

Market Liquidity and Heterogeneity in the Investor Decision Cycle

Richard Bookstaber

Office of Financial Research

richard.bookstaber@treasury.gov

Michael D. Foley

University of Vermont

mikefoley37@gmail.com

Brian F. Tivnan

The MITRE Corporation

btivnan@mitre.org

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

Market Liquidity and Heterogeneity in the Investor Decision Cycle

Richard Bookstaber · Michael D. Foley ·
Brian F. Tivnan

March 10, 2015

Abstract During liquidity shocks such as occur when margin calls force the liquidation of leveraged positions, there is a widening disparity between the reaction speed of the liquidity demanders and the liquidity providers. Those who are forced to sell typically must take action within the span of a day, while those who are providing liquidity do not face similar urgency. Indeed, the flurry of activity and increased volatility of prices during the liquidity shocks might actually reduce the speed with which many liquidity providers come to the market. To analyze these dynamics, we build upon previous agent-based models of financial markets to develop an order-book model with heterogeneity in trader decision cycles. The model demonstrates an adherence to important stylized facts such as a leptokurtic distribution of returns, decay of autocorrelations over moderate to long time lags, and clustering volatility. We show that the heterogeneity in decision cycles can increase the severity of market shocks, and even absent a shock can have notable effects on the stochastic properties of market prices.

1 Introduction

Just as the day-to-day movement of prices gives little insight into the potential for the dislocations that emerge during a fire sale, so also the day-to-day liquidity of the market gives little indication for the drying up of liquidity that often occurs during those events. Contrary to the tenets of demand theory,

R. Bookstaber
Office of Financial Research, U.S. Treasury
E-mail: richard.bookstaber@treasury.gov

M. Foley
University of Vermont E-mail: mikefoley37@gmail.com

B. Tivnan
MITRE Corporation E-mail: btivnan@mitre.org

supply may not be immediately forthcoming at any price. Indeed, it often seems to be the contrary; the rapid drop in prices as those who face the urgent need to liquidate try to elicit the other side of the trade actually reduces the supply in the market.

The key feature of the market that leads to illiquidity during times of urgent demand is that the frequency with which agents arrive at the market is heterogeneous among market participants. Some agents, such as high frequency and statistical arbitrage traders, are continuously engaged in the market and can take action very quickly based on changes in prices. Others, such as longer-term fundamental investors, will not enter the market immediately in reaction to changes in the market price because they have a slower decision cycle. Their decision cycle might involve consulting with others in their firm or reformulating a broader investment plan before taking action. Furthermore, as a practical matter, these longer-term investors tend to be the ones with the deepest pockets in terms of liquidity supply.

Because of the heterogeneous arrivals, and more to the point because the liquidity demand has a greater frequency of arrival than does the deep-pocket supply, liquidity will appear to have dried up during periods of high immediacy. Impatience reduces the effectiveness of price as a signal. What is more, the very attempt to use prices as a signal can lead to perverse results. When those with a lower frequency finally do arrive, the price might have dropped to the point that they are not willing to supply liquidity without regrouping, reducing their frequency of participation as they huddle with their analysts to assess the market action and become more timid with the liquidity they do supply.

The heterogeneous decision cycle dynamic was evident during the October 19, 1987 market crash (Presidential Task Force on Market Mechanisms 1988). Early on that Monday morning, computer-generated sell orders from portfolio insurance programs hit the futures market. Program traders then bought these positions in the index and simultaneously put in orders to sell the underlying stocks, leading to a flood of sell orders into the New York Stock Exchange. The specialists, with limited capital to inventory these orders, rapidly dropped the prices of the stocks in order to elicit buying. But many of the large investors, focused on longer-term positions rather than intra-day trading, were not as quick to buy as the computer programs were to sell. Though the specialists signaled in classical economic fashion by reducing prices, these fell, so to speak, on deaf ears. By the time many of the stock investors got around to their start-of-week portfolio activities, prices had dropped through successive specialist price reductions to the point that the liquidity that might have otherwise been forthcoming was forestalled.

The heterogeneity of decision cycles led to a similar result during the Flash Crash on May 6, 2010 (CFTC and SEC 2010). As confirmed by Kirilenko et al. (2011), the replenishment of the order book could not keep up with the essentially instantaneously triggered market stop loss orders, even with the microsecond speed of the high frequency traders. And, once again, the heterogeneity of the decision cycle came with a one-two punch: first prices dropped precipitously as market stop loss orders ran through the order book,

and then, with the prices in disarray, many liquidity providers pulled away from the market to recalibrate their market activity, leading prices to continue their cascade, in some cases to absurd levels. As Kirilenko et (2011) concluded, while heterogeneity in investors' decision cycles did not trigger the Flash Crash, that heterogeneity certainly exacerbated market volatility.

Duffie (2010) presents a model highlighting the impact of inattentive investors with particular interest on its implications for the 2008 financial crisis. The effect of heterogeneous decision cycles is like that of the inattentive investors in Duffie's model; in either case price signals will fail to transmit to a subset of market participants who are not actively in the market at the moment the price signal is made. As Duffie shows, this will have implications for price dynamics even during typical day-to-day market activity. But it will have particularly adverse effects during times of forced selling. If the liquidity demander takes the lack of response to his price concessions as being an indication that the price level is still too high rather than understanding that the lack of response is due to a lack of attentiveness or a lag in the ability of the other side of the trade to make a decision, it will lead to a further drop in prices that is both unnecessary and ineffective because those who might be willing to take on the other side of the trade might literally be out to lunch. As observed by Bookstaber (2007), referring to the 1987 market crash, "Supply dried up because of the difference in time frames between the demanders and suppliers. By the time equity investors could have reacted to the prices and done some bargain hunting the specialists had moved prices so precipitously that these potential liquidity suppliers were scared away. The key culprit was the difference in liquidity because of the different trading time frames between the demanders and the suppliers in the two markets."

The objective of this paper is to detail the development and testing of a model which can capture essential characteristics of asset price dynamics through the use of heterogeneous decision cycles that range along the spectrum from intra-day traders to the long-term, fundamental investors. These market dynamics fit a number of important stylized facts and are particularly important in understanding the drying up of liquidity and resulting cascades that occur during market breaks and crises along the lines of events like the 1987 Crash and the 2010 Flash Crash, and similar events such as the Long-Term Capital Management-related fire sale in late 1998, and the August, 2007 Quant Quake.

2 Model Overview: An Agent-based approach to heterogeneous decision cycles

We model the heterogeneity of agent arrival by separating the agents into three groups: short-term liquidity takers, short-term liquidity providers, and long-term liquidity providers. For some agents, time matters more than price, while for others, price matters more than time; in a fire sale scenario, the first group includes those who are facing margin calls.

Each agent arrives at the market only periodically. That is, each agent visits the market with an arrival frequency, and the frequency for those in the first two groups is higher than for those in the last group. In particular, agents arrive in the market based on a Poisson process. The agents then buy or sell an amount dictated by a draw from a normal distribution, and specific trader class factors which we discuss in the model description.

The market consists of an order book that posts a given amount to be bought or sold at every price. Market clearing occurs in each period by matching buyers and sellers at the last-period price, and then moving up or down the price to extract more orders from the order book. The price at which the buying and selling is finally equilibrated becomes the current posted price for that period.

Even if the frequency of arrival and the quantity demanded are not affected by market conditions, the demand for immediacy might overwhelm the supply that arrives at a slower pace, and arrives at the slower pace no matter what the price change (due to a longer decision-making and trade approval process). In this situation, the liquidity demanders might sell to drop prices, not see enough supply forthcoming, and so drop prices further, resulting in even more demand. In this situation the price might not be an effective signal to the liquidity suppliers who have a slower arrival rate. Changes in market conditions can accentuate the market cascades. The market might recover from a market shock, or might go into a free fall depending on the relative effect of market conditions on the frequency of arrival on the one hand and the reduction of the quantity offered on the other. For example, it might be that once the lower-frequency liquidity demanders arrive, they see the change in market conditions and, uncertain what is causing this price drop, elect to tread into the market more gingerly.

In the process of developing and testing the model: (a) we replicated and extended a leading market model while synthesizing with other leading models; (b) during the development and testing of our model, we employed the highest level criteria for validation testing, while also thoroughly sampling the relevant region of the parameter space of our model; (c) we applied our model to illuminate interdependencies between liquidity and trader decision cycles and their effect on market stability; and (d) we demonstrated that our model can be used to run controlled policy experiments (e.g., stress tests). The organization of the remainder of the paper begins with a review the foundational literature underlying the development and testing of the model, a detailed description of the model and its evolution, our design of experiments, the analysis of those experiments and our conclusions.

2.1 Review of Foundational Literature

In this study, we build upon three distinct bases of literature: modeling methodologies, empirical findings for financial markets known as stylized facts, and the evolving set of agent-based models of financial markets.

In the general modeling literature, Balci (1997) defines an iterative approach for verification (i.e., tests to confirm whether we built the model correctly) and validation (i.e., tests to confirm that we built the correct model). As part of verification testing, Axtell et. al. (1996) call for researchers to contribute to the cumulative foundation provided by previous models from previous studies. Therefore, the successful researcher first replicates the model from a previous study to ensure an independent verification of that study and its findings. Only then should the researcher undertake the extension of the replicated model to address the model characteristics required of the new study. Lastly, the extended model must be subjected to a rigorous set of validation and calibration tests (LeBaron 2001a) such as the Axtell and Epstein (1994) framework for evaluating empirical relevance of agent-based models. See Tivnan et. al. (2011) for an example of a study that follows the above approach.

The empirical targets for the above validation tests are known as stylized facts. While many researchers have conducted empirical analyses of financial markets, Mandelbrot (1963) was one of the first to identify that price returns do not necessarily follow a Gaussian random walk. Many have built upon Mandelbrot's foundational work and Cont (2001) provides a comprehensive review of these empirical analyses which identify several stylized facts of financial markets. These stylized facts have been shown to be robust across the price time series of various markets over varying time intervals. While many of these stylized facts have been debated in recent years, three have largely been agreed upon and applied by the quantitative finance community. These accepted stylized facts include decaying autocorrelations over moderate time lags, clustering volatility, and fat tails in the distribution of returns.

Put simply, decaying autocorrelations indicate very small correlations in price movements over moderate to long time lags, clustering volatility indicates that price returns of similar absolute size tend to follow one another, and fat tails indicate that returns are much more peaked (that is, have higher kurtosis and fatter tails) than a Gaussian distribution would predict. As described by Cont (2001: 224), these stylized facts are so constraining that it is not easy to exhibit even an (ad hoc) stochastic process which possesses the same set of properties and one has to go to great lengths to reproduce them with a model. Stated differently, this set of stylized facts provides empirical evidence to guide the comprehensive, validation testing of market models.

While there has been much discussion of the potential for agent-based models to advance the study of financial markets (Bookstaber 2012), it is important to note the evolution of agent-based models of financial markets. Consistent with two of the prevailing design issues (Lebaron 2001a), the evolution of the set of agent-based models of financial markets follow two distinct paths, one focused on the representation of the agents and the other focused on the representation of the trading mechanism. The former trajectory began in the early 1990s with the Santa Fe Institute (SFI) market model (Palmer et. al 1994). Largely predicated on Holland's (1977) genetic algorithms, the SFI market model was well received as a novel departure from equilibrium models, largely

based on its qualitative depiction of agreement with empirical observations of market dynamics such as bubbles and crashes. Following the SFI market model, Lux et. al. (1999) introduced a model with a single trade type (i.e., market orders) that was the first to demonstrate clustered volatility, one of the stylized facts common to many markets. Similar to the Lux market with market orders, LeBaron (2001b) introduced a market model with agents that learn based on a neural network. In the LeBaron model, agents decided how much of their total wealth to invest; therefore successful agents can have a large impact in the market. Cont and colleagues (Ghoulmie et al. 2005) developed a model that qualitatively reproduces the three prevailing stylized facts described above. Of note, the traders in the Cont models have heterogeneous trading thresholds, and the traders, many of whom trade rather infrequently, adapt their thresholds based upon performance feedback.

The trajectory of models focused on trading mechanisms began with Maslov (2000) nearly a decade after the SFI market model. Maslov's model was quickly extended by Darley et al. (2001) and Farmer et al. (2005). The Farmer model (subsequently referred to as the ZIM for Zero-Intelligence Model) built a model of zero-intelligence traders active within the structure of a continuous, double auction placing two types of orders: market orders and limit orders. Market orders are orders that enter the market with an intent to buy or sell a certain number of shares and do not specify a particular price. Limit orders, on the other hand, enter the market with both a specified price and quantity of shares. As market orders do not specify a price they are executed immediately upon entering the market at the best available price. Limit orders, however, will accumulate in the market until their specified price is met or they are cancelled. The accumulation takes place in a prioritized queue by price and arrival time. This accumulation of limit orders is called the limit order book. It is this accumulation of limit orders that creates liquidity (i.e., the ability for market orders to be executed) in the market. In the ZIM, both market and limit orders arrive and are cancelled according to a Poisson process. While not rich enough to replicate the dynamics of market events such as the 2010 Flash Crash, the ZIM is useful in exploring structural aspects of market dynamics.

2.2 The Preis Model

A follower of the cumulative approach espoused by Axtell et al. (1996), Preis et. al. (2006, 2007) first replicated and then extended the ZIM. Their model reproduced a leptokurtic distribution of returns, which the ZIM did not. Because of this, we use the Preis model as a building block from which we extend to explore market-relevant aspects of heterogeneity that are not contained within the Preis model.

As with the ZIM, Preis' market consists of a continuous double auction order book in one asset. Both market and limit orders can be placed in the order book. When a sell (buy) market order is placed, the order is executed immediately against a buy (sell) limit order with the specified highest (lowest)

price and earliest time-stamp. We will refer to the highest buy price as the best bid ($p_b(t)$), and the lowest sell price as the best ask ($p_a(t)$).

The Preis model consists of two sets of agents, patient and impatient agents. Impatient agents place only market orders with rate μ orders per trader per time step. When an impatient agent is activated, they buy with probability q_{taker} and sell with probability $1 - q_{taker}$, where q_{taker} is a mean reverting random walk which starts at 0.5, takes steps of size .001, and is bounded between 0.45 and 0.55. In this light, we refer to q_{taker} as a trending variable, as prices will tend to drop when $q_{taker} < 0.5$ and rise when $q_{taker} > 0.5$. Patient agents place only limit orders with rate α . Patient agents place buy and sell orders with equal probability. Limit orders are priced away from the spread according to

$$\lambda(t) = \lambda_0 \left(1 + 10 \frac{|q_{taker}(t) - .5|}{\sqrt{\langle (q_{taker}(t) - .5)^2 \rangle}} \right) \quad (1)$$

where $\lambda_0 = 100$. Sell orders are placed at price $p_b(t) + X$ and buy orders are placed at price $p_a(t) - X$, where X is drawn from an exponential distribution. Note that X is independently drawn from the exponential distribution for each limit order placed. This equation ensures that as q_{taker} deviates from its mean of 0.5, traders will tend to place orders further from the spread. All traders place only orders of size σ , which is kept constant at one throughout Preis' experiments. Limit orders are cancelled randomly with rate δ per order per time step. Figure 1 provides a graphical depiction of the limit order book from Preis' model.

3 Methods

Our methodological approach to this study unfolded in two distinct phases: the development of the model and its subsequent application to explore the effects of heterogeneous decision cycles. The model development proceeded according to the following iterative cycle:

- Update the formulation of the conceptual model as necessary
- Verify the implementation of the agent-based model against the formulation
- Construct and execute of a Design of Experiment (DOE)
 - Employ a comprehensive and efficient sampling of the parameter space known as an Orthogonal, Latin Hypercube (Kleijnen et al. 2005)
- Validate the results of the DOE using a set of stylized facts beyond those demonstrated in the ZIM and Preis models

Our first iteration through this cycle addressed the replication of the Preis model. Subsequent iterations introduced the following extensions to the Preis model, in succession:

- “Agentizing” the Preis model with individual traders rather than representative trader classes

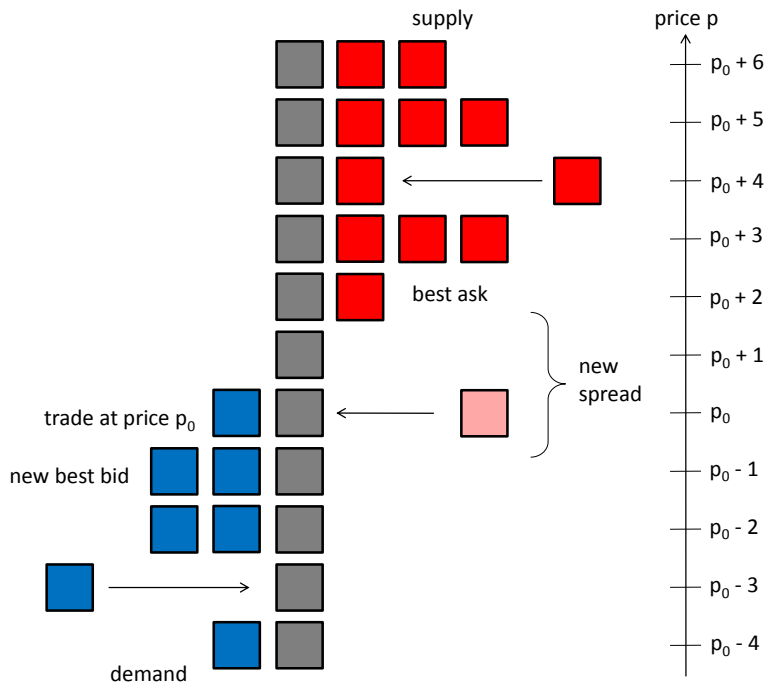


Fig. 1: Fig. 1: Example of a limit order book, adapted from Preis et al (2006). The dark red blocks are limit sell orders while the blue blocks are limit buy orders. The light red block is a sell market order, which is executed against the best (highest) priced limit buy order. The grey blocks represent a notional demarcation between supply and demand.

- Injecting heterogeneity in order rates and sizes
- Implementing an exogenous mechanism for liquidity shocks
- Introducing large firms with price sensitivity
- Introducing capital constraints for each trader

The replication of the Preis model, our extensions and application are described in greater detail below. But first, we define our methodology for validation testing using the stylized facts.

3.1 Application of the Stylized Facts

In our application of the stylized facts, we draw from Cont (2001) and require that kurtosis is greater than 5, and that autocorrelation of absolute returns is at least 5 times greater than autocorrelation of raw returns. Another important factor to consider when measuring stylized facts are what we refer to as “market collapse” cases. When a market order is placed and there are no limit orders on the opposite side of the book to fill, we call this a collapse. While

collapses occur in real markets, they are rare, and in a valid market model we would not expect to see a collapse under normal conditions. In total, our use of this set of stylized facts represents a more restrictive approach to validation testing than previously applied to agent-based models of limit order books.

Thus for design point p , where a design point is the combination of levels for all model parameters and also known as a point in the parameter space of the model (Kleijnen et al. 2005), the stylized facts pass rate R_p is calculated as:

$$R_p = \frac{1}{n} \sum_{i=1}^n \left((K(r_{p_i}) > 5) * (CC(r_{p_i}) > 5) * t(ACF) > .05) \right) \quad (2)$$

$$\text{where } CC(r_{p_i}) = \frac{\sum_{j=1}^{20} ACF(|r_{p_i}|)}{\sum_{j=1}^{20} ACF(r_{p_i})}$$

and n is the number of random seeds, ACF is the autocorrelation function, CC is the clustering volatility constant, K is the kurtosis function, $t(ACF)$ refers to the p -value of a Student t-test of the ACF, and r_{p_i} is a vector of returns for design point p and random seed i . If all 3 of the above conditions are true, the randomly seeded run is given a value of 1, and if any of the conditions are false, the run is given a value of 0. Any run which collapses is assigned a score of 0.

Replication

The first step in our model progression was to replicate Preis and his results. Following Preis' procedure for allowing the model to run for sufficient time such that it achieves an initial stationarity, we were then able to show that the parameter set described in Preis passed each of our three stylized facts (see Figure 2) Given the importance of liquidity to this study, we needed an approach to capture the underlying order book dynamics, so we adopted Paddrik et al. (2013) approach to visualizing order books. Figure 3 provides examples of order book density plots thereby enabling a direct inspection of order book liquidity. The yellow to red spectrum represents buy orders, while the purple-blue spectrum represents sell orders. Therefore, the color within the density plot reflects the amount of liquidity available at a specific price and time.

We then generated a set of 500 design points using a comprehensive and efficient search of the parameter space, known as nearly orthogonal, Latin-hypercube sampling, for the α , δ , and μ parameters. The μ and δ values were drawn uniformly from $[0, 0.2]$, while α values were drawn uniformly from $[0, 0.5]$. We discarded all design points for which $\mu > \alpha(1 - \delta)$, since any market with more market orders than limit orders will quickly collapse. This left 204

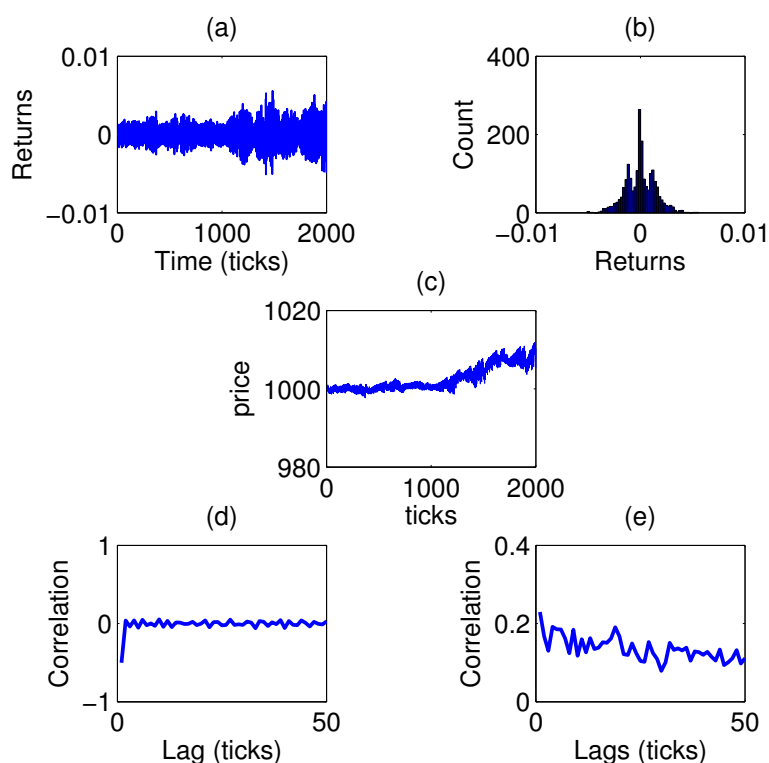


Fig. 2: Here we see an example of some of the stylized facts plots from the original Preis model. (a) Shows the time series of returns, (b) shows the histogram of returns, (c) is the raw price series, (d) shows the decay in autocorrelations past small time lags, and (e) shows the persistence of autocorrelations over moderate time lags.

design points for us to validate against the stylized facts. For each design point, we instantiated 200 unique replicates, and we ran each replicate for 10,000 time steps. Each replicate generated a price series, which we analyzed according to the stylized facts described above. Lastly, for each design point, we assigned a pass rate score according to equation (2), aggregated from the 200 replicates for that design point. The results from this validation testing are as follows:

Out of the 204 design points, 29 of them passed the stylized facts with a rate of 95% or better while also avoiding any market collapses. The next few extensions to Preis were tested against the stylized facts on only these 29 design points.

3.2 Extensions of Preis

Agentizing

Next, in order to keep better track of individual trader behavior, we agentized the Preis model. By agentized, we intend the following: rather than a single representative agent for an entire class of agents as is common in many neo-classical, economic models, we instantiate each member of the class, thereby infusing the possibility of heterogeneity within the class (Guerrero and Axtell 2011). Instead of applying the rates α and μ to the trader classes, each individual trader is assigned a time between trades, T_{Δ} drawn from a Poisson distribution with mean α , for patient traders, T_{Δ} drawn from a Poisson distribution with mean μ for impatient traders. Our stylized facts were mostly unaffected by this change, and we actually saw improvement in the pass rates of some design points due to random effects. We will refer to patient traders as short-term liquidity providers, and impatient traders as short-term liquidity demanders throughout the rest of the paper.

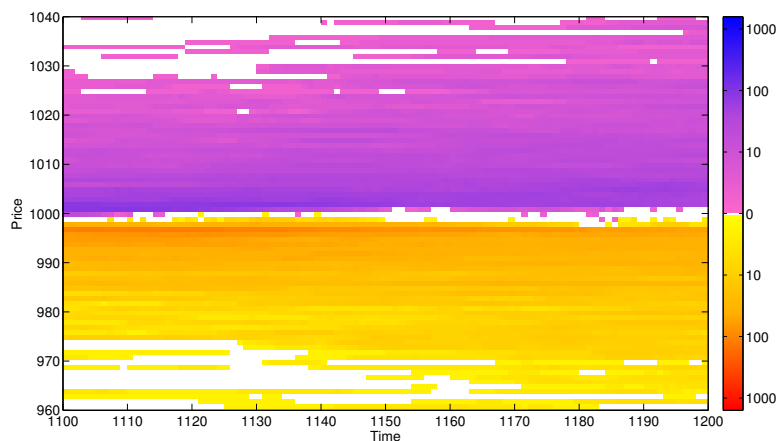
Heterogeneity in Order Rates and Sizes

We next let liquidity providers place orders of size σ , while liquidity demanders continue to place orders of size 1. We then varied σ over five values: 1, 5, 10, 25, and 50. We use these parameter values to show our consistency with the base Preis model (i.e., $\sigma = 1$) and our meaningful divergence from the Preis model. To keep liquidity stationary, we adjusted the T_{Δ} of liquidity providers so that if a trader placed an order of size σ , they would trade $\frac{1}{\sigma}$ as often. To fill orders which are heterogeneous in size, we follow conventional, electronic order book queuing. If the size of the market sell (buy) order is larger than the best priced limit buy (sell) order, the next highest (lowest) priced limit order with the earliest time-stamp is filled, and this process continues until the entire market order is filled. If only a piece of a limit order is executed, the rest of the limit order remains in the order book with its original time-stamp until more market orders arrive and the entire limit order is filled.


Exogenous Shocks

We next introduced a shock in the market price as typically sets off a cascade, thus initiating a market event, and studied effects of this shock as σ varied. As shown in Figure 3, a shock occurs when a large number of market sell orders (or a few large orders) are placed in succession over a short time. We implement a shock as follows: the total number of limit buy shares is summed. The shock size is then calculated as a percentage of the total buy side liquidity. A market sell order of this size is placed at time t_{shock} , and q_{taker} is subsequently set to 0.47 and the mean reverting, random walk to 0.5 begins anew. Note that such an exogenous manipulation of q_{taker} serves as an abstraction for volume-adjusted measure of order flow toxicity, a measure developed by Easley et al. (2011).

(a)



(b)



(c)




Fig. 3: Examples of order book density plots. (a) is an instance of our replication of the Preis model. (b) is an example of a market shock, where the price drops and recovers rather quickly, indicating that an influx of sell market orders occurred from time 520 to 620, and a greater number of buy market orders occurred from time step 620- 700. (c) is a limit order book density plot for a $\sigma = 50$ experiment of our finalized model. Note the two dark red buy orders near the spread at timesteps 585 and 660. Since they are close to the spread, they are quickly filled. Contrast this when the dark red buy order placed further from the spread at timestep 650, which stays in the book for over 20 timesteps and is eventually cancelled.

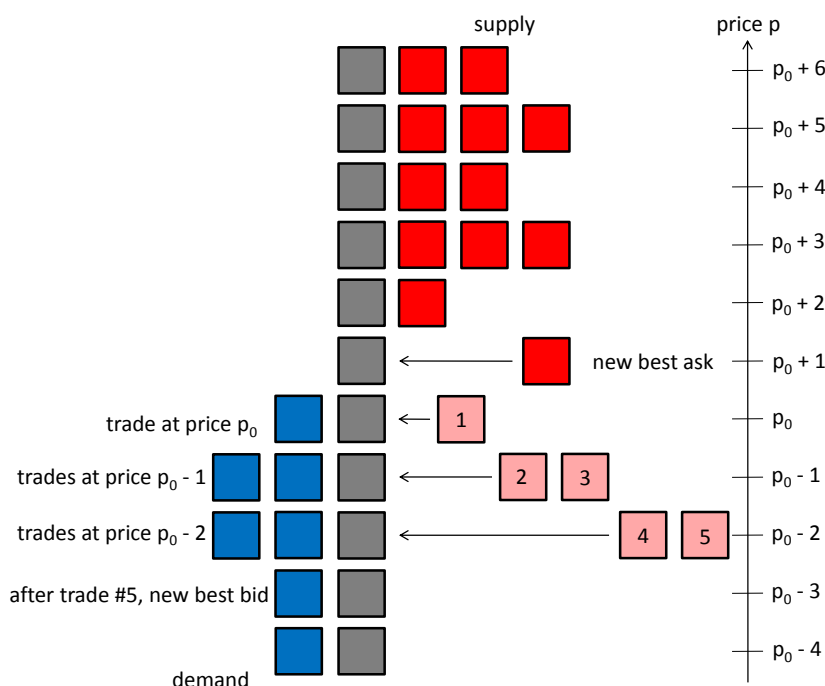


Fig. 4: Example of a limit order book with a shock. When multiple market sell orders (light red blocks) are placed at the same time, the buy limit orders closest to the spread are each filled, causing the price to quickly drop.

Firms

We then introduced a third trader class, the firms (or long-term liquidity provider). Firms place orders of size σ (with T_Δ increased by a factor of σ), whereas the short-term liquidity providers and short-term liquidity demanders place orders of size 1, and trade as they would in the base Preis model. We introduce N firms at this stage, and reduce the number of short-term liquidity providers by N so that liquidity maintains its stationarity. See Figure 3c for a density plot of the order book depicting the presence of firms in the market. Once again stylized facts results were upheld, and we see minimum price results which are comparable to our previous step.

Figure 3c gives insight into the effect of adding the heterogeneity of order sizes and decisions cycles. In the order book density plot of Preis (Figure 3a), a modest amount of liquidity accumulates at each edge of the spread, whereas Fig 3c illustrates a significant accumulation of liquidity further away from the spread as firms try to maximize price discovery, as depicted by the contrast between dark red blocks indicating buy orders placed by firms, and yellow-orange and purple-blue blocks indicating buy and sell orders placed by short-term liquidity providers, respectively. The effect of the long-term, liquidity

provider is evident at Time 660 when a large buy order becomes the best bid. With the arrival of the firm's order, the series of negative returns ends, the spread decreases and a series of positive returns begins. We hypothesize that the timing and size of these orders from the long-term, liquidity providers are important to market stability, a point which we will further expand on in Section 4.

Capital Constraints

Lastly, for added realism we impose capital constraints on all trader classes. Our market has a market capitalization of M , split as follows. 50% of the market capitalization is off-book, and not being actively traded. Initial trader wealth is allocated as follows:

- N firms are given a total of $\$0.125M$ in cash and $\frac{0.125}{p_0}M$ shares; such that each firm receives $\frac{\$0.125M}{N}$ in cash and $\frac{.125M}{p_0N}$ shares.
- 125 short-term providers are given $\$.125M$ in cash ($\$.01M$ each) with $\frac{.125}{p_0M}$ shares ($\frac{.01}{p_0M}$ each).
- 250 short-term demanders are allocated $\$.25M$ in cash ($\$.01M$ each) with $\frac{.25}{p_0M}$ shares ($\frac{.01}{p_0M}$ each).

With capital constraints, the firms trade larger limit orders in the following, price-sensitive manner: Each firm maintains a perception of the fundamental value of the traded asset, p^* , drawn from a normal distribution with mean $p^*_0 = 1,000.00$ and standard deviation of 5. The size of the order of firm i at time t is

$$O_i(t) = \begin{cases} \text{Buy} : \frac{\theta\sigma(p_t - p^*)}{100p_t} c_i(t) & : (p_t - p^*) < 0 \text{ and } |p_t - p^*| \leq 100 \\ \text{Buy} : \frac{\theta\sigma}{p_t} c_i(t) & : (p_t - p^*) < 0 \text{ and } |p_t - p^*| > 100 \\ \text{Sell} : \frac{\theta\sigma(p_t - p^*)}{100} s_i(t) & : (p_t - p^*) > 0 \text{ and } |p_t - p^*| \leq 100 \\ \text{Sell} : \theta_i\sigma s_i(t) & : (p_t - p^*) > 0 \text{ and } |p_t - p^*| > 100 \end{cases} \quad (3)$$

Where $\theta \in [0, 1]$ is the aggression of the trader (Thurner et al. 2013) and σ is a parameter which ties the firm's arrival rate to the size of the firm's orders. For $\sigma = 2$, firms place orders of double the size half as often as they would if $\sigma = 1$. Length between trades for firms is taken from a Poisson distribution with mean $T_\Delta = \frac{\sigma}{\alpha}$. Firms also specify the price of their limit orders according to equation (1), described above.

The sizes of the buy orders for the short term traders (both liquidity providers and demanders) are now determined by $O_i(t) = \theta_i * c_i(t)$, and $c_i(t)$ is the cash of trader i at time t . Sell orders from these traders are of size $O_i(t) = \theta * s_i(t)$, where $s_i(t)$ is shares on hand (not currently being traded in the order book) for trader i at time t .




Fig. 5: A limit order book snapshot showing the introduction of an order placed by a firm.

Stylized Facts Results for Final Version of the Model

Due to the possibility for qualitative differences in market dynamics resulting from the introduction of capital constraints, we reran the original 204 design points using the same DOE as before. Using the same test of the stylized facts, we found 18 design points with a pass-rate of 90% or better. However, by only slightly relaxing our stylized facts constraints and accepting runs with kurtosis greater than 4.5 (rather than 5), a clustering coefficient greater than 4 (rather than 5), and using the same autocorrelation test, 36 design points pass the stylized facts with a rate of 90% or better. In fact, 29 design points passed the stylized facts in more than 95% of the replicates. Due to the fact that these relaxed standards are still qualitatively consistent with Cont’s thresholds for the stylized facts, we ran the experiments which follow on the 36 design points with a pass rate of at least 90%. Figure 6 illustrates just how much of an effect our initial (standard) stylized facts test had on the pass rates of the design points. Figure 7 shows an example of the stylized facts in an instance of our final model.

	Pass Rate		
	80%	90%	95%
Relaxed	48	36	29
Standard	32	18	13

Fig. 6: The number of design points which passed our “relaxed” stylized facts test by percentage of replicates compared to the number of design points which passed our standard stylized facts test. For example, 18 design points had no collapses, kurtosis ≥ 5 , a clustering coefficient ≥ 4 , and no significant autocorrelations of returns in at least 90% of the randomly seeded replicates.

3.3 Outline for Final Version of the Model

We have built up the model step-by-step in order to illuminate the important phases of its development. For the final model, on which we run controlled experiments and analyze in Section 4, the following simulation procedure and parameter settings are used.

Final Simulation Procedure

- At setup (t_0), traders are endowed with the capital and rules described above. The initial price is set to p_0 . Each trader is also assigned a random initial trading time which is drawn from a uniform distribution in $[1, T_\Delta]$.
- For the first 30 time steps, only liquidity providers are active, placing limit orders in the order book to ensure that the order book has an adequate amount of liquidity for regular market behavior to occur. After this initialization period, liquidity demanders become active and can begin to place market orders.
- At each time step t , the set of agents is randomly permuted, and they act in this sequence. Only traders which last traded T_Δ time steps ago are given the opportunity to trade, based on the trader’s class rules. If trader i chooses to buy, they first attempt to buy as many shares as allowed by the equations described above, rounded down to an integer number of shares. The same behavior happens if the trader chooses to sell. At the end of the time step, a number of shares $= \delta * L_{Total}$ are cancelled, where δ is the order cancellation rate, and L_{Total} is the total number of shares available in the book as liquidity. This is done by randomly cancelling orders until this number is met.
- When firms enter the market, they determine the mispricing by looking at their view of price and where the current price p_t is in regards to that. They then determine an order size with equation 3 described above.
- Wealth and shares for all traders must stay greater than or equal to 0.

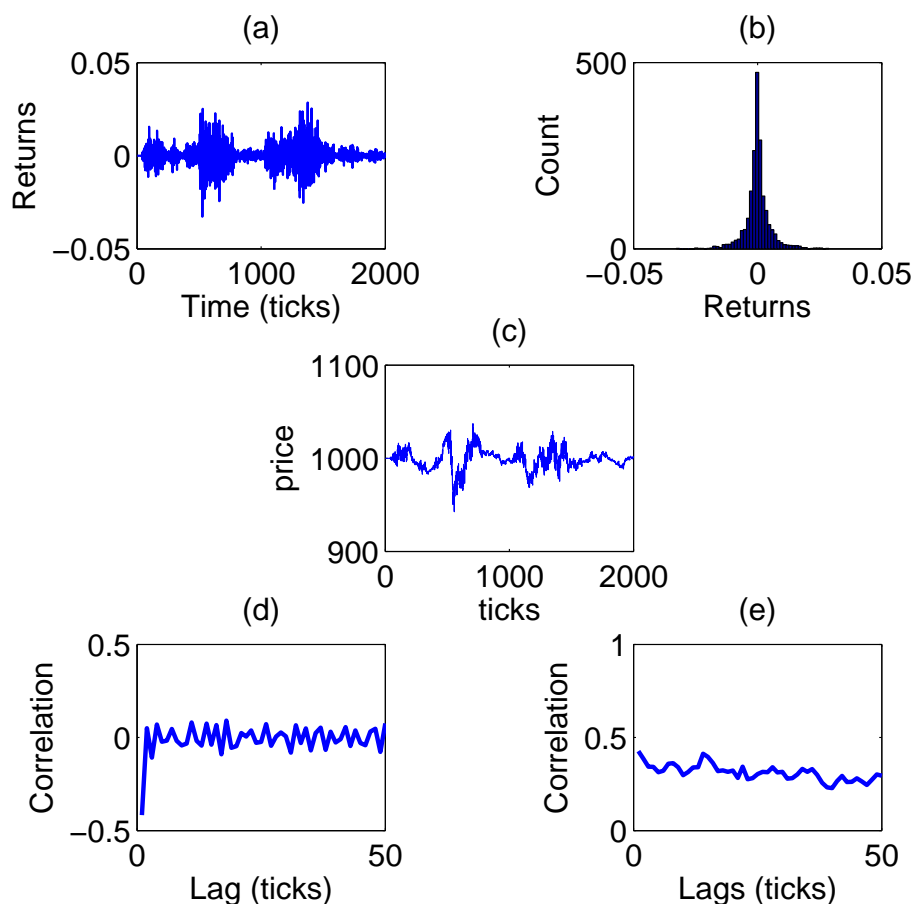


Fig. 7: An example of a stylized facts from the Design Point 7 of our final model version. (a) shows how returns tend to cluster in size over time (clustering volatility), while (e) displays the autocorrelations of absolute returns over small lags. (b) shows the distribution of returns over the entire run. (c) is a price series, and (d) shows autocorrelations of raw returns, which tend to decay quickly and stay near 0.

Model Parameter Specifications

The market has a total market capitalization of \$100 million dollars. The capital is split between the traders as follows: Each liquidity demander receives \$100,000 and 100 shares. There are 250 demanders, so as a trader class, they own \$25 million and 25,000 shares. There are 5 firms, each of which owns \$2.5 million and 2500 shares, for a class total of \$12.5 million and 12,500 shares. For the remaining liquidity providers, each of them receives \$100,000 and 100 shares. There are 125 liquidity providers, so the trader class owns \$12.5 million and 12,500 shares. Half of the market capital, \$50 million, is held by traders

who do not regularly trade, therefore these traders are not explicitly reflected in our model. Theta is set to 0.25 for all traders.

4 Analysis

Experiments

In the following section, we describe our controlled experiments to investigate the impact of heterogeneous decision cycles on market stability. In the first experiment, we vary σ to investigate the impact of heterogeneous order sizes and rates, which are attributes of classes of traders (i.e., short-term liquidity takers, short-term liquidity providers and long-term liquidity providers). In the second experiment, we vary the firms' perception of the asset's fundamental value, which may be an attribute of the individual firm. In our last experiment, we vary σ as well as introduce an exogenous, liquidity shock to investigate the effect of heterogeneous decision cycles during market crises. Together, these experiments demonstrate the ability to investigate dynamics at three levels: individual traders, classes of traders, and the market level where liquidity accumulates and prices emerge.

All of the experiments which are analyzed in this section are tested on the sample of 36 design points described above, each of which passed the relaxed stylized facts in at least 90% of the replicates.

As before, we first varied the value of σ , across five values: 1, 5, 10, 25, and 50. Figure 8 illustrates the stability of the stylized facts across various values of σ , where pass rate is calculated according to our stricter standard test.

We focused on the average minimum price over a set of simulations in a given experiment. That is, for each run of the experiment, we looked at the lowest price after the price shock occurs, i.e., the trough of the resulting market event, and averaged this across all of the runs. Depicting experiments for two design points with results representative of the remaining design points, Figure 9 shows σ has large effects on the average minimum price; the average minimum price is lower as we increase σ , though its effects start to dampen as it reaches values of 25 and 50. It is also apparent that in the worst-case scenarios, the value of σ has major effects on minimum price, but when the median of the minimum price is taken across all random seeds, the effects are almost negligible. This highlights an important aspect for financial markets, and complex systems in general, which is that in any fat-tailed distribution, the mean and median are not sufficiently reliable statistics to give accurate insight into the behavior of a system. From Figure 8, we note that the average kurtosis does not increase significantly as σ increases, even though the range of minimum prices does increase. This is significant because as we have hypothesized, the effect of σ on average minimum price is not due to qualitatively different sizes of individual negative returns (i.e., larger, individual price drops) but due to the inability of short-term liquidity providers to meet the immediate needs of short-term liquidity takers between the arrival of long-term, liquidity providers.




Fig. 8: Each plot is a histogram of data taken from 36 design points. (a), (b), and (c) plots are stylized facts (determined by the standard test) data for $\sigma = 1$. (d), (e), and (f) are data from $\sigma = 10$. (g), (h), and (i) are taken from $\sigma = 50$. The first column depicts the overall average pass rate across all replicates. The second column shows average kurtosis, and the third column shows the distribution of clustering volatility.

The nonlinear effects of σ on minimum price as σ increases seem to be due to our cancellation methodology, which we derived from both Farmer and Preis. Small orders are more likely to be cancelled than larger orders. Even if a firm's order were randomly selected to be cancelled, if the cancellation of the firm's order would cause the aggregation of cancelled shares to exceed the cancellation level determined by the order cancellation rate, then the firm's order would not be cancelled and another order would be randomly selected. An area for future research and a potentially meaningful extension to the model would be a behavioral rule for order cancellation which is grounded in empirical research, hence the first demonstration of Level 3 Empirical Relevance in the Axtell and Epstein (1994) framework for an agent-based model of an order book.

For the next experiment, we tested the impact of firms' perception of fundamental value on price. We considered three implementations of firms' fundamental view of price: Constant and homogeneous, constant and heterogeneous

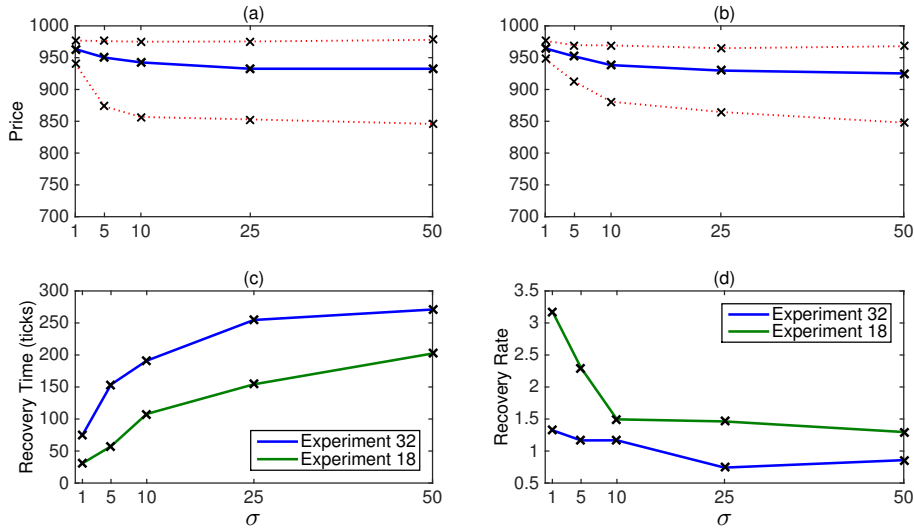


Fig. 9: (a) and (b) are minimum price plots for Experiment 18 and 32 respectively. The blue line represents the median minimum price across all replicates for each of our σ experiments. The red lines represent the 5th and 95th percentile minimum prices. (c) Displays recovery time in both experiments as a function of σ . (d) Shows how the rate of recovery, calculated as $\frac{\text{Price Change}}{\text{Recovery Time}}$, where Recovery Time is how long it takes the price to swing back to p_{200} , which we determined was the end of the transient period. The price change is calculated as $p_{200} - p_{Min}$.

(mean p_0 , standard deviation 5), and a heterogeneous random walk (starting at p_0 , step size 0.5). We do not see qualitatively different results across these three design points, which suggests that our results are robust to the firms' independent views of fundamental price. Thus, for the sake of simplicity, we used the constant and heterogeneous fundamental prices for our remaining DOE.

Lastly, we re-ran the σ experiment with a downward price shock to test how each market would respond (i.e., a stress test). Here we see similar, yet more pronounced results than the ones described above. Interestingly, we see less of a 'dampening' effect as σ increases to 25 and 50 than we do in the un-shocked cases. Although the absolute value of the slope does decrease in both experiments, it decreases more slowly in the shocked case, lending more support for our hypothesis that markets are especially time-sensitive during fire-sale scenarios, and sometimes the only traders with sufficient capital to stabilize the market are the traders who have the longest deliberation times.

For these shocked experiments, we analyzed the effect of σ on market collapses, as depicted in Fig 11. The rate of collapses increases monotonically up to $\sigma = 25$, and then decreases at $\sigma = 50$. This positive correlation between σ




Fig. 10: Here we see the mean over all design points of the 5th percentile of minimum prices across all random seeds. The blue line represents unshocked cases, while the red line represents model runs with a 15% liquidity shock.

and collapse rate provides further evidence that in the market crisis scenarios, the long-term, liquidity providers were simply not responsive enough to provide sufficient liquidity to stabilize the market.

Overall, this set of results comports with our hypothesis that the timeliness of large firms is very important to market stability.

To summarize, we first ran experiments to vary the size of σ and found that increasing levels of σ contribute to lower prices in a given market, longer times for prices to recover and at slower rates. We then tested for potential biases to our results from particular representations of firms' perception of fundamental value, but found no such biases. Lastly, we demonstrated the model as a stress testing platform by introducing a liquidity shock and analyzing the effect of σ on minimum prices and collapses. Again, we found evidence to support our hypothesis that markets are especially time-sensitive during crises, and sometimes the only traders with sufficient capital to stabilize the market are the traders who have the longest decision cycles.

5 Conclusion

To explore the dynamics of heterogeneous decision cycles and their effect on liquidity and subsequently prices, we develop a robust and valid model by replicating and extending a leading market model while synthesizing with other leading models. During the development and testing, we demonstrated that




Fig. 11: Rate of collapse across all design points and all replicates. Experiments with $\sigma = 1, 5, 10$ have 200 replicates, while experiments with $\sigma = 25, 50$ have 1000 replicates to better control random effects.

the model can reproduce those stylized facts for realistic price dynamics which are commonly applied to market models (i.e., the levels of leptokurtic returns, decay of autocorrelations over moderate time lags, and clustering of volatility). While not yet accepted as a "stylized fact," we incorporated an aspect of order flow to reflect recent empirical findings on order flow toxicity. We demonstrated the robustness of the model results by comprehensively sampling the relevant region of the parameter space while preserving the reproduction of the stylized facts. We applied the model to illuminate interdependencies between liquidity and trader decision cycles and their effect on market stability; and lastly, we demonstrated that the model can be used to run controlled policy experiments (e.g., stress tests).

The model also gives a realistic demand dynamic for periods of market dislocation, highlighting how a market shock can lead to a further downward market cascade, and even lead to a market collapse. We show that heterogeneity in the investment decision cycle can extend the extremes for market shocks, leading to an accentuation of these fat-tailed events. From a policy standpoint, only models that have met stringent tests against the stylized facts should be used to inform policy decisions. The demonstration that larger degrees of heterogeneity increase the average severity of a market event gives support to the use of circuit breakers, regulator-imposed pauses in trading, because these circuit breakers in effect reduce the speed of the short-term liquidity demanders, thereby reducing the heterogeneity of the decision cycles.

In conclusion, even without the impact of adverse market conditions, the demand for immediacy might overwhelm the supply that arrives at a slower

pace. And if that pace of the liquidity supply is governed by institutional characteristics such as longer decision-making or trade approval process, then the supply might not be quickly altered through price changes. If the liquidity demanders act unaware of these characteristics, dropping prices in order to elicit supply, and then, not see enough supply immediately forthcoming, dropping prices further, then the price will not be an effective signal to the liquidity suppliers.

Acknowledgements For their helpful comments and feedback, we wish to thank Lashon Booker, Chris Danforth, Peter Dodds, Darrell Duffie, Doyne Farmer, Neil Johnson, Matthew Koehler, Adam McLeod, Matthew McMahon, Mark Paddrik, Tobias Preis and the anonymous reviewers in the Office of Financial Research reviewing process. All remaining errors are the sole responsibility of the authors.

References

1. Axtell, R., Axelrod, R., Epstein, J., and Cohen, M.D. "Aligning Simulation Models: A Case Study and Results," *Computational and Mathematical Organization Theory*, 1 (2), pp: 123-141 (1996).
2. Axtell, R. L. and J. M. Epstein "Agent-Based Models: Understanding Our Creations", *Bulletin of the Santa Fe Institute*. (1994)
3. Balci, O. "Verification, Validation and Accreditation of Simulation Models", in *Proceedings of the 1997 Winter Simulation Conference*, editors. S. Andradhottir, K. J. Healy, D. H. Withers, and B. L. Nelson. (1997)
4. Bookstaber R. *A Demon of Our Own Design*, Wiley Publishing (2007).
5. Bookstaber, R. "Using agent-based models for analyzing threats to financial stability," Working Paper No. 3, U.S. Office of Financial Research. (2012)
6. Cont, R. (2001), "Empirical Properties Of Asset Returns: Stylized Facts And Statistical Issues," *Quantitative Finance*, 1, pp: 223-236.
7. Darley, V., A. Outkin, T. Plate, and F. Gao. "Learning, evolution and tick size effects in a simulation of the NASDAQ stock market," presented at *Proceedings of the 5th World Multi-Conference on Systemics, Cybernetics and Informatics*, Orlando, FL. (2001)
8. Darrell Duffie, "Presidential Address: Asset Price Dynamics with Slow-Moving Capital", *Journal of Finance*, 65:4, August 2010., pp. 1237 - 1267.
9. Easley, David, MM Lopez De Prado, and Maureen O'Hara. "The microstructure of the flash crash: Flow toxicity, liquidity crashes and the probability of informed trading." *Journal of Portfolio Management* 37.2 (2011): 118-128.
10. Farmer, J.D., P. Patelli, and I. I. Zovko, "The Predictive Power of Zero Intelligence in Financial Markets," *Proceedings of the National Academy of Science*, vol. 102, pp. 2254-2259, (2005).
11. Ghoulmie, F., R. Cont, and J-P Nadal. "Heterogeneity and Feedback in an Agent-based Market Model," *Journal of Physics: Condensed Matter*, vol. 17, pp. 1259-1268. (2005)

12. Guerrero, Omar A. and Robert Axtell. 2011. Using Agentization for Exploring Firm and Labor Dynamics. *Emergent Results of Artificial Economies: Lecture Notes in Economics and Mathematical Systems* 652(4):139-150.
13. Holland, John H., and Judith S. Reitman. "Cognitive systems based on adaptive algorithms." *ACM SIGART Bulletin* 63: 49. (1977)
14. Iori, G., M. G. Daniels, J. D. Farmer, L. Gillemot, S. Krishnamurthy, and E. Smith, "An Analysis of Price Impact Function in Order-Driven Markets," *Physica A*, vol. 324, pp. 146-151, (2003)
15. Kirilenko, Andrei A. and Kyle, Albert S. and Samadi, Mehrdad and Tuzun, Tugkan, *The Flash Crash: The Impact of High Frequency Trading on an Electronic Market* (May 5, 2014). Available at SSRN:<http://ssrn.com/abstract=1686004> or <http://dx.doi.org/10.2139/ssrn.1686004>
16. Jack P. C. Kleijnen, Susan M. Sanchez, Thomas W. Lucas, Thomas M. Cioppa. "A User's Guide to the Brave New World of Designing Simulation Experiments," *INFORMS Journal on Computing*, 17 (3), pp. 263-289. (2005)
17. LeBaron, B. "A Builder's Guide To Agent-Based Financial Markets," *Quantitative Finance*, 1, pp:254261 (2001a)
18. LeBaron, B. "Evolution of Time Horizons in an Agent-Based Stock Market," *Macroeconomic Dynamics*, 5, pp. 225-254. (2001b)
19. Lux, T. and M. Marchesi. "Scaling and Criticality in a Stochastic, Multi-Agent Model of a Financial Market." *Nature* 397: 498-500. (1999)
20. Mandelbrot, B. "Variation of Certain Speculative Prices," *Journal of Business*, 36:4, pp: 294-419. (1963)
21. Maslov, S. "Simple Model Of A Limit Order-Driven Market," *Physica A*, 278, pp: 571-578. (2000)
22. Paddrik, Mark E, Haynes, Richard, Todd, Andrew, Scherer, William and Beling, Peter. "Visualizations for Financial Market Regulation." <http://dx.doi.org/10.2139/ssrn.2252598> (2013)
23. Palmer, R.G., W. Brian Arthur, John H. Holland, Blake LeBaron, Paul Talyer. "Artificial Economic Life: A Simple Model Of A Stock Market," *Physica D*, 75, pp: 264-274. (1994)
24. Preis, T., S. Golke, W. Paul and J. J. Schneider. "Multi-AgentBased Order Book Model Of Financial Markets," *Europhysics Letters*, 75 (3), pp. 510516. (2006)
25. Preis, T., Sebastian Golke, Wolfgang Paul, and Johannes J. Schneider. "Statistical Analysis Of Financial Returns For A Multi-agent Order Book Model Of Asset Trading," *Physical Review E* 76, 016108. (2007)
26. Presidential Task Force on Market Mechanisms: Report of the Presidential Task Force on Market Mechanisms, "The Brady Report", Nicholas F. Brady, (Chairman), U.S. Government Printing Office. (1988)
27. Thurner, S., J. Doynne Farmer and John Geanakoplos. "Leverage causes fat tails and clustered volatility," *Quantitative Finance*, 12:5, 695-707 (2012)

28. Tivnan, B., Matthew Koehler, Matthew McMahon, Matthew Olson, Neal Rothleder, Rajani Shenoy. "Adding to the Regulator's Toolbox: Integration and Extension of Two Leading Market Models," ArXiv e-prints, arxiv.org/abs/1105.5439v1. (2011)
29. U.S. Commodity Futures Trading Commission and U.S. Securities & Exchange Commission, Findings Regarding The Market Events Of May 6, 2010, Report Of The Staffs Of The CFTC and SEC To The Joint Advisory Committee On Emerging Regulatory Issues, September 30, 2010, <http://www.sec.gov/news/studies/2010/marketevents-report.pdf>. (2010)