From Discrete-Time Models to Continuous-Time, Asynchronous Models of Financial Markets

Katalin Boer, Uzay Kaymak and Jaap Spiering

ERIM REPORT SERIES *RESEARCH IN MANAGEMENT*	
ERIM Report Series reference number	ERS-2006-009-LIS
Publication	February 2006
Number of pages	31
Persistent paper URL	
Email address corresponding author	kaymak@few.eur.nl
Address	Erasmus Research Institute of Management (ERIM) RSM Erasmus University / Erasmus School of Economics Erasmus Universiteit Rotterdam P.O.Box 1738 3000 DR Rotterdam, The Netherlands Phone: + 31 10 408 1182 Fax: + 31 10 408 9640 Email: info@erim.eur.nl Internet: www.erim.eur.nl

Bibliographic data and classifications of all the ERIM reports are also available on the ERIM website:
www.erim.eur.nl

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS	
Abstract	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 modelling of financial markets. We study the behaviour 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. Since 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 modelling.
Free Keywords	Agent-Based Computational Finance, Artificial Stock Markets, Market Microstructure, Glosten and Milgrom Model, Informational Asymmetry, Continuous Trading, Autonomous Behaviour.
Availability	The ERIM Report Series is distributed through the following platforms: Academic Repository at Erasmus University (DEAR), DEAR ERIM Series Portal Social Science Research Network (SSRN), SSRN ERIM Series Webpage Research Papers in Economics (REPEC), REPEC ERIM Series Webpage
Classifications	The electronic versions of the papers in the ERIM report Series contain bibliographic metadata by the following classification systems: Library of Congress Classification, (LCC) LCC Webpage Journal of Economic Literature, (JEL), JEL Webpage ACM Computing Classification System CCS Webpage Inspec Classification scheme (ICS), ICS Webpage

From discrete-time models to continuous-time, asynchronous models of financial markets

Katalin Boer Uzay Kaymak Jaap Spiering

Erasmus School of Economics, Econometric Institute

P.O. Box 1738, 3000 DR Rotterdam, The Netherlands

Email: {kboer, kaymak}@few.eur.nl

Abstract

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 modelling of financial markets. We study the behaviour 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. Since 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 modelling.

Keywords

Agent-based computational finance, artificial stock markets, market microstructure, Glosten and Milgrom model, informational asymmetry, continuous trading, autonomous behaviour.

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 [20]. Usually, individuals are represented by (software) agents interacting in an artificial environment. By using agents for studying market dynamics, heterogeneous, boundedly rational, and adaptive behaviour 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 [2]. 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. [15, 6, 5, 13, 17]). Although call markets are encountered regularly, continuous trading sessions and quote-driven execution mechanisms are more common in real financial markets [10, 12]. This has prompted some groups to look beyond the call market structure. Examples of such models have been described in [14, 7, 19, 18]. 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 [18], where traders "sleep" after actions, and they "wake up" at predefined times, or as a result of certain events. Other studies [14, 17] model continuous sessions through discrete time simulation, in which the trader whose decision is carried out during the next trading session is centrally (usually randomly) selected. Further, 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 markets 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 [4], which implements a continuous, asynchronous agent-based artificial stock market on top of JADE (Java Agent Development Environment) [1]. 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 behaviour model of JADE provides for the execution of concurrent (asynchronous) agent activities. Continuity is modelled 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 one hand, a lot of information is available to the agents in turn-based models, since 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. Since 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 in order to assess the limitations of the current modelling practice. In this paper, we consider this question by studying the behaviour 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. Further, we consider what additional considerations are needed in order to extend the turn-based model into the continuous, asynchronous model.

We consider an information-based model, since they provide insights into the adjustment process of prices that we are interested in [16]. We study the behaviour of a market maker in order to focus the scope of the investigation. The research presented in this paper is based on the learning market maker from [9]. This model extends the information-based model in [11], 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 rate of informed and uninformed traders. In [8] and [9] 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. In the model discrete time simulation is applied and a probabilistic representation of order flows is considered. We implement this model 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 artificial stock market 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 [9]. 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 modelling 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 analyzed. 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 in order to study market dynamics is based on an extended version of the information-based Glosten and Milgrom model proposed in [9]. 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 [3] 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 more 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. The fundamental value can be thought of as a quantification of news with respect to the stock traded. The jump process is modelled 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 [9] 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 represented: informed traders and uninformed traders. Depending on the simulation parameters, the informed traders may be *perfectly informed* or *noisy informed*. Perfectly informed traders observe the correct fundamental value (V_t), while noisy informed investors observe a distorted fundamental value $W_t = V_t + \tilde{\eta}\ (0, \sigma_W)$. Here, $\tilde{\eta}\ (0, \sigma_W^2)$ 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 ($V_t > A_t$ in case of perfectly informed traders and $W_t > A_t$ in case of noisy informed traders). He will sell if the fundamental value that he observes is below the bid price ($V_t < B_t$ or $W_t < B_t$). He will place no order if the observed fundamental value is within the bid-ask spread ($B_t \leq V_t \leq A_t$ or $B_t \leq W_t \leq A_t$). 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. In order to discover the fundamental value, Bayesian learning is applied, based on the extended Glosten and Milgrom model [9].

The Glosten and Milgrom model [11] 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 arrived, and the ask price to the expectation of the true value given that a buy order arrived. In order to compute the expectations in the extended model in [9], 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 ($\overset{\sim}{\omega}\ (0, \sigma)$);
- whether a change in the fundamental value occurred;
- the distribution function of the noise process ($\overset{\sim}{\eta}\ (0, \sigma_W^2)$).

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 probably lower than the current bid price, and a series of buy orders might indicate that the fundamental value is probably higher than the current ask price. However, the market maker will have to take into account the noise incorporated by the orders of the noisy informed traders, and the noise implied by the orders submitted by uninformed traders.

2.4.1 Adjusting the bid and ask quotes

The market maker tries to set 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 [9]:

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

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

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 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 noisy informed traders, the probabilities for sell and buy orders are determined by the following equations:

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

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

$$\Pr(\text{Buy}|V = V_i, V_i \leq A) = (1-\alpha)\eta + \alpha \Pr(\widetilde{\eta}\ (0, \sigma_W^2) > (A - V_i)), \tag{7}$$

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

The second term in (5) reflects the probability that a noisy 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 (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 noisy 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 noisy 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 $V_{max} = V_0 + 4\sigma - 1$ (rounded off to an integer value in cents) to contain the prior value probabilities. The values are initialized by setting the i-th 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, so 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

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

where the Action represents a buy order, a sell order or noorder. 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, in order to ensure market liquidity and encourage trading.

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

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

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

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

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

$$\begin{aligned} \Pr(\text{No order}|V = V_i,&B \leq V_i \leq A) = (1-\alpha)(1-2\eta) + \\ &+ \alpha\Big(\Pr(\widetilde{\eta}\ (0, \sigma_W^2) > (B - V_i)) + \Pr(\widetilde{\eta}\ (0, \sigma_W^2) > (V_i - A))\Big), \end{aligned} \tag{13}$$

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

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 behaviour 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 further 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 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). In order 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 an 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. In order 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 behaviour of an investor in

[2, 3]). 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. Further, 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 [9] so that we can verify our implementation and compare the results. Some values or processes needed to be converted, however, from the discrete time model into the continuous time model, as described in Section 3. Further, additional parameters need to be specified related to the individual, autonomous representation of the investors.

Fundamental value related settings.

- The initial value (V_0) is 7500 cents in accordance with Das.
- The standard deviation (σ) of the jump process is set to 50 cents.
- In [9], every trading period a jump in the fundamental value occurs with a small probability (1 in 1000). 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).

Investor related settings.

- The fraction of perfectly (and respectively noisy) informed traders (α) varies across experiments, taking the values of 0.33, 0.5, 0.75 and respectively 1. In the latter 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 to be 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: 1 informed - 2 uninformed for $\alpha = 0.33$, and 1 informed - 1 uninformed for $\alpha = 0.5$, 3 informed - 1 uninformed for $\alpha = 0.75$, and 1 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.

Market maker related settings.

- The market maker knows the fraction of uninformed traders and the probability with which they trade.

- A variable indicates whether informed traders are perfectly informed or noisy informed in the experiments.

- The time before the market maker counts a ”no order” (the *supported inactivity time*) is set to 1500 milliseconds, after some experimentation. If no order arrives after this period of time, the market maker updates his probability estimate and bid ask spread accordingly.

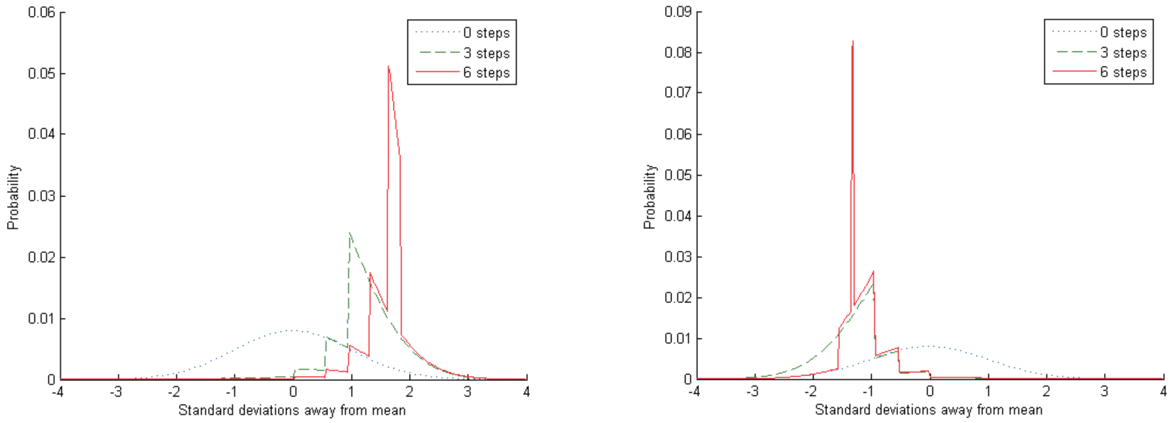

(a) (b)

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

In order to verify wether our implementation is correct, we validate it against the model described in [9]. 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 at respectively 0, 3 and 6 update steps. Similar images can be found in [9]. Each step represents one trade or ”no trade” event and one update round including the normalization of the probabilities. Step 0 coincides with either the start of the market or the moment that a jump in the fundamental value occurred. In either case, the market maker's probability estimate has just been initialized or re-centered.

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 1a, a situation is shown where most of the mass shifted to the right a couple of times, which results in an upwards 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, so 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 upwards, as he counted a ”no trade”, increasing the certainty that the actual fundamental value is indeed within this range of values.

Figure 1b 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 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 noisy informed traders

With noisy 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 vulnerable when it comes to noise, since the area where most of the probabilities are in is more restricted than is the case with noisy 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 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 2a shows a situation where the noisy informed traders have submitted a series of buy

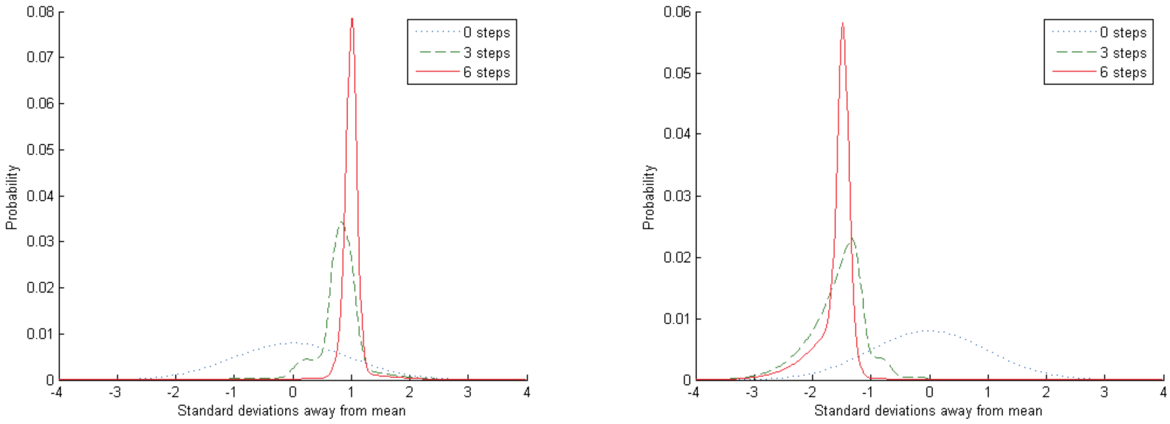

(a) (b)

Figure 2: Two different paths of evolution for the market maker's probability density estimate with noisy informed traders.

orders, making the market maker shift the probability mass to the right. The last two steps no more orders were received, so the probabilities for the possible fundamental values that lie within the current bid-ask spread are shifted upwards, while those outside the spread are shifted downwards. Since 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 noisy informed trader.

Figure 2b shows a situation that is similar to that in Figure 2a, 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 2a, i.e. the market maker is less certain about the fundamental value. The market maker received a buy order after the first three orders, which increased 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

Since 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 typical outcome of experiments with the discrete-time market maker model of [9]. Figure 3 shows

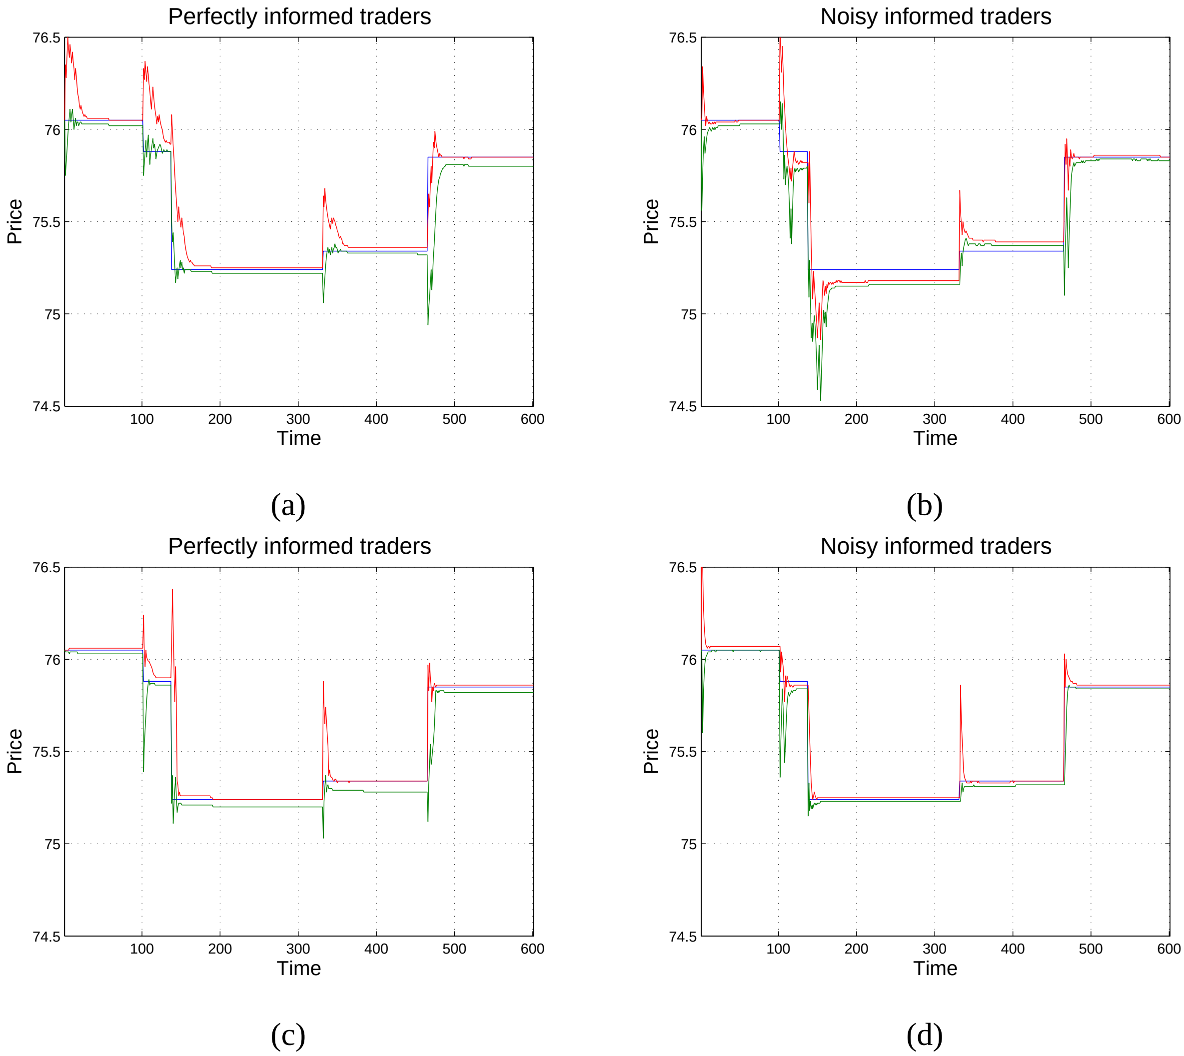

Figure 3: Market prices in the discrete-time model of [9]. (a) 50% perfectly informed traders, (b) 50% noisy informed traders, (c) 75% perfectly informed traders, (d) 75 % noisy informed traders.

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 3a and Figure 3b, the fraction of informed traders is 0.5. Figure 3a 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 noisy informed, the market can still learn the underlying fundamental value, but he can be making small errors sometimes, as seen in Figure 3b.

In Figure 3c and Figure 3d, the fraction of informed traders is 0.75. In this case, the market

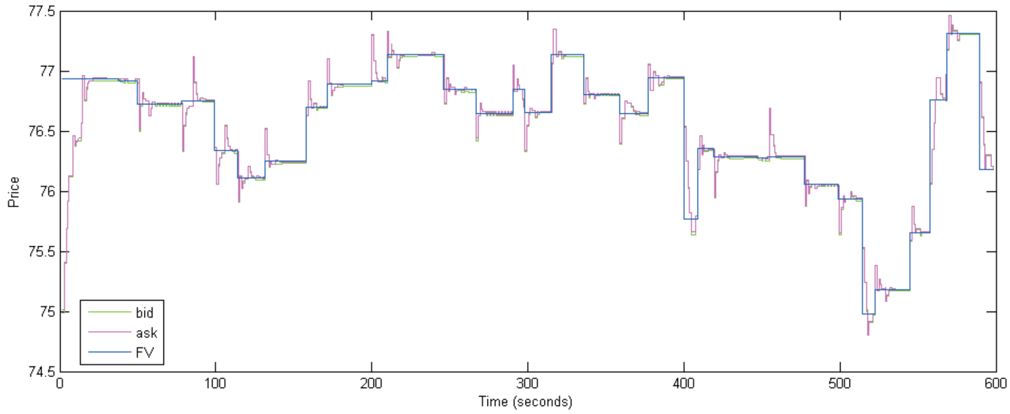

Figure 4: Tracking the fundamental value with one single perfectly informed trader.

maker is able to learn the underlying fundamental value more quickly. This is to be expected, since 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 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. In [9], the author has shown this to be the case in the discrete-time model. Our results in Section 4.3.1 also confirm this conclusion.

In the following, we first consider the case with perfectly informed traders, and then we examine the influence of the noisy informed traders.

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 shows that the market maker is able to track the fundamental value almost perfectly without any large delays when

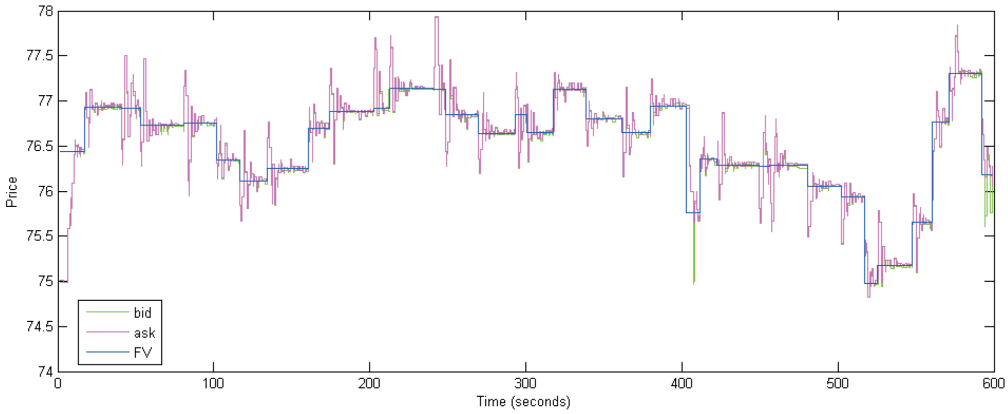

Figure 5: Tracking the fundamental value with 75% perfectly informed traders.

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. Most of the time the bid and ask prices are equal or differ just one cent, following from the feature that the market maker is aware of the fact that traders are informed. This is similar to the results from the market maker in the discrete model.

As the fraction of informed traders decreases it takes the market maker more and more time to approximate the fundamental value. In case of 75% of traders being perfectly informed (3 informed and 1 uninformed) the market maker still tracks the fundamental value very well (Figure 5). A kind of overreaction can be observed, however. 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 probably caused by the larger amount of traders in the market: three informed and one uninformed, as opposed to one single trader in the earlier validation experiment. 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 of the discrete model. With the discrete model, an overreaction is visible 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.

By increasing the fraction of the uninformed traders, it becomes more and more difficult to the market maker to learn the new fundamental value (Figure 6). He is still able to track it, but it takes him longer to adjust to the actual fundamental value. As we expect, large jumps, such

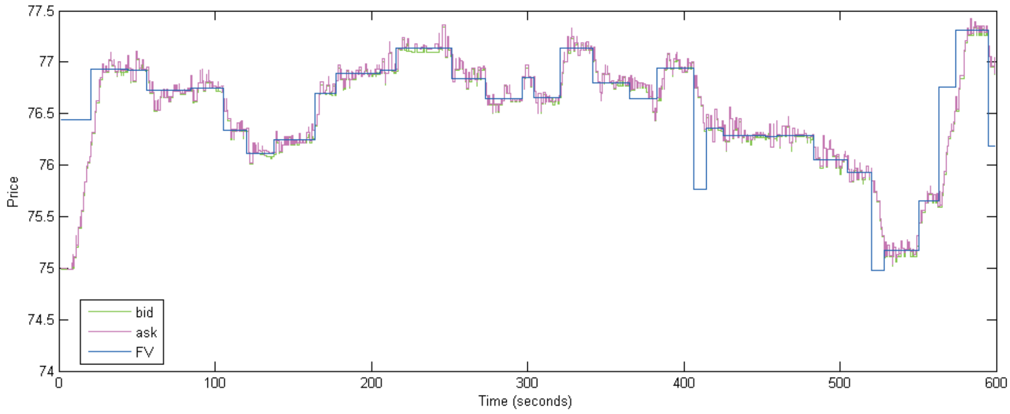

Figure 6: Tracking the fundamental value with 50% perfectly informed traders.

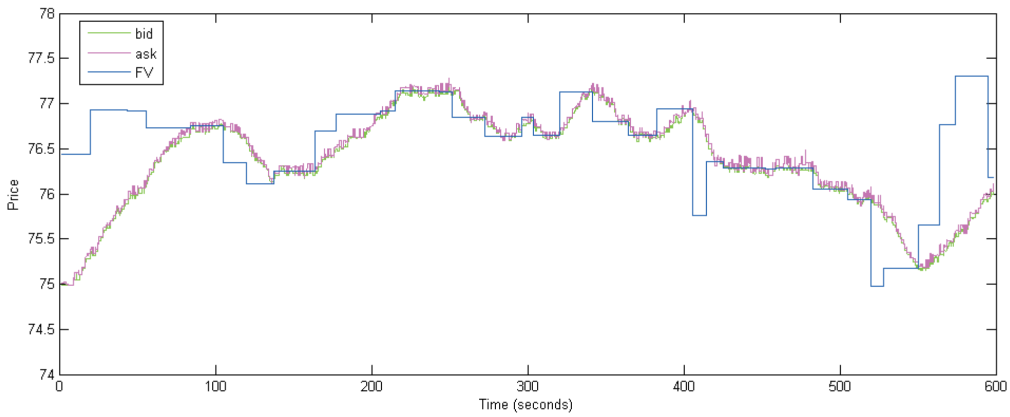

Figure 7: Tracking the fundamental value with 33% perfectly informed traders.

as the one around the 400th second in Figure 6 are more difficult to track.

If there are more uninformed traders than informed traders, the market maker is still able to track the value relatively quickly when the jumps in the actual fundamental value are not too large . However, in a situation where there is a sharp movement of the value in one direction, such as just after the 500th second in Figure 7, it takes the market maker a long time to adjust the quotes to the right value, because of the relatively high amount of undirected trades. This only occurs after a series of jumps in the same direction. If the fundamental value fluctuates without going strongly in one direction, such as around the 300th second in the simulation, the market maker catches on the current value as soon as it moves back in the direction of the current bid and ask prices. This shows that, as expected, when the proportion of uninformed traders increases, the movement of the market maker's prices towards the fundamental value is

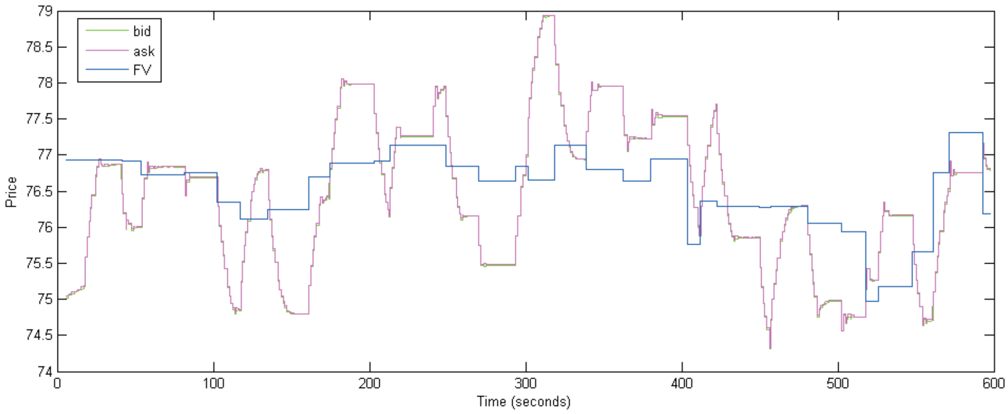

Figure 8: Tracking the fundamental value with one noisy informed trader.

slowed down, i.e. uninformed traders decrease the observability of the fundamental value for the market maker.

Performance with noisy informed traders

Let us now analyze how noisy information diffuses into prices. For this case, we consider noisy informed traders instead of the perfectly informed ones in the experiments presented below.

If the trading crowd consists of a single noisy informed trader, the market maker is able to follow the value from the informed trader, although he is at times rather far from the actual fundamental value due to the additional noise. The additional noise might, for example, first push the observed fundamental value upwards and, after the next jump in the actual value, downwards (or vice versa), resulting in a large change in the *observed* value. In an extreme case, the observed fundamental value might lie outside of the range of possible fundamental values that is considered by the market maker, making it impossible for the market maker to track this value under the current circumstances, until the next jump occurs and the range is re-normalized.

If next to noisy informed traders, uninformed traders interact, the bid and ask prices can be far away from the actual fundamental value because of the influence of both the additional noise in the noisy informed traders' observations and the noise from the uninformed traders' orders. The simulation result in Figure 9 looks more extreme compared to the corresponding simulation with perfectly informed traders as each noisy informed trader observes a different fundamental value. Each of the traders bases his orders on this observed value, and the market

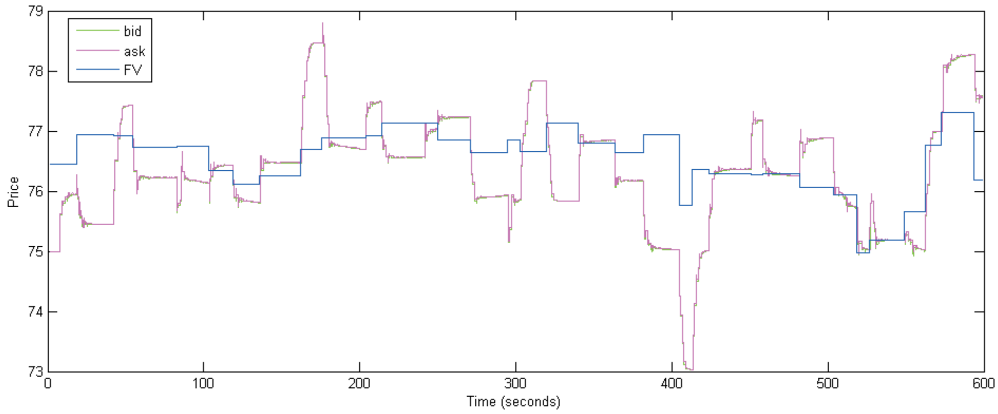

Figure 9: Tracking the fundamental value with 75% noisy informed traders.

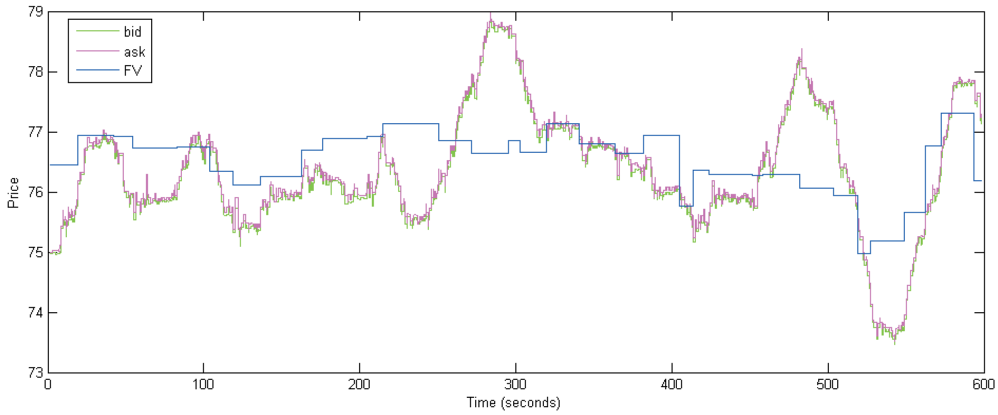

Figure 10: Tracking the fundamental value with 33% noisy informed traders.

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, so all investors will keep submitting orders almost continuously.

By increasing the rate of uninformed traders in the experiments, the observed (nearly) constant bid-ask values, that we observed in case of 100% noisy investors tend to disappear bid-ask prices fluctuating more often (see Figure 10). Further, because of the increased uncertainty the bid-ask spread size increases as well.

Table 1: Samples of bid-ask spread sizes (in cents)

100% perfectly informed	0.93
75% perfectly informed	1.65
50% perfectly informed	1.68
33% perfectly informed	3.52
100% noisy informed	1.33
75% noisy informed	1.65
50% noisy informed	3.19
33% noisy informed	6.73

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, as opposed to a large spread size, it is less likely that the actual fundamental value falls within this range. Informed traders will not submit orders if the market maker's bid and ask prices are set around their observed fundamental value. Thus, since 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. The table shows that the spread size increases as the fraction of informed traders decreases, i.e. the uncertainty increases. Especially the presence of noise in the fundamental value that is observed by the informed trader contributes to a larger spread size.

It is shown in [9] 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 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,

Table 2: Sample cash balance and inventory of the market maker.

Simulation	Cash balance	Inventory	Total wealth	Average per trade
100% perfectly informed traders	2805.72	-36	63.24	0.17
75% perfectly informed traders	-829.49	11	9.48	0.01
50% perfectly informed traders	1721.47	-22	30.11	0.04
33% perfectly informed traders	3945.23	-51	63.62	0.05
100% noisy informed traders	7009.69	-91	55.47	0.15
75% noisy informed traders	309.13	-1	232.65	0.13
50% noisy informed traders	918.35	-11	78.94	0.10
33% noisy informed traders	2370.32	-30	56.12	0.05

which is represented, in our implementation, by a minimum and maximum for the 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

In this model, when quoting bid and ask prices, the market maker satisfies 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 (as noted in [9]). An interesting question is to how large the difference is.

In the experiments, there are no restrictions made on the wealth position and inventory of the market maker, nor on those 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. Further, 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 [9]. We did not study, however, the consequences of such a setting for the purpose 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. Further, 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 behaviour 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 3 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 uninformed investors. 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. In order to answer this question, we analyze situations with populations of more purely informed traders.

In the experiments presented the number of traders is kept relatively small, while in the models in [11] and [9] 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 (so, 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 are either zero or one orders to be processed.

In our continuous model, in which investors are represented as individuals and exhibit autonomous, asynchronous behaviour, 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 noisy informed traders, as perfectly informed traders all observe the same fundamental value, while noisy informed traders

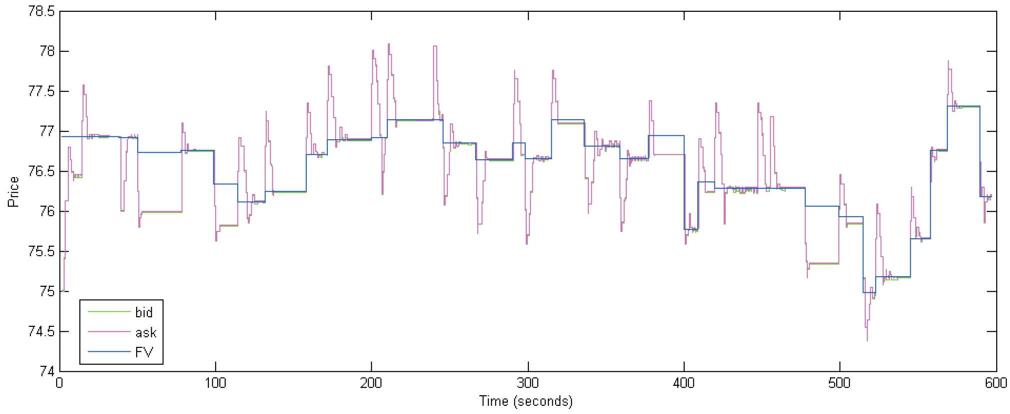

Figure 11: Tracking the fundamental value with four perfectly informed traders.

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. It takes the market maker more time to set the prices close to the fundamental value, as can be seen in Figure 11. Whenever there is a jump in the fundamental value, all four perfectly informed traders immediately respond to this change, and, since they all use the same information, they perform the same action. Because the market maker bases his bid and ask prices on the received orders, this causes an overreaction to any change in the fundamental value. Even a slight change, such as around the 450th second in the simulation, causes a large peak in the direction of the jump because of the immediate response of all traders. Again, such effects are not visible in discrete-time models.

Another side effect is that, because of the large amount of orders received after a change in the fundamental value, the probability density estimate is quickly concentrated in a very small area. This leads to situations where, because the estimate is restricted to such a small area, newly arriving trades no longer have an effect on the probability density estimate, so there is no change in the market prices until the next jump in the fundamental value. This can be observed at several moments in Figure 11, when the bid and ask prices are just below the fundamental value, but following a straight line.

When increasing the number of noisy informed traders in the experiments with 100% informed traders, the bid and ask quotes could move far away, but also closer to the real value (see Figure 12). At some points, the market maker is much less influenced by the noise, be-

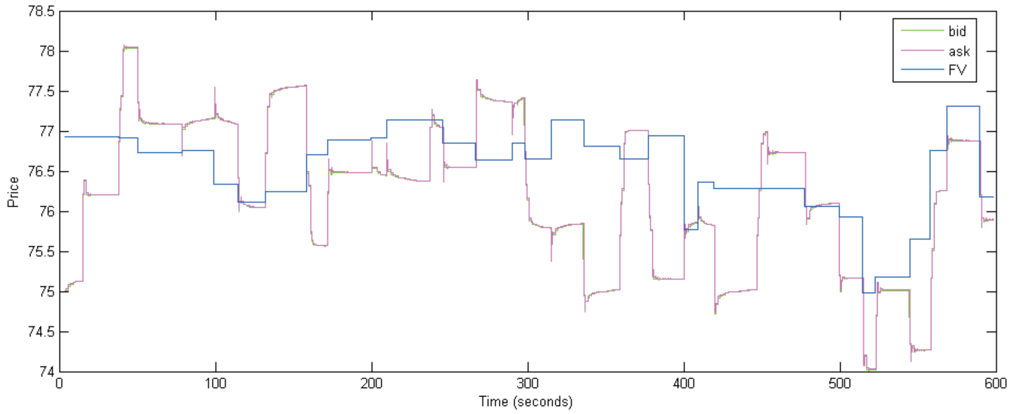

Figure 12: Tracking the fundamental value with four noisy informed traders.

cause the average observed fundamental value with additional (normally distributed) noise of the four market participants is near the actual fundamental value (such as at the end of this simulation run), but at some points the noise drives all four traders in the same direction, causing the market maker to set prices far from the actual fundamental value. This can be observed at around the 350th second in Figure 12. Accordingly, when there is a jump in the fundamental value, in most cases a peak in the resulting bid and ask prices can be observed, similarly to the experiments with a population of four perfectly informed traders. However, the effect is much less pronounced, because a peak appears only if the observed fundamental value jumps in the same direction for all noisy informed traders.

In simulations with a relatively high amount of informed traders order queues cause large peaks appeared in the bid and ask prices that are set by the market maker, whenever a jump in the fundamental value occurs. Figure 13a shows a part of a simulation run with ten traders, of which eight are perfectly informed and Figure 13b shows the size of the order queue in the same part of the simulation. It can be observed that a jump in the fundamental value is followed by a queue of eight pending orders. Since 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 is repeated until the market maker sets

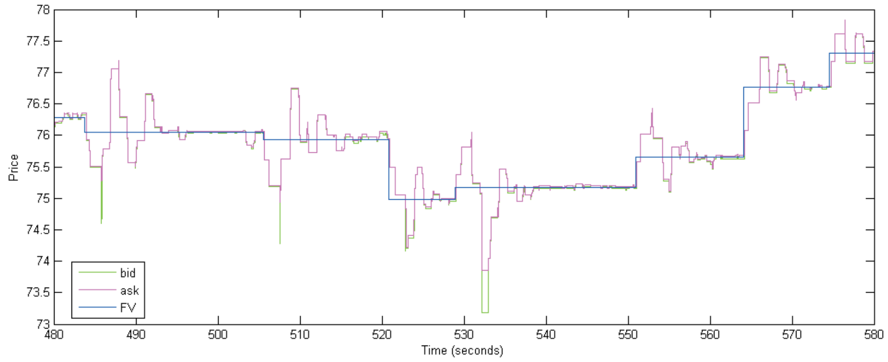

(a)

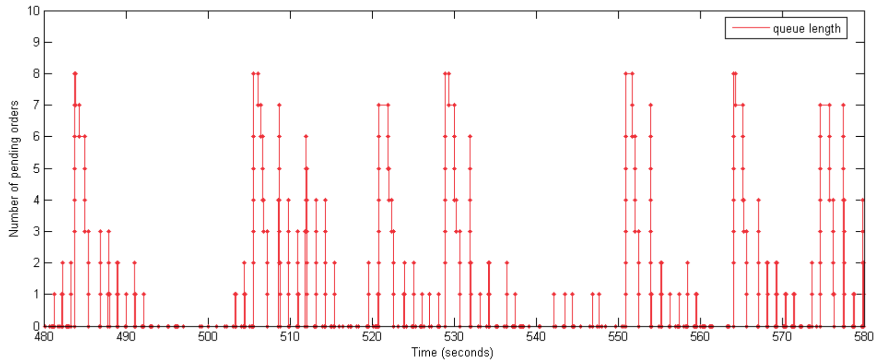

(b)

Figure 13: Market prices (a) and the pending order queue (b) over time, in a market with ten traders, of which eight are perfectly informed.

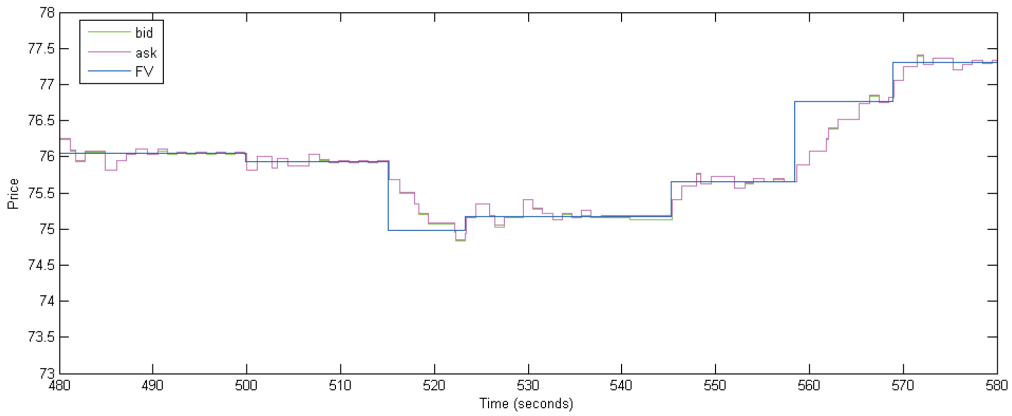

Figure 14: Market prices in the market of [9] with 70% informed traders

his bid and ask prices around the fundamental value. These *overshoots*, caused by the herd-like behaviour 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 model of [9], 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.

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 is why the amplitude of fluctuations increases with the increasing number of informed investors. In order to improve the learning algorithm, the market making algorithm should 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 behaviour and homogeneous setting of informed traders is to introduce more variation in the order placement behaviour of the traders. Introducing different reaction times for the investors, for example, would reduce the length of the queue that arises after a change in the fundamental value of the risky asset occurs. Different reaction times can be the result of different news-sources, location or time needed to handle and analyze the news, etc.

6 Conclusions

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 diverging market dynamics.

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' behaviour influence the market dynamics. Given that most financial markets apply continuous trading sessions, these features should thus be taken into account during modelling.

We found the main difference in the outcomes of the discrete-time and continuous-time simulation in 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 behaviour 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 one hand, informed traders do not necessarily decide to trade all at the same time in reality. We propose to reduce the size of order queues 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. Further, 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 behaviour of the market maker.

References

[1] F. Bellifemine, G. Caire, T. Trucco, and G. Rimassa. JADE programmer's guide, 2002.

[2] K. Boer, A. de Bruin, and U. Kaymak. On the design of artificial stock markets. ERIM report series Research in Management, Erasmus Research Institute of Management, Eras-

mus University Rotterdam, the Netherlands, Jan. 2005. ERS-2005-001-LIS.

[3] K. Boer, U. Kaymak, and A. de Bruin. A modular agent-based environment for studying stock markets. ERIM report series Research in Management, Erasmus Research Institute of Management, Erasmus University Rotterdam, the Netherlands, Apr. 2005. ERS-2005-017-LIS.

[4] K. Boer, M. Polman, A. de Bruin, and U. Kaymak. An agent-based framework for artificial stock markets. In *16th Belgian-Dutch Conference on Artificial Intelligence (BNAIC'04)*, pages 83–90, 21-22 October, Groningen, The Netherlands 2004.

[5] W. A. Brock and C. H. Hommes. A rational route to randomness. *Econometrica*, 65(5):1059–1095, Sep. 1997. ARED-cobweb.

[6] D. Challet, M. Marsili, and Y.-C. Zhang. *Minority Games: Interacting Agents in Financial Markets*. Oxford University Press, 2005.

[7] N. T. Chan and C. Shelton. An electronic market maker. Working Paper AI Memo 2001-005, Massachusetts Institute of Technology, 2001.

[8] S. Das. An agent-based model of dealership markets. In *International Workshop on Complex Agent-based Dynamic Networks, Oxford,*, October 2003.

[9] S. Das. A learning market-maker in the Glosten-Milgrom model. *Quantitative Finance*, forthcoming, 2005.

[10] M. Demarchi and T. Foucault. Equity trading systems in europe: A survey of recent changes. *Annales D'Economie et de Statistique*, 60:73–115, October–December 2000. It briefly describes the structure of the major stock markets in Europe!

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

[12] L. Harris. *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press, 2003.

[13] B. LeBaron. Building the Santa Fe artificial stock market. Wrorking paper, Brandeis University, 2002.

[14] O. Loistl and O. Vetter. *KapSyn Computer Modelled Stock Exchanges, User Manual. Ver. 3.02.*, 2000.

[15] M. Marsili. Toy models of markets with heterogeneous interacting agents. In A. Kirman and J.-B. Zimmermann, editors, *Economics with Heterogeneous Interacting Agents*, number 503 in Lecture Notes in Economics and Mathematical Systems, pages 161–182. Springer, Berlin, 2001.

[16] M. O'Hara. *Market microstructure theory*. Blackwell Publisher, 2002.

[17] M. Raberto, S. Cincotti, S. Focardi, and M. Marchesi. Agent-based simulation of a financial market. *Physica A*, 299(1–2):319–327, 2001.

[18] M. Shatner, L. Mushnik, M. Leshno, and S. Solomon. A continuous time asynchronous model of the stock market. *arXiv:cond-mat*, 2000.

[19] E. Smith, J. D. Farmer, L. Gillemot, and S. Krishnamurthy. Statistical theory of the continuous double auction. Working Paper 02-10-057, Santa Fe Institute, 2002.

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

Publications in the Report Series Research* in Management

ERIM Research Program: "Business Processes, Logistics and Information Systems"

2006

Smart Business Networks Design and Business Genetics
L-F Pau
ERS-2006-002-LIS
http://hdl.handle.net/1765/7319

Designing and Evaluating Sustainable Logistics Networks
J. Quariguasi Frota Neto, J.M. Bloemhof-Ruwaard, J.A.E.E. van Nunen and H.W.G.M. van Heck
ERS-2006-003-LIS
http://hdl.handle.net/1765/7320

Design and Control of Warehouse Order Picking: a literature review
René de Koster, Tho Le-Duc and Kees Jan Roodbergen
ERS-2006-005-LIS
http://hdl.handle.net/1765/7322

A Theoretical Analysis of Cooperative Behavior in Multi-Agent Q-learning
Ludo Waltman and Uzay Kaymak
ERS-2006-006-LIS
http://hdl.handle.net/1765/7323

Supply-Chain Culture Clashes in Europe. Pitfalls in Japanese Service Operations
M.B.M. de Koster and M. Shinohara
ERS-2006-007-LIS
http://hdl.handle.net/1765/7330

From Discrete-Time Models to Continuous-Time, Asynchronous Models of Financial Markets
Katalin Boer, Uzay Kaymak and Jaap Spiering
ERS-2006-009-LIS

---

* A complete overview of the ERIM Report Series Research in Management:
https://ep.eur.nl/handle/1765/1

ERIM Research Programs:
LIS Business Processes, Logistics and Information Systems
ORG Organizing for Performance
MKT Marketing
F&A Finance and Accounting
STR Strategy and Entrepreneurship