Joint Discussion Paper
Series in Economics

by the Universities of

Aachen · Gießen · Göttingen
Kassel · Marburg · Siegen

ISSN 1867-3678

No. 15-2015

Alexandru Mandes

Impact of inventory-based electronic liquidity providers within a high-frequency event- and agent-based modeling framework

This paper can be downloaded from
http://www.uni-marburg.de/fb02/makro/forschung/magkspapers/index_html%28magks%29

Coordination: Bernd Hayo • Philipps-University Marburg
School of Business and Economics • Universitätsstraße 24, D-35032 Marburg
Tel: +49-6421-2823091, Fax: +49-6421-2823088, e-mail: hayo@wiwi.uni-marburg.de

Impact of inventory-based electronic liquidity providers within a high-frequency event- and agent-based modeling framework

Alexandru Mandes*
University of Giessen
Licher Str. 64, 35394 Giessen, Germany
Alexandru.Mandes@wirtschaft.uni-giessen.de

June 4, 2015

Abstract. This contribution addresses the impact of high-frequency electronic liquidity provision strategies on financial markets' intraday dynamics, by evaluating the interaction between multiple trading strategies within a computer laboratory, i.e. an artificial stock market. Initially, a realistic base-line model is set up around a continuous double auction market, with trading being pursued only by four types of low-frequency market participants. Sequentially, the high-frequency agents are added to the model and the corresponding changes related to various measures of market quality and market systemic risk are analyzed, under both regular and market stress conditions, such as when the order flow balance is suddenly disrupted by a large volume-in-line sell program. A detailed intraday analysis of a flash crash emergence is also conducted. Finally, possible regulatory policies such as minimum holding or quote resting time and financial-transaction taxes are assessed.

Key Words: agent-based modeling, continuous double auction, high-frequency trading, electronic liquidity provision, market quality, systemic risk, flash crash, regulatory policies

JEL classification: C63, G17, G28

*The author would like to thank Deutscher Akademischer Austauschdienst (DAAD) for the awarded PhD scholarship. Also, the author is grateful to WEHIA 2015 participants for their comments and suggestions.

Introduction

The recent technological advances and the development of electronic financial markets have led to an unprecedented level of market automation, whose consequences are still a heated subject of media debate. According to a TABB group estimate presented in November 2009, the market volume share accounted for by high-frequency trading (HFT) in the USA has reached 61%.[1] A more recent report published by the ESMA, covering a sample of 100 stocks from nine EU countries for May 2013, shows that the HFT activity accounted for between 24% or 43% of the total traded value, depending on the estimation method, i.e. based on the primary business of firms (direct approach) or on the lifetime of orders (indirect approach). For the number of trades the corresponding numbers for HFT activity were 30% and 49%, and for the number of orders 58% and 76% respectively.[2]

In 2012 the UK Government Office for Science has commissioned a foresight project to address the complex issue of computer generated trading in financial markets and to provide policy strategic options (*Foresight: The Future of Computer Trading in Financial Markets* (2012)). The scientific evidence shows that, on one side, the general market quality has been improved, but on the other side there might be a greater risk of periodic illiquidity. Still, many specific questions have remained unanswered due to the significant challenges in the empirical evaluation of HFT, determined either by a lack of proper data identification (Friederich and Payne (2011)) or by endogeneity issues – HFT growth coincides with the 2008–2009 market turmoil (Brogaard (2010)). For example, one open problem refers to the high-frequency traders' (HFTs) liquidity providing behavior during regular market conditions versus situations of market stress. Friederich and Payne (2011) question the robustness of a system relying on the liquidity provided by HFTs without any obligations of staying in the market, as well as the economic value of predatory strategies which profit at the expense of traditional investors. Nevertheless, even if passive HFTs would not flee during the periods of high order toxicity, it might be that their trading strategies could temporarily amplify some liquidity issues, even if they are not their direct cause. Regarding the causal relation between HFT and volatility, Brogaard (2010) and Linton (2011) show that it is not clear whether larger volatility is the effect of HFT or just a favorable condition which stimulates HFT activity. Given that groups of stable systems could interact in highly unstable ways ("fallacy of composition"), Danielsson and Zer (2012) suggest that the systemic risk of computer-based strategies is currently not well understood. One factor which might contribute to an increase of the systemic risk is a large homongeneity of strategies, e.g., when large clusters

[1] "High Frequency Trading – What Is It & Should I Be Worried?", Larry Tabb, World Federation of Exchanges, 2009

[2] "High-frequency trading activity in EU equity markets", European Securities and Markets Authority, Economic Report Number 1, 2014

of market participants are following exactly the same strategies. Farmer and Skouras (2011) and Sornette and Von Der Becke (2011) underline the chaotic properties of the financial system, determined by the complex interactions between the market participants, which can sometimes exhibit very different behaviors as a result of very small initial state changes. Also, it is not clear to what extent do HFTs contribute to these non-linear dynamics and reinforcing feedback loops. Regarding the price discovery mechanism, Brogaard (2010) and Brogaard et al. (2012) conclude that by trading in the direction of the incoming information, HFTs positively contribute to price efficiency. However, since HFTs put more emphasis on short term rather than fundamental information, the arrival of information – concentrated or diffused – is also important. Finally, with respect to policy making, further analysis for a better understanding of the potential impact of the various regulatory options is recommended in order to prevent any unexpected and detrimental market reactions.

Given the challenges faced by the empirical approach, a viable alternative for testing the various hypothesis regarding the impact of HFT, as well as the potential effect of policy regulations, is to run simulations within an artificial stock market by means of agent based modeling. A list of prominent milestones of designing intraday financial market models should start with Chiarella and Iori (2004), Chiarella et al. (2009) who have developed a continuous double auction model with stylized agents corresponding to human strategies. The authors have shown that, by including chartist strategies, the simulated price time series exhibit realistic stylized facts such as the fat tails of returns, volatility clustering and the fat-tailed distribution of limit order placement relative to the midpoint. Gsell (2008) has extended the previous model and made a first step towards modeling computer-based trading by adding two different implementations of an algorithmic trader. More recently, Vuorenmaa and Wang (2013) has included market-makers with a tight inventory control and simulated the 2010 Flash Crash (May 6, 2010) scenario, where a sudden market drop was triggered by a volume-inline execution algorithm. Their results show that the probability of a flash crash increases with the number of high-frequency (HF) agents, tightness of inventory sizes and smaller tick size, while quality metrics such as spread-at-touch and volatility depend diametrically on the previous parameters.

However, the previous intraday market designs suffer from a series of severe limitations. For once, the time dimension is not properly operationalized – being reduced to an atemporal sequence of simulation steps – and therefore these models do not function as expected when agents are supposed to act at different time frames, as detailed in Section 2. Secondly, low-frequency (LF) agents place their trading orders randomly and do no react to the changing microstructure conditions. In reality, the order book is not just a trading instrument, but also an important source of endogenous information, which influences to a large extent the way HFTs behave. The importance of building a realistic and robust model, which could then be successfully used as a laboratory for policy simulation, has been underlined by Prof. Alan Kirman during the WEHIA 2015 round table: "If we wish to build a model which will be useful for studying the effect of policy measures we should base it on realistic assumptions derived from observed empirical behavior. However, we should also make sure that simulations of that model do not depend on a very limited set of parameter values for these assumptions. We want a model which is robust in the sense that small modifications in the assumptions do not change radically the nature of the results."

The current contribution proposes a higher fidelity intraday market model with a proper implementation of time which allows for simulating trading strategies at different temporal frequencies. Additionally, a microstructure-based order submission mechanism as in Mandes (2014) is included, which allows LF agents to react to the current liquidity and volatility conditions. Based on the proposed model, a series of simulations is run in order to answer several research questions regarding the impact of HFT on market quality and systemic risk, as well as the effects and side-effects of two policy measures. Section 1 presents a list of market quality indicators which helps in assessing the various market scenarios. Section 2 introduces the underlying framework of the agent-based models, as well as the various LF and HF market participants. In Section 3, the set of simulated scenarios are described, covering different market configurations, agent parameterizations and possible market states. The main outcomes are interpreted in Subsection 3.2 and the flash crash event is analyzed distinctly in Subsection 3.3. Two potential regulatory policies, i.e. imposing a financial-transaction tax and minimum holding/ quote resting times, are evaluated in Section 4 and afterwards the final conclusions are drawn.

1 Market quality measures

In Brogaard (2010) and *Foresight: The Future of Computer Trading in Financial Markets* (2012) market quality is assessed along three different dimensions: (i) liquidity and transaction costs, (ii) price efficiency and price discovery, and (iii) volatility and financial market stability.

Liquidity and transaction costs

According to Harris (2002), liquidity is the most important characteristic of well-functioning markets, reflecting the ability of executing large orders in a short period of time and with low trading costs. Besides the explicit cost components related to commissions and other fees, which are constant and fixed by brokers or stock exchanges, the implicit costs are highly dependent on the variable market conditions and often have a greater impact on the overall trading performance (up to 80% of overall transaction costs).[3] For example, if the spread is wide, the execution of a market order is expensive, i.e. the cost of immediacy is high. Furthermore, if the order size is large, one has to take into account also the available total quantity, as well as the dispersion of resting limit orders, which are to be sequentially matched until the entire market order is filled ("walking up the book").

Harris (2002) summarizes the different attributes of liquidity as: immediacy, market width, market depth and market resiliency. The immediacy, also known as tightness or spread-at-touch, represents the cost of liquidity for small orders. It is reflected by the bid-ask difference or, more common, by the relative spread computed as the ratio of the bid-ask spread to mid-price. The market width or breadth captures the market impact of executing larger orders and can be measured as the cost of filling a given order size. Alternatively, the market width could be sampled by recording the potential market impact of all orders submitted by fundamental traders, which corresponds to an unconditional market impact function. If only actual executions are taken into consideration, the average effective spread can be computed as the share-weighted average of the double amount of the signed difference between the order arrival price and the actual execution price. Additionally, the displayed market depth represents the order size that can be executed at a given cost. For example we will measure the average market depth, at one-minute frequency, as total volumes available at the best quotes, as well as within a 1% and 2% range around the mid-price level. Finally, market resilience refers to the speed with which the order book is replenished and former market equilibrium is restored after a series of orders executed by uninformed traders have moved the prices away. However, even if highly important for informed traders, this last property will not be tackled in the current paper.[4]

Price efficiency and price discovery

Prices are considered to be efficient when they correctly reflect the asset's underlying fundamental values. According to Hendershott (2011), variance ratios are standard ways to measure price efficiency. For example, the variance of weekly returns can be compared with (five times) the variance of daily returns. If the variance ratio is one, markets are considered to be efficient, while lower variance ratios imply that prices tend to overreact, whereas

[3] "Market quality: The new benchmark for trading venues", The Trade, Issue 35, Jan–Mar 2013, `http://www.xetra.com/blob/1208528/2fe40dfdebf6ba7c9cff52a879894d17/data/Xetra-Market-Quality-Article-in-The-Trade.pdf`

[4] *Foresight: The Future of Computer Trading in Financial Markets* (2012) reviews also other liquidity measures, such as the Amihud measure of illiquidity, which measures the amount of price change per unit of trading volume.

higher variance ratios imply that prices tend to follow short-term trends. In order to assess the impact of high-frequency traders we will compute the following variance ratios: 10 to 1 minute variance and 1 minute to 10 seconds variance.

Volatility and financial market stability

Harris (2002) describes volatility as the tendency for prices to change unexpectedly, either in response to new fundamental information (fundamental volatility) or to impatient traders' demand for liquidity (transient volatility). In the current implementation, the underlying fundamental value basically remains unchanged during the entire trading session. However, the flash crash scenario described in Section 3 can also be interpreted as the case when a single market participant, in possession of some new information about the fundamental values, trades accordingly, even if in a highly aggressive manner. On the other side, the transitory volatility is well designed in our model by a rich heterogenous ecology of order flow generating agents and by the realistic implementation of the microstructure-based order submission mechanism. Volatility is not constant through time, and episodes of low volatility alternate with high volatility periods (volatility clustering). From the perspective of financial market stability, the episodes of extreme volatility are of interest. For example, rare events when large price changes happen in a short period of time, just to recover as fast to their previous levels, can be interpreted as market failures.

In this paper, the high-frequency volatility is estimated in various ways. One first measure is the volatility range, computed as the percentage ratio of the highest over the lowest price, within some period of time. Another estimate is the daily standard deviation of one-minute returns, given by the realized variance equal to the sum of squared price changes at the selected frequency. Subsampling at a frequency lower than every tick reduces the effect of microstructure noise, e.g., the bid-ask bounce. Finally, for analyzing the patterns of the one-minute price series, we compute the Higuchi fractal dimension (Higuchi (1988)).

2 The market model

The most common price discovery mechanism implemented by regulated exchanges all over the world is the continuous double auction. Basically, market participants send either limit or market orders, which are sequentially matched against the outstanding limit orders previously stored in a double queued order book. Technically, the incoming order flow is processed by a matching engine which generates trades based on a price and time priority.

For example, buy market orders and crossing limit orders, i.e. with a limit price higher than the lowest price of the outstanding sell orders, are executed at once under the available size constraint, while non-crossing limit orders are placed in the buy side of the order book, waiting to be matched by a future opposite counterparty order.

An important key-word used in the above description of the continuous double auction mechanism is "sequentially", which means that the order of incoming trading orders is essential. Most intraday agent-based models (ABMs) deal with this asynchrony issue by activating the agents one at a time. For example, Chiarella and Iori (2002), Chiarella et al. (2009) propose a random polling model where individual agents are randomly picked at each simulation step, and with a fixed probability may enter the market. In Gsell (2008) two distinct frequency-classes of agents (stylized and algorithmic traders) are reflected by different probabilities of market participation as expressed by a latency multiplier. Similar, Vuorenmaa and Wang (2013) model both low- and high-frequency agents and introduce a dynamic random polling mechanism, where the latency factor (activity frequency) of a HFT agent is a function of its inventory, i.e. the higher its asset holdings, the more active the HF trader is.

However, the previous activation mechanisms distort the historical price time series, which is essential for the investment decision of LF traders. An example of mixing various time-frequencies is given in Figure 1. The upper timeline illustrates the activity of a LF trader, who follows a τ frequency, as represented by the main ticks on the horizontal (time) axis. The second from the top timeline shows the potential activity of a HF agent, which is able to trade at a higher granularity than τ. For example, while the LF trader "rests" between $\tau = 5$ and $\tau = 6$, the HF trader executes two trades. The third timeline reflects the outcome of random polling which assures the same activation sequence given by the combined previous LF and HF agents. In this case, the time span between two consecutive LF events is widened (but could also be contracted) depending on the number of in-between HF events. For example, the dynamic random polling in Vuorenmaa and Wang (2013) leads to intense activity clustering when HF traders are responsible for almost the entire market activity. If the time frame of LF traders is not decoupled from the fluctuation in number of market events, i.e. time is not independently represented from order flow frequency, LF traders will end up miss-behaving by applying their LF strategies at a HF time frame. Furthermore, no market statistics which are to be computed as related to time can be computed.

The bottom timeline in Figure 1 represents the correct output where the LF time-stamps are not shifted in time anymore and which can be achieved by implementing an event-driven model as described in Muchnik et al. (2006) and Daniel (2006). Under this event-driven framework, agents decide in an

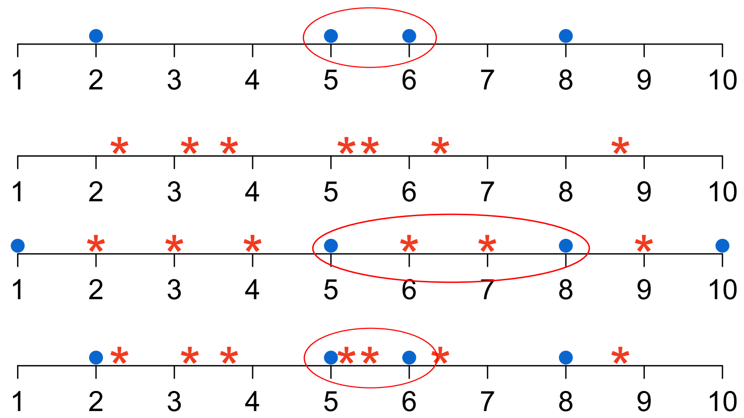

Figure 1: Mixing low and high time-frequencies

autonomous manner when to become active, depending on their 'inner' clock (time frame) or external market events, allowing for a more realistic implementation of time. Each model component has its own, independent timeline which intersects now and then with the timelines of the other processes. The central polling component is thus replaced with dispersed and autonomous agents' trading requests, which are queued and executed sequentially based on their arriving time priority.

The model proposed in the current paper builds on the previously mentioned event-based model and adds a proper implementation of various price time series, conditioned on the required time frame, as a necessary base for developing trading strategies which operate at different frequencies. Basically, besides agents' wake-up requests, time series update events are added to a master queue, from where all events are dispatched to their corresponding handlers in a timely order. Therefore, a low-frequency time series (different from a high-frequency tick-by-tick time series) can be updated at fixed time intervals, e.g., every minute, independent of the number of trades occurring within each time interval. Explicitly, each agent type has its own trading frequency and its activation follows a Poissonian process, where the inter-wake-up durations are random draws from an exponential distribution with the rate parameter $\lambda = 1/\beta_{\mathtt{tf}}$, where $\beta_{\mathtt{tf}}$ is equal to the individual characteristic frequency. Additionally, three price time series at tick-by-tick, one minute and daily frequencies are recorded. The daily price series is initialized for the time-period preceding the simulation start by means of a random walk process with a daily price drift $\mu_{\mathtt{eod}}$ and a daily price volatility $\sigma^2_{\mathtt{eod}}$.[5] The initial fundamental value is set and is kept constant during the trading session – the fundamental time series is allowed to change only at the end of the trading day. It might be worth mentioning here that all actions related to the market matching engine are also modeled as independent events, e.g., new orders, order removal/expiration, trade and quote change notifications.

[5]By this, we set up an initial trend for the end-of-day low frequency chartists, as detailed in Subsection 2.1.

This allows for an effective representation of latency differences which could lead to different market quotes from the moment of creating and sending the order until the order actually "hits" the market.

2.1 Low frequency agents

The agents' design starts from the classification in (Kirilenko et al.; 2011) where six categories were identified to participate in the E-mini S&P 500 stock index futures market between 3-6 May, 2010: intermediaries, HFTs, fundamental buyers and sellers, small traders and opportunistic traders. The first two categories correspond to whats is known as high-frequency trading (will be described in Subsection 2.2), while the latter four can be considered low-frequency trading.

- Low-frequency strategies

We define a pool of low-frequency strategies S which can be employed by the LF traders, as follows: fundamental, trend following, mean reversion and noise trading. Fundamentalists' expectations are equal to their own evaluations of the fundamental value and can be defined as in (2.1), where $\epsilon^{\texttt{fund}} \sim U(0.99, 1.01)$ is a noise term, contributing to the heterogeneity of the model:

$$E_t[p^{\texttt{fund}}] = f_t\,\epsilon^{\texttt{fund}} \tag{2.1}$$

The two chartist strategies can be applied both at daily (end-of-day), as well as at intraday frequencies. The future price expectation of trend followers at time t, $E_t[p^{\texttt{trend}}]$, is based on extrapolating the weighted average of the average past return over some previous time-span τ and the current return r_t (see equation (2.2)).[6] End-of-day time spans are random deviates from a uniform distribution $\tau_{\texttt{eod}} \sim U(1, 9)$, as well as intraday time spans $\tau_{\texttt{id}} \sim U(10, 60)$. On the opposite, reversal strategies expect the price to return to a long-term moving average (see equation (2.3)). In the case of the daily time frame, the p_t time series refers to end-of-day (closing) prices, while for intraday time frames it depends on the selected frequency, e.g., 5 minutes, 1 minute or tick-by-tick.

$$E_t[p^{\texttt{trend(eod/id)}}] = p_t * \left(1 + \left(0.95\,\frac{1}{\tau}\sum_{i=1}^{\tau} r_{t-i} + 0.05\,r_t\right)\,\tau\right) \tag{2.2}$$

[6]The current return is based on the most recent intraday trading price relative to the previous daily closing price or relative to the previous time stamp for intraday time series, respectively.

$$E_t[p^{\texttt{rev(eod/id)}}] = \frac{1}{\tau} \sum_{i=1}^{\tau} p_{t-i} \tag{2.3}$$

Finally, the expectations of the noisy trading components are build around the current mid-price m_t as in equation (2.4), with the noise term $\epsilon^{\texttt{noise}} \sim U(0.96, 1.04)$.

$$E_t[p^{\texttt{noise}}] = m_t \, \epsilon^{\texttt{noise}} \tag{2.4}$$

The implementation of a low frequency trader can mix up several strategies, as in Chiarella and Iori (2002), Chiarella et al. (2009), given individual weights allocated to each individual strategy, i.e. fundamentalist $w_{\texttt{fund}}$, end-of-day/ intraday trend following $w_{\texttt{trend(eod)}}$/ $w_{\texttt{trend(id)}}$, end-of-day/ intraday mean reversion $w_{\texttt{rev(eod)}}$/ $w_{\texttt{rev(id)}}$ and noise strategy $w_{\texttt{noise}}$.

$$E_t[p] = \frac{\sum\limits_{s \in S} w_s \, E_t[p^s]}{\sum\limits_{s \in S} w_s} \tag{2.5}$$

- Agent activation

The trade frequency of each trader type, which proxies the activity of traders, is considered to be proportional to the market participation (% of trades) computed in Kirilenko et al. (2011) – even if some orders are not executed at all, while others are counterpart to multiple trades. We compute the time frame $\beta_{\texttt{tf}}$ of an agent cluster, i.e. the mean time duration between successive activity wake-ups, as the total length of a trading session expressed in microseconds divided by the number of trades for which the trader type is accounted for.

- Order size

The actual order sizes are randomly generated from an upper-truncated exponential distribution. The mean and upper bound of these exponential distributions are specific for each trader type and are set to match the summary statistics of trader categories in Kirilenko et al. (2011). Thus, given the desired order size mean, an appropriate pair of rate $\lambda = 1/\beta_{\texttt{size}}$ and right-side cutoff parameter $b_{\texttt{size}}$ – corresponding to the maximum order size – of the exponential distribution modeling order sizes can be chosen.[7] Moreover, all random draws are rounded to the closest integer, except for values smaller than one which are only upper rounded (ceiling function) in order to avoid null order sizes.

- Order submission

The order placement mechanism is built around an optimization problem which minimizes the risk adjusted execution cost, taking into consideration a set of relevant micro-structure factors – such as order book liquidity, order flow proxied by the order book imbalance and transient volatility –, as well as intrinsic agents characteristics – such as the sense of urgency. The entire procedure is described in detail in Mandes (2014) and is reproduced without any changes in the current implementation.

- Order cancellation

Since the low-frequency traders are not modeled as individuals, but as clusters, we rely on a probabilistic order cancellation process as in Cui and Brabazon (2012). At each activation, the agent cluster makes a random choice between removing its oldest outstanding limit order with probability λ_c and trading, i.e. submitting a new order, with the complementary probability $1 - \lambda_c$.

Fundamental buyers and sellers

The actual low-frequency trader types match the previous trading strategies with the trader categories in Kirilenko et al. (2011). Thus, we consider fundamental buyers and sellers to form price expectations by employing a mixture of end-of-day strategies such as fundamental, trend following and mean reversion. However, the decision of buying, respectively selling, is constant and hard-coded for each of the two types, similar to Cui and Brabazon (2012) – the main feature of these two traders is that they accumulate directional positions to be held for longer periods of time. In the current implementation, their future price expectations do not influence their trading decision, but only their order size distribution by tilting its mean $\tilde{\beta}_{\texttt{size}} = (E_t[p]/p_t)\,\beta_{\texttt{size}}$.

[7] The mean of the truncated exponential distribution $f_X(x|\beta, b) = \frac{\frac{1}{\beta}\exp(-x/\beta)}{1-\exp(-b/\beta)}$, for $0 < x \leq b$ where $\beta > 0$, is $E[X] = \beta\,\frac{1-(k+1)\exp(-k)}{1-\exp(-k)}$, where $k = b/\beta$ (see Olive (2008, Chapter 4)). Therefore, by knowing $E[X]$ and b, one can solve the previous equation for β. However, if k is large enough (e.g., $k \geq 7$) then $\beta \approx E[X]$.

Opportunistic/ day traders

Opportunistic traders are considered to be day traders, which follow a mix of various strategies applied both at daily as well as intraday time frames. To account for heterogeneity within the cluster, at each trader activation, the weights associated with the individual strategy components are randomly shuffled according to an uniform distribution $U[0,1)$. The trading decision to buy or sell a specific asset quantity is generated by comparing the expected price $E_t[p]$ with the current midpoint or market price p_t.

Small traders/ noise traders

Finally, small traders, which trade very few contracts per day, are modeled as zero-intelligence/ liquidity traders, i.e random buying or selling. In other words, in order to fit to the standardization described at the beginning of the current subsection, small traders' "strategies" are based only on the noisy trading component.

2.2 High frequency agents

Market professionals consider high-frequency traders (HFT) to be not a strategy by itself, but a technology which allows for the automation and for the application of a wide spectrum of trading strategies at a "drastically compressed time scale",[8] where human traders are not able to react. However, a set of common denominators for the entire range of HFT-based strategies has been identified both by academic researchers (see Aldridge (2009), Brogaard (2010)), as well as by market regulators across the world.[9] The list of common attributes could be summarized as follows: (1) computer-based non-discretionary/ automated strategies, (2) proprietary trading (as opposed to agency activity), (3) use of low latency technologies (e.g., co-location services, proximity hosting, direct market access, individual data feeds offered by exchanges), (4) real-time tick-by-tick data processing, (5) high amount of intraday messages (orders, quotes or cancellations) and trades, (6) small margins per trade, (7) high capital turnover, (8) flat over-night positions (no

[8]NYSE Euronext comments on the Committee of European Securities Regulators' (CESR) call for evidence on micro-structural issues of the European equity markets.

[9]BaFin (German HFT Bill – Hochfrequenzhandelsgesetz), European Comission (MiFid 2 – Directive 2014/65/EU of the European Parliament and of the Council of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU), European Securities and Markets Authority (consultation paper "Guidelines on systems and controls in a highly automated trading environment for trading platforms, investment firms and competent authorities" – Reference ESMA/2011/224), Securities and Exchange Commission ("Concept release on equity market structure" – Release No. 34-61458, File No. S7-02-10, page 45).

positions or fully hedged). A comparison to lower-frequency strategies shows that HFT-based strategies address very short time frames for generating entry and exit signals (associated with small expected returns), hold positions for very short periods of time (in terms of seconds or milliseconds), trade very often and need a smaller risk capital, which they are able to relocate several times on a short-time basis, leading to their overall profitability (high volume).

Both categories from the classification of Kirilenko et al. (2011) which were not treated previously, i.e. intermediaries and HFTs, stand for short horizon investors which participate in a large number of transactions without accumulating a significant net position. For example, intermediaries sum up to less than 2% of all traders, but account for around 10% of all trades and volume, while their total net holding is not more than 2000 contracts. The case of HFTs is even more extreme in the sense that, while representing only 0.1% of all traders, they account for almost one third of the total daily trading volume. Moreover, their net position tends to quickly revert to a mean of about zero (an estimated half-life of the inventory holding period of 137 seconds) and rarely exceeds more than 3000 contracts. Moreover, the average order sizes of intermediaries are almost half of those of HFTs, and intermediaries tend to reduce their inventories to their target levels at a slower rate than HFTs do it (inventory half-life is 173 seconds), by primarily using marketable orders. Overall, for both HFTs and intermediaries, almost 80% of their orders are entirely passively executed, even if these stand for only 50% of their traded volume, which still corresponds to lower aggressiveness ratios than those of LFTs.

Besides the previous quantitative statistics which might be given by intrinsic features such as latency, trading frequency, order sizes, target inventories, there is also more substantial behavioral discrepancy between the two types, which also leads to a difference in their profitability. According to Kirilenko et al. (2011), HFTs seem to correctly anticipate price changes, open positions aggressively in the direction of the current price (contemporaneous and of the past four seconds), and try to close their positions by submitting resting limit orders in the direction of the anticipated price move. However, if their risk and inventory management objective requires it, HFTs can also offset their positions by submitting marketable orders.[10] On the other side, intermediaries trade in the opposite direction of the immediate prices and their first two lags, but in the direction of the past three to eight seconds.

As in the previous section we try to model and implement strategies matching the empirically observed behaviors. According to EUREX, the main HFT-based strategies are liquidity provision, (statistical) arbitrage,

[10] After about 10 to 20 seconds, HFTs appear to reverse the direction of their trading.

short term momentum trading and liquidity detection.[11] Except for the statistical arbitrage, which deals with pairs of assets rather than a single asset, the other identified strategies can be classified in two main groups: (i) passive (uninformed) strategies which try to cash in the bid-ask spread when the order flow is balanced, and (ii) active (informed) strategies which try to predict the short-term market direction due to new information or just due to participants' order flow. Strategy-wise, we consider these two classes to correspond to the intermediaries and HFT types described in Kirilenko et al. (2011), i.e. we propose an electronic liquidity provision (ELP) strategy for intermediaries and a predatory strategy for HFTs. In the current contribution, we will deal only with the first type and leave the latter for further research.

Electronic liquidity provision (ELP) – liquidity traders

ELP is a high-frequency quasi market making strategy, which seeks to capture both the bid-ask spread and the rebates paid by the trading venues as incentives for posting liquidity, with no obligation of participating in the market at all times. ELP agents are poorly informed about fundamental values and therefore try to minimize the risk of liquidity provision by quickly offsetting even small inventory imbalances or by trying to identify toxic order flow which would ultimately lead them to opening positions against the future market move.[12] These tight inventory management objectives make them use marketable orders more often than what would be expected in the case of traditional market makers. However, even if their round-trip gains and inventory sizes are smaller than the ones of better informed (human) market makers, their overall profitability comes from a larger number of round-trips due to quoting more often at the best bid and ask, which is equivalent to cashing in a higher cumulative spread.

Our implementation of the ELP agents reproduces to a large extent the model proposed in Vuorenmaa and Wang (2013), which is further based on the model introduced by Avellaneda and Stoikov (2008) – a representative model for inventory-based market making. In more detail, the optimal bid and ask quotes are identified following a two-step procedure. First, a reservation price (a personal indifference valuation for the asset) is computed as a deviation from the mid-price m_t, given the current inventory imbalance I_t

[11] "High-frequency trading – a discussion of relevant issues", London, 8 May 2013, `http://www.eurexchange.com/blob/exchange-en/455384/490346/6/data/presentation_hft_media_workshop_lon_en.pdf`

[12] In a literature review, Aldridge (2013) identifies two main types of automated market making models: inventory- and information-based models. The first class is concerned only with the effective management of inventory, without any opinion on the drift or any autocorrelation structure, while the latter tries to identify the true value from the market itself, rather than relying on external fundamental information.

and the current instant volatility σ_t^2 (see equation(2.6)).

$$r_t(I_t) = m_t - \gamma^{\mathtt{elp}}\, I_t\, \sigma_t^2 \tag{2.6}$$

where $\gamma^{\mathtt{elp}}$ stands for the risk aversion coefficient.

In the second step, the optimal spread size around the reservation price r_t is determined. For this, we rely on the simplified version proposed by Vuorenmaa and Wang (2013), where the market maker's spread is arbitrary equal to $s_t + 2\,\exp(|\xi|)$, where s_t is the (current) bid-ask spread and ξ is a normal distributed noise with a standard deviation $\sigma_\xi = 0.01$, contributing to some heterogeneity. Most of the heterogeneity in quotes is induced by the different inventory holdings and by the effective activation and requoting time. Furthermore, we consider the bid and ask quote sizes to be constant and equal to $\beta^{\mathtt{ELP}}_{\mathtt{size}}$ – if one of the quotes is partially or completely filled, it will be replaced with a new full-sized quote at the next agent activation. The ELP agents' wake-up clock is fixed at constant time intervals $\beta^{\mathtt{ELP}}_{\mathtt{tf}}$. In order to desynchronize their activation processes, only the first wake-up event is drawn from an exponential distribution with $\lambda = 1/\beta^{\mathtt{ELP}}_{\mathtt{tf}}$. However, not every wake-up coincides with quotes being updating, since requoting is set up as a probabilistic event constrained by the current inventory imbalance relative to the maximum allowed inventory threshold $I^{\mathtt{ELP}}_{\mathtt{max}}$ – the higher the inventory imbalance, the more probable a requoting event is, and when the inventory imbalance exceeds the maximum threshold requoting becomes certain (see equation(2.7)).

$$\mathtt{Prob}_t(I_t) = \max(1, (I_t/I^{\mathtt{ELP}}_{\mathtt{max}})^2) \tag{2.7}$$

3 Simulations setup and results

3.1 Model parameters and calibration

Simulations are run over one individual trading day which lasts for 6 hours and 45 minutes (from 08:30 to 15:15 CT), summing up to 243E05 microseconds.[13] Before recording the actual trading session, each simulation is warmed-up for one hour. The underlying fundamental value $f = 1200.00$ is considered constant throughout the entire simulation – since during the trading day there are no exogenous fundamental shocks, the entire intraday dynamics is only due to endogenous micro-structure factors. The daily time series of the

[13] This period corresponds to the open hours of the components of S&P 500 index, which is the underlying asset of the CME S&P 500 E-mini futures.

previous ten closing prices is initialized using a random walk process with an initial closing price $p^{\texttt{eod}}_{-10} = f$, a daily drift $\mu^{\texttt{eod}} = -0.005\, f$ and a daily volatility $\sigma^{\texttt{eod}} = 0.005\, f$, corresponding to an overall negative daily trend.

As already mentioned, the LF agents are represented by a cluster-type implementation, i.e. one single instance per type, since their trading decisions are not influenced by the current individual portfolios. The parameters of the clusters refer to the following dimensions: time frame $\beta_{\texttt{tf}}$, strategy weights (i.e. fundamentalist $w_{\texttt{fund}}$, end-of-day/ intraday trend following $w_{\texttt{trend(eod)}}/w_{\texttt{trend(id)}}$, end-of-day/ intraday mean reversion $w_{\texttt{rev(eod)}}/w_{\texttt{rev(id)}}$, and noise strategy $w_{\texttt{noise}}$ respectively), order size distribution mean $\beta_{\texttt{size}}$ and upper cutoff $b_{\texttt{size}}$. The time frame, the order size distribution mean and upper cutoff are computed as described in Subsection 2.1, in order to reproduce the traders category descriptive statistics from (Kirilenko et al.; 2011). These setup parameters are marked as bold in Table 1. The other parameters represent ad-hoc choices, so that the emergent market properties reproduce a set of target statistics, reflecting a wide range of market and trader categories characteristics such as: price, volume, trading frequency, quotes and order flow, agents' aggressiveness, market share and net inventories. A second set of parameters deals with the order submission mechanism and their values are set as in Mandes (2014): order book imbalance base $\mu = 2$, book depth levels $N = 3$, size penalty exponent $\eta = 0.8$, volatility bands $\sigma_{is} = 7\,\sigma$, $\sigma_{dyn} = 7.5\,\sigma$, and the weights of the adjusted risk cost function $\alpha_0 = 0.1$, $\alpha_1 = 0.5$, $\alpha_2 = 0.25$, $\beta = 2.5$. The cancellation probability for the oldest order sent by a LF cluster is $\lambda_c = 0.4771$ as in Cui and Brabazon (2012), where it was determined from empirical data.

	BUY	SELL	OPP.	SMALL
$\beta_{\texttt{tf}}$ (ms)	**268**	**268**	**81**	**1226**
$\beta_{\texttt{size}}$	**14.09**	**14.20**	**8.80**	**1.23**
$b_{\texttt{size}}$	**200**	**200**	**100**	**9**
$w_{\texttt{fund}}$	0.9	0	–	0
$w_{\texttt{trend(eod)}}$	0.2	0.9	–	0
$w_{\texttt{trend(id)}}$	0	0	–	0
$w_{\texttt{rev(eod)}}$	0	0.2	–	0
$w_{\texttt{rev(id)}}$	0	0	–	0
$w_{\texttt{noise}}$	0	0	–	1

Table 1: Calibration of the LFT types

The implementation of the ELP agents is realized in more detail than in the LFT case. Instead of building clusters of trading agents, each ELP is individually instantiated, with its own timeline and portfolio. The strategy time frame is assumed to be constant $\beta^{\texttt{ELP}}_{\texttt{tf}} = 1$ second – dependent on the agents' technological capacities, i.e. fixed activation time, and its value corresponds to the average trading frequency documented in Kirilenko et al. (2011). Furthermore, the risk aversion coefficient is common for all ELP agents $\gamma^{\texttt{ELP}} = 1.0$, as well as their default quote sizes $\beta^{\texttt{ELP}}_{\texttt{size}} = 7$. The two missing parameters related to ELP, i.e. the number of agents $N^{\mathtt{ELP}}$ and their maximum inventory thresholds $I^{\mathtt{ELP}}_{\mathtt{max}}$, are described together with the simulation scenarios in Subsection 3.2.

3.2 Simulation scenarios

In this section, several scenarios trying to capture various market configurations and global outcome possibilities are defined. In order to assess the impact of HFTs, a base-line scenario with LFTs only is set up (marked with "LFT" in Table 2 and following tables). One comparative HF scenario marked with "LFT w. hom. ELPs" extends the previous scenario by adding $N^{\mathtt{ELP}} = 100$ electronic liquidity provision agents with a tight and homogeneous maximum inventory threshold $I^{\mathtt{ELP}}_{\mathtt{max}} = 10$. The homogeneity assumption is useful for the grid-analysis of the relation between various pairs of factors and associated outcomes, and could be explained by agents applying a set of common popular algorithms and by using the same short-time window high-frequency data as input.[14] The homogeneous ELPs scenario is further divided into two sub-scenarios, one without and one with financial transaction taxes, as explained in Section 4. As an alternative to the homogeneous setup, we also run a heterogenous scenario marked "LFT w. het. ELPs", where the quoted sizes are uniformly distributed $\beta^{\mathtt{ELP}}_{\mathtt{size}} \sim U(1,12)$ and $I^{\mathtt{ELP}}_{\mathtt{max}} = k\,\beta^{\mathtt{ELP}}_{\mathtt{size}}$ are multiples of the quoted sizes, where $k \sim U(1,10)$ – the ELPs' strategies remain common, but they have different parameterizations.

The previous market configurations are run under two different market states, representing regular and extreme ("stressed") market conditions. The latter correspond to the "Flash Crash of 2:45" scenario, as described in Kirilenko et al. (2011). Their investigation results show that the flash crash was triggered by a large sell program of 75,000 contracts, which has used an execution algorithm based on a volume-in-line strategy (also known as percentage of volume), targeting an execution rate of 9% of the previous minute trading volume. Therefore, in the flash crash scenario an additional agent (marked as "POV"), reproducing the sell program described earlier, is added to the implementation. The agent is activated at 13:30 CT and executes market sell orders of a size equal to 9% of the previous minute volume without its own execution share, at a fixed $\beta^{\mathtt{POV}}_{\mathtt{tf}} = 1$ minute frequency, until the program is finished or until the end of the trading session, i.e. 15:15 CT.

Each of the scenarios is run 100 times with different random seeds and both the mean and the standard deviation of the individual key statistics are reported. Tables 3 and 4 present a list of descriptive statistics reflecting various facets of market quality, corresponding to the previously introduced

[14] Actually, this is one of the concerns expressed by Danielsson and Zer (2012) – a too large homogeneity of strategies could contribute to an increasing systemic risk.

		Transaction tax	Market state	
			regular	stressed
Agents	LFT only	no		
	LFT w. hom. ELPs	no		
		yes		
	LFT w. het. ELPs	no		

Table 2: Simulation scenarios configurations

experiment setups. Moreover, Tables 5 and 6 capture the agents' trading statistics under the various scenarios. The results show that under both market conditions, the total trading volume increases substantially (six to eight times) when ELPs agents are introduced. When no extreme event occurs (such as a flash crash), ELPs account for about 45% of the total number of trades and total trading volume – the extreme event outcome will be treated separately in Subsection 3.3. The overall liquidity is also highly improved in the simulations including ELPs. All spread measures are four times smaller when the ELPs are homogenous and three times smaller with heterogenous ELPs, as compared to the base-line scenario. The virtual market impact of executing a market order with a size of 100 is also three to four times lower. The market depth at quotes and within a 25 bps range around the mid-price are about two times larger, while the market depths within 50 and 100 bps are similar, and the total market depth within 200 bps is actually slightly smaller. This means that the total order book liquidity is lower and concentrated more towards the best quotes (thinner tails) in the scenarios with ELPs. On one side, this is due to the ELPs strategy of placing quotes close to the best bid and asks, but also because of LFTs preferring to send four times more market orders, and corespondent fewer limit orders, due to the lower immediacy costs. The same aspect can also be identified in the increasing trade and volume aggressiveness ratio of LFTs, i.e. the share of orders and proportion of orders consuming order book liquidity (in other words, executed against passive/ outstanding limit orders). It is to be noted that, in our implementation, the LFTs' order submission behavior is reactive to the market microstructure factors. As a consequence, LFTs might also compete with ELPs for a better execution chance by placing more aggressive limit orders. The improved liquidity can also be observed in the case of the POV sell program, which is able to execute a larger share of its initial order and with smaller execution costs. By looking at the percentage of market depth accounted by ELPs, one can see that ELPs contribute with 80% of the total volume at the best quotes and their contribution decreases for further away price levels. The book gaps measures have a high standard deviation and we cannot conclude for any significant differences. However, the logic could argue that the sum of the first gaps is smaller for heterogenous ELPs due to their strategy differences.

Both range volatility and the realized standard deviation at 1-minute frequency show a substantial decrease – the standard deviation of 1-minute returns is almost half than in the base-line model and the daily range volatility gets close to zero. Consequently, also the daily return is close to zero while it is negative in the LFT-only case.[15] The differences are even more pronounced when the market stress scenario is evaluated. In the base-line case the market drop is one percent larger and the range volatility is higher, while these statistics remain unchanged when the ELPs are present and no extreme event occurs. Basically, for the HF set ups, no major differences between the regular and stress market states can be identified by looking at the general market statistics. A deeper analysis of the flash crash events, looking at the inter-agents feed-back loops, is conducted in the next section. Finally, regarding price efficiency, both considered variance ratios do no capture any significant difference, but in the end there is no new exogenous information incoming into the model during the trading day.

The differences between the base-line and the high-frequency setups can also be assessed by visually inspecting the plots 3, 4 and 5, where a mixed scenario is recorded – in the first two and a half hours of the trading day both LFTs and ELPs are active, while after 11:00 CT ELPs are switched off and the market is made by LFTs only – actually, for about an hour, some remaining ELP limit orders are still in the market, but are not refreshed after their execution. These plots reflect the intraday dynamics of various market and agent measures for a single run. Regimes switches at the moment when ELPs are disabled can be observed in the intraday time series for most of the individual plots, and the changes are also in line with the average results previously described in this section.

3.3 The anatomy of a flash crash

A flash crash is a market event of extreme volatility, where the price experience a sudden drop, followed by a rapid rebound. The most notorious flash crash is the one which has occurred on May 6, 2010 on the E-mini S&P 500 stock index futures market.[16] However, since 2007 more than 15,000 different events similar to a flash crash have been identified in several individual

[15]The negative day return in the LFT-only set up is due to the previous descending daily trend and due to the selected mix of strategies applied by the opportunistic traders, which puts a larger weight on chartist components – there are four chartist versus one fundamentalist component.

[16]On that day, according to Kirilenko et al. (2011), a large sell program started sometime after 13:30 CT. During a 13 minute period, between 13:32:00 and 13:45:27 CT, the E-mini contract declined with 5.1%. Over the course of the next second, a cascade of executed orders caused the price of the E-mini to drop another 1.3% and the following transaction would have triggered a drop in price of another 6.5 index points. This has triggered the CME Globex Stop Logic Functionality, which pauses all transactions for 5 seconds even if the market remains open for orders to be submitted, modified or canceled. At 13:45:33, the E-mini exited the Reserve State, the market resumed trading and after fluctuating for a few seconds, the price began a rapid ascent until it reached the levels before the drop.

stocks.[17]

In our artificial environment, similar sudden price drops can be replicated for HF set ups including ELPs and a POV sell program. The probability of a flash crash depends however on the chosen set of model parameters. For example, in order to test the influence of the number of ELP agents $N^{\mathtt{ELP}}$ and of the maximum inventory threshold $I^{\mathtt{ELP}}_{\mathtt{max}}$, under the homogeneity assumption, we run a grid simulation over the Cartesian product for two sets of values, one for each factor, and plot the results in the upper heatmap of Figure 2.[18] It seems that the critical number of agents, which can coordinate in a way leading to the emergence of a flash crash, is between 10 and 50. Moreover, the tighter the inventory control, the higher is the probability of a flash crash. Even if seldom, extreme events occur also for larger populations of ELPs when their maximum inventory thresholds are small. The parameters-wise heterogenous ELPs scenario is also tested and its results with respect to different levels of activity time frame and agents number are represented in the bottom plot of Figure 2. Compared to the homogenous set up, the flash crash probabilities for the time frame (1000 milliseconds) and same average inventory threshold (approx. 30) become zero, even within the critical agents number range.

In order to analyze what happens during a flash crash, we select the market configuration "LFT w. hom. ELPs" described in the previous subsection, i.e. number of ELP agents $N^{\mathtt{ELP}} = 100$ and maximum inventory threshold $I^{\mathtt{ELP}}_{\mathtt{max}} = 10$. This corresponds to a flash crash probability of 6% which makes it a rare event. The overall market quality measures in Table 4 and the agents' activity statistics in Table 6 show that the total trading volume increases fivefold – on May 6, 2010 a twofold increase in activity has been reported. Individually, the ELPs carry the main responsibility with a factor of ten increase in activity. The fundamental buyers, the day traders and the POV sell program also contribute by doubling to tripling their traded volumes. However, the general market quality is worse than in the non-flash crash scenario, under all aspects – the order book liquidity is thinner, the volatility skyrockets and the variance ratios show an increasing price overreaction. Another important remark is that the aggressiveness ratios of ELPs jumps from 6% to 50%, which reflects a very hectic inventory management activity. On the other side, the LF market participants find the opportunity of executing their trades in a more favorable passive manner. A striking difference can be observed with respect to the ELPs total profitability – their steady profits made during normal market conditions or during market stress, but without a flash crash, conditions become very volatile when a flash crash occurs and can turn even into huge losses. The large standard deviation fully reflects this uncertainty.

[17] `http://www.nanex.net/FlashCrash/OngoingResearch.html`

[18] The other ELP-related parameters are fixed as follows: activation time frame $\beta^{\mathtt{ELP}}_{\mathtt{tf}} = 1$ second, risk aversion coefficient $\gamma^{\mathtt{ELP}} = 1.0$ and default quote sizes $\beta^{\mathtt{ELP}}_{\mathtt{size}} = 7$.

Next, we propose a "radiography" of the intraday dynamics along a large set of indicators computed at a one-minute frequency. Two runs with the same calibration as previously described, but with different random seeds and opposite outcomes with respect to the flash crash occurrence, are chosen and depicted in two sets of figures: 6, 7, 8 and 9, 10, 11 respectively. From the "ELP long holdings" time series (the third plot from the top) one can observe that, when the POV sell program starts, the long positions of ELPs rise to levels two times higher than at any previous moment during the day. In the no flash outcome, these positions are offset quickly just by quoting more aggressive sell limit orders – the "Spread" time series shows how the bid-ask spread tightens. The main consequence is that ELP trading is conducted with LFTs – there is little ELP to ELP volume as indicated by the "ELP-ELP volume share" time series. Afterwards and until the end of the trading session, ELPs mostly carry a total net short position caused by LF opportunistic traders' activity, which is compensated and complementary to the continued POV sell activity. Actually, after start of the sell program we observe an increased market quality reflected by the lowest daily values of instant volatility, higher depth at quotes and steady bid-ask spreads.

On the other side, in the flash crash case, the initial shock to the ELPs' long positions induced by the beginning of the POV sell program is offset aggressively against other ELPs which open other long positions instead ("hot-potato" or "pass-the-parcel" effect) – observable in the "ELP-ELP volume share" time series of Figure 9. This activity is associated with a one-minute volume spike which triggers larger POV sell market orders, and a volume feed-back loop is created. In repeatedly trying to close their refreshed and increasingly larger long positions, ELPs get more and more aggressive, and consume the order book liquidity on their way (see all market depth time series). The LFTs seem not to be capable of significantly refilling the bid side of the order book in such a short period of time, and the ELPs are left to be the main liquidity providers at quotes and within the 25 and 50 bps ranges around the mid-price (see the correspondent "ELP depth share" time series). Eventually, the order book liquidity deteriorates and no market participant can provide the necessary liquidity required by the "hyperinflated" POV sell program.

4 Testing regulatory policies

After a crisis or after extreme events like the May 6, 2010 Flash Crash, regulators are faced with a large amount of public pressure to take action and correct what went wrong. This is not always an easy task, because an unambiguous historical precedent from which to draw conclusions might not exist, or just because financial markets constitute such complex systems whose reaction might be very tricky to be assessed. Moreover, regulators have to take into consideration the risk that the consequences of their actions could exacerbate or generate a new bigger problem than the one which is supposed to be tackled. A variety of possible policies targeted to improving the overall market performance and to reducing the risks of market failure, in the context of large market automation, have been proposed in *Foresight: The Future of Computer Trading in Financial Markets* (2012). These include notification of algorithms, circuit breakers, minimum tick size requirements, market maker obligations, minimum resting times, minimum order-to-execution ratios, periodic auctions, maker-taker pricing, order priority rules and internalisation of agency order flow. In Europe there are also heated debates around the opportunity of imposing financial-transaction taxes, and countries like France and Italy have already implemented such taxes for specific asset types.

Agent-based modeling frameworks are attractive because they allow for the virtual experimentation of innovative policy options and for simulating a large range of scenarios. If the model assumptions are realistic enough, the model parameters are calibrated to the specific problem at hand, and the models are able to reproduce realistic market behaviors in a robust manner, reliable policy making could also be developed within such artificial laboratories. The current contribution tries to assess the impact of imposing minimum resting times and financial-transaction taxes. Harris (2002) suggests that transaction taxes penalize HF strategies more than longer-term investors, since the former imply a larger number of trades. In our implementation, the resting time can be simulated by controlling the activity time frame of the ELPs, while the introduction of a transaction tax is considered to increase the minimum quoted spread size of the ELPs, in order to cover the extra trading cost and remain profitable.[19]

In order to experiment the previously mentioned policy options, we set up a model with the most unfavorable parametrization corresponding to a 100% flash crash probability under market stress condition, i.e. $N^{\texttt{ELP}} = 30$ homogenous ELPs with $I^{\texttt{ELP}}_{\texttt{max}} = 10$ maximum inventory threshold. Two ranges of values for the two policy factors are proposed and 100 runs for each of the possible combinations are recorded. The results are presented in the form of a heatmap in Figure 2 and show that the flash crash probability reacts more to varying the minimum spread than to the agents' trading frequency. The main explanation is that by quoting larger spreads, ELPs diminish their share of volume at best quotes with about 30%, and thus the "hot-potato" effect can be avoided.

Nonetheless, imposing a financial transaction tax, which completely elim-

[19] Just for completion, all other simulations in the paper consider a no minimum quoted spread, which is equivalent to a minimum quoted spread size of zero.

inates the risk of a flash crash, has also an impact on the agents' trading activity and on the overall market quality. HFTs' market share drops from 46% to around 20% and the majority of most market quality measures take values between the LFT-only and HFT with no transaction tax market configurations, but more closer to the LFT-only base-line set up.

5 Conclusions and outlook

We have designed and implemented an event-driven agent based model of a continuous double auction financial market. This design allows for a realistic representation of time and for the independent activity of all its components, with no central control. As a consequence, the agents included in our model can be calibrated according to various time-related trading statistics which are found in the literature or computed from empirical data. Furthermore, the model output can be compared directly with real time series. The four types of low-frequency agents identified in the literature – fundamental buyers and seller, opportunistic day traders, small noisy traders – have been implemented, matching their empirical trading statistics. The order submission mechanism is a key component of order book trading and we have included one of the most realistic existing models, which is built around an optimisation problem minimizing the risk-adjusted execution cost and taking into account a relevant set of microstructure factors. This makes the LF agents reactive to their trading environment, which is ultimately influenced also by HF market participants. Finally, one type of HFTs corresponding to the electronic liquidity provision strategies has also been included into the model.

Next, a sequence of different scenarios has been simulated in order to try to answer some of the open questions regarding the impact of HFTs on current financial markets. We find that in general ELPs have a positive contribution to the market quality, as assessed by a large range of market quality indicators, grouped along three main lines: (i) liquidity and transaction costs, (ii) price efficiency and price discovery, (iii) volatility and financial market stability. These results are also in line with the empirical findings. Furthermore, at least for ELP strategies, there is no causality between ELP activity and increased volatility. We can also confirm that a large homogeneity with respect to HFT strategies can contribute to the systemic risk – when heterogeneity within parameters is simulated, even for a common strategy, the market becomes more robust under stress conditions.

Moreover, we have also reproduced the chaotic properties of financial systems and showed how rare events such as flash crashes can emerge by conducting a detailed inspection. Thus, we can conclude that even if ELPs are not the primary cause of a flash crash, they can amplify the total trading volume while trying to aggressively offset their inventory imbalances. When ELPs discard their positions to other ELPs, a "hot-potato" effect emerges which is associated with an increased traded volume. If a POV sell program, targeting a fixed trading rate linked to the previous market volumes, does not implement also some limits in order to account for extraordinary and unexpected scenarios, we have the sufficient ingredients for keeping alive a volume feed-back loop. One useful side-note for execution algo developers is that volume alone should not be mistaken for liquidity. Finally, two possible policy options regarding minimum resting times and financial-transaction taxes have been evaluated within the developed artificial environment. We have shown that imposing transaction taxes helps by lowering the probability of flash crashes, but they also negatively impact the ELPs activity and the overall market quality.

Future research suggestions include the assessment of other policy measures such as minimum order-to-execution ratios, notification of algorithms accompanied by volume throttling, minimum tick size requirements and maker-taker pricing. Additionally, the heterogeneity of ELPs could be extended to other strategy parameters such as the sensitivity coefficient $\gamma^{\texttt{elp}}$ or even to implementing different ELP strategies such as those from the information-based class. An example of a market making strategy relying on Bayesian learning to combine new information with prior beliefs about the true market value is provided by Das (2005, 2008). Also the currently implemented ELP and POV strategies could be upgraded to more intelligent algorithms, including also a risk management component. For example, the ELPs could monitor an flow toxicity metric such as the Volume-Synchronized Probability of Informed Trading introduced by Easley et al. (2012) and switch off when order flow becomes toxic. On the other side, the POV execution algorithm could also include a price cap and stop selling when the price drops below a given limit. Finally, a further extension should implement also the second class of HFT described in Kirilenko et al. (2011), i.e. predatory or order anticipating strategies.

References

Aldridge, I. (2009). *High-frequency trading: a practical guide to algorithmic strategies and trading systems*, Vol. 459, Wiley.

Aldridge, I. (2013). *High-frequency trading: a practical guide to algorithmic strategies and trading systems, Second Edition*, Wiley.

Avellaneda, M. and Stoikov, S. (2008). High-frequency trading in a limit order book, *Quantitative Finance* **8**(3): 217–224.

Brogaard, J. (2010). High frequency trading and its impact on market quality.

Brogaard, J., Hendershott, T. and Riordan, R. (2012). High frequency trading and price discovery.

Chiarella, C. and Iori, G. (2002). A simulation analysis of the microstructure of double auction markets, *Quantitative Finance* **2**(5): 346–353.

Chiarella, C. and Iori, G. (2004). The impact of heterogeneous trading rules on the limit order book and order flows, *Quantitative Finance Research Centre Research Paper* (152).

Chiarella, C., Iori, G. and Perelló, J. (2009). The impact of heterogeneous trading rules on the limit order book and order flows, *Journal of Economic Dynamics and Control* **33**(3): 525–537.

Cui, W. and Brabazon, A. (2012). An agent-based modeling approach to study price impact, *Computational Intelligence for Financial Engineering & Economics (CIFEr), 2012 IEEE Conference on [proceedings]*, IEEE Press, pp. 1–8.

Daniel, G. (2006). *Asynchronous simulations of a limit order book*, PhD thesis, University of Manchester.

Danielsson, J. and Zer, I. (2012). Systemic risk arising from computer based trading and connections to the empirical literature on systemic risk, *UK Government Office for Science, Foresight Driver Review - The Future of Computer Trading in Financial Markets* .

Das, S. (2005). A learning market-maker in the glosten–milgrom model, *Quantitative Finance* **5**(2): 169–180.

Das, S. (2008). The effects of market-making on price dynamics, *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 2*, International Foundation for Autonomous Agents and Multiagent Systems, pp. 887–894.

Easley, D., de Prado, M. and O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world, *Review of Financial Studies* **25**(5): 1457–1493.

Farmer, D. and Skouras, S. (2011). An ecological perspective on the future of computer trading, *UK Government Office for Science, Foresight Driver Review - The Future of Computer Trading in Financial Markets* .

Foresight: The Future of Computer Trading in Financial Markets (2012). *Technical report*, UK Government Office for Science.

Friederich, S. and Payne, R. (2011). Computer based trading, liquidity and trading costs, *UK Government Office for Science, Foresight Driver Review - The Future of Computer Trading in Financial Markets* .

Gsell, M. (2008). Assessing the impact of algorithmic trading on markets: A simulation approach, *Technical Report 2008/49*, Center for Financial Studies, Frankfurt, Main.
URL: *http://hdl.handle.net/10419/43250*

Harris, L. (2002). *Trading and exchanges: Market microstructure for practitioners*, Oxford University Press, USA.

Hendershott, T. (2011). High frequency trading and price efficiency, *UK Government Office for Science, Foresight Driver Review - The Future of Computer Trading in Financial Markets* .

Higuchi, T. (1988). Approach to an irregular time series on the basis of the fractal theory, *Physica D: Nonlinear Phenomena* **31**(2): 277–283.

Kirilenko, A., Kyle, A., Samadi, M. and Tuzun, T. (2011). The flash crash: The impact of high frequency trading on an electronic market, *Available at SSRN 1686004* .

Linton, O. (2011). What has happened to uk equity market quality in the last decade? an analysis of the daily data, *UK Government Office for Science, Foresight Driver Review - The Future of Computer Trading in Financial Markets* .

Mandes, A. (2014). Order placement in a continuous double auction agent based model, *Discussion Paper 43-2014*, MAGKS Joint Discussion Paper Series in Economics.

Muchnik, L., Louzoun, Y. and Solomon, S. (2006). Agent based simulation design principles – applications to stock market, *Practical Fruits of Econophysics*, Springer, pp. 183–188.

Olive, D. J. (2008). Applied robust statistics, *Preprint M-02-006., http://www.math.siu.edu* .

Sornette, D. and Von Der Becke, S. (2011). Crashes and high frequency trading, *UK Government Office for Science, Foresight Driver Review - The Future of Computer Trading in Financial Markets* .

Vuorenmaa, T. A. and Wang, L. (2013). An agent-based model of the flash crash of may 6, 2010, with policy implications, *Available at SSRN 2336772* .

	Regular market conditions			
	LFT only	LFT w. hom. ELPs		LFT w. het. ELPs
		non transaction tax	transaction tax	
Trading volume	134,670 (8,982)	1,039,163 (136,308)	356,577 (48,998)	866,909 (75,366)
Average trade size	4.39 (0.037)	4.47 (0.058)	4.58 (0.037)	4.19 (0.15)
1-min trading frequency	75.80 (4.92)	573.83 (67.56)	192.14 (25.11)	511.00 (50.85)
Percentage return (%)	-1.54 (0.52)	-0.013 (0.26)	0.59 (0.8)	-0.39 (0.15)
Market orders share (%)	4.69 (0.35)	23.07 (1.91)	11.91 (1.77)	20.97 (1.55)
Market impact exponent	0.76 (0.12)	0.87 (0.17)	0.77 (0.075)	0.84 (0.11)
Spread-at-touch	2.67 (0.19)	0.64 (0.069)	1.83 (0.03)	0.81 (0.071)
Spread-at-touch (bps)	23.68 (1.90)	5.66 (0.60)	16.04 (0.43)	7.14 (0.63)
Relative effective spread (bps)	22.24 (2.08)	6.00 (0.39)	16.48 (0.36)	7.41 (0.41)
Impact for 100 units (bps)	14.77 (2.15)	3.71 (0.55)	8.53 (0.22)	4.36 (0.41)
Market depth (at quotes)	9.93 (0.44)	24.41 (8.38)	13.86 (0.8)	20.42 (1.79)
Market depth ($\leq$ 25 bps)	1,059 (113)	2,242 (144)	1,485 (119)	2,158 (158)
Market depth ($\leq$ 50 bps)	3,054 (214)	4,093 (289)	2,879 (281)	4,079 (325)
Market depth ($\leq$ 100 bps)	6,781 (519)	7,068 (561)	5,648 (512)	7,195 (611)
Market depth ($\leq$ 200 bps)	12,057 (1,292)	10,670 (972)	9,699 (913)	11,202 (1,108)
ELP depth share (at quotes)		85.89 (2.04)	52.51 (4.1)	79.22 (2.70)
ELP depth share ($\leq$ 25 bps)		29.56 (4.17)	49.48 (4.55)	29.36 (2.79)
ELP depth share ($\leq$ 50 bps)		17.34 (1.40)	26.74 (3.02)	16.16 (1.67)
ELP depth share ($\leq$ 100 bps)		10.11 (0.81)	12.85 (1.09)	9.21 (1.01)
ELP depth share ($\leq$ 200 bps)		6.74 (0.64)	7.54 (0.66)	5.98 (0.76)
Book gaps ($\leq$ 25 bps)	49.34 (5.63)	59.17 (19.87)	35.4 (5.25)	35.35 (11.92)
Book gaps ($\leq$ 50 bps)	166.12 (10.86)	212.26 (13.34)	126.43 (17.99)	201.98 (16.28)
Book gaps ($\leq$ 100 bps)	471.53 (23.63)	469.44 (31.56)	556.37 (39.51)	477.45 (30.89)
Book gaps ($\leq$ 200 bps)	1,169.90 (62.81)	955.31 (69.85)	1,169.46 (95.55)	979.01 (76.22)
10-min var. / 1-min var.	0.15 (0.043)	0.12 (0.029)	0.42 (0.14)	0.12 (0.022)
1-min var. / 10-sec var.	0.23 (0.028)	0.18 (0.017)	0.22 (0.025)	0.17 (0.012)
Range volatility (%)	1.97 (0.52)	0.37 (0.16)	1.54 (0.49)	0.47 (0.11)
Realized 1-min std. dev. (%)	0.023 (0.002)	0.010 (0.001)	0.021 (0.002)	0.012 (0.001)
Flash crash probability	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)	0.0 (0.0)
Higuchi fractal dimension	1.86 (0.033)	1.94 (0.028)	1.64 (0.054)	1.94 (0.020)

Table 3: Market key metrics under normal conditions

The market width reflects the cost per unit of liquidity and we compute it in terms of market impact to trade a size of 100, i.e. the cost to trade compared to the mid-point. The effective spread is twice the midpoint benchmarked transaction cost and the relative effective spread represents the ratio of the effective spread to the current midpoint. A basis point (abbreviated as 'bps') is a unit that is equal to 1/100th of 1%. Market depth is computed in number of shares. The values in parenthesis represent the standard deviations of the descriptive statistics computed over 30 runs.

	LFT only	LFT w. hom. ELPs			LFT w. het. ELPs
		non transaction tax		transaction tax	
		non Flash Crash	Flash Crash		
Trading volume	131,107 (9,327)	1,057,907 (139,312)	5,015,989 (469,487)	358,526 (48,101)	870,114 (74,634)
Average trade size	4.45 (0.048)	4.54 (0.056)	4.58 (0.096)	4.61 (0.035)	4.19 (0.15)
1-min trading frequency	72.78 (5.02)	574.91 (68.91)	2,706.17 (267.56)	191.86 (24.63)	513.02 (50.74)
Percentage return (%)	-2.45 (0.74)	0.009 (0.26)	-100.00 (0.000)	-0.81 (0.54)	-0.43 (0.14)
Market orders share (%)	4.50 (0.35)	23.07 (1.98)	4.81 (0.89)	11.55 (1.62)	20.68 (1.52)
Market impact exponent	0.83 (0.22)	0.86 (0.18)	–	0.75 (0.077)	0.83 (0.11)
Spread-at-touch	2.71 (0.20)	0.77 (0.073)	1,662.26 (1,576.71)	1.83 (0.028)	0.92 (0.068)
Spread-at-touch (bps)	24.11 (1.98)	6.75 (0.64)	16,583 (5,736.48)	16.07 (0.40)	8.11 (0.59)
Relative effective spread (bps)	22.35 (1.97)	6.66 (0.43)	63.12 (255.95)	16.52 (0.35)	8.09 (0.39)
Impact for 100 units (bps)	15.12 (2.23)	4.06 (0.61)	4.11 (1.38)	8.57 (0.22)	4.69 (0.41)
Market depth (at quotes)	10.50 (0.46)	31.80 (9.12)	19.55 (4.19)	16.59 (1.89)	24.12 (1.80)
Market depth ($\leq$ 25 bps)	982.36 (100.75)	2,188.12 (144.19)	1,578.31 (99.07)	1,367.05 (96.81)	2,117.71 (155.95)
Market depth ($\leq$ 50 bps)	2,861.51 (201.60)	4,042.24 (295.41)	2,803.55 (171.29)	2,671.88 (244.92)	4,047.83 (321.25)
Market depth ($\leq$ 100 bps)	6,498.26 (496.14)	7,021.64 (574.38)	4,908.25 (298.98)	5,472.88 (485.78)	7,174.66 (606.44)
Market depth ($\leq$ 200 bps)	11,883 (1,196.50)	10,627 (1,004.22)	7,328.69 (540.50)	9,806.25 (904.01)	11,181 (1,104.94)
ELP depth share (at quotes)		87.86 (2.06)	66.22 (10.47)	58.38 (3.81)	81.27 (2.28)
ELP depth share ($\leq$ 25 bps)		28.99 (3.99)	24.70 (1.59)	53.69 (4.40)	28.89 (2.84)
ELP depth share ($\leq$ 50 bps)		16.90 (1.44)	14.13 (0.60)	29.12 (3.19)	15.80 (1.67)
ELP depth share ($\leq$ 100 bps)		9.81 (0.83)	8.09 (0.45)	13.10 (1.13)	8.98 (1.01)
ELP depth share ($\leq$ 200 bps)		6.54 (0.65)	5.47 (0.29)	7.40 (0.66)	5.83 (0.75)
Book gaps ($\leq$ 25 bps)	50.37 (5.46)	56.71 (18.17)	43.75 (27.59)	38.58 (4.86)	34.90 (11.58)
Book gaps ($\leq$ 50 bps)	170.99 (10.30)	209.30 (12.93)	164.72 (16.28)	141.68 (19.39)	204.67 (15.65)
Book gaps ($\leq$ 100 bps)	479.40 (22.82)	462.87 (31.00)	369.38 (30.65)	568.10 (37.65)	472.22 (30.82)
Book gaps ($\leq$ 200 bps)	1,172.18 (59.87)	929.85 (68.36)	784.98 (34.10)	1,179.81 (90.89)	960.61 (75.45)
POV program completed (%)	3.53 (0.59)	30.37 (2.49)	81.05 (16.09)	11.48 (2.25)	24.58 (1.50)
POV implem. shortfall	-6.82 (4.22)	-0.62 (0.55)	-1,072.02 (386.20)	-7.36 (2.97)	-1.49 (0.33)
POV implem. shortfall (%)	-0.6 (0.4)	-0.1 (0.0)	-93.0 (32.0)	-0.64 (0.26)	-0.1 (0.0)
10-min var. / 1-min var.	0.20 (0.061)	0.14 (0.037)	0.025 (0.38)	0.49 (0.16)	0.13 (0.027)
1-min var. / 10-sec var.	0.18 (0.059)	0.12 (0.016)	0.008 (0.012)	0.19 (0.029)	0.11 (0.009)
Range volatility (%)	2.76 (0.81)	0.42 (0.18)	254,020,459 (208,446,679)	1.72 (0.27)	0.47 (0.10)
Realized 1-min std. dev. (%)	0.023 (0.004)	0.009 (0.001)	332,281 (362,787)	0.021 (0.002)	0.010 (0.001)
Flash crash probability (%)	0.0 (0.0)	6.0 (0.0)		0.0 (0.0)	0.0 (0.0)
Higuchi fractal dimension	1.79 (0.044)	1.89 (0.035)	1.65 (0.29)	1.61 (0.049)	1.91 (0.023)

Table 4: Key metrics under the market stress scenario (including POV)

The implementation shortfall of the POV Sell Program is computed as the difference between the volume weighted average execution price (VWAP) and the arrival price.

		LFT only	LFT w. hom. ELPs		LFT w. het. ELPs
			non transaction tax	transaction tax	
BUY	Trading volume	56,879 (3,985.51)	205,933 (20,129)	113,164 (23,300)	183,807 (13,637)
	Trading market share (%)	17.92 (0.43)	8.86 (0.70)	13.63 (0.93)	9.33 (0.75)
	Volume market share (%)	20.56 (0.67)	9.48 (0.92)	15.27 (1.14)	10.13 (0.93)
	Trade aggressiveness ratio (%)	74.73 (5.87)	96.33 (1.22)	94.89 (1.34)	94.92 (0.34)
	Volume aggressiveness ratio (%)	72.20 (7.26)	95.55 (1.45)	94.63 (1.60)	93.41 (0.56)
	e-o-d inventory	56,879 (3,985.51)	205,933 (20,129)	113,164 (23,300)	183,807 (13,637)
	e-o-d profit (x1,000)	-63,107 (21,509)	-233,577 (26,049)	-130,016 (30,289)	-209,424 (16,921)
SELL	Trading volume	70,779 (4,067.87)	412,251 (25,340)	154,899 (10,991)	372,255 (13,829)
	Trading market share (%)	23.16 (1.47)	18.31 (1.23)	19.57 (1.50)	19.51 (1.34)
	Volume market share (%)	25.80 (1.80)	19.54 (1.56)	21.42 (1.65)	21.03 (1.35)
	Trade aggressiveness ratio (%)	25.99 (7.45)	84.94 (1.92)	39.16 (6.30)	80.97 (1.70)
	Volume aggressiveness ratio (%)	27.76 (8.10)	83.24 (2.23)	37.91 (6.46)	77.70 (2.23)
	e-o-d inventory	-70,779 (4,067.87)	-412,251 (25,340)	-154,899 (10,991)	-372,255 (13,829)
	e-o-d profit (x1,000)	86,371 (181.43)	468,669 (34,071)	173,531 (12,095)	424,222 (22,549)
OPP.	Trading volume	136,816 (13,559)	473,244 (48,648)	308,692 (70,192)	419,660 (37,375)
	Trading market share (%)	52.73 (1.95)	24.57 (2.17)	44.71 (4.03)	25.24 (2.31)
	Volume market share (%)	50.17 (2.20)	22.44 (2.28)	42.36 (4.03)	23.79 (2.19)
	Trade aggressiveness ratio (%)	56.59 (0.93)	89.31 (1.33)	64.25 (3.97)	87.56 (1.64)
	Volume aggressiveness ratio (%)	53.50 (1.06)	88.27 (1.36)	62.17 (3.95)	85.94 (1.87)
	e-o-d inventory	17,864 (3,996.89)	210,273 (13,078)	46,866 (16,330)	192,315 (13,563)
	e-o-d profit (x1,000)	-15,131 (20,585)	-238,340 (22,222)	-53,412 (24,033)	-218,918 (18,395)
SMALL	Trading volume	4,865.88 (240.29)	7,914.18 (338.46)	6,609.13 (518.92)	7,447.67 (275.15)
	Trading market share (%)	6.19 (0.37)	1.34 (0.11)	3.32 (0.24)	1.41 (0.11)
	Volume market share (%)	1.03 (0.17)	0.000 (0.000)	0.14 (0.35)	0.000 (0.000)
	Trade aggressiveness ratio (%)	10.74 (0.73)	70.71 (1.62)	22.70 (2.09)	64.69 (2.48)
	Volume aggressiveness ratio (%)	10.65 (0.79)	70.98 (1.67)	22.85 (2.14)	64.99 (2.51)
	e-o-d inventory	-3,963.62 (222.06)	-3,839.48 (167.83)	-4,558.83 (310.61)	-3,748.55 (135.36)
	e-o-d profit (x1,000)	27.59 (9.56)	-3.98 (3.81)	-22.72 (31.20)	2.28 (2.14)
ELP	Trading volume		978,983 (232,839)	129,791 (18,141)	750,648 (135,985)
	Trading market share (%)		46.91 (3.90)	18.76 (3.55)	44.51 (4.09)
	Volume market share (%)		46.19 (4.40)	18.05 (3.65)	42.57 (3.89)
	Trade aggressiveness ratio (%)		6.01 (5.59)	0.0 (0.0)	4.80 (5.63)
	Volume aggressiveness ratio (%)		6.53 (6.03)	0.0 (0.0)	4.57 (5.31)
	e-o-d inventory		-115.39 (115.46)	-571.24 (427.14)	-117.48 (77.90)
	e-o-d profit (x1,000)		244.78 (31.16)	40.95 (26.26)	252.21 (13.81)

Table 5: Agent statistics under normal conditions

The end-of-day profits reflect the intraday cash flow associated with trading, as well as the virtual profit from marking to market the end-of-day inventory.

		LFT only	LFT w. hom. ELPs			LFT w. het. ELPs
			non transaction tax		transaction tax	
			non Flash Crash	Flash Crash		
BUY	Trading volume	56,909 (4,314.64)	213,599 (19,825)	570,361 (124,216)	108,842 (20,273)	188,964 (13,808)
	Trading market share (%)	18.28 (0.45)	9.06 (0.73)	4.21 (0.62)	13.14 (0.70)	9.57 (0.76)
	Volume market share (%)	21.18 (0.70)	9.68 (0.93)	4.98 (1.00)	14.57 (0.90)	10.42 (0.91)
	Trade aggressiveness ratio (%)	71.23 (6.03)	96.21 (0.96)	31.79 (10.32)	92.83 (2.32)	94.85 (0.36)
	Volume aggressiveness ratio (%)	67.50 (7.26)	95.41 (1.14)	25.62 (9.59)	91.96 (2.78)	93.31 (0.54)
	e-o-d inventory	56,909 (4,314.64)	213,599 (19,825)	570,361 (124,216)	108,842 (20,273)	188,964 (13,808)
	e-o-d profit (x1,000)	-60,187 (20,872)	-239,806 (25,477)	-177,382 (10,959)	-126,393 (27,919)	-214,248 (13,655)
SELL	Trading volume	68,727 (4,652.86)	405,934 (23,576)	460,437 (31,010)	153,227 (12,541)	366,416 (12,565)
	Trading market share (%)	23.31 (1.47)	17.80 (1.25)	3.93 (0.61)	19.39 (1.31)	19.07 (1.31)
	Volume market share (%)	25.75 (1.82)	18.84 (1.55)	4.07 (0.61)	21.01 (1.43)	20.67 (1.39)
	Trade aggressiveness ratio (%)	28.06 (7.51)	81.48 (1.61)	62.04 (6.67)	44.88 (5.28)	76.51 (1.64)
	Volume aggressiveness ratio (%)	30.02 (8.15)	79.54 (1.84)	55.58 (7.52)	43.82 (5.45)	72.56 (2.19)
	e-o-d inventory	-68,727 (4,652.86)	-405,934 (23,576)	-460,437 (31,010)	-153,227 (12,541)	-366,416 (12,565)
	e-o-d profit (x1,000)	86,138 (5,935.74)	460,849 (30,647)	680,872 (273,446)	175,667 (13,478)	413,676 (22,653)
OPP.	Trading volume	129,448 (12,991)	488,102 (48,153)	841,034 (146,760)	283,657 (59,819)	429,625 (37,752)
	Trading market share (%)	51.91 (1.96)	25.05 (2.23)	7.58 (0.40)	41.44 (3.14)	25.76 (2.31)
	Volume market share (%)	48.80 (2.21)	22.77 (2.29)	7.95 (0.61)	38.78 (3.09)	24.27 (2.22)
	Trade aggressiveness ratio (%)	56.10 (0.93)	88.10 (1.22)	46.19 (10.58)	66.54 (3.18)	86.52 (1.65)
	Volume aggressiveness ratio (%)	53.06 (1.10)	87.10 (1.26)	37.69 (10.85)	64.49 (3.17)	84.86 (1.83)
	e-o-d inventory	18,063 (3,822.30)	219,497 (12,746)	-46,513 (131,786)	56,869 (11,687)	199,980 (12,981)
	e-o-d profit (x1,000)	-14,860 (20,465)	-250,441 (21,147)	719,348 (556,539)	-65,715 (22,545)	-224,610 (19,857)
SMALL	Trading volume	4,480.59 (264.29)	8,043.18 (315.32)	12,899 (1,620.17)	6,068.33 (462.99)	7,535.43 (267.90)
	Trading market share (%)	5.94 (0.35)	1.36 (0.12)	0.35 (0.028)	3.06 (0.20)	1.43 (0.11)
	Volume market share (%)	1.00 (0.000)	0.000 (0.000)	0.000 (0.000)	0.010 (0.10)	0.000 (0.000)
	Trade aggressiveness ratio (%)	11.30 (0.82)	67.39 (1.76)	42.57 (6.74)	25.08 (2.01)	60.77 (2.12)
	Volume aggressiveness ratio (%)	11.23 (0.87)	67.76 (1.82)	33.74 (9.78)	25.24 (2.05)	61.09 (2.14)
	e-o-d inventory	-3,594.67 (231.26)	-3,822.23 (180.33)	-2,961.76 (1,101.73)	-4,065.67 (244.97)	-3,734.43 (125.28)
	e-o-d profit (x1,000)	61.20 (21.72)	-4.59 (3.59)	9,074.94 (1,214.68)	38.29 (14.17)	4.31 (1.23)
ELP	Trading volume		977,359 (239,819)	8,086,457 (694,688)	156,649 (14,827)	729,256 (134,853)
	Trading market share (%)		46.01 (4.09)	83.62 (0.66)	22.17 (2.66)	43.31 (4.12)
	Volume market share (%)		45.20 (4.48)	80.10 (1.53)	21.61 (2.71)	41.19 (3.88)
	Trade aggressiveness ratio (%)		6.23 (5.83)	50.56 (1.89)	0.0 (0.0)	4.89 (5.71)
	Volume aggressiveness ratio (%)		6.66 (6.23)	52.36 (2.28)	0.0 (0.0)	4.68 (5.41)
	e-o-d inventory		-563.23 (261.11)	341.23 (176.55)	190.61 (706.07)	-362.36 (136.30)
	e-o-d profit (x1,000)		268.21 (29.75)	-979,729 (2,007,815)	39.51 (27.44)	265.18 (14.73)
POV	Trading volume	2,649 (440.93)	22,777 (1,869.79)	60,790 (12,066)	8,608 (1,684)	18,432 (1,124.90)
	Trading market share (%)	0.57 (0.062)	0.72 (0.066)	0.30 (0.034)	0.81 (0.074)	0.85 (0.083)
	Volume market share (%)	0.53 (0.50)	0.85 (0.36)	0.000 (0.000)	0.97 (0.17)	0.85 (0.36)
	Trade aggressiveness ratio (%)	100.00 (0.000)	100.00 (0.000)	100.00 (0.000)	100.00 (0.000)	100.00 (0.000)
	Volume aggressiveness ratio (%)	100.00 (0.000)	100.00 (0.000)	100.00 (0.000)	100.00 (0.000)	100.00 (0.000)
	e-o-d inventory	-2,650 (441)	-22,777 (1,869.79)	-60,790 (12,066)	-8,608 (1,684)	-18,432 (1,124.90)
	e-o-d profit (x1,000)	14.57 (10.14)	40,603 (9,697.04)	4,613.39 (17,151)	42.60 (21.76)	12,886 (19,781)

Table 6: Agent statistics under stress market conditions (including POV)

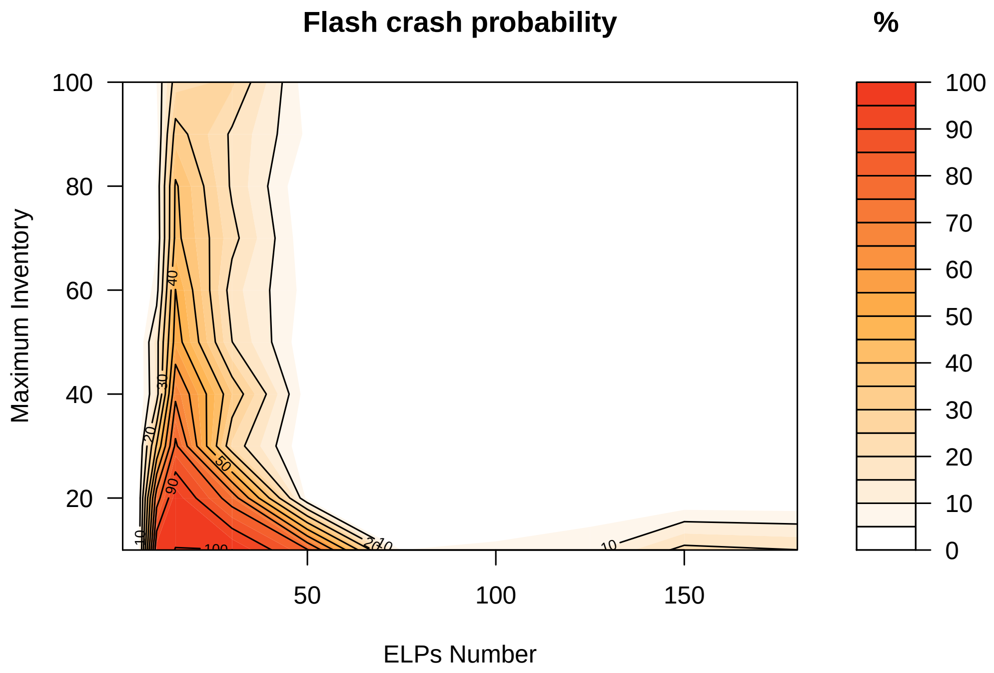

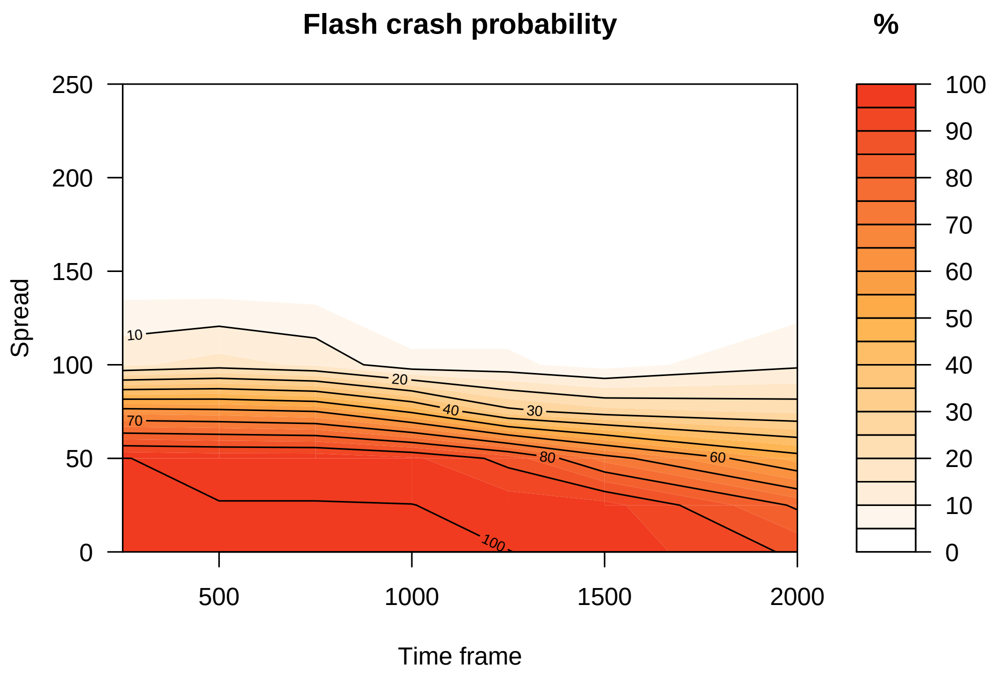

Figure 2: Flash crash probabilities for various model configurations and policy options

The upper plot reflects the outcomes of different number of homogenous ELP agents vs. their maximum threshold inventory levels (their activity time frame is fixed to one second and their quote size to seven). The lower plot shows the impact of varying the activity time frames (milliseconds) vs. the minimum spread (ticks) on the flash crash probability.

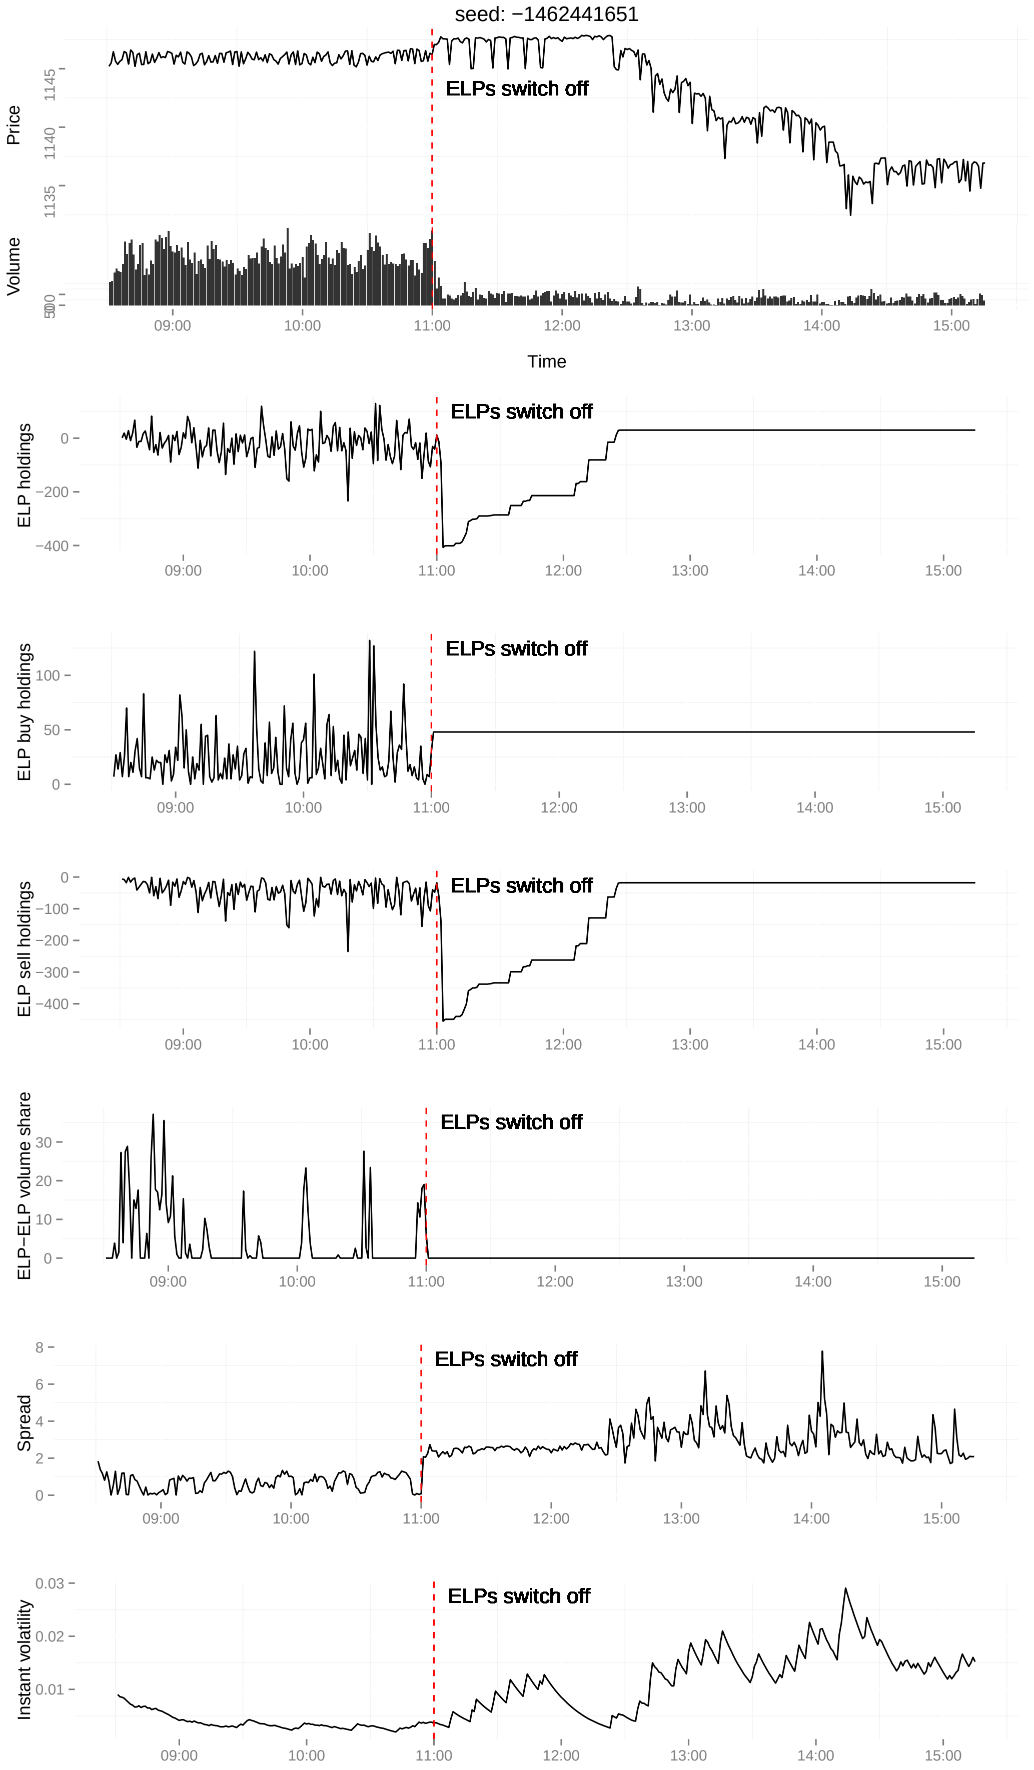

Figure 3: Price, volume and market quality time series – LFT w. hom. ELPs until 11:00

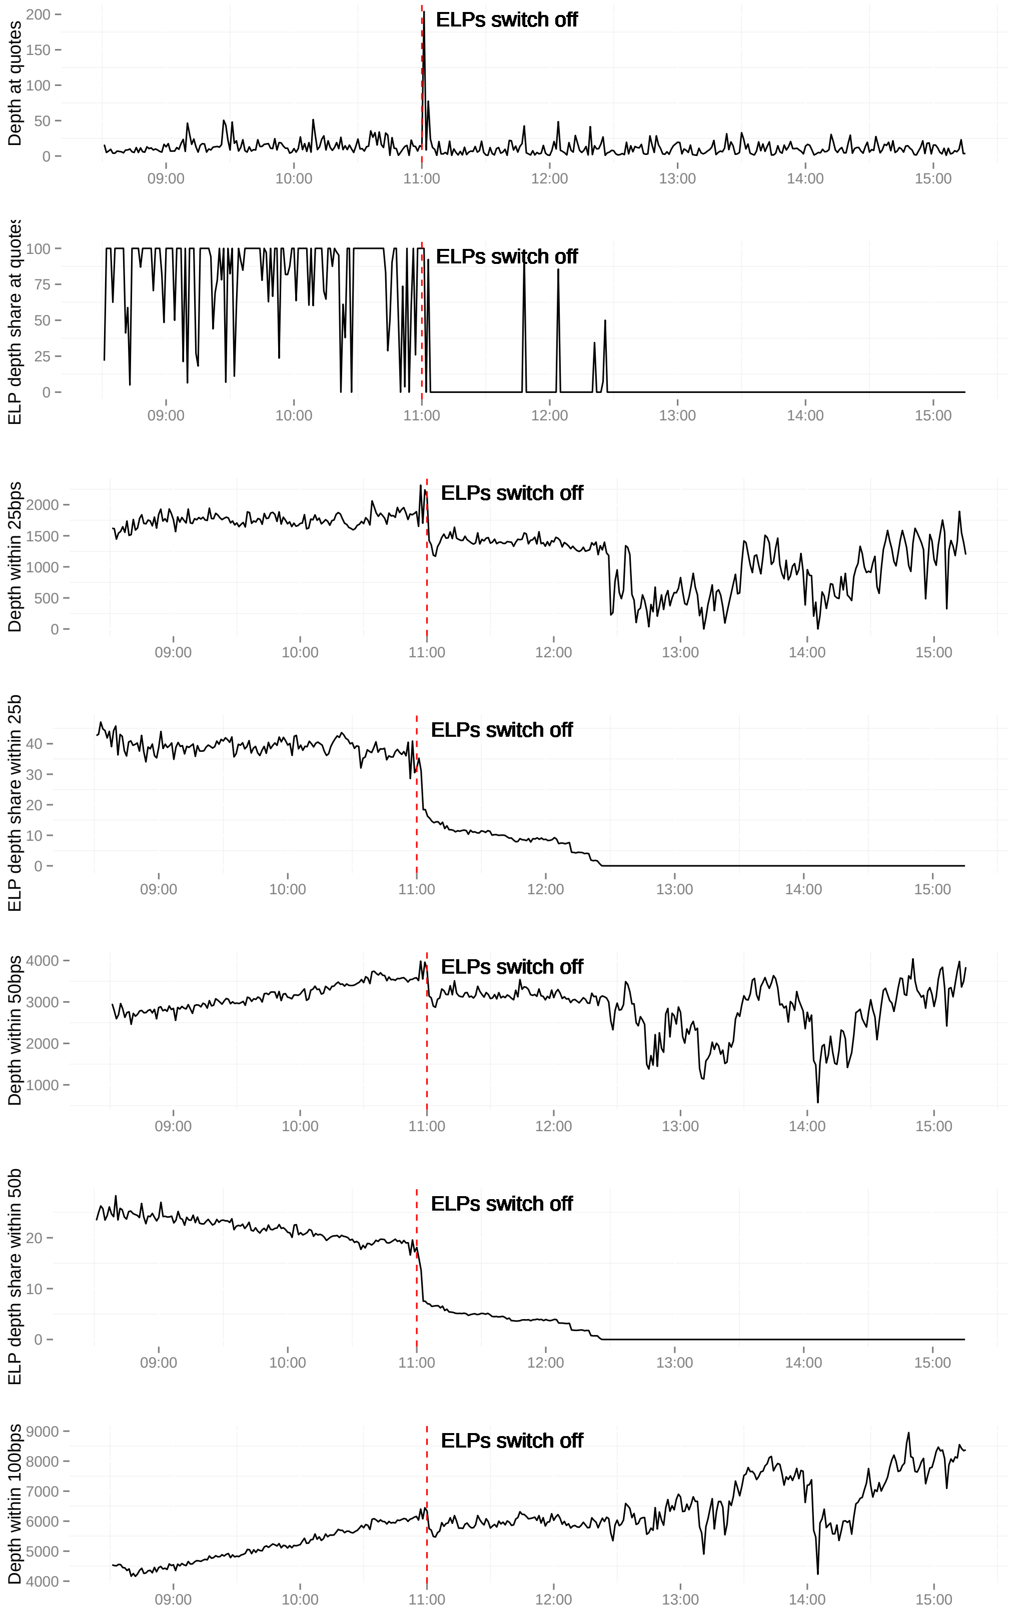

Figure 4: Other market quality time series (ctd.) – LFT w. hom. ELPs until 11:00

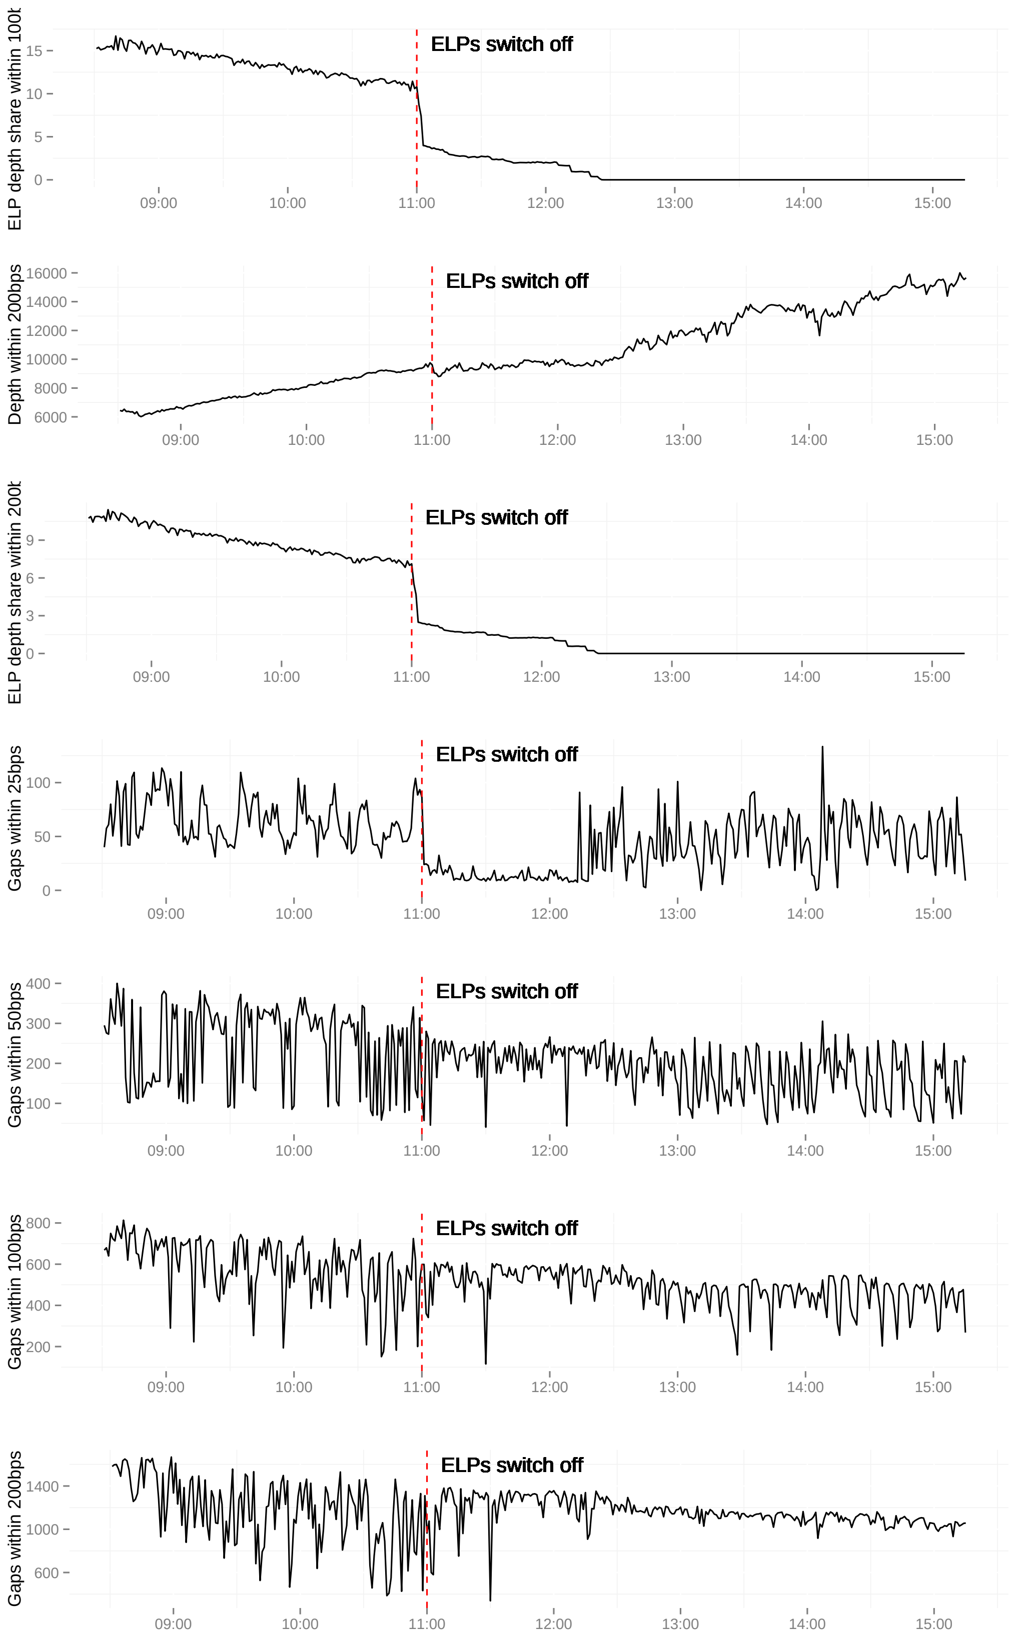

Figure 5: Other market quality time series (ctd.) – LFT w. hom. ELPs until 11:00

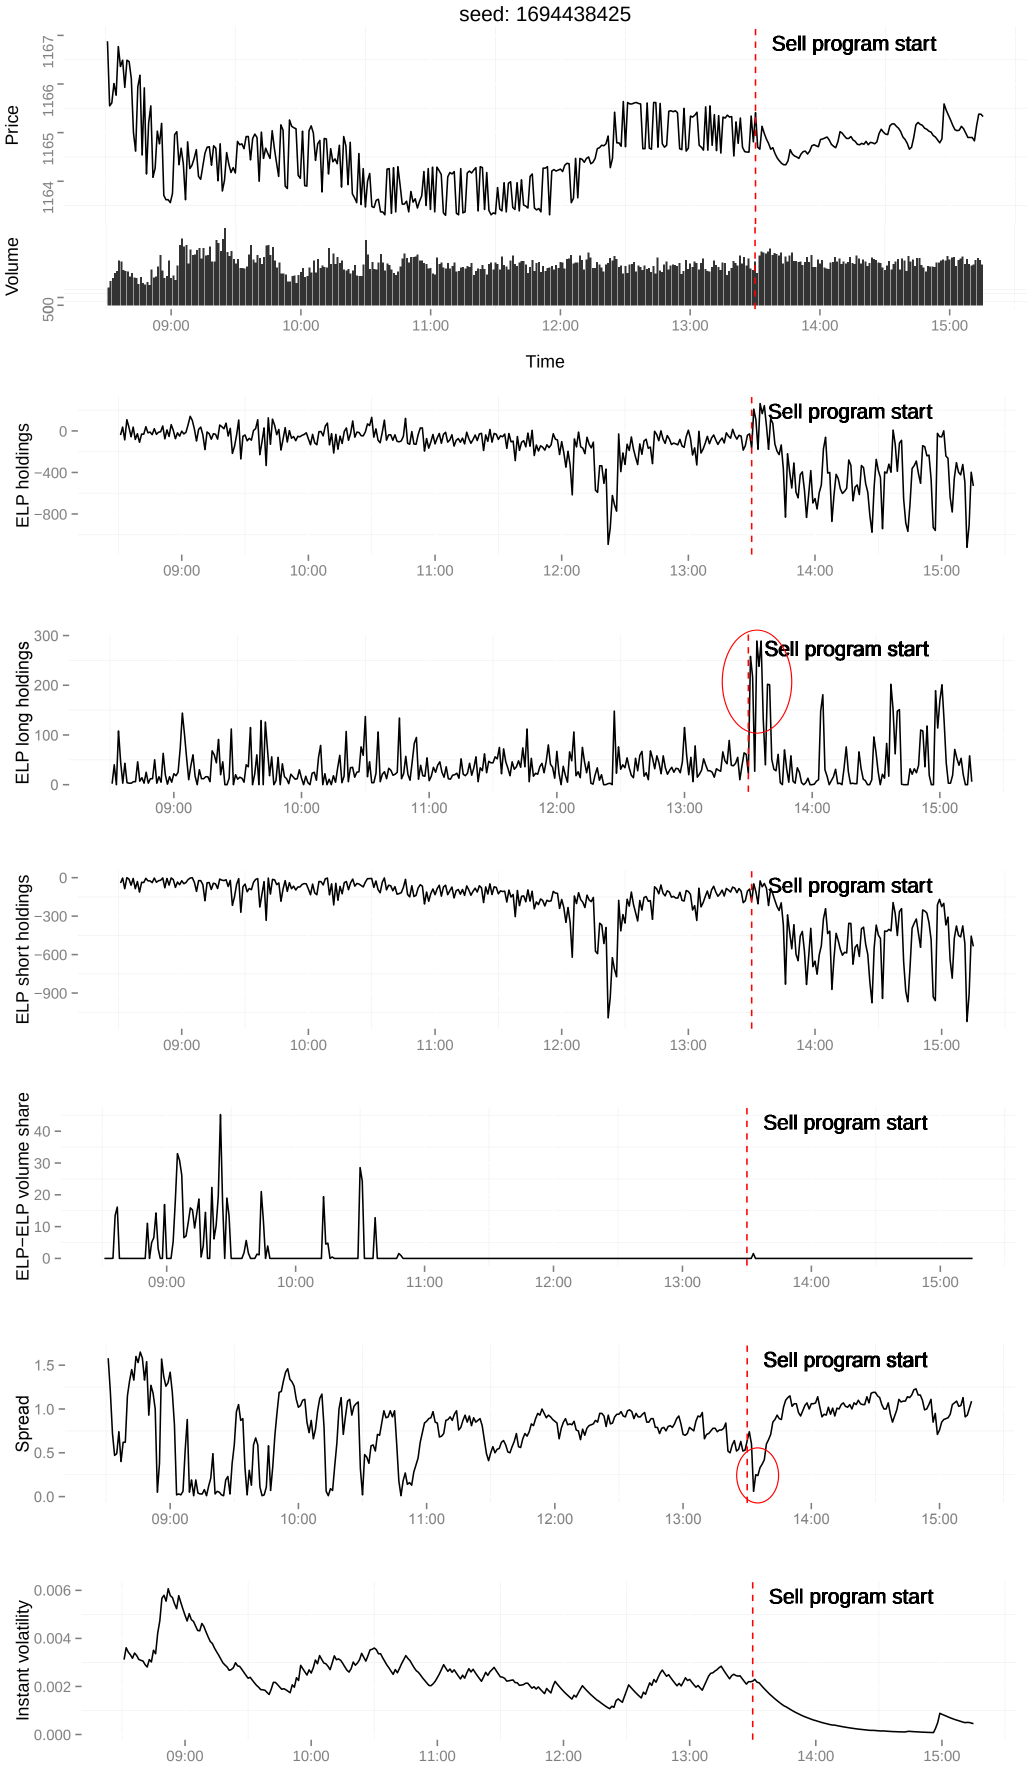

Figure 6: Price, volume and market quality time series – LFT w. hom. ELPs, POV (no flash crash)

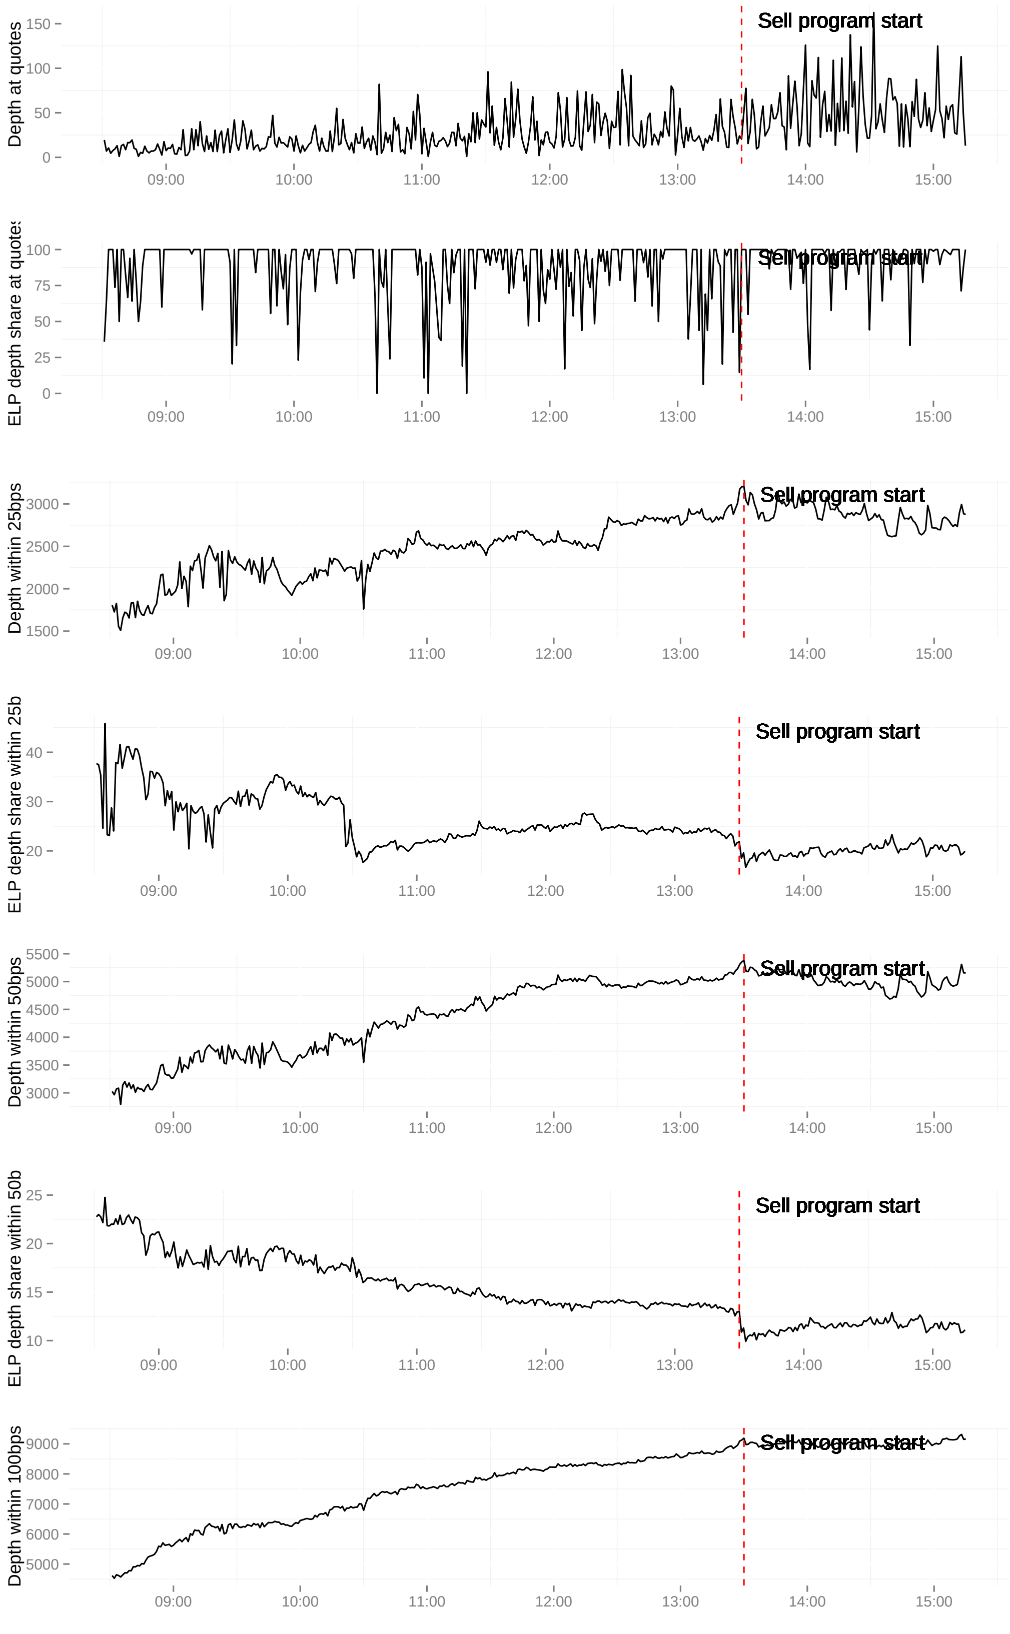

Figure 7: Other market quality time series (ctd.) – LFT w. hom. ELPs, POV (no flash crash)

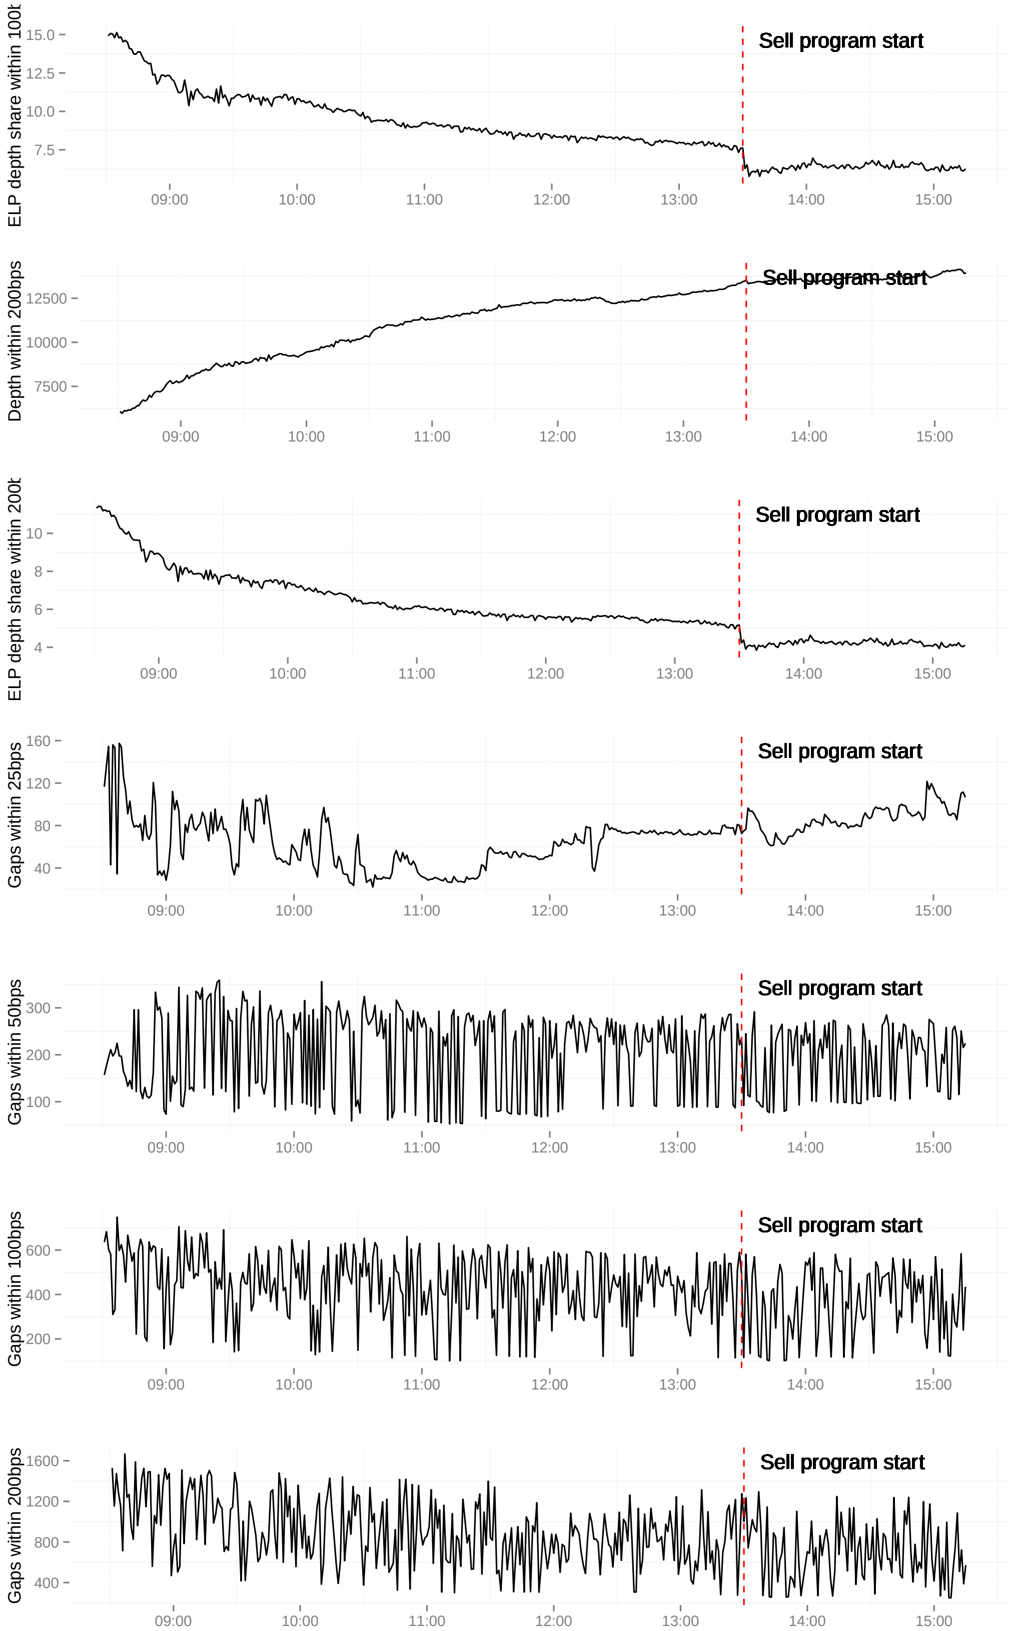

Figure 8: Other market quality time series (ctd.) – LFT w. hom. ELPs, POV (no flash crash)

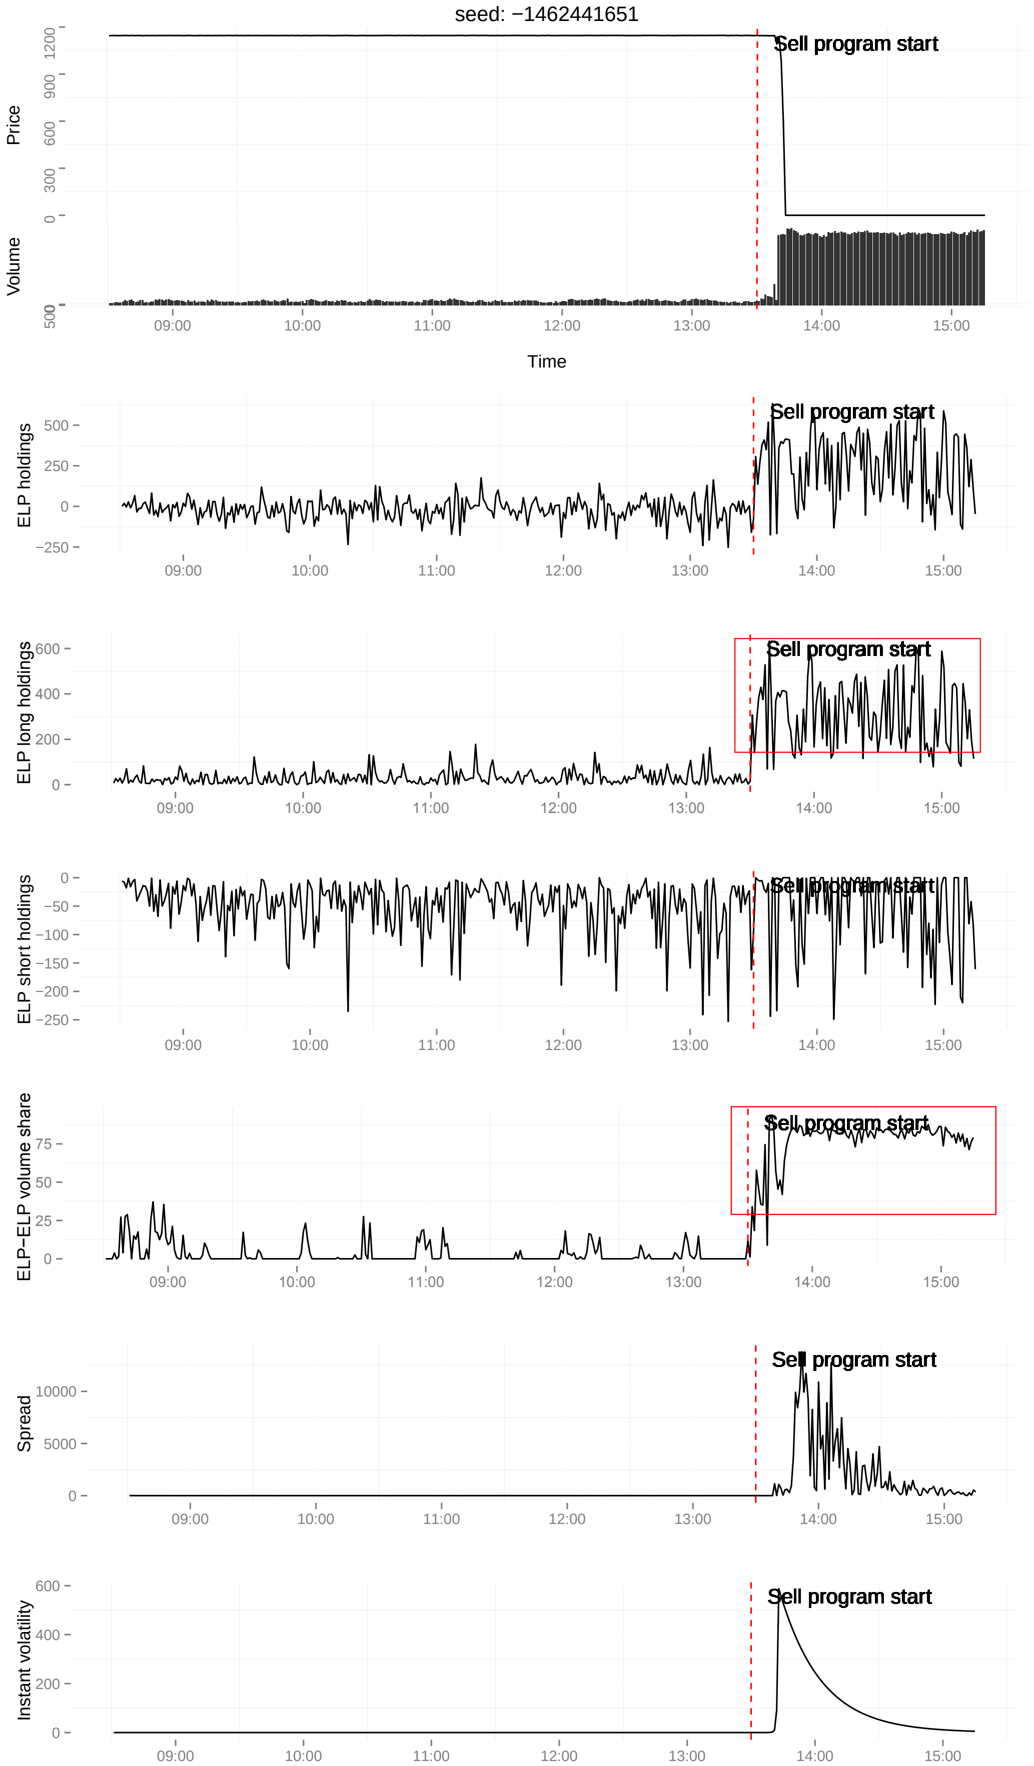

Figure 9: Price, volume and market quality time series – LFT w. hom. ELPs, POV (flash crash)

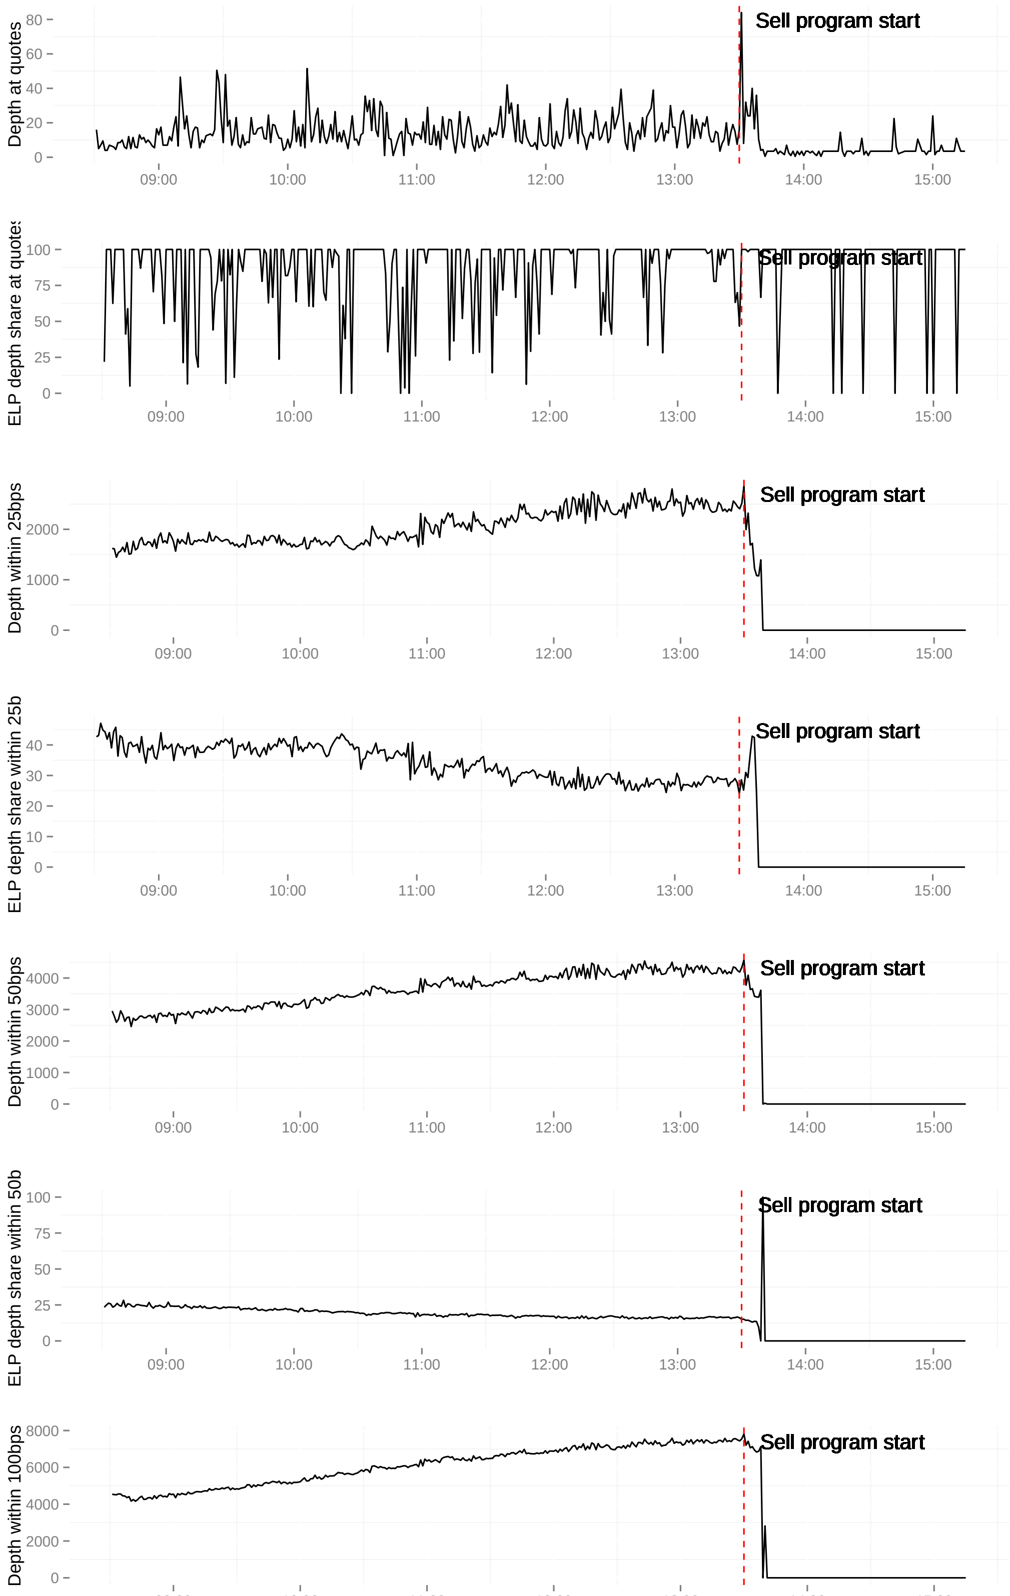

Figure 10: Other market quality time series (ctd.) – LFT w. hom. ELPs, POV (flash crash)

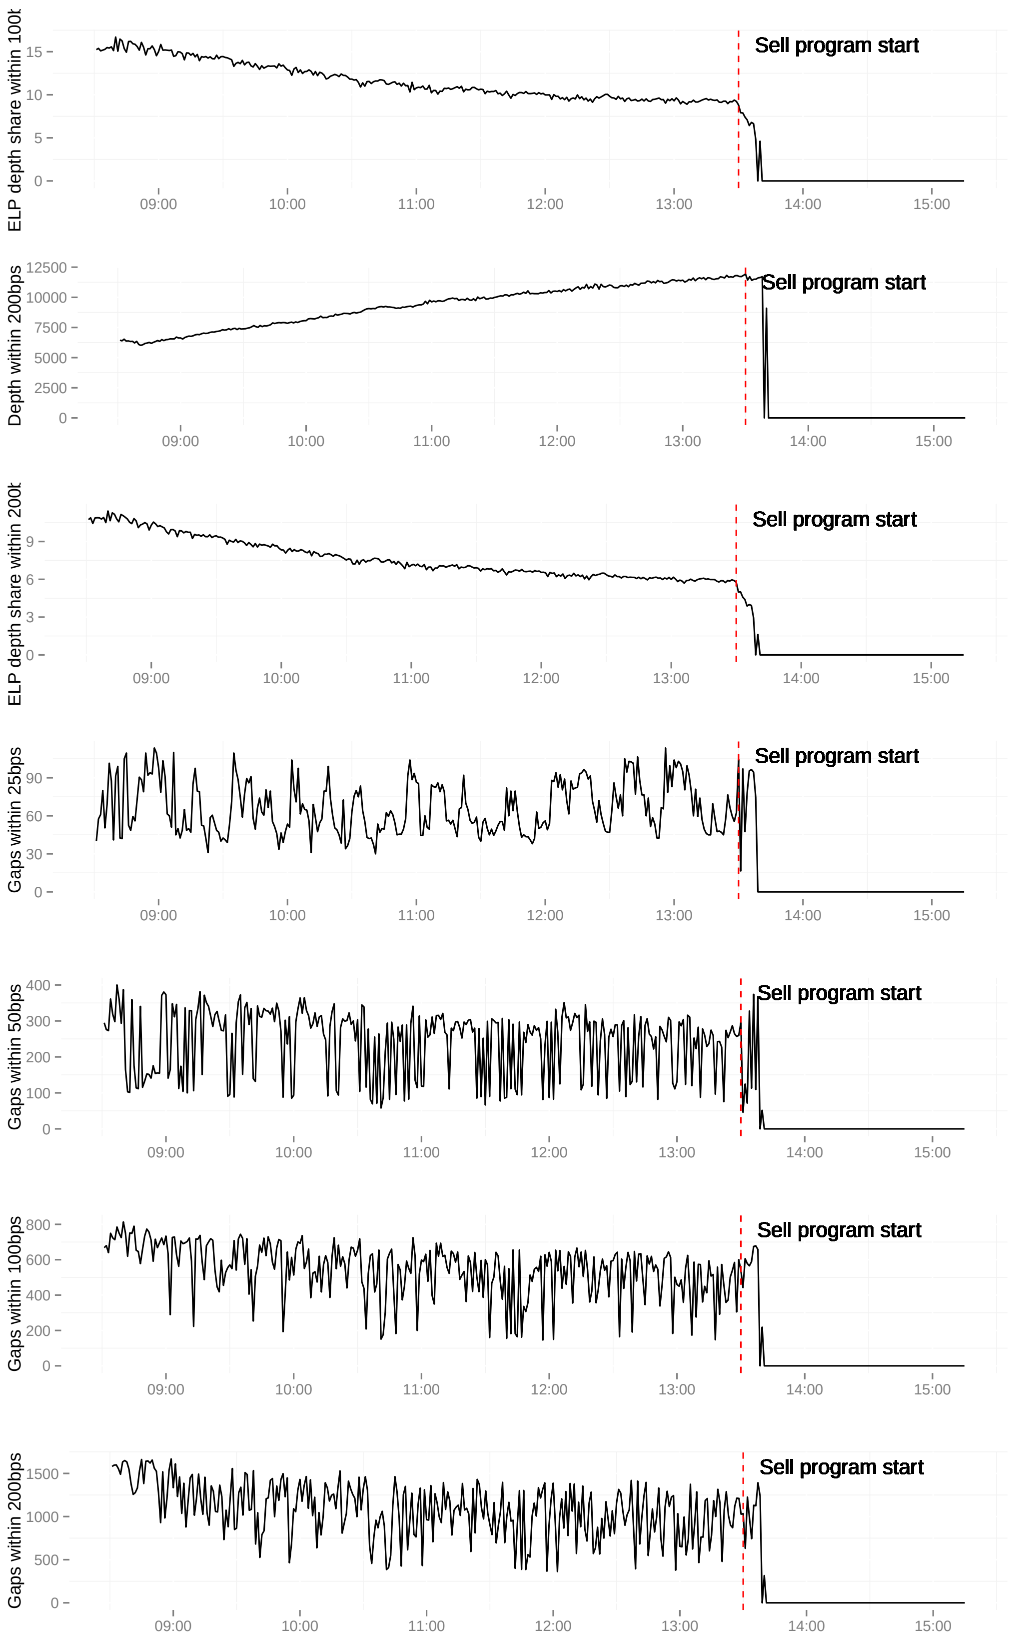

Figure 11: Other market quality time series (ctd.) – LFT w. hom. ELPs, POV (flash crash)