the weekend should approximately equal out each other. The dataset contains ETFs, stocks listed on NASDAQ, and stocks listed on the NYSE. We choose 36 stocks that are listed on NASDAQ with the highest limit order activity and matched them with 36 ETFs that have a similar number of added orders to keep the results for the different, comparable securities. The list of stocks and ETFs used for this analysis is given in Table 2.

The structure of the data enables us to measure limit order lifetimes without any noise. Each new limit order that arrives on the NASDAQ receives a day-unique order reference number, and each change or deletion of it is marked with it. That means that once a limit order is placed, its lifetime can be determined by looking for its deletion time. The lifetime is then $L = t_d - t_p$, where L is the lifetime, t_d the deletion time and t_p the placement time.

The other two modes are subject to noise. ITCH data is anonymous and does not carry any identifier of the market participant that places the limit order. Thus, we cannot detect every structured approach in these two proxies, because any other market participant can add a limit order before the one who we are looking at does. However, because of the high speed of high-frequency trading, high frequency limit order strategies indeed leave visible footprints.

3 Empirical Results

As Hasbrouck and Saar (2011) show, the majority of limit orders of highly liquid stocks are often deleted within a matter of a few milliseconds. One of their aims was to show the basic structure of the order dynamics within the time frame of fleeting orders as defined by Hasbrouck and Saar (2001, 2009). Hasbrouck and Saar (2011) use a dataset with a timestamp precision of a millisecond, which is sufficient for the microstructure of order dynamics. As a consequence, all limit orders with lifetimes of under one millisecond were shown as one millisecond. In 2010, however, NASDAQ's timestamp precision level was increased to nanoseconds. This improvement enables us to 'zoom in' and see what really happens in the atomic regions of limit order data.

We will perform a limit order book analysis based on three different modes: first, we look at the limit order lifetime, which measures the time between the placement of a limit order and its deletion. Second, we analyse the limit order revision time. We measure the time that passes between the deletion of a limit order until the next placement of a limit order. The third mode is the inter-order placement time. This is the time the order book for a specific stock does not receive new limit orders.

Because we are interested in the 'algorithmic nano level' using nanosecond timestamps we 'zoom in' the few-millisecond lifetimes. After a short recapitulation of the evidence in the macro-level of order lifetimes, we examine the regions that are obviously the most interesting for many high-frequency algorithms: the sub-millisecond lifetimes.

3.1 Limit Order Lifetimes

Limit order lifetimes are fundamentally non-normally distributed. Rather, they can be described with the Weibull distribution, which defines a survival probability S of

$$\mathbb{P}(T > t) = S(t) = \exp(-\exp(\beta_0)t^p), \tag{1}$$

where β_0 is the scale parameter and p is the shape parameter (Cleves et al., 2002, p. 212). The estimated β_0 in our dataset is constantly smaller than one and greater than zero, and p averages around 0.31, which yields a hyperbolic distribution. Many limit orders are deleted before their execution within a handful of milliseconds. The more limit order placement activity there is for an equity, the shorter the limit order lifetime becomes. For the most active stocks or ETFs, it is not uncommon that more than 80 per cent of limit orders are deleted without execution within less than one second.

Following Prix et al. (2007, 2008), we look for irregularities in the densities of limit order lifetimes. Whereas Prix et al. (2007, 2008) find multiple peaks in the kernel density estimations of their Xetra datasets at 250 milliseconds, two seconds and multiples of 30 seconds, we only find one peak at 100 milliseconds in our NASDAQ dataset. While some stocks or ETFs show an additional peak at one second, only the peak at 100 milliseconds is statistically significant.

Many limit orders, as already mentioned, are deleted after one or two milliseconds. With a dataset with an unmatched timestamp precision, we can examine the 'nanostructure' of today's stock markets and capture ultra-high-frequency high-frequency trade activity. The result is surprising. We only find one significant peak at the micro level with millisecond timestamps. With the more exact timestamps, it becomes clear that the entire density of limit order lifetimes consists of peaks, that are invisible at a lower 'resolution.' The ultra-short lifetimes show peaks at multiples of 50 microseconds. Figure 1 illustrates this with four exemplary histograms, which depict two NASDAQ-listed stocks, and two ETFs. The individual assets were chosen randomly, and the day presented is 22 February. The histograms show the frequency of limit order lifetimes in the interval of $[0, 2]$ milliseconds. The solid line shows the cumulated share of limit orders deleted until time t relative to the total amount of limit orders for the security on the day. For example, observe that more than 15 per cent of the limit orders placed during the day on CENX were deleted within two milliseconds.

Figure 2 shows the cumulative average share of limit orders that have been deleted within n seconds. The share is calculated relative to the total number of limit orders that have been inserted on the same day. The figure shows two curves: one shows the average limit order lifetime of 36 ETFs, the other one that of 36 NASDAQ stocks covering five trading days, i.e., one trading week. The main underlying function—disregarding the jump at 100 milliseconds—is concave and approaches unity.

On average, limit orders are deleted within a very short time. In the sample of five trading days, more than 50 per cent of inserted limit orders have been deleted within approximately 0.6 seconds. It is noticeable that the graphs of the stocks and ETFs differ a little. The basic structure of limit order deletion times is in both cases

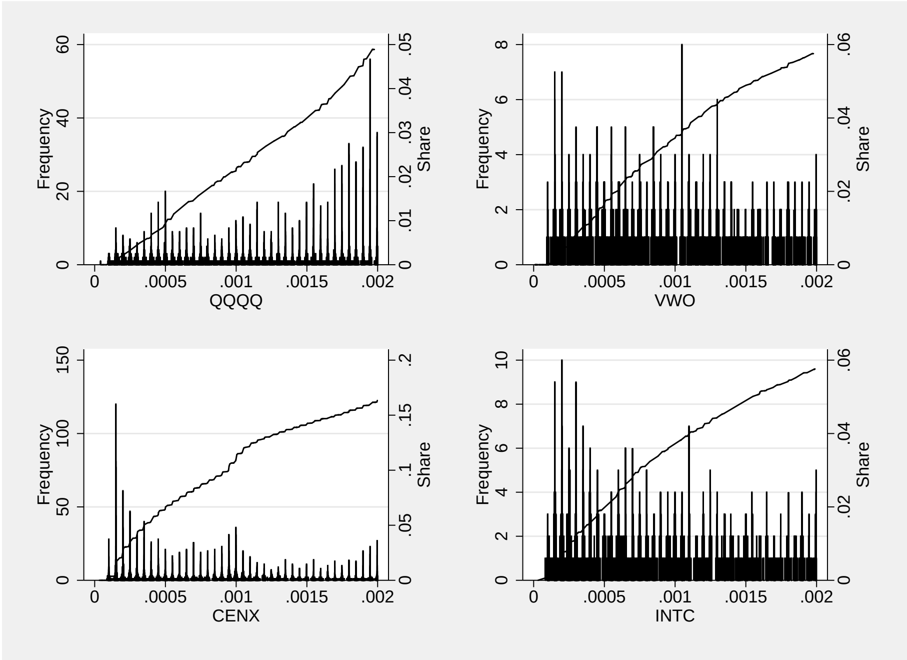

Figure 1: Exemplary histograms of order lifetimes. Each bar represents a timeframe of 100 nanoseconds from 0 to two milliseconds. Depicted are the two ETFs QQQQ (Powershares QQQ, 125,680 observations) and VWO (Vanguard Emerging Markets ETF, 39,121) and two stocks listed on NASDAQ: CENX (Century Aluminum Co, 57,783), and INTC (Intel Corp, 48,977). Solid line: cumulated share of limit order lifetime relative to the total amount of limit orders on the day in per cent.

a concave function with a jump at 100 milliseconds. Stocks, however, tend to have a greater proportion of limit orders with a lifetime of below 100 milliseconds, but the slope of the function of limit order lifetimes decreases more rapidly for stocks than for ETFs.

The functions in Figure 2 can possibly be regarded as limit order risk functions $f(\cdot)$. The probability that a limit order does not optimally fit market conditions increases over time, $\partial f(\cdot)/\partial t \geq 0$. To fit the limit orders to the market, traders constantly have to adjust price and/or quantity tags in order not to be exposed to the two main risks of limit orders: non-execution risk and undesired execution risk (Handa and Schwartz, 1996).

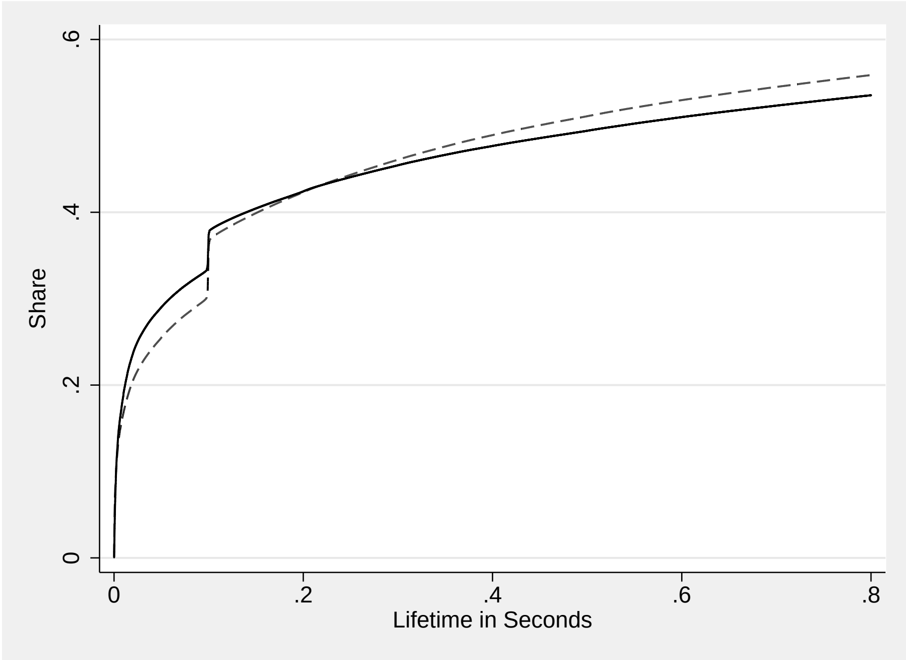

Figure 2: Cumulative average limit order lifetimes. Solid line: NASDAQ listed stocks, dashed line: ETFs. X-axis: limit order lifetime in seconds. Y-axis: share of limit orders thathave been added and deleted on the same day.

Every order starts with a risk of non-optimality greater than zero, because no model can perfectly reflect reality, nor can it predict the future accurately. No model, however complex, can say with absolute certainty that the properties of any limit order are impartially optimal, because models are always less complex than reality. At any point in time, some external factors change and the trading system deletes the limit order in favour of a new one. This possibly leads to the clustering of observations at values close to zero in all the proxies. This 'baseline risk' is reflected by the rapid increase of the function especially in the regions close to zero. The more insecure the perception of optimality of the calculated limit order is, the more likely it is that the limit order will be deleted after a very short time. Because high-frequency traders do not mind acting within microseconds or any other speed that hardware and software allow, this risk level can be adjusted with almost infinite precision.

There exists a wide range of possible factors that influence the suitability of limit orders: bid and offer price, order book depth, order book change rate, ad hoc news, position of the order in the order book, (implied) volatility, each of these factors for some benchmark or a correlated asset, and many more. The value of the function and perhaps even the function itself changes in time, and they can change rather rapidly. The more limit orders there are, the quicker market conditions change and limit orders have to be adjusted more rapidly. Obviously, with a growing share of high-frequency trading engines on the market, this results in a self-nurturing process, because the algorithms change the very factors they observe.

We now turn to the scale parameter β_0 of the Weibull function and regress it against the log of the number of limit orders. This univariate regression shows that the more limit order activity there is for a security, the shorter-lived limit orders become. As mentioned earlier, the Weibull distribution fits the survival probabilities of limit orders remarkably well. However, the peak cannot be replicated by the smooth parameterised curve and distorts the parameter estimation. Hence, we left lifetimes of around 100 milliseconds out of consideration. Specifically, we ruled out lifetimes if $.095s \leq L_t \leq 0.105s$. We choose this range because the peak is not only at exactly 100 milliseconds but at some lifetimes t around it, probably due to effects of the IT infrastructure. We made an individual regression for ETFs and common stocks. The estimation results are shown in Table 1.

The two types of assets are, therefore, broadly comparable in terms of limit order lifetime changes in relation to limit order activity. The constant is a little smaller for stocks than for ETFs, but the slope of the regression curve is less steep. In both cases, however, the slope is positive. This means that the more limit orders arrive at the stock exchange, the higher the probability of very short-lived limit orders with a lifetime of only a handful milliseconds becomes.

Table 1: Regression results of $y = \alpha + \beta\mathbf{x} + \epsilon$, where the scale parameter β_0 is the dependent variable and the natural logarithm of the number of limit orders that have been added and deleted on the same day is the independent variable.

	Coeff.	Est. Result	Std. Err.	95% Conf. Int.	N	Adj. R^2	Avg. p
ETF	α	-3.182	0.530	-4.229, -2.135	175	.138	.321
	β	0.280	0.052	0.177, 0.382			
NASDAQ	α	-2.776	0.560	-3.880, -1.671	175	.099	.317
	β	0.246	0.056	0.138, 0.354			

With the addition of a dummy variable δ, which is 1 when $t \geq 0.1s$ and 0 otherwise, the Weibull distribution fits the actual data best. The fitted Weibull curve function is then

$$\mathbb{P}(T > t) = S(t) = \exp(-\exp(\beta_0)t^p) + \delta d,$$

where d is the value for the jump in the deletion probability. The value for d is on average 0.045 for ETFs and 0.058 for NASDAQ stocks. This indicates that, on average, around five per cent of all inserted limit orders without execution are deleted after precisely 100 milliseconds.

Limit order lifetimes over a trading day usually show a distinct break between market hours and non-market hours. Active trading algorithms that generate many short-lived limit orders and place and remove limit orders are most active during market hours when the limit order activity is highest. Figure 3 shows the results of an aggregation of limit order lifetimes and number of limit orders over the five trading days in February 2010. The figures were created by splitting the trading day into five-second intervals.

While traders place limit orders for both ETFs and NASDAQ-listed stocks in pre-market hours, the data only contain limit orders for ETFs in after-market hours.

In addition to the different average limit order lifetimes during market and non-market hours, the lifetimes during market hours show a rough reverse smile. The lifetimes shortly after market opening and before market close are shorter than at noon. The lower two figures show the total amount of limit orders arriving in each five-second interval. They show the well-known smile effect with more limit order activity at the beginning and at the end of the trading day than around noon (see for example, Jain and Joh (1988), Foster and Viswanathan (1993), or Biais et al. (1995)). This supports the hypothesis that high-frequency trade strategies are most active when many limit orders are in the market, perhaps in order to hide themselves from sniffer algorithms that seek to reverse-engineer other algorithms' strategies to frontrun them.

This concave pattern of average limit order lifetimes over the trading day fits smoothly into the framework of a limit order risk function. With a rapidly changing order book structure, it becomes more likely that a limit order placed at time t_0 becomes non-optimal for the market at some time $t_0 + x$. At the beginning and at the end of a trading day, trading activity is commonly at its highest level over the day. After market opening, market participants trade to find a consensus on the fair price of an asset and process the information of the night and from other markets. Before market close, traders often close their position in order to avoid overnight risk or take the position their portfolio manager has ordered. Of course, this leads to a lot of trading action, with many order insertions and -deletions. These activities increase the slope of the limit order risk function for lifetimes close to zero.

Fast markets with a lot of volatility should prove to be a good environment to test the concept of a limit order risk function as well as analyse order dynamics in turbulent conditions. Especially interesting is the so-called 'flash crash' of 6 May 2010, in which HFT was involved; see, for example CFTC and SEC (2010, pp. 45–

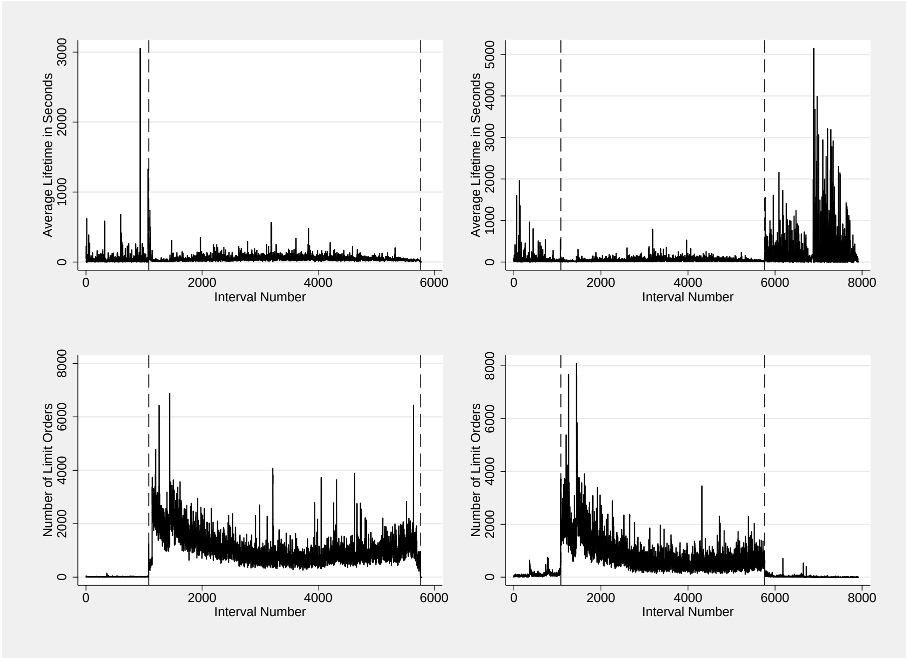

Figure 3: Top: average limit order lifetimes of limit orders for NASDAQ-listed stocks (left) and ETFs (right). Bottom: number of limit orders for the same period and assets. The x-axis represents the number of the five-second intervals, starting at 0 and representing the interval 08:00.00 - 08:00.04 AM, when the market opens. The dashed vertical lines show the beginning of market hours (09:30.00 AM, or the 1,080th interval, and 04:00.00 PM, or the 5,760th interval).

57). By investigating the trading day using the limit order lifetime proxy, we receive a further indication that it is likely that the proxy correlates with the intensity of high-frequency trading.

As Figure 4 shows, the average lifetime of limit orders on 6 May behaves as the average limit order lifetimes of the February lifetimes, only with a more elevated variance of the lifetimes over time than in Figure 3. At the time of the flash crash, the average lifetime of limit orders plummets from values ranging from around two to thirteen seconds to average lifetimes of around one second and lower. This indicates that during the flash crash, the share of high-frequency trading increased, because humans are unlikely to systematically insert and delete orders within much less than a second or so.

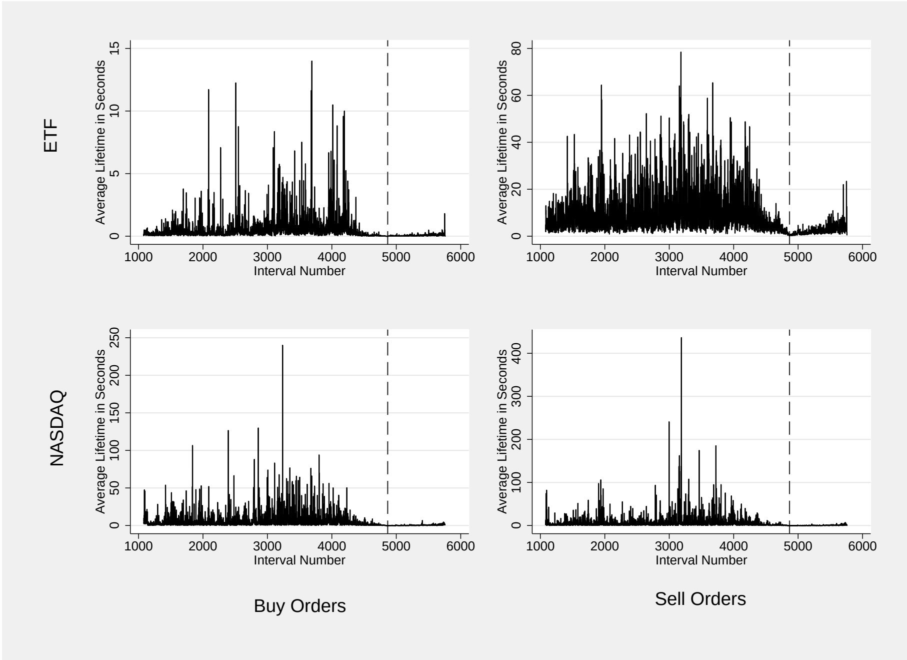

Figure 4: Upper graphs: average Lifetimes of limit buy (left) and sell (right) orders on ETFs. Lower graphs: average lifetimes of buy (left) and sell (right) orders of NASDAQ-listed stocks. Five-second-intervals, starting at 9:30 AM (interval number 1080). All graphs show the time 09:30 AM to 04:00 PM, i.e., the regular trading hours, of 6 May 2010. The dashed line highlights 2:45.30 (or interval number 4,866), when the Dow-Jones index reached its minimum on that day with a little more than nine per cent losses for a short period of time.

In distressed markets, there are two possibilities for high-frequency traders: the first possibility is to 'pull the plug' because the algorithm's model depends on 'normal' markets. Because the market changes very rapidly, the model does not yield limit orders with risk levels below the threshold. This causes the algorithm to temporarily halt trading. The second possibility is the opposite: to trade with a higher frequency because the market changes more rapidly. The faster the market changes, the faster increases the risk that the limit order does not fit to the market and needs to be replaced. This is the case if the algorithm's risk level for new orders is still below the risk threshold that would prevent it from trading.

As is visible in all lifetime graphs in Figure 4, the average lifetime decreases and reaches its minimum at around 2:45 PM, when the Dow-Jones hit the minimum value of the day with a minus of approximately nine per cent. This signifies that at the time of the market turmoil, high-frequency algorithms were very active, placing and deleting limit orders rapidly. It would be too much, however, to draw any definitive conclusion from this indicator.

The different structures of limit order lifetimes for ETFs and common stocks enable us to create the concept of limit order risk functions as perceived by market participants. If a trader deletes a limit order quickly after its placement, he or she gives up the time priority in exchange for not being exposed to the two main risks of limit orders, non-execution and adverse execution. In the next section, we will analyse the proxy order-revision time, which shows how long the trader waits after a deletion before he or she places a new order for the same stock.

3.2 Order Revision Times

Order revisions are a fundamental feature of the price formation process on order-driven stock markets. With the adjustment of limit order properties, traders can adjust their perception of the supply/demand curves for securities and contribute to the price-formation process. Even though most stock exchanges provide a built-in routine for replacing limit orders, it seems common to manually delete limit orders and manually insert a new one. Because limit order revisions are a fundamental part of active trading, it serves as the second proxy.

Human traders and computers alike constantly compute their optimal limit orders—according to their market perception. That means that the time between the deletion of a non-optimal limit order and the next insertion of a limit order with adjusted properties can be very small or even zero. The revision time depends on the trader's decision to wait or not to wait for a market reaction on the deletion of the limit order. However, with a timestamp precision of a nanosecond, it is theoretically possible but very unlikely to observe revision times of zero: the new order would have to be placed with less than a nanosecond delay, and the network technology would have to be perfectly constant, both of which is rather unlikely.

The slope of the limit order risk function for small values of t is usually very steep. To keep non-optimality risk low, traders adjust their limit orders very quickly in order to adapt to changing market conditions. We show that many securities show a distributional peak at 100 milliseconds in the distribution of order-revision times comparable to the one of limit order lifetimes. This is in line with the findings of Hasbrouck and Saar (2011), who employ two older datasets with more peaks in the hazard function. There, the lifetime distribution becomes smoother, a process that seems to have continued further. We expect a similar shape for average cumulated revision times as we find for order lifetimes. We assume it is more likely that a trader keeps trading in the same stock with a higher probability than foregoing it, so revision times should resemble lifetimes.

It is possible but unlikely that revision times are often shorter than a few milliseconds through the random placement of different market participants. The information of an order deletion has to arrive at the market participant, which even on very advanced platforms takes some 100 to 150 microseconds or so even for users with their trade engines co-located next to the exchange's servers. Adding the processing time for the now-changed order book and the time it takes to send a limit order back to the stock exchange, it is likely that at least around 0.3 to 0.5 milliseconds pass.

The market participant who deletes the limit order does not have to wait for the information of the deletion (as he or she generates it). As trivial as it sounds, he or she knows about the change of the order book caused by his or her deletion before all other market participants do. Thus, not considering the unlikely coincidental placement or deletion of a limit order by another market participant in such small timeframes, the market participant who is about to delete his or her limit order knows the structure of the limit order book before anyone else does. This enables him or her to calculate the optimal limit order for the then-changed order book. For the deleting market participant, the order-revision time can be infinitesimally small or even negative, if the new order is routed more quickly than the deletion message. Of course, only high-frequency engines can perform actions in such short intervals. Even though human traders might know exactly the properties of their next limit order they plan to place after they delete their old one, they will not be able to place it in a matter of milliseconds, probably not even in a tenth of a second.

Figure 5 shows histograms of order-revision times for the same two stocks and two ETFs used in Figure 1. The data reveals that very short order-revision times of only a few milliseconds are very common. In this set of two ETFs (VWO and QQQQ) and two stocks (CENX and INTC), order-revision times are lower than or equal to two milliseconds in some 40 per cent of the cases. It is noticeable that similar to limit order lifetimes, limit order revision times tend to occur every 50 microseconds. This explains the step-wise shape of the cumulated average limit order revision share.

To verify that the large shares of revision times of under two milliseconds in Figure 5 are not just special cases, we take the same securities as we did in Section 3.1 and calculate the average cumulated share of order-revision times. Again, the function describes a concave curve and resembles the one we discovered with the

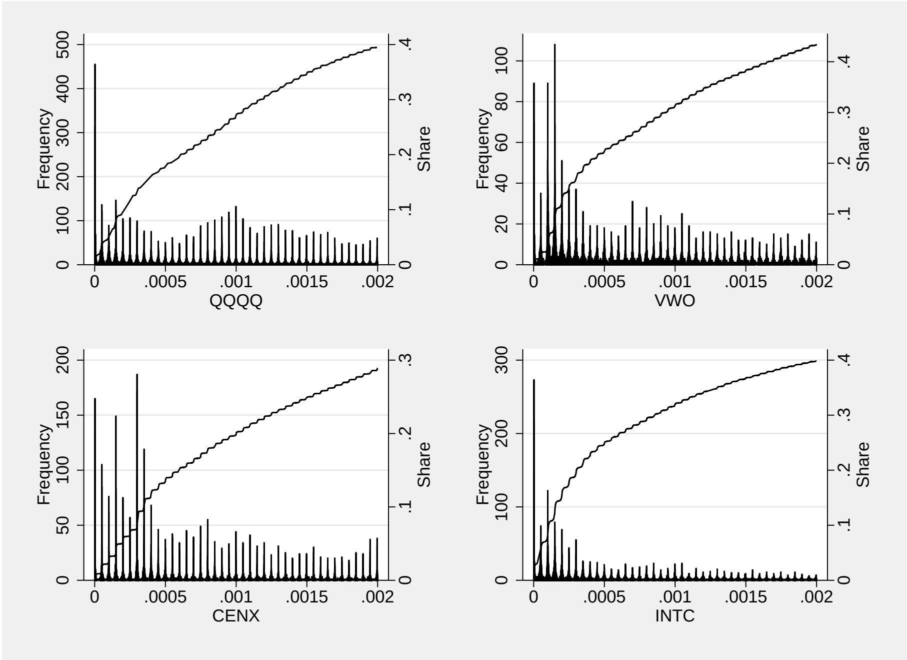

Figure 5: Histograms and cumulated proportions of limit order revision times on 22 February 2010 of VWO and QQQQ (ETFs) and CENX and INTC (NASDAQ-listed stocks). In each subfigure, the left y-axis shows the frequency of each limit order revision time (with increments of 0.2 microseconds), and the right y-axis shows the cumulated share of the order-revision time t with respect to all limit order revisions for the security on the same day.

limit order lifetime proxy. It is much steeper, which are shown by the exemplary figures in Figure 5.

The cluster of order-revision times of a few milliseconds may be the effect of the often-made argument that traders seek for liquidity by placing ultra-short-termed limit orders in quick succession. In the context of a limit order risk function, this repeated insertion of limit orders also makes sense for the traders. Whereas the deleted limit order climbed the limit order risk function over time, the trader can delete that limit order and add a new one with adjusted properties to start at a risk of zero that it does not fit the market—at least within the trader's model framework. Because computers can calculate optimal limit orders continuously, there is no necessity to wait after a deletion to place the new order. In some cases, e.g., if there is uncertainty, it may be advisable to wait a few fractions of a second for the market's reaction to the deletion. However, high-frequency traders do not seem to wait very often, as can be deduced from the distribution of the cumulative average revision times.

Figure 6 shows average cumulative limit order revision times for NASDAQ stocks and ETFs. It shows that orders for ETFs are revised differently than stocks. ETFs tend to have shorter revision times than common stocks, indicating that high-frequency traders are more active on those structured products than on plain vanilla equity.

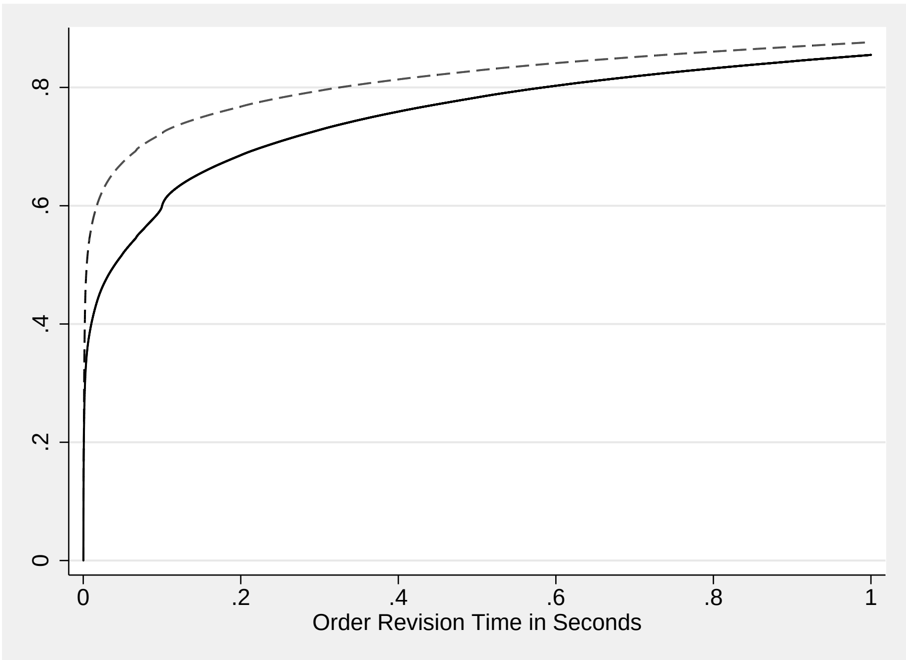

Figure 6: Average limit order revision times for ETFs (dashed line) and stocks listed on NASDAQ (solid line). The lines show the average cumulated share of order-revision times t relative to all observations of the message sequence [Delete–Add] per day per security.

The cause is probably the structured nature of ETFs. Market participants—traders, and especially in the case of ETFs, market makers—know the constituents of the ETF and can easily calculate its net asset value (NAV). This makes the risk that the limit order is not optimally suited for the current market, lower than for common stocks, where traders always face the problem that they do not know its fundamental value. Because this knowledge is relatively easily available in the case of ETFs, the only edge traders have is speed, which makes the use of algorithms pivotal. Notice that the curve for stocks shows a small jump at around 0.1 seconds, which does not exist for ETFs.

Figure 7 shows that order-revision times over the day behave as one would assume. As the order flow decreases around noon as shown in the lower half of Figure 3, the average time that passes between the deletion of a limit order and the arrival of a new limit order for the same stock increases. Before market opening, and especially after market close, order revisions become rather scarce, which leads to much longer order-revision times than during market hours. The informational content of non-market hours is only limited, so we do not include them in this figure.

To see how active trading changes in distressed markets, we employ the data of 6 May 2010 as for the limit order lifetime. We create several order-revision time figures with the technique employed in Figure 6. According to CFTC and SEC (2010, p. 57), the share of high-frequency trading on total market volume hovered a little over 40 per cent over the day and peaked at 50per cent at 2:45 PM, when the Dow-Jones index hit its minimum. Their figures refer to 17 HFT firms, i.e., high-frequency trading firms. They split up the day into 15-minute intervals, and for each the share of the 17 HFT firms on the overall market is given. They are thus minimum values, as other algorithmic market participants were not taken into the dataset. The cumulative share of limit order revisions on 6 May for four 15-minute intervals is given in Figure 8.

It becomes apparent that the eagerness to place new orders after the deletion of a previous limit order is greater for ETFs than for NASDAQ stocks. This hints

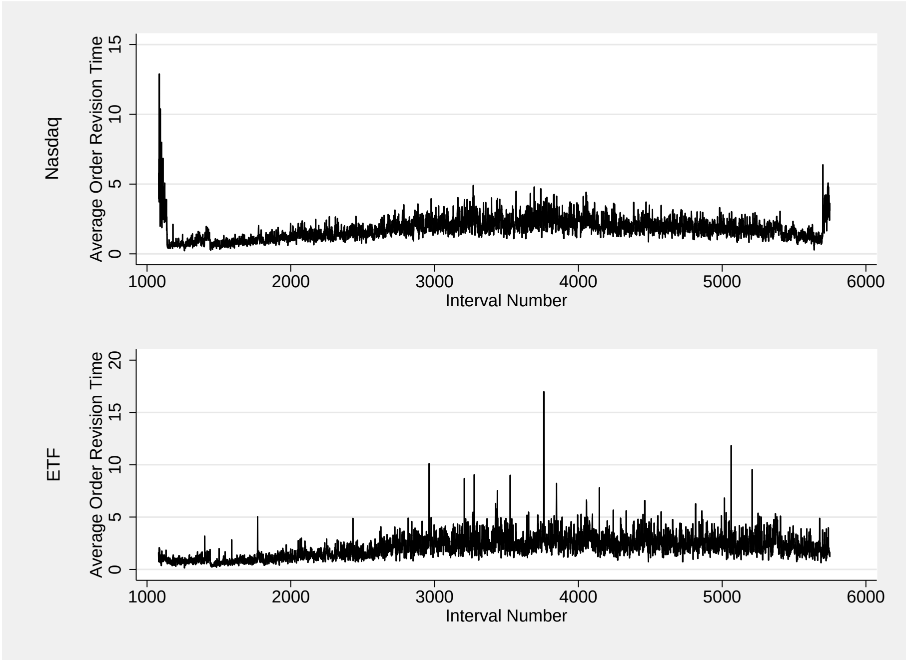

Figure 7: Average order-revision times over the trading day. The x-axis is the interval number; each interval represents five seconds of trading. The figures show the time 9:30 AM to 4:00 PM. Each dataset consists of 36 stocks or ETFs, respectively. The trading week is 22–26 February, i.e., five trading days.

at a greater share of algorithmic price formation for ETFs than for common stocks. At 10:00 AM as well as at 12:00 PM on 6 May, the figures are comparable to a very quiet market, given in Figure 6. At the time of the flash crash, however, which happened between around 2:30 and 3:00 PM, the speed of limit order revisions increases rapidly for both limit orders to sell and limit orders to buy.

This proxy alone probably does not yield a relative share of high-frequency trading in the market. However, a decrease of limit order revision times indicates a higher proportion of high-frequency trading if the value is adjusted for a faster market, which automatically brings down limit order revision times with the noise it generates. Ultra-fast revision times of only a few milliseconds can possibly serve as an approximate indicator for high-frequent limit order activity. However, with-

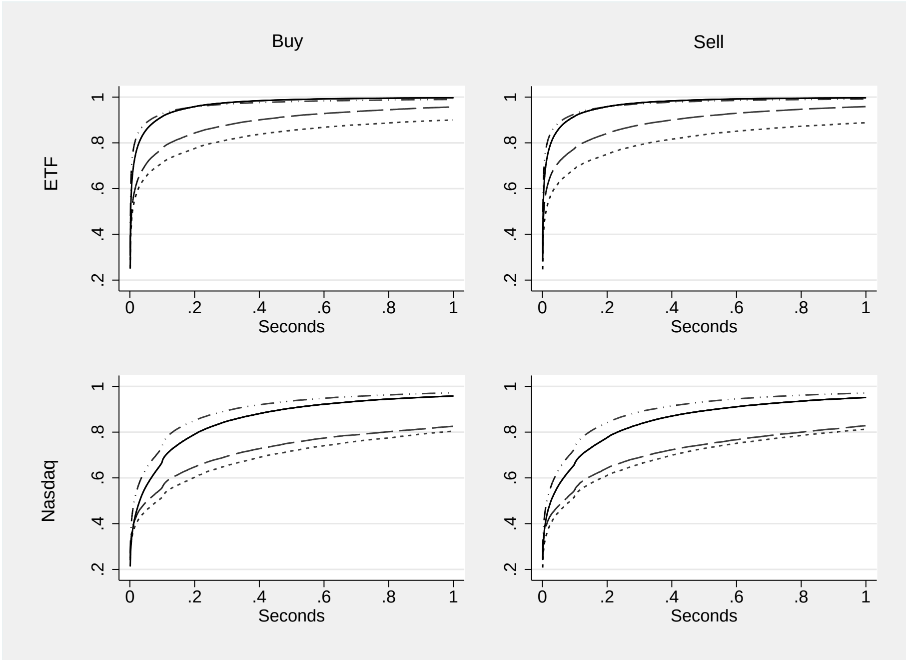

Figure 8: Average share of order-revision times t within four 15-minute intervals on 6 May 2010. —— interval from 2:30 PM to 2:45 PM; – · · – interval from 2:45 PM to 3:00 PM; — — interval from 10:00 AM to 10:15 AM; - - - interval from 12:00 AM to 12:15 PM. For example, of all the observed occurrences of the message flow of the form [Delete–Add] that were sell orders on ETFs (upper right figure), in around 80 per cent of the cases, a new order was placed within 0.4 seconds from 12:00 AM to 12:15 PM.

out a reference dataset with an indicator for algorithmic orders, it is impossible to construct a measure.

3.3 Inter-Order Placement Times

Although some research papers analyse inter-trade or inter-transaction durations (e.g., Engle and Russell (1998), Ivanov et al. (2004)), to the best of our knowledge, there are no scientific papers on the inter-order placement duration, i.e., the time that passes from the placement of one limit order until the next limit order arrives.

Limit order insertions in quick succession can have various origins. For example, the much criticised so-called 'quote-stuffing' works this way. If one market participant places many limit orders at once for one stock, the algorithms of other market participants have to read and process them, which consumes calculation time. They face a (possibly small) time disadvantage. The market participant who placed the limit orders does not have to react to the new limit orders, he or she knew the structure of the limit orders beforehand simply because he or she placed them. Due to the fact that the dataset is anonymous, we cannot say if the limit orders we analyse are part of a quote-stuffing attempt or not. We can, however, say with a high probability that ultra-short inter-order placement times originate in high-frequency trading.

Inter-order placement times tend to be very short by nature. If there are many limit order insertions, even without order clustering the time between the arrival of limit order insertions decreases. In addition to that, it is well-known that liquidity attracts liquidity, causing inter-order placement times to further diminish at times of much limit-order activity. But for the ultra-short inter-order placement times that we observe, this explanation alone does not suffice. Figure 9 shows the inter-order placement times of four securities for up to two milliseconds and the cumulated share relative to all placed limit orders on the day (22 February).

In the case of QQQQ and INTC, with a probability of 50 per cent a trader would only have to wait two milliseconds after the insertion of a limit order to see the next one to be placed. Perfectly distributed over the trading day with 23,400 seconds, we would expect inter-order placement times distributed around average values of around 0.2 seconds for QQQQ (125,680 orders) or 0.5 seconds for INTC (489,777 orders). For the securities VWO and CENX, the probability of a new limit order insertion within two milliseconds after the last one lies between 25 and 30 per cent.

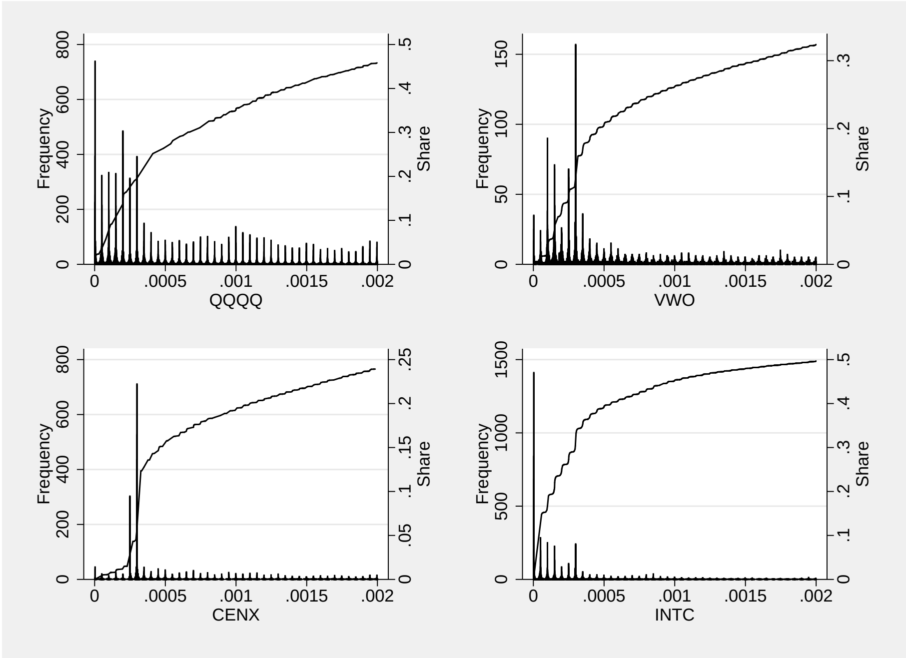

Figure 9: Inter-order placement times of QQQQ and VWO (two ETFs) and CENX and INTC (two stocks listed on NASDAQ). The bars show the frequency of inter-order placement times t, the solid line shows the cumulated share of limit orders at t relative to all limit orders placed on that day for the individual security.

The quickest inter-order placement time is two microseconds, i.e., 0.000002 seconds—which is clearly not coincidental. The bulk of inter-order placement times is placed within less than a millisecond. Note the different shapes of the histograms given in Figure 9. Every security shows the clustering around multiples of 1/20,000 of a second. While 3 per cent of all limit orders of CENX, for example, arrive earlier than 0.0002 seconds after the last limit order, it shows a massive clustering of limit orders at 0.00025 and 0.0003 seconds—almost 12 per cent of limit orders are placed with exact that speeds. The limit orders of QQQQ, in comparison, come in at a higher rate, around 20 per cent of the limit orders are inserted within 0.0004 seconds. As for the limit order revision time, is not possible to say with certainty that the two successive orders generating such short inter-order placement durations come from the same trader, because the dataset is anonymous and the realised latency of the stock market and the market participants to process the order is unknown.

To compare ETFs and common stocks with each other, we calculate the average shares of inter-order placement times t relative to all limit orders of the individual securities. The results are given in Figure 10. The figures show five curves, representing the inter-order placement durations for the intervals (0, 1/10,000), (0, 1/1,000), (0, 1/100), (0, 1/10), and (0, 1) seconds. For example, for the sample of 36 ETFs and five trading days, with a probability of 40 per cent the average duration between two successive limit order insertions is only $0.6 \times 1/100s = 0.006$ seconds. This indicates that limit orders are clustered. NASDAQ-listed stocks are quite comparable to ETFs regarding this proxy, except the jump at 100ms.

Observe that at 3/10,000 of a second or 300 microseconds, there is a peak in the distribution of inter-order placement times, as can be seen from the curve labelled 1/1,000 of a second. In addition to that, the solid line of NASDAQ stocks shows a jump at 0.1 seconds that are possibly part of strategic runs as found by Prix et al. (2007, 2008), and Hasbrouck and Saar (2011). The jumps indicate that there are relatively static algorithms in the market that wait 300 microseconds or 100 milliseconds after a limit order has been added and place a new one.

The often much shorter inter-order placement durations for more active stocks and ETFs indicate that algorithms take up speed with an increased order flow. This behaviour can also be explained with the concept of a limit order risk function. The more limit orders arrive, the faster increases the risk that a limit order placed in the market some time earlier turns non-optimal. This leads the algorithm to delete the old order in favour of a new one, leading to decreased order revision- and decreased inter-order placement times.

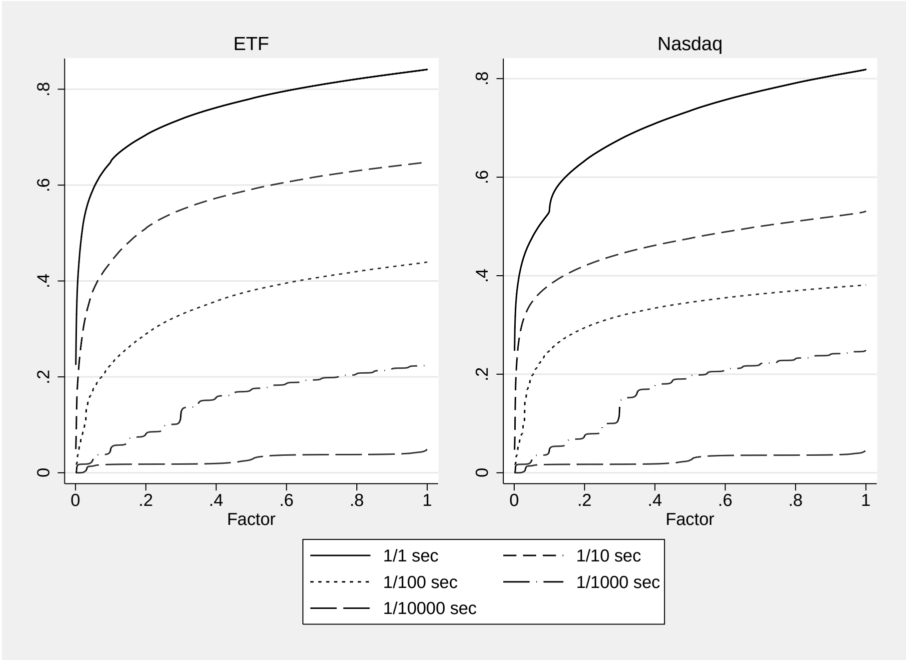

Figure 10: The lines show the different cumulated probabilities of order-revision time as being equal to t. Each curve represents different intervals, which can be calculated by multiplying the factor given at the x-axis with the corresponding time scale given in the legend. For example, the solid line (1/1 sec) shows the interval (0, 1) seconds, the dashed line directly below it the interval (0, 1/10) seconds and so forth. For example, with a probability of 80 per cent, a new order arrives within one second after the last insertion for both ETFs and stocks.

4 Conclusion

We analyse raw order book message data for traces of high-frequency trading. From the analysis of the microstructure of order dynamics in time frames that are still perceivable by humans (as described above), we know that a great deal of limit order activity occurs in a few milliseconds' time. Until recently, the timestamp precision of most datasets 'only' reached milliseconds, prohibiting a thorough and detailed analysis of ultra-high frequency algorithms. This paper aims at helping close this gap.

We operate with an order book protocol from the US stock exchange NASDAQ. It logs all order book events, such as limit order insertions, deletions, executions, etc. In order to give traders and other market participants an idea of the way high-frequency trading works on a modern order-driven market, we analyse the structure of limit order lifetimes, limit order revision times, and inter-order placement times. All three proxies show a clustering of observations at a few milliseconds, which makes an analysis of pure high-frequency trading behaviour impossible with timestamp precisions of a millisecond or worse. The dataset has timestamps which are exact to the nanosecond. This enables us to perform analyses at a greater accuracy than ever before, which is necessary to observe high-frequency traders that currently operate at microsecond levels.

The limit order lifetime, i.e., the time from insertion to the deletion of a limit order, is the first proxy we analyse. It is also the most exact one; the data does not produce any noise, because every order is equipped with a day-unique order reference number. Many limit orders for common stocks and ETFs are only active for a few milliseconds, often only a few microseconds, before they are deleted. This proxy shows clusterings of observations at multiples of 50 microseconds. Order dynamics differ for ETFs and common stocks. A greater proportion of limit orders for common stocks than for ETFs is deleted within less than 100 milliseconds. During the flash crash on 6 May 2010, the average limit order lifetimes plummeted to very small values both for buy orders and for sell orders.

Order revision times, i.e., the time that passes between a deletion of a limit order and the next insertion of a limit order, are very small. Their density resembles the one of limit order lifetimes, they also show peaks at multiples of 50 microseconds. As in the case of limit order lifetimes, revision times for ETFs and for common stocks are different. This becomes apparent in the average cumulated share of revision times;

the slope is much steeper for ETFs for values close to zero than for common stocks. During the flash crash on 6 May 2010, revision times for both ETFs and common stocks decreased, indicating that the share of high-frequency trading increased as compared to ‘normal’ markets.

The inter-order placement time measures the time that passes between two successive limit order insertions for a stock. A great share of inter-order insertion durations is very short. On average, in 50 per cent of the cases a new order is placed within less than 100 milliseconds. This proxy shows the peaks at multiples of 50 microseconds that can be observed for lifetimes and revision times. For very active stocks or ETFs, this time decreases to two or three milliseconds. The differences between ETFs and stocks are rather negligible in comparison with the two other proxies revision times and lifetimes. As one would expect, inter-order placement times decreased for both asset classes during the flash crash.

The structures of the limit order proxies can possibly be explained by a limit order risk function. The function describes the increasing risk of a limit order not to be optimal for the market any more. It is an increasing function that can depend on various factors that influence both the optimal order strategy and tactics, such as, for example, depth of the order book, volatility, implied volatility, depth of correlated securities, and any other factor the trader deems important for an optimal order strategy.

The short order-revision times and lifetimes could be a direct result of this. To illustrate this for limit order revision times: if a trader deletes an existing limit order, a newly inserted order restarts at a risk of non-optimality of zero. In the case of order lifetimes: if a previously inserted limit order reaches a threshold level of inappropriateness, the algorithm or trader deletes the order and inserts a new one that is optimal according to the employed model. High-frequency algorithms can do both of these things very quickly. They rapidly read and process large amounts of market information. They feed this information to their order calculation models, which act accordingly. The more powerful information and communication technology becomes, the faster the value of the limit order risk function changes, which results in decreasing values for the proxies. At the time of the flash crash, the order-revision times decreased significantly, because the market itself generated a lot of new information through large price movements.

ETFs are treated differently than common stocks. Their limit order lifetimes are on average longer, their revision times shorter, and their inter-order placement times are also shorter. This can be explained by the different structure of ETFs compared to common stocks. Because ETFs are usually diversified portfolios with a known inner structure, their fundamental value can relatively easily be calculated via its net asset value. This lowers the risk that an inserted order is not optimal for the current market conditions. This has a positive effect on limit order lifetimes, because the suitability of the order for the market decreases more slowly, especially in the regions close to zero. For the order-revision time, the effect is negative, because the trader can calculate the optimal order with less uncertainty for ETFs than common stocks, so a new order can be placed with no delay. It goes without saying that this leads to clusters of observations close to zero, making a very accurate timestamp precision necessary.

Future research will include modelling and testing a limit order risk function and the connection of the proxies to the structure of the order book. It will be interesting to see if the ultra-short lifetimes, order-revision times or inter-order placement times happen within, at or slightly away from the BBO.

Acknowledgement

We are grateful to NASDAQ OMX for generously providing this order message dataset.

A Stocks and ETFs in the Dataset

Table 2: Ticker symbol, type, name, and number of limit orders of the stocks and ETFs used in the empirical analysis.

No.	Ticker Symbol	Type	Company Name	Limit Orders
1	IYR	ETF	iShares Dow-Jones US Real Estate Index Fund	313,998
2	SPXU	ETF	ProShares UltraPro Short S&P 500	302,766
3	IAU	ETF	iShares Gold Trust	278,115
4	VWO	ETF	Vanguard MSCI Emerging Markets ETF	243,793
5	SMDD	ETF	ProShares UltraPro Short MidCap400	229,215
6	DUG	ETF	ProShares UltraShort Oil & Gas	224,297
7	EDC	ETF	Direxion Daily Emerging Markets Bull 3X Shares	210,257
8	SH	ETF	ProShares Short S&P500	189,223
9	TLT	ETF	iShares Barclays 20+ Year Treasury Bond Fund	184,564
10	IXG	ETF	iShares S&P Global Financials Sector Index Fund	183,516
11	UPRO	ETF	ProShares UltraPro S&P 500	170,674
12	USD	ETF	ProShares Ultra Semiconductors	149,717
13	XHB	ETF	SPDR S&P Homebuilders ETF	147,519
14	TBT	ETF	ProShares UltraShort 20+ Year Treasury	146,791
15	IEO	ETF	iShares Dow-Jones US Oil & Gas Exploration & Production Index Fund	146,481
16	SMH	ETF	Semiconductor HOLDRs Trust	142,775
17	SLV	ETF	iShares Silver Trust	14,2620
18	XOP	ETF	SPDR S&P Oil & Gas Exploration & Production ETF	141,408
19	FPL	ETF	Futura Polyesters Ltd	138,815
20	KRE	ETF	SPDR KBW Regional Banking ETF	132,532
21	XLK	ETF	Technology Select Sector SPDR Fund	132,232
22	TWM	ETF	ProShares UltraShort Russell2000	131,086
23	DAI (now DUST)	ETF	Direxion Gold Miners Bull 2X Shares	129,203

Continued on next page

Table 2 – continued from previous page

No.	Ticker Symbol	Type	Company Name	Limit Orders
24	EWC	ETF	iShares MSCI Canada Index Fund	125,131
25	ERY	ETF	Direxion Daily Energy Bear 3X Shares	121,007
26	XLI	ETF	Industrial Select Sector SPDR Fund	120,967
27	OEF	ETF	iShares S&P 100 Index Fund	119,704
28	EPP	ETF	iShares MSCI Pacific ex-Japan Index Fund	117,567
29	ICF	ETF	iShares Cohen & Steers Realty Majors Index Fund	116,879
30	XLV	ETF	Health Care Select Sector SPDR Fund	114,838
31	TRA	ETF	Terra Industries Inc	112,987
32	IWF	ETF	iShares Russell 1000 Growth Index Fund	111,749
33	DOG	ETF	ProShares Short Dow30	110,545
34	EEV	ETF	ProShares UltraShort MSCI Emerging Markets	107,322
35	LHB	ETF	Direxion Daily Latin America Bear 3X Shares	104,625
36	AGQ	ETF	ProShares Ultra Silver	102,882
1	QCOM	NASDAQ	QUALCOMM Inc	314,698
2	INTC	NASDAQ	Intel Corp	307,902
3	CENX	NASDAQ	Century Aluminum Co	283,047
4	BRCM	NASDAQ	Broadcom Corp	244,583
5	STLD	NASDAQ	Steel Dynamics Inc	229,428
6	MSFT	NASDAQ	Microsoft Corp	225,705
7	STX	NASDAQ	Seagate Technology PLC	211443
8	BRCD	NASDAQ	Brocade Communications Systems Inc	189,830
9	CSCO	NASDAQ	Cisco Systems Inc	184,703
10	WFMI (now WFM)	NASDAQ	Whole Foods Market Inc	181,518
11	ORCL	NASDAQ	Oracle Corp	170,411

Continued on next page

Table 2 – continued from previous page

No.	Ticker Symbol	Type	Company Name	Limit Orders
12	INTU	NASDAQ	Intuit Inc	149,600
13	NTAP	NASDAQ	NetApp Inc	148,801
14	CTXS	NASDAQ	Citrix Systems Inc	147,367
15	CTSH	NASDAQ	Cognizant Technology Solutions Corp	146,316
16	SNDK	NASDAQ	SanDisk Corp	143,250
17	UAUA (now UAL)	NASDAQ	United Continental Holdings Inc	142,098
18	MRVL	NASDAQ	Marvell Technology Group Ltd	141,903
19	NIHD	NASDAQ	NII Holdings Inc	138,778
20	ASML	NASDAQ	ASML Holding NV	132,701
21	DTV	NASDAQ	DIRECTV	132,340
22	MU	NASDAQ	Micron Technology Inc	131,756
23	XLNX	NASDAQ	Xilinx Inc	130,376
24	GILD	NASDAQ	Gilead Sciences Inc	126,173
25	LLTC	NASDAQ	Linear Technology Corp	122,328
26	APOL	NASDAQ	Apollo Group Inc	120,823
27	TEVA	NASDAQ	Teva Pharmaceutical Industries Ltd	119,679
28	MCHP	NASDAQ	Microchip Technology Inc	118,128
29	PTEN	NASDAQ	Patterson-UTI Energy Inc	116,396
30	NVDA	NASDAQ	NVIDIA Corp	113,321
31	KLAC	NASDAQ	KLA-Tencor Corp	112,373
32	AMLN	NASDAQ	Amylin Pharmaceuticals Inc	111,876
33	CMCSA	NASDAQ	Comcast Corp	111,031
34	HCBK	NASDAQ	Hudson City Bancorp Inc	104,100
35	LBTYA	NASDAQ	Liberty Global Inc	104,095
36	SBUX	NASDAQ	Starbucks Corp	102,938

References

Biais, B., Hillion, P., and Spatt, C. (1995). An empirical analysis of the limit order book and the order flow in the paris bourse. *The Journal of Finance*, 50(5):1655–1689.

CFTC and SEC (2010). Findings regarding the market events of May 6, 2010. Technical report, U.S. Commodity Futures Trading Commision, U.S. Securities & Exchange Commission.

Cleves, M. A., Gould, W. W., and Gutierrez, R. G. (2002). *An introduction to survival analysis using Stata*. Stata Press, College Station, Texas.

Engle, R. and Russell, J. (1998). Autoregressive conditional duration: A new model for irregularly spaced transaction data. *Econometrica*, 66(5):1127–1162.

Foster, F. D. and Viswanathan, S. (1993). Variations in trading volume, return volatility, and trading costs: Evidence on recent price formation models. *The Journal of Finance*, 48(1):187–211.

Groth, S. (2009). Algorithmic trading engines and liquidity contribution: The blurring of "traditional" definitions. In Godart, C., Gronau, N., Sharma, S., and Canals, G., editors, *Software Services for e-Business and e-Society*, volume 305, pages 210–224. Springer Boston.

Gsell, M. (2009). Algorithmic activity on Xetra. *The Journal of Trading*, 4(3):74–86.

Handa, P. and Schwartz, R. A. (1996). Limit order trading. *Journal of Finance*, 51(5):1835–1861.

Hasbrouck, J. and Saar, G. (2001). Limit orders and volatility in a hybrid market: The Island ECN. Working paper, NYU.

Hasbrouck, J. and Saar, G. (2009). Technology and liquidity provision: The blurring of traditional definitions. *Journal of Financial Markets*, 12(2):143–172.

Hasbrouck, J. and Saar, G. (2011). Low-latency trading. Working paper, Stern School of Business (NYU) and Cornell University.

Hendershott, T., Jones, C. M., and Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *Journal of Finance*, 66(1):1–33.

Hendershott, T. and Riordan, R. (2009). Algorithmic trading and information. Unpublished Working Paper #09-08, NET Institute.

Ivanov, P. C., Yuen, A., Podobnik, B., and Lee, Y. (2004). Common scaling patterns in intertrade times of U. S. stocks. Quantitative finance papers, arXiv.org.

Jain, P. and Joh, G. (1988). The dependence between hourly prices and trading volume. *Journal of Financial and Quantitative Analysis*, 23(3):269–283.

Maurer, K.-O. and Schäfer, C. (2011). Analysis of binary trading patterns in Xetra. *The Journal of Trading*, 6(1):46–60.

Muthuswamy, J., Palmer, J., Richie, N., and Webb, R. (2011). High-frequency-trading: Implications for markets, regulators, and efficiency. *The Journal of Trading*, 6(1):87–97.

Prix, J., Loistl, O., and Hütl, M. (2007). Algorithmic trading patterns in Xetra orders. *The European Journal of Finance*, 13(8):717–739.

Prix, J., Loistl, O., and Hütl, M. (2008). Chain-structures in Xetra order data. Working paper, Vienna University of Economics and Business Administration.

Storkenmaier, A. and Wagener, M. (2010). Research on competition and regulation in European equities trading: An executive summary. Technical report, Karlsruhe Institute of Technology.