Computational Intelligence, Volume 23, Number 2, 2007

FROM DISCRETE-TIME MODELS TO CONTINUOUS-TIME, ASYNCHRONOUS MODELING OF FINANCIAL MARKETS

KATALIN BOER, UZAY KAYMAK, AND JAAP SPIERING

Erasmus School of Economics, Econometric Institute, Rotterdam, The Netherlands

Most agent-based simulation models of financial markets are discrete-time in nature. In this paper, we investigate to what degree such models are extensible to continuous-time, asynchronous modeling of financial markets. We study the behavior of a learning market maker in a market with information asymmetry, and investigate the difference caused in the market dynamics between the discrete-time simulation and continuous-time, asynchronous simulation. We show that the characteristics of the market prices are different in the two cases, and observe that additional information is being revealed in the continuous-time, asynchronous models, which can be acted upon by the agents in such models. Because most financial markets are continuous and asynchronous in nature, our results indicate that explicit consideration of this fundamental characteristic of financial markets cannot be ignored in their agent-based modeling.

Key words: artificial stock markets, agent-based computational economics, continuous-time simulation, information asymmetry, market maker.

1. INTRODUCTION

Within agent-based computational economics (ACE), artificial stock markets (ASM) are studied extensively to assess how global regularities arise from individual interactions of market participants (Tesfatsion 2001). Usually, individuals are represented by (software) agents interacting in an artificial environment. By using agents for studying market dynamics, heterogeneous, boundedly rational, and adaptive behavior of market participants can be represented and its impact on market dynamics can be assessed.

Most agent-based simulation models of financial markets are discrete-time in nature. At a high level of abstraction, the discrete-time models of financial markets are turn-based games, in which the market participants take turns to execute their actions. The specific characteristics of the simulation models determine how the players take turns and how they arrive at their decisions. Because these models are turn-based, there are explicit mechanisms in the simulation to ensure that the players coordinate their actions by waiting for their turn. Such players remain passive in the market until it is their turn (i.e., they are selected) to take decisions and perform actions.

By far the most common models in the literature represent a call market, in which orders are aggregated at discrete points in time and market price is set at equilibrium (Boer, de Bruin, and Kaymak 2005a). At each round of the game, the investors are asked to submit their orders, and the market price is determined by aggregating supply and demand (e.g., Brock and Hommes 1997; Marsili 2001; Raberto et al. 2001; LeBaron 2002; Challet, Marsili, and Zhang 2005). Although call markets are encountered regularly, continuous trading sessions and quote-driven execution mechanisms are more common in real financial markets (Demarchi and Foucault 2000; Harris 2003). This has prompted some groups to look beyond the call market structure. Examples of such models have been described in (Loistl and Vetter 2000; Chan and Shelton 2001; Smith et al. 2002; Shatner et al. 2000). These markets try to implement continuous order matching and even asynchronous decision making in a turn-based simulation model. A continuous trading model is proposed in Shatner et al. (2000), where traders "sleep" after actions, and they "wake up" at predefined times, or as a result of certain events. Other studies (Loistl and Vetter 2000; Raberto et al. 2001) model continuous sessions through discrete-time simulation, in which the trader whose decision is carried out during the next trading session is centrally selected. In these models new orders are automatically matched with pending ones if possible.

The disadvantage of selecting agents centrally and randomly, such that their actions are carried out, is that this turn-based execution mechanism assumes implicitly that the market participants coordinate their actions, in the sense that they wait for their turn. In reality, however, traders make decisions asynchronously, i.e., each investor may be carrying out a different task at the same point in time. This property of asynchronous behavior is not captured in turn-based models. The developments in agent technology provide mechanisms for going beyond the turn-based models, by implementing ASM in agent-based development environments. One example of such an ASM is the ABSTRACTE framework proposed in Boer et al. (2004), which implements a continuous, asynchronous agent-based ASM on top of JADE (java agent development environment) (Bellifemine et al. 2005). Market participants with various roles such as investors, brokers, and market makers are implemented in ABSTRACTE as autonomous agents using the generic agent features offered by JADE. The behavior model of JADE provides for the execution of concurrent (asynchronous) agent activities. Continuity is modeled by concurrent execution of agent actions (in Java threads), which interact by asynchronous message passing.

There are some important differences in the nature of available information in turn-based models and continuous, asynchronous agent-based models like ABSTRACTE. On the one hand, a lot of information is available to the agents in turn-based models, because each agent can observe the consequences of the previous decisions (e.g., prices that have been formed as a result of other agents' trading decisions). In continuous, asynchronous models, there is uncertainty regarding this information, because the agents take decisions based on available information at some point in time, but the market state may change between the placement of an order and its execution. On the other hand, additional information might be revealed in the continuous, asynchronous models due to the race conditions, which is not available in turn-based models. For example, a sudden increase in the number of entries in the order book might entail actionable information, while it is not available in turn-based models.

Given the difference in the nature of information between the turn-based models and continuous, asynchronous models, an important question is to what degree the models developed in turn-based simulations are extensible to continuous, asynchronous simulations. Because most financial markets are continuous with asynchronously interacting traders, while the agent-based models are often turn-based, this is an important question to address to assess the limitations of the current modeling practice. In this paper, we consider this question by studying the behavior of a learning market maker in a market with information asymmetry. We study the characteristics of the market prices that arise in continuous, asynchronous simulations, and compare it to the characteristics of the prices in the turn-based models. Furthermore, we consider what additional considerations are needed to extend the turn-based model into the continuous, asynchronous model.

We consider an information-based model, because they provide insights into the adjustment process of prices that we are interested in (O'Hara 2002). We study the behavior of a market maker to focus the scope of the investigation. The research presented in this paper is based on the learning market maker from Das (2005). This model extends the Glosten and Milgrom (1985) information-based model, which was proposed to show the influence of informational asymmetry on the bid-ask spread in financial markets. In this model, the market maker tries to discover the fundamental value of a stock by means of Bayesian learning. He determines the bid and ask quotes based on his expectation of the real value, the order flow, and his prior knowledge regarding the ratio of informed and uninformed traders. In Das (2003, 2005) a nonparametric density estimation technique is proposed for maintaining a probability distribution over the true value that the market-maker can use to set prices. Discrete-time simulation is applied in the model, and a probabilistic representation of order flows is considered. In this paper, we implement the model of Das (2005) in ABSTRACTE for continuous simulation with individual investors interacting asynchronously, and we report our results.

The outline of the paper is as follows. Section 2 describes the ASM and the learning market maker model that we consider in this paper. Most of this section is based on the learning market maker model of Das (2005). In Section 3, we discuss how the market maker model of the previous section can be applied in a continuous, asynchronous simulation. We mention the additional parameters whose values must be determined, and describe the general modeling approach in the continuous, asynchronous setting. Our experimental design and the results we have obtained are described in Section 4. Here the successfulness of the price discovery mechanism, the bid ask spread and the market maker's wealth are studied. A discussion regarding the implications of continuous, asynchronous representation is given in Section 5, while Section 6 presents the conclusions of the paper.

2. THE ARTIFICIAL STOCK MARKET

The organization of the artificial market that we use to study market dynamics is based on an extended version of the information-based Glosten and Milgrom model proposed in Das (2005). We combine the learning market maker, which Das describes in a turn-based model, with investors that interact asynchronously and autonomously. We implement our model on top of ABSTRACTE (Boer, de Bruin, and Kaymak 2005b) that applies continuous-time simulation instead of discrete-time simulation. In this section, we describe the characteristics of the represented market.

2.1. The Marketplace

In the studied market model trading sessions are continuous and the execution system is quote-driven. There is one stock traded. One market maker and multiple investors are represented. Investors place market orders for one single share. The market maker is responsible for the liquidity of the stocks and the execution of orders. The market price is thus formed at the bid or ask quote of the market maker, depending on whether a sell or a buy order is matched against it. The stock does not pay dividends. It is assumed that the stock has an underlying *fundamental value*, which is generated exogenously to the market.

2.2. The Fundamental Value

The underlying fundamental value of the stock at time t is V_t. This value is exogenous to the model and follows a jump process. In other words, V_t is constant most of the time, and its value changes occasionally. This process can be thought of as a simple representation of the value of a stock during intraday trading: its value does not change for some periods of time and is modified when news arrives regarding the traded stock. The jump process is modeled as a random process following $V_{t+1} = V_t + \tilde{\omega}(0, \sigma)$, where $\tilde{\omega}(0, \sigma)$ represents a sample from a normal distribution with mean zero and variance σ^2. In Das (2005) a jump occurs with some probability at every trading period, that is at every discrete point in time. In our model, continuous-time simulation is applied, and accordingly, a jump in the fundamental value will occur at randomly drawn times, as explained in Section 3.

2.3. The Investors

Investors are differentiated based on the information they receive regarding the fundamental value. There are two types of investors: informed traders and uninformed traders.

Depending on the simulation parameters, the informed traders may be *perfectly informed* or *noisily informed*. Perfectly informed traders observe the correct fundamental value (V_t), while noisily informed investors observe a distorted fundamental value $W_t = V_t + \tilde{\psi}(0, \sigma_W)$. Here, $\tilde{\psi}(0, \sigma_W)$ represents a sample from a normal distribution with mean zero and variance σ_W^2. Finally, uninformed traders do not know what the underlying fundamental value is, and they trade randomly.

Informed traders decide whether to trade or not, based on their observation of the fundamental value. An informed trader will buy if the fundamental value that he observes is higher than the market maker's ask price, i.e., if $V_t > A_t$ in the case of perfectly informed traders, and $W_t > A_t$ in the case of noisily informed traders. He will sell if the fundamental value that he observes is below the bid price, i.e., if $V_t < B_t$ or $W_t < B_t$. He will place no order if the observed fundamental value is within the bid-ask spread, i.e., $B_t \le V_t \le A_t$ or $B_t \le W_t \le A_t$). Uninformed traders place buy and sell orders with equal probability (η). They can also decide not to place orders for a while with probability $1 - 2\eta$.

2.4. The Learning Market-Maker

The market maker aims to set bid/ask prices to capture the underlying fundamental value of the stock unknown to him. To discover the fundamental value, Bayesian learning is applied, based on the extended Glosten and Milgrom model (Das 2005).

The Glosten and Milgram (1985) model was proposed to show the influence of informational asymmetry on the bid-ask spread. The market maker sets the bid price to the expectation of the true value given that a sell order has arrived, and the ask price to the expectation of the true value given that a buy order has arrived. To compute the expectations in the extended model in Das (2005), the market maker keeps a probability density estimate over a whole range of possible values for the stock, which is updated after the arrival of an order. The bid and ask prices will in turn be based on the expected value given this estimate. The market maker tries to learn in this way the fundamental value known by informed investors.

The market-maker knows the following information:

- the fraction of informed traders (α) and uninformed traders ($1 - \alpha$) in the market;
- the probability for an uninformed trader to trade (η);
- the initial fundamental value V_0;
- the distribution function of the jump process ($\tilde{\omega}(0, \sigma)$);
- whether a change in the fundamental value has occurred;
- the distribution function of the noise process ($\tilde{\psi}(0, \sigma_W)$).

The market maker carries out the following tasks:

1. receive and execute orders;
2. update the probability density estimate based on the received orders;
3. calculate the expected value of the stock based on the updated probability values;
4. adjust the bid and ask quotes according to the changes in the expected value.

Steps 2 to 4 are also carried out if no orders are placed.

In this market, information regarding the fundamental value of the stock diffuses from the informed traders to the market maker. The private information is contained in the submitted orders from the traders. A series of sell orders might indicate that the fundamental value is lower than the current bid price, and a series of buy orders might indicate that the fundamental value is higher than the current ask price. However, the market maker will have to take into account the noise incorporated by the orders of the noisily informed traders, and the noise implied by the orders submitted by the uninformed traders.

2.4.1. Adjusting the Bid and Ask Quotes. The market maker tries to set the bid and ask prices such that these reflect the fundamental value of the stock. He ensures market efficiency in this way, as he tries to incorporate the information into the market price by learning from the orders. The bid (B) and ask (A) quotes are calculated according to the learning algorithm described in Das (2005):

$$B = \frac{1}{P_{\text{Sell}}} \sum_{V_i = V_{\min}}^{V_{\max}} V_i \Pr(\text{Sell} \,|\, V = V_i) \Pr(V = V_i), \tag{1}$$

$$A = \frac{1}{P_{\text{Buy}}} \sum_{V_i = V_{\min}}^{V_{\max}} V_i \Pr(\text{Buy} \,|\, V = V_i) \Pr(V = V_i), \tag{2}$$

where P_{Sell} is the a priori probability of a sell order, and P_{Buy} is the a priori probability of a buy order.

In a market with perfectly informed traders, the probability for a sell order or a buy order depends on the fraction of various traders and the probability they will trade. Accordingly, the market maker bases his estimates on the expectation that (rational) informed traders will always buy if the perceived fundamental value is above the ask price, will always sell if the perceived value is below the bid price, and will not trade otherwise. Then,

$$\Pr(\text{Sell} \,|\, V = V_i; V_i < B) = \Pr(\text{Buy} \,|\, V = V_i; V_i > A) = \alpha + (1 - \alpha)\eta, \tag{3}$$

$$\Pr(\text{Sell} \,|\, V = V_i; V_i \geq B) = \Pr(\text{Buy} \,|\, V = V_i; V_i \leq A) = (1 - \alpha)\eta. \tag{4}$$

In case of models with noisily informed traders, the probabilities for sell and buy orders are determined by the following equations:

$$\Pr(\text{Sell} \,|\, V = V_i, V_i < B) = (1 - \alpha)\eta + \alpha \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) < (B - V_i)\right), \tag{5}$$

$$\Pr(\text{Sell} \,|\, V = V_i, V_i \geq B) = (1 - \alpha)\eta + \alpha \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) \geq (V_i - B)\right), \tag{6}$$

$$\Pr(\text{Buy} \,|\, V = V_i, V_i \leq A) = (1 - \alpha)\eta + \alpha \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) \geq (A - V_i)\right), \tag{7}$$

$$\Pr(\text{Buy} \,|\, V = V_i, V_i > A) = (1 - \alpha)\eta + \alpha \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) < (V_i - A)\right). \tag{8}$$

The second term in equation (5) reflects the probability that a noisily informed trader sells if the fundamental value V is below the current bid price B. This occurs if the observed fundamental value, including the noise, is below the bid price. Similarly, the second term in equation (6) reflects the same probability, under the assumption that the fundamental value V is equal to or greater than the bid price B. That means, that if the noise in the noisily informed trader's observation is greater than the difference between the fundamental value and the bid price, the trader will submit a sell order. Although a perfectly informed trader would not sell in this case, the additional noise can cause a noisily informed trader to make different decisions.

2.4.2. Adjusting the Probability Density Estimate. To solve equations (1) and (2), the market maker needs to know the values $\Pr(V = V_i)$ for the approximation of the probability density estimate. The market maker updates these estimates as trading continues. To do this, the market maker uses an array from $V_{min} = V_0 - 4\sigma$ to $V_{max} = V_0 + 4\sigma - 1$ to contain the prior value probabilities. The values are initialized by setting the ith value in the array to $\int_{-4\sigma+i}^{-4\sigma+i+1} \mathcal{N}(0, \sigma)\, dx$. Here, $\mathcal{N}$ is the normal density function in x with mean zero and standard deviation σ. The array is kept in a normalized state at all times, thus the entire probability mass for V lies within it.

When an order arrives, the market maker updates the probabilities for V_i by scaling the distribution, based on the type of order. The values are updated using Bayes' rule according to

$$\Pr(V = V_i \mid \text{Action}) = \frac{\Pr(\text{Action} \mid V = V_i) * \Pr(V = V_i)}{\Pr(\text{Action})}, \tag{9}$$

where the Action represents a buy order, a sell order, or no order. The prior probability $\Pr(V = V_i)$ is known from the existing probability estimates. The probabilities and the conditional probabilities for buy and sell orders can be computed using the equations introduced in the previous section.

In addition to receiving buy or sell orders, it is also possible that the market maker does not get any orders at some point in time. The prior probability for no order $P_{\text{No order}}$ is equal to $1 - (P_{\text{Sell}} + P_{\text{Buy}})$. The fact that there are no (informed) traders who want to trade, given the current bid and ask prices and the current fundamental value, suggests that the bid and ask prices are currently set around the fundamental value. By adjusting the estimated probabilities, the market maker can make the bid-ask spread smaller, to ensure market liquidity and to encourage trading.

If the market contains perfectly informed traders, the following updates are being made to determine the new probability estimates using equation (9):

$$\Pr(\text{No order} \mid V = V_i, B \leq V_i \leq A) = (1 - \alpha)(1 - 2\eta) + \alpha, \tag{10}$$

$$\begin{aligned}\Pr(\text{No order} \mid V = V_i, V_i < B) &= \Pr(\text{No order} \mid V = V_i, V_i > A)\\ &= 1 - \alpha - (1 - \alpha)(1 - 2\eta).\end{aligned} \tag{11}$$

If the market contains noisily informed traders, the following updates are being made:

$$\Pr(\text{No order} \mid V = V_i, V_i < B) = (1 - \alpha)(1 - 2\eta) + \alpha \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) > (B - V_i)\right), \tag{12}$$

$$\begin{aligned}\Pr(\text{No order} \mid V = V_i, B \leq V_i \leq A) &= (1 - \alpha)(1 - 2\eta)\\ &+ \alpha\Big(\Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) \geq (B - V_i)\right) + \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) \geq (V_i - A)\right)\Big),\end{aligned} \tag{13}$$

$$\Pr(\text{No order} \mid V = V_i, V_i > A) = (1 - \alpha)(1 - 2\eta) + \alpha \Pr\left(\tilde{\psi}\left(0, \sigma_W^2\right) > (V_i - A)\right). \tag{14}$$

After all probabilities have been updated, the array is normalized again. When a jump in the fundamental value occurs the market maker re-centers the estimates around the last expected value.

3. CONTINUOUS, ASYNCHRONOUS IMPLEMENTATION OF THE MARKET MODEL

We have implemented the model described in Section 2 within ABSTRACTE. In addition to the market maker, we have also implemented individual investors, where they interact asynchronously and autonomously, deciding for themselves when to trade, and their decision is taken into account at all times. We emphasize that, in such a setting, the investors are most probably carrying out different tasks, at the same moment. One of the agents, for example, might just listen to news, while another one is analyzing the market, and a third one is waiting for his order to be executed, and all this as a consequence of their autonomous feature and not coordinated by some central system.

The continuous-time implementation of the market and the autonomous, asynchronous implementation of the traders, implies other kinds of behavior of the participants than in discrete-time simulation with some order generating process or centrally selected traders. This feature in turn necessitates a different representation of some of the parameters and the specification of additional parameters for the market and the participants in the studied model. In the following, we consider these differences. In particular, we consider the jump process, the no-order condition, and some time-related factors that need to be taken into account.

The jump process. In discrete-time simulation of the model described in Section 2, the underlying fundamental value of the risky asset is constant for most of the time. Changes (jumps) occur occasionally at various moments. A jump in the fundamental value occurs with some probability at every discrete time step (trading round). To model this process in the continuous setting, we draw the lengths of the periods in which the fundamental value does not change from a uniform distribution in an interval. Jumps occur at the end of these periods.

The no-order condition. In the discrete situation, the "no-order condition" occurs when a buy or sell order is not placed in a trading round. During continuous simulation, however, orders are not placed at fixed times. In this case, the market maker must determine when to update his bid-ask spread in case no orders arrive for some time. At the moment, we have selected a fixed interval whose size is determined after some experimentation.

The number of traders. The individual representation of the investors entails that next to the rate of different types of traders the number for each type has to be specified. While a small number of investors might be not representative, a lot of traders could overload the market maker with orders, if they place orders faster than the market maker can handle them. To avoid this situation, our investors can wait for a small time after their order is executed. The length of the waiting period is currently the same for all the investors.

Time horizon. Investors analyze market conditions, place orders, and wait for news or for the execution of a placed order (see, for example, the typical behavior of an investor in Boer, de Bruin, and Kaymak (2005a, 2005b)). If an investor has no orders placed, because it was not worth for him, and no news arrive for a while, he could reconsider the possibility to trade. In the discrete implementation, investors decide whether to trade or not whenever they are selected to be processed by the central simulation manager. In a continuous setting, a lot of things can happen in the market while the investor is waiting for the execution of his order. Furthermore, after his order is executed, he could analyze the possibility to trade again immediately or could wait for a while before trading again. For how long he waits before examining the market conditions again depends on his time horizon. This setting might represent the different time horizon of the investors, respectively. In our experiments, investors analyze market conditions either when news arrives or when their individual time horizon "expires."

4. EXPERIMENTS AND RESULTS

In this section we describe the experimental setup for our study and report the results obtained. We first consider the experimental settings. Then, we validate our implementation against the original model. Finally, we present some additional results and discuss some implications of the asynchronous, continuous setting.

4.1. Experimental Settings

Most of the settings correspond to the ones from Das (2005) so that we can verify our implementation and compare the results. Some values or processes need to be converted, however, from the discrete-time model into the continuous-time model, as described in Section 3. Furthermore, additional parameters need to be specified related to the individual, autonomous representation of the investors.

4.1.1. Fundamental Value-Related Settings.

- The initial value (V_0) is 7,500 cents in accordance with Das.
- The standard deviation (σ) of the jump process is 50 cents.
- In Das (2005), every trading period a jump in the fundamental value occurs with a small probability (1 in 1,000). In our model, which applies continuous-time simulation, the fundamental value jumps randomly every 5 to 30 seconds in real time.
- The noise process has a mean of 0 and a standard deviation of 0.05 (5 cents).

4.1.2. Investor-Related Settings.

- The fraction of perfectly informed traders (α) varies across experiments, taking the values of 0.5, 0.75, and 1, respectively. In the last case, the market maker's α is set to $1 - 1 \cdot 10^{-6}$ to prevent the collapse of the update equations (i.e., to prevent all updated values being multiplied with 0).
- Given the individual representation of the investors in our implementation, the number of each type of investor needs to be specified to each fraction applied. For the purpose of this paper, we conduct experiments with the minimum necessary number of investors: one informed—one uninformed for $\alpha = 0.5$, three informed—one uninformed for $\alpha = 0.75$, and one informed for $\alpha = 1$. Additionally we also study the implications of considering more interacting investors on the market dynamics when more investors interact.
- The probability (η) that uninformed traders place a buy order (and respectively sell order) is set to 0.3. Consequently, the probability that uninformed traders do not trade is set to 0.4.
- The time horizon of the investors is set to 0 in the experiments presented here, meaning that they trade continuously.
- All investors place market orders for one quantity of the risky stock. The investors do not withdraw their order once it is submitted.

4.1.3. Market Maker-Related Settings.

- The market maker knows the fraction of uninformed traders and the probability with which they trade.
- The market maker also knows whether informed traders are perfectly informed or noisily informed in the experiments.
- The time before the market maker counts a "no order" (the *supported inactivity time*) is set to 120 msec after some experimentation. If no order arrives after this period of time, the market maker updates his probability estimate and bid-ask spread accordingly.

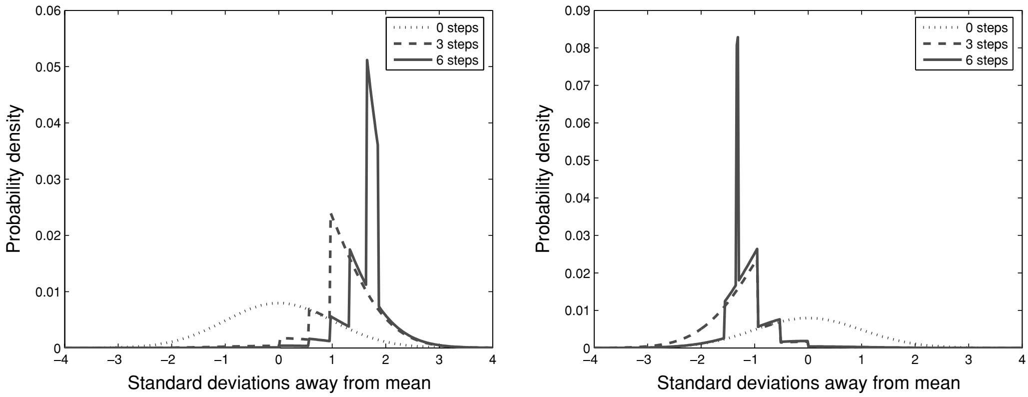

(a) (b)

FIGURE 1. Two different paths of evolution for the market maker's probability density estimate with perfectly informed traders.

Simulations last 10 minutes in real time. The number of transactions within this time span depends on the number and type of traders in the market. The number of transactions increases with the increasing number of investors. Short time horizon values for the investors imply more transactions in an experiment.

4.2. Validation and Verification

To verify whether our implementation is correct, we validate it against the model described in Das (2005). For this reason the properties of the probability density estimates are compared. After the arrival of an order, the market maker adjusts his current probabilities for all possible fundamental values in an attempt to track the actual fundamental value. Figure 1 shows the evolution of the market maker's density estimate after, respectively, 0, 3, and 6 update steps. Similar images can be found in Das (2005). Each step represents one trade or "no-trade" event and one update round including the normalization of the probabilities. Step 0 coincides with the moment that a jump in the fundamental value occurred and the market maker's probability estimate has just been initialized or recentered.

4.2.1. The Case of Perfectly Informed Traders. Figure 1 shows the evolution of the market maker's probability density estimate at, respectively 0, 3, and 6 steps after a (re)initialization in an arbitrary market run, with 70% of informed traders in the market. There are many possible variations with the same settings, and these two figures show two possible scenarios. In Figure 1(a), a situation is shown where most of the mass shift to the right a couple of times, which results in an upward shift of the probabilities over a smaller range. This corresponds to a situation where investors submit a series of buy orders up to a point where the (perfectly informed) submitters stop submitting orders, when the market maker's bid-ask spread is around the actual fundamental value, thus it is not worthwhile for them to trade any longer. The last few steps then resemble steps where the market maker shifted the probabilities for the region within the bid-ask quote upward, as he counted a "no-trade," increasing the certainty that the actual fundamental value is indeed within this range of values.

Figure 1(b) shows a different scenario. Here, there is three times a shift from the center to the left, corresponding to a series of sell orders. The fact that a sell order is submitted

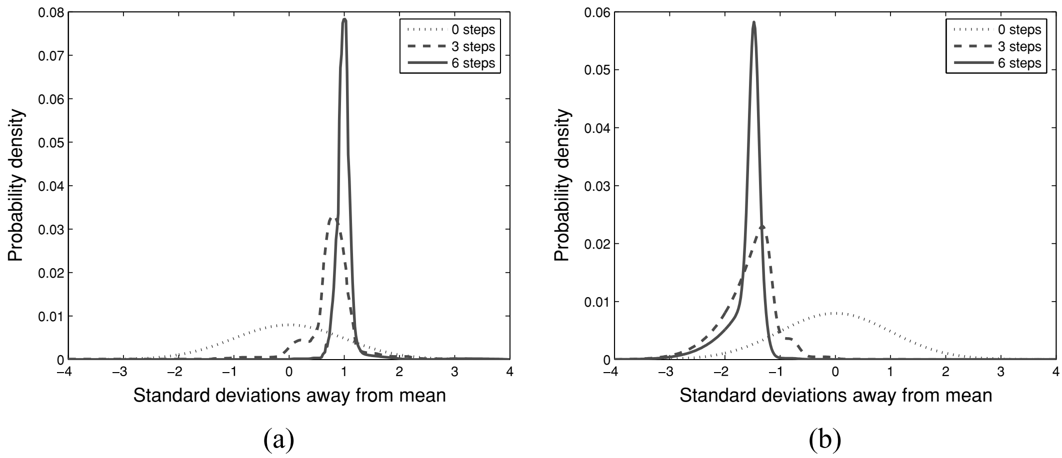

FIGURE 2. Two different paths of evolution for the market maker's probability density estimate with noisily informed traders.

is interpreted by the market maker as the current bid-ask spread being probably above the actual fundamental value. After this shift the market maker's bid-ask spread is set around the actual fundamental value (as in the previous example) so the probabilities within the area between the bid and ask price are increased while the others are decreased, creating a high peak around the actual fundamental value (in the market maker's opinion), until another trade arrives.

4.2.2. The Case of Noisily Informed Traders. With noisily informed traders, the updating algorithm results in a much smoother probability distribution. This happens because for every update by the market maker, the probabilities for a whole range of values are taken into account because of the additional noise in the traders' decisions. With perfectly informed traders, all probabilities above the current price are moved in one direction, and below the current price in the other direction. This leads to a distribution with several peaks, which is more sensitive to noise, because the area where most of the probability mass is located is more restricted than the case with noisily informed traders.

Figure 2 shows two situations after a jump in the fundamental value occurs, in a market setting where the informed traders received the information with additional noise added. The two scenarios shown are somewhat similar to the two that have appeared in the situation with perfectly informed traders. Again, these are just two scenarios out of many possible variations for the probability density estimate after 0, 3 and 6 steps, of which some differ only very slightly. Figure 2(a) shows a situation where the noisily informed traders have submitted a series of buy orders, making the market maker shift the probability mass to the right. For last two steps no more orders were received, thus the probabilities for the possible fundamental values that lie within the current bid-ask spread are shifted upward, while those outside the spread are shifted downward. Because the market maker has to rely on noisy information, the expected value most likely does not exactly follow the fundamental value, but the fundamental value with additional noise that is received by the noisily informed trader.

Figure 2(b) shows a situation that is similar to that in Figure 2(a), but in this case there is a shift from the center to the left, corresponding to a series of sell orders. However, the peak of the distribution is not as high as the peak in the distribution of Figure 2(a), i.e., the market

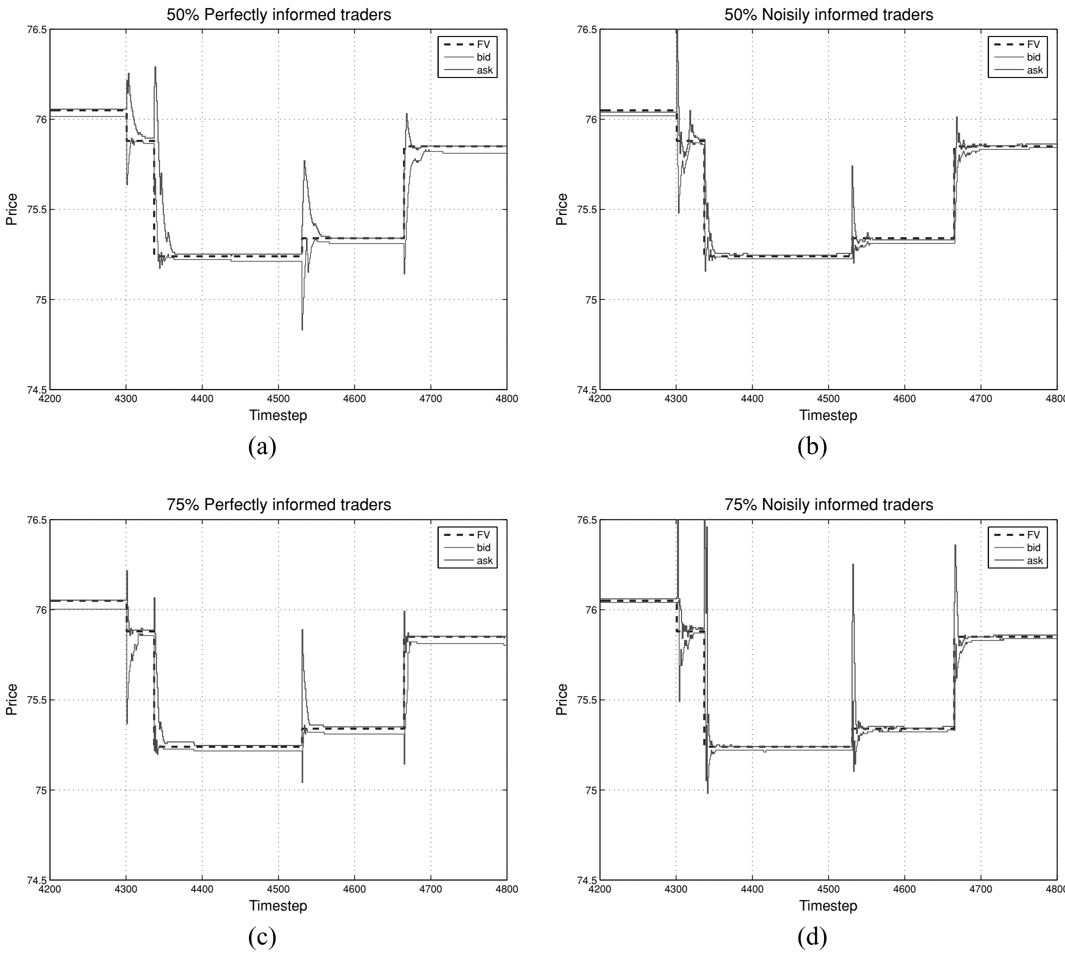

FIGURE 3. Market prices in the discrete-time (Das 2005) model. (a) 50% perfectly informed traders, (b) 50% noisily informed traders, (c) 75% perfectly informed traders, (d) 75% noisily informed traders.

maker is less certain about the fundamental value. The market maker receives a buy order after the first three orders, which increases the uncertainty (as the market maker is receiving mixed signals) and results in higher probabilities on the left of the peak in the distribution.

4.3. Findings

4.3.1. Discrete Model. Because we are interested in understanding how a market maker model developed for discrete-time simulation performs in a continuous, asynchronous simulation, we discuss in this section only typical outcome of experiments with the discrete-time market maker model of Das (2005). Figure 3 shows how the market maker sets his bid-ask spread trying to follow a given profile of fundamental value. For comparison purposes, the parameters in the experiments have been set as described in Section 4.1.

In Figure 3(a) and Figure 3(b), the fraction of informed traders is 0.50. Figure 3(a) shows the result of a typical experiment for perfectly informed traders. In this case, the market maker is able to learn the underlying fundamental value of the asset fairly quickly from the trades of the informed traders. When the traders are noisily informed, the market can still learn the underlying fundamental value, but he can be making small errors sometimes, as seen in Figure 3(b).

In Figures 3(c) and 3(d), the fraction of informed traders is 0.75. In this case, the market maker is able to learn the underlying fundamental value slightly more quickly. This is to be expected, because the market maker learns primarily from the trades of the informed traders. Hence, he can learn more quickly when there are more informed traders in the market. Note also the increased uncertainty after news arrives in the market, i.e., after a jump in the fundamental value. The bid-ask spread is initially large, but the spread is reduced gradually as time passes. The learning progresses fairly regularly, and the prices evolve without much fluctuation. This indicates stable learning on the part of the market maker.

4.3.2. Price Discovery in the Continuous Model. In this section we analyze the performance of the market maker in tracking the fundamental value in the asynchronous simulation. We examine whether the market maker is able to react timely to changes in the fundamental value and whether the fundamental value is correctly reflected in the market prices in our continuous, asynchronous simulations. Das (2005) has shown this to be the case in the discrete-time model. Our results in Section 4.3.1 also confirm this conclusion for 50% and 75% informed traders in the market.

In the following, we first consider the case with perfectly informed traders, and then we examine the influence of the noisily informed traders. We have run the simulation experiments for 10 minutes, but we display in the figures only a portion of the simulation results to increase the clarity of the graphics.

4.3.2.1. Performance with Perfectly Informed Traders. With 100% of the traders in the market being informed, the market maker should be able to quickly adjust to any changes in the fundamental value. Figure 4(a) shows that the market maker is able to track the fundamental value almost perfectly when receiving the orders of one single perfectly informed trader. The only points where the values are slightly inaccurate are directly after a jump during which the market maker is learning the new fundamental value. After the market maker learns the fundamental value, the bid and ask prices are equal to the fundamental value or differ just one cent from it. Note that the speed of adjustment depends on the value of the "no-order" condition. In case of a 100% perfectly informed trader, the market maker quickly captures the fundamental value. The bid-ask spread is, however, quite large immediately after a jump, and hence the perfectly informed trader does not trade. Consequently, the bid-ask is not updated based on new orders, but only based on the no-order condition. The smaller the value of the no-order condition, the faster the convergence of the bid and ask values to the fundamental value.

As the fraction of informed traders decreases, the market maker takes more time to approximate the fundamental value. Both with 75% of the traders being perfectly informed (Figure 4(c)), and with 50% of the traders being perfectly informed (Figure 4(e)), the market maker learns the fundamental value very well. Large jumps, such as the one around the 410th second are more difficult to track. An interesting feature that we observe is that the market maker seems to be more uncertain when working with 75% informed traders as opposed to 50% informed traders. In the market with 50% informed traders, the uncertainty can be exclusively ascribed to the presence of the uninformed trader. In Figure 4(c) a kind of overreaction can be observed. After the fundamental value jumps, the market maker's bid and ask prices begin to fluctuate with a decreasing amplitude, before approximating the desired value. This phenomenon is caused by the larger amount of traders in the market with 75% perfectly informed traders. Here, three informed and one uninformed investors

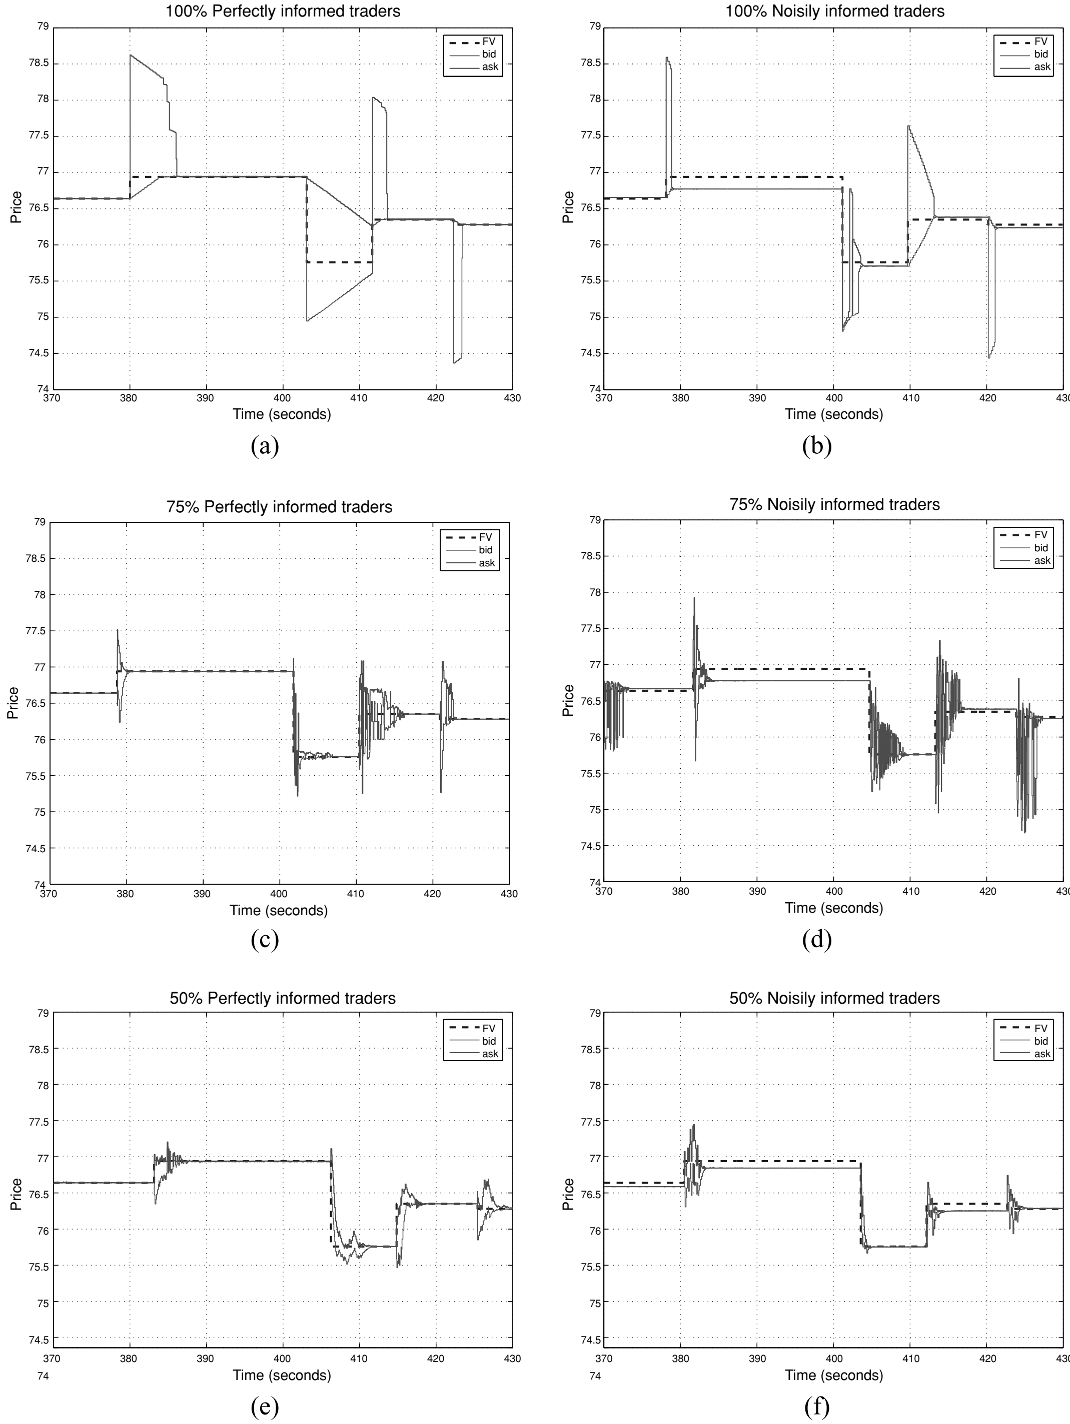

FIGURE 4. Tracking the fundamental value with perfectly informed ((a), (c), (e)) and noisily informed ((b), (d), (f)) traders.

are interacting, as opposed to one informed and one uninformed trader in the experiments with 50% informed traders. The informed traders respond all immediately to the change in the fundamental value by placing the same kind of order, thus causing the peaks in the resulting bid and ask prices. Note that the repeated peaks are not visible in Figure 3(c) of the discrete model. With the discrete model, the market maker is uncertain about the correct bid-ask prices after a jump in the fundamental value, but the market maker quickly learns the new fundamental value, especially when the fraction of informed traders in the market is high. There are no other large oscillations.

4.3.2.2. Performance with Noisily Informed Traders. Let us now analyze how noisy information diffuses into prices. For this case, we consider noisily informed traders instead of the perfectly informed ones. If the trading crowd consists of a single noisily informed trader (Figure 4(b)), the market maker is able to follow the value from the informed trader, although most of the time he over- or underestimates slightly the actual fundamental value due to the additional noise. The additional noise might, for example, first push the observed fundamental value upward and, after the next jump in the actual value, downward (or vice versa), resulting in a small change in the *observed* value.

If uninformed traders are also present next to noisily informed traders, the bid and ask prices can be further away from the actual fundamental value because of the influence of both the additional noise in the noisily informed traders' observations and the noise from the uninformed traders' orders. Compared to the corresponding simulation with perfectly informed traders, there are more fluctuations in the simulation results shown in Figures 4(d) and 4(f). Each noisily informed trader observes a different fundamental value. Each of the traders bases his orders on this observed value, and the market maker attempts to infer the actual fundamental value based on these trades, being aware of the presence of noise and knowing the noise function. The market maker attempts to set his bid and ask prices around the mean price of the different observed values. However, this value is most likely not the observed fundamental value of any of the traders in the crowd, and thus all investors will submit more orders, leading to additional fluctuations. This phenomenon is more pronounced when the fraction of uninformed traders increases.

4.3.3. Bid-Ask Spread Size. The bid-ask spread is the result of the market maker's price setting process. With a small spread size, it is less likely that the actual fundamental value falls within the bid-ask range. Informed traders will not submit orders if the market maker's bid and ask prices are set around their observed fundamental value. Because one of the market maker's tasks is to ensure market liquidity, he has to set the bid and ask prices sufficiently small to make sure there are investors who are willing to trade. However, with a spread that is too small, the market maker will constantly lose money to well-informed traders. Consequently, the market maker needs to find a balance between his profitability and market liquidity.

A sample of the average spread sizes (in cents) from the experiments we have conducted can be found in Table 1. Relatively high spread values can be observed in the experiments with one single perfectly informed trader. This large spread is a result of the speed of adjustment of the bid-ask spread, which in turn depends on the "no-order" condition as we discussed above. After a jump in the fundamental value, the informed trader does not submit orders as long as the fundamental value is within the spread. The market maker can only update the bid-ask spread under the "no-order" condition, whose size determines how fast the convergence takes place. It could thus be argued that the market maker could update the bid-ask prices faster in the experiments with an 100% perfectly informed trader.

Das (2005) shows that the spread size is also related to the jump probability and the standard deviation of the jump process. The market maker is uncertain about the fundamental

TABLE 1. Samples of Bid-Ask Spread Sizes*

Fraction	Number of Informed Traders	Number of Uninformed Traders	Perfectly Informed Case	Noisily Informed Case
100%	1	0	33.14	7.99
100%	4	0	10.95	1.39
75%	3	1	4.16	7.68
50%	1	1	6.70	3.93

*In cents.

TABLE 2. Sample Cash Balance and Inventory of the Market Maker

Simulation	Cash Balance	Inventory	Total Wealth	Average per Trade
100% perfectly informed traders	23,175	−305	10.02	0.00
75% perfectly informed traders	−1,252	18	118.97	0.01
50% perfectly informed traders	1,575	−19	129.11	0.02
100% noisily informed traders	5,720	−75	0.05	0.00
75% noisily informed traders	29,597	−383	438.49	0.03
50% noisily informed traders	4,053	−52	91.47	0.01

value whenever a jump in the fundamental value has just occurred. Whenever this happens, the market maker is only informed of the fact that a jump has occurred, but not of the size or direction of the jump, which will be more extreme in cases with a higher standard deviation for the jump process. Only after a period of trading the market maker knows in which direction the prices should be moved, and it takes time to adjust his spread size. With a lower jump probability, which is represented in our implementation by a longer time between two consecutive changes, the market maker can stick with a small spread size for a longer period.

4.3.4. Wealth and Inventory. When quoting bid and ask prices, the market maker aims to satisfy a zero-profit condition. Given, however, that the market making algorithm that is used is an approximation, it is unlikely the method will be exactly zero-profit. An interesting question is how large the difference is.

In our experiments, there are no restrictions on the cash and inventory of the market maker. There are also no restrictions on the cash and inventory of the investors. This allows them to go short, i.e., sell more assets than they currently have in possession, or to "loan" money, i.e., buy more assets than they have money for. Table 2 shows the cash balance (initially zero) and inventory (initially zero shares) at the end of the simulation run for the simulation runs displayed in the previous section. Furthermore, the market maker's total wealth at the end of the simulation and the average wealth contribution of a single trade are presented. These results show that the market maker, in general, makes a small profit or loss, corresponding to the zero-profit condition. The market maker could make more profit by enlarging the spread, as proposed in Das (2005). We did not study, however, the consequences of such a setting for the purposes of this paper.

5. IMPLICATIONS OF CONTINUOUS, ASYNCHRONOUS REPRESENTATION

In Section 3, we have pointed out that continuous implementation of markets and the individual, autonomous representation of traders carrying out tasks asynchronously requires the specification of additional parameters. For example, the no-order condition in the continuous setting can take other values than in the discrete setting. Furthermore, attention should be paid to the number of different types of investors, in addition to their fraction, when traders are individually represented. Individual, autonomous representation, together with the asynchronous behavior of the traders can influence market dynamics in an interesting way.

Consider, for example, the experiments with 75% informed traders. Here, the number of informed investors is three as opposed to the experiments with the population consisting of a single trader. The difference thus, between the two experiments, is not only the fraction of the uninformed investors. The number of players in the market is also different. The question is now, whether the observed fluctuations are caused only by the random traders, or whether the increased number of informed investors (in absolute terms, and hence not in terms of the fraction) also plays a role. To answer this question, we analyze situations with populations of multiple informed traders.

In our experiments, the number of traders is kept relatively small, while in the models in Glosten and Milgrom (1985) and Das (2005) the number of traders is not relevant (only the fractions of various types of traders matter). This has to do with the typical feature that in the original models investors are either not represented individually (thus, just the order flow is generated) or one single investor is centrally selected to trade at each trading period. A result of this assumption is that, in each period, there is either zero or one order to be processed.

In the continuous model, in which investors are represented as individuals and exhibit autonomous, asynchronous behavior, every trader has the opportunity to submit an order whenever the trader determines it is worthwhile to do so. This feature implies that if more investors interact on the market, it can happen that some of them decide to place orders at the same time. For example, when all informed traders observe the same jump in the fundamental value, and it is worthwhile to submit an order, they will all do so. Given that the market maker is able to process only one order at a time, this homogeneous, simultaneous decision will result in a queue of orders for him. If there are more informed traders in the market, there is also a larger queue. This effect is stronger with perfectly informed traders than with noisily informed traders, as perfectly informed traders all observe the same fundamental value, while noisily informed traders observe this value with additional noise, which might lead them to take different decisions.

In one of the experiments, we have increased the number of perfectly informed traders from one to four in the experiments with 100% informed traders. Whenever there is a jump in the fundamental value, all four perfectly informed traders respond to this change in the same way because they all use the same information. Hence, they perform the same action. In this way, the market maker can sometimes learn the underlying fundamental value faster. For comparison, consider the change around 380th second with one perfectly informed trader in Figure 4(a). Toward the end of the learning period, the ask price is updated several times before the market maker learns the true fundamental value. In Figure 5(a), it can be seen that the last learning steps occur much faster, thereby reducing the total learning time for the market maker. However, the arrival of multiple orders with the same actions can also cause an overreaction to any change in the fundamental value. An example of this can be seen after the jump around 415th second in Figure 5(a).

The influence of the overreaction is more prominent in the case of four noisily informed traders as shown in Figure 5(b). In this case, the market maker is not always able to learn the fundamental value accurately, although the probability density estimate may quickly become concentrated in a very small area due to the large amount of orders received after a change in the fundamental value. Because different traders perceive different values of the fundamental,

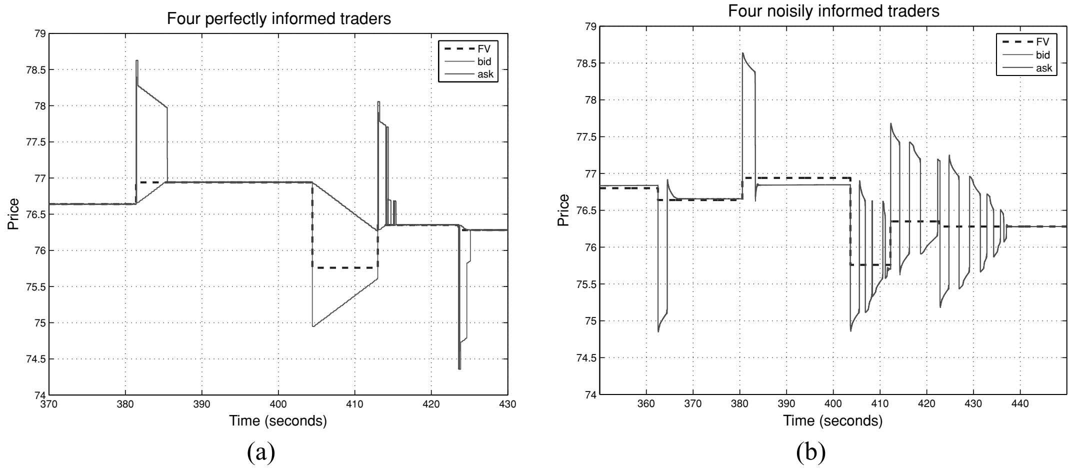

FIGURE 5. Tracking the fundamental value (a) with four perfectly informed and (b) with four noisily informed traders.

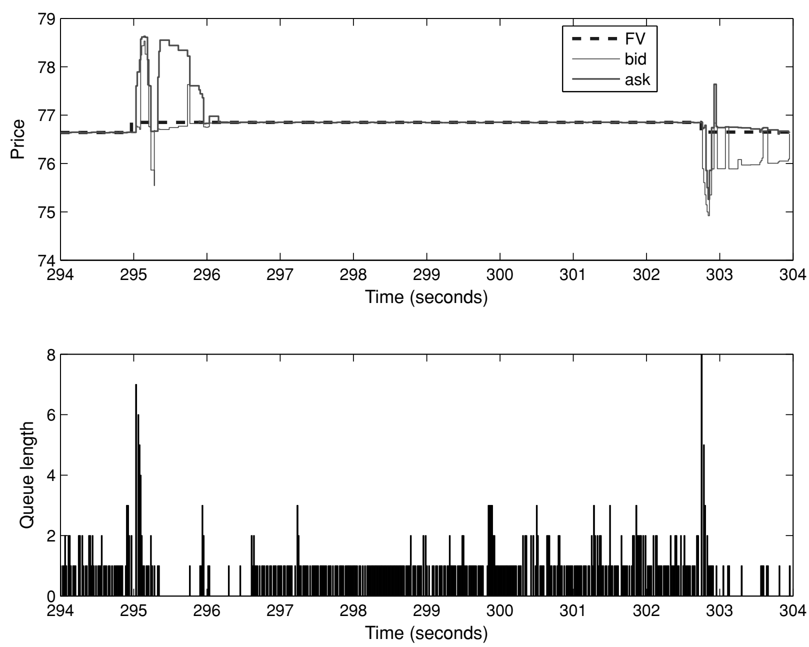

FIGURE 6. Market prices and the pending order queue over time, in a market with 10 traders, of which eight are perfectly informed.

they may trade at different times, leading to oscillations in the market maker's estimate of the fundamental value. Such effects are not visible in the discrete-time model.

In simulations with a relatively high amount of informed traders, order queues cause large peaks in the bid and ask prices that are set by the market maker, whenever a jump in the fundamental value occurs. Figure 6 shows a part of a simulation run with 10 traders, of which eight are perfectly informed. Both the bid-ask prices and the order queue are displayed in the figure.

It can be observed that the queue length of pending orders increases significantly after a jump in the fundamental value. Because all informed traders submit an order of the same type, each individual order pushes the market maker's bid and ask prices in the same direction, causing a strong movement in one direction. As soon as these orders are processed, the bid and ask prices are typically over- or underestimating the fundamental value. Given that the informed traders know the actual fundamental value and see the results of this overreaction, they will again submit an order, driving the prices back in the direction of the fundamental value and even further away in the other direction. This process may be repeated several times until the market maker sets his bid and ask prices around the fundamental value. These *overshoots*, caused by the herd-like behavior of informed traders who immediately react to a change in the fundamental value, is the consequence of the autonomous nature of the traders. In the original (Das 2005) model, this effect is not present, because only one trader can submit an order at a time, and thus, when the next trader enters the market, he observes the new bid and ask prices, in which the new information is already (partly) processed.

When a queue of pending orders arises, the market maker will aim to handle them one by one, as seen with the queue peak after the 295th second. Market orders that need to wait for other orders to be executed before they are handled will not be cleared at the price they were placed for. The market maker, however, processes all the orders as if they would have been placed one after each other, taking into account the effects of transacting the earlier arrived orders. This implies that the amplitude of fluctuations increases with the increasing number of informed investors. In order to improve the learning algorithm, the market making algorithm could be adjusted to handle this specific situation. The market maker could learn from the queue as this indicates the arrival of some news.

Another way to deal with this specific situation caused by the asynchronous behavior and homogeneous setting of informed traders is to introduce more variation in the order placement behavior of the traders. Introducing different reaction times for the investors, for example, would reduce the length of the queue that arises when a change in the fundamental value of the risky asset occurs. Different reaction times could be the result of different news-sources, location, or time needed to analyze the news.

6. CONCLUSIONS AND FUTURE RESEARCH

In this paper we have presented a continuous-time simulation of an information-based market model with a learning market maker. We have implemented the individual investors as asynchronously interacting autonomous agents. We have compared the model against its discrete simulation counterpart and have shown that continuous, asynchronous setting can entail significantly different market dynamics compared to the discrete simulation.

Continuous asynchronous implementation implies that additional parameters (such as the number of traders and the time-horizon for traders) need to be defined and determined. Additionally, some other parameters have to be represented in a different way (such as the probability of a jump, and the no-order condition in our model). The results of our experiments point out that the continuous simulation of continuous markets, and individual asynchronous representation of traders' behavior influence the market dynamics. Given that most financial markets apply continuous trading sessions, these features should thus be taken into account during modeling.

We found the main difference in the outcomes of the discrete-time and continuous-time simulation is the fluctuation of bid-ask quotes (and consequently prices). Namely, prices tend to fluctuate more often and with a larger amplitude in the continuous, asynchronous setting with investors represented individually, before they converge to the fundamental value or settle down to a somewhat stable value. The amplitude of the fluctuations tends to increase with the increasing number of investors. The explanation we found behind this phenomenon is the herd-like behavior of informed investors, who react all at the same time to changes in the fundamental value causing the market maker to overreact to changes, causing in turn overshoots in the bid-ask quotes.

Receiving and reacting to news simultaneously causes increased order queues in the continuous setting. We propose two ways to deal with this specific situation. One the one hand, informed traders do not necessarily decide to trade all at the same time in reality. The size of the order queues can be reduced by introducing a variation for the reaction times of the informed traders to the available news. On the other hand, order queues contain valuable information from which the market maker could learn. They indicate changes in the fundamental value and also the fact that the stock is probably mispriced. The market maker could try to avoid overreaction by, for example, modifying his expectations more smoothly, when he observes a queue of orders on the same trading side. Furthermore, he could also temporarily modify the fraction of informed traders when using the learning algorithms, assuming that queue consists mainly of the orders of informed traders who react to mispriced stocks. In our future research we intend to investigate these extensions to the learning behavior of the market maker.

Besides continuous-time specific settings, there are several other interesting situations to experiment with in both the discrete-time and the continuous-time framework. It would be interesting, for example, to vary the jump process and investigate the market maker's learning behavior under different circumstances. Large jumps could be introduced in this sense, or historical time series could be used to model the jump process. Another challenging task is to adapt the learning algorithm of the market maker to account for both perfectly informed and noisily informed traders.

It might seem that the asynchronous, continuous-time framework just makes the dynamics, and the interpretation of the results, more intricate. It has, however, several advantages compared to the original model. While in the original framework, there is no attention paid to the way orders arise, in the model presented in this paper there is special attention paid to the behavior of individual investors. Investors in the continuous setting are more realistically represented, being able to make decisions autonomously. Furthermore, the agents may have more information available in the continuous-time framework, for example, via order queues. We expect that using this additional information would enable the market maker to learn several quantities that would be unknown to the market makers on actual markets (such as the fraction of informed traders and the distribution of the jump process).

REFERENCES

BELLIFEMINE, F., G. CAIRE, T. TRUCCO, and G. RIMASSA. 2005. JADE Programmer's Guide. Available at http://jade.tilab.com/doc/programmersguide.pdf.

BOER, K., M. POLMAN, A DE. BRUIN, and U. KAYMAK. 2004. An agent-based framework for artificial stock markets. *In* 16th Belgian-Dutch Conference on Artificial Intelligence (BNAIC'04), pp. 83–90.

BOER, K., A. DE BRUIN, and U. KAYMAK. 2005a. On the design of artificial stock markets. ERIM Report Series Research in Management, ERS-2005-001-LIS. Available at http://hdl.handle.net/1765/1900.

BOER, K., A. DE BRUIN, and U. KAYMAK. 2005b. A modular agent-based environment for studying stock markets. ERIM Report Series Research in Management, ERS-2005-017-LIS. Available at http://hdl.handle.net/1765/1929.

BROCK, W. A., and C. H. HOMMES. 1997. A rational route to randomness. Econometrica, **65**(5): 1059–1095.

CHALLET, D., M. MARSILI, and Y.-C. ZHANG. 2005. Minority games: Interacting Agents in Financial Markets. Oxford University Press, Oxford, UK.

CHAN, N. T., and C., SHELTON. 2001. An electronic market maker. AI Memo 2001–005, Massachusetts Institute of Technology, Cambridge, MA.

DAS, S. 2003. An agent-based model of dealership markets. *In* International Workshop on Complex Agent-based Dynamic Networks, University of Oxford, UK.

DAS, S. 2005. A learning market-maker in the Glosten-Milgrom model. Quantitative Finance, **5**(2): 169–180.

DEMARCHI, M., and T. FOUCAULT. 2000. Equity trading systems in Europe: A survey of recent changes. Annales d'Economie et de Statistique, **60**: 73–115.

GLOSTEN, L. R., and P. R. MILGROM. 1985. Bid, ask and transaction prices in a specialist market with heterogenously informed traders. Journal of Financial Economics, **14**: 71–100.

HARRIS, L. 2003. Trading and Exchanges: Market Microstructure for Practitioners. Oxford University Press, New York.

LEBARON, B. 2002. Building the Santa Fe artificial stock market. Working Paper, Brandeis University.

LOISTL, O., and O. VETTER. 2000. KapSyn Computer Modelled Stock Exchanges, User Manual. Ver. 3.02. University of Economics and Business Administration, Vienna, Austria.

MARSILI, M. 2001. Toy models of markets with heterogeneous interacting agents. *In* Economics with Heterogeneous Interacting Agents, Lecture Notes in Economics and Mathematical Systems, *Edited by* A. Kirman, and J.-B. Zimmermann, vol. 503. Springer, Berlin, pp. 161–182.

O'HARA, M. 2002. Market Microstructure Theory. Blackwell Publishing, Oxford, UK.

RABERTO, M., S. CINCOTTI, S. FOCARDI, and M. MARCHESI. 2001. Agent-based simulation of a financial market. Physica A, **299**: 319–327.

SHATNER, M., L. MUSHNIK, M. LESHNO, and S. SOLOMON. 2000. A continuous time asynchronous model of the stock market. arXiv:cond-mat. Available at http://arxiv.org/PS-cache/cond-mat/pdf/0005/0005430.pdf.

SMITH, E., J. D. FARMER, L. GILLEMOT, and S. KRISHNAMURTHY. 2002. Statistical theory of the continuous double auction. Working Paper 02-10-057, Santa Fe Institute.

TESFATSION, L. 2001. Introduction to the special issue on agent-based computational economics. Journal of Economic Dynamics and Control, **25**: 281–293.