A Modular Agent-Based Environment for Studying Stock Markets

KATALIN BOER, UZAY KAYMAK AND ARIE DE BRUIN

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT

REPORT SERIES
RESEARCH IN MANAGEMENT

ABSTRACT AND KEYWORDS	
Abstract	Artificial stock markets are built with diffuse priors in mind regarding trading strategies and price formation mechanisms. Diffuse priors are a natural consequence of the unknown relation between the various elements that drive market dynamics and the large variety of market organizations, findings, however, might hold only within the specific market settings. In this paper we propose a framework for building agent-based artificial stock markets. We present the mechanism of the framework based on a previously identified list of organizational and behavioural aspects. Within the framework experiments with arbitrary many trading strategies, acting in various market organizations can be conducted in a flexible way, without changing its architecture. In this way experiments of other artificial stock markets, as well as theoretical models can be replicated and their findings compared. Comparisons of the different experimental results might indicate whether findings are due to traders' behaviour or to the chosen market structure and could suggest how to improve market quality.
Free Keywords	Computational economics, agent-based modelling, artificial stock markets, behavioural finance, market microstructure, price formation, model verification and validation.

A Modular Agent-Based Environment for Studying Stock Markets

Katalin Boer   Uzay Kaymak   Arie de Bruin

Erasmus University Rotterdam, Faculty of Economics

P.O. Box 1738, 3000 DR, Rotterdam, the Netherlands
e-mail: {kboer,adebruin,kaymak}@few.eur.nl

Abstract

Artificial stock markets are built with diffuse priors in mind regarding trading strategies and price formation mechanisms. Diffuse priors are a natural consequence of the unknown relation between the various elements that drive market dynamics and the large variety of market organizations, findings, however, might hold only within the specific market settings. In this paper we propose a framework for building agent-based artificial stock markets. We present the mechanism of the framework based on a previously identified list of organizational and behavioural aspects. Within the framework experiments with arbitrary many trading strategies, acting in various market organizations can be conducted in a flexible way, without changing its architecture. In this way experiments of other artificial stock markets, as well as theoretical models can be replicated and their findings compared. Comparisons of the different experimental results might indicate whether findings are due to traders' behaviour or to the chosen market structure and could suggest how to improve market quality.

Keywords

Computational economics, agent-based modelling, artificial stock markets, behavioural finance, market microstructure, price formation, model verification and validation.

1 Introduction

Artificial stock markets (ASM's) are primarily designed to help us understand market dynamics. Several ASM's provide a tool to test theoretical hypothesis, these being commonly the efficient market hypotheses and rational expectations hypotheses (Chen and Yeh, 2001; Brock and Hommes, 1997; Loistl and Vetter, 2000). Further, several ASM's try to study what kind of structures and behaviours can reproduce time series with similar characteristics as on real markets (Hommes, 2001; Gaunersdorfer et al., 2001). The question is, however, what are the properties of the time series on the market? The answer is not univocal. Some people argue that prices follow a random walk, others seem to find patterns ("stylized facts") in series. Are those there by chance or is there some predictability in the prices?

In the ASM's, the successfulness of different traders can be studied since traders make different predictions regarding future values. Most of the studies compare successfulness of fundamentalists versus technical analysts, or of random trading agents against learning agents and, indeed, find differences in their wealth outcomes (Hommes, 2001; LeBaron, 2002). Of its own trying to find prosperous strategies is not a scientific goal. However, repeated successfulness of technical analysts over fundamentalists indicates predictability in the time series which reflects inefficiency in the markets. Experiments with inefficient markets might indicate what kind of organizational changes should be made in order to make them more efficient and in this way improve market quality.

Artificial stock markets are however, as LeBaron (2001) states "only toy models that represent a complicated social situation in a highly stylized fashion". Given the difficulty to analyze how strong the connection of market models (theoretical, empirical, experimental or agent-based) to the real markets is, opinions regarding their validity are extremely diversified. As LeBaron notices, the benefits of agent-based ASM's lie in the "large number of parameters for which our priors are extremely diffuse", however, this peculiarity makes them also difficult to validate.

Two important features of ASM's in relation to other market models and real stock markets are emphasized: our **diffuse priors** and the ability to handle **large number of parameters**. Taking into account a large number of parameters, on the one hand, turns market models to contain more representative elements of the de facto microstructure of stock markets, strengthening their validity. Given the changing and large variety of market organizations and the several

"hidden" features, such as details behind price formation mechanisms and traders' decision, on the other hand, imply broadly diffuse prior assumptions. Consequently, the dilemma that one has to deal with, when representing markets is, for which prior to choose for, how valid is the assumption, and how does the choice influence the dynamics.

In ASM's in the literature a common choice is to study call markets and investor type of behaviour. Regarding the price formation mechanisms and trading strategies behind this commons structure even, several interpretations and implementations are given. The diverse possibilities to represent markets and traders, the implementation difficulties and the several unconsidered market features by the current ASM's discussed in (Boer et al., 2005) motivated us to design an ASM framework that aims to be more representative and to cover the missing features.

In this article we introduce the framework that provides a tool for representing several types of markets and an arbitrary number of trading strategies. We design and describe the framework based on the list of organizational and behavioural critical factors that we have identified in a previous article, and on the analysis of several artificial stock markets presented regarding these factors (see Boer et al. (2005)). By designing this framework we aim to achieve multiple objectives. First of all, we would like to experiment with rarely studied market types and behaviours (such as: continuous trading sessions, brokers, market maker based price settings, autonomously and asynchronously acting agents) by incorporating these features in the framework or allowing for their flexible implementation. Further, we aim to replicate, test and validate some of the existing artificial markets on the top of the framework. In this way, the introduced framework can help us to study whether findings of experiments within different market models can globally explain some market dynamics or are the results of the chosen settings.

The outline of this article is as follows. In Section 2 we present briefly the structure of the framework. In Section 3 we elaborate on the organizational features of the framework, pointing to the various market setting possibilities. Then, in Section 4 we discuss the generic behaviour of the different market participants and the possibility to flexibly plug-in various strategy implementations. A short evaluation of the framework regarding its accuracy is provided in Section 5 through some simple implementations and analysis of the experimental results. Discussion and future research finalize the paper in Section 6.

2 The structure of the framework

Trading in a market is organized around three main components: the instruments that are traded, the market participants who handle and the institutional structure behind price formation. The institutional structure, that is market microstructure, directly drives the price formation mechanism and defines the type of a market. Instruments, namely stocks are issued by certain companies, representing part of ownership, and their value, or better said valuation depends therefore on the performance of the company. Being the object of exchange however, they are more or less part of the market where they are traded, their value being influenced as by issuer specific news, as well as by the market organization and by the decision of the market participants. Regarding the market participants, we differentiate two types of them (see Boer and Kaymak (2003) and Boer et al. (2005)), financial traders, like brokers or market makers, who have specific well-defined roles, and investors, who are in fact, not internal part of the market organization, but can observe it, make trading decisions and send their orders to participants on the market place.

According to the discussion above, the framework consists of three main components (Figure 1):

- the marketplace;

- the set of investors;

- an information source.

The existence of brokers and market makers, and their behaviour is strongly related to the market organization: they have special roles and tasks, and their activities can be regulated by trading rules, that are part of the market structure. Investors, however do not have a special role within the market, and therefore function more or less independently. As a consequence, the marketplace represents the institutional structure behind price formation, and includes subcomponents of the financial traders such as brokers and market makers. Further, market specific features of stocks that can be traded on a specific market are also defined here. Market participants who act in fact from outside the institutional structure, that is investors are represented separately. As a consequence, they can trade in multiple markets if they wish so. The information source component is introduced to represent information coming from the issuer and is designed with the aim to generate news related to stocks traded on the market, such as dividends and information regarding the fundamental value.
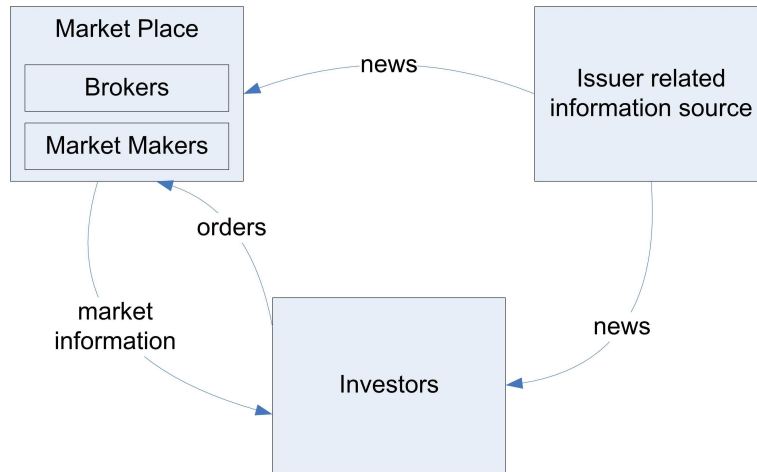
Figure 1: The components of the framework

3 Organizational aspects

Given the various market organizations and trading strategies, we decided not to focus on a specific market structure, but to allow for experiments in multiple trading environments. Therefore we design a framework that can be configured to experiment with various market types and traders. During the design phase we take into account the institutional organization of stock markets and several market models. The list of organizational and behavioural factors identified by Boer et al. (2005) drive the presentation of the framework. When designing the framework we primarily focused on the representation of the perceived ignored features of stock markets, like continuous trading sessions. However, we intended to incorporate the structure of ASM's from the literature in order to replicate and test their results. In this section we discuss how the organizational factors: instruments, orders, market participants, trading sessions, execution systems, protocols and transparency; are represented in the framework, such that it can reflect multiple market structures. Behavioural aspects are focused in the next section.

- **Traded instruments**

  A number of risky stocks can be traded within the framework, all identified with a unique symbol. However, like the majority of the ASM's, for the time being, we focus on experiments with trading one risky stock. Cash supplies are available to express the risk free choice of traders. The information source component that represents the issuer can

generate dividends and fundamental values if required, and inform traders about the new values. The algorithms behind news-generation can be flexibly varied and extended.

- **Orders**

  Both limit and market orders can be placed by traders acting within the framework. Orders are characterized by a stock name, size, side, quoted price and time-stamp.

- **Market participants**

  In contrast to existing artificial stock markets, featuring only investors and (rarely) market makers, in the introduced framework, next to investors and market makers, brokers can interact. The presence of brokers is required, for example, in markets where continuous double-auctions can be conducted, like the NYSE. If necessary brokers can be excluded from experiments, for instance, in order to replicate experiments of ASM's that do not implement them. There is at least one market-maker assigned to each stock, responsible for the liquidity of the stock. If an implementation on top of the framework allows for more market makers taking responsibility for the same stock, the structure reflects dealer-markets. The detailed generic behaviour of the market participants within the framework is discussed in the next section.

- **Trading sessions**

  Since continuous trading sessions are very common on stock markets (Harris, 2003; Demarchi and Foucault, 2000) but are rarely encountered in current ASM's, we focus on implementing a continuous trading mechanism. However, since most artificial markets implement call-auctions, experiments with markets where trading occurs at discrete points in time can be also conducted within the framework.

- **Execution systems**

  We support the representation of two types of continuous market structures: with and without continuous double-auctions. A market with continuous double-auction represents, for example, the trading mechanism on the NYSE. In addition the framework provides the possibility to implement hybrid markets, like continuous markets with call-auctions to determine the opening and the closing prices, given that the organization of markets in this way is very common (Demarchi and Foucault, 2000).

We consider continuous structures as quote-driven markets, where market makers quote bids and asks. Experiments with several continuous execution systems can be conducted within the framework, depending on the implemented and configured bid-ask quoting strategy of the market makers.

Call auctions within the framework are designed to represent price-driven markets. Several call-auction related execution systems within the framework for setting market prices can be flexibly implemented again on top of the framework.

- **Protocols**

  Protocols are related to the type of the trading session and specific execution mechanism applied within the framework. In general, based on the specified market type, they might refer to the tick-size, the upper and lower limit of a bid-ask spread, the number of call-auctions that should be conducted, the time interval between two call-auctions, the specification of events that can trigger call-auctions on hybrid markets (e.g. the intensity of change in prices), the (simulated) length of a trading day, etc.

- **Transparency**

  Regarding the publication of a-priori information: historical price series, the bid-ask quotes of the market maker, and eventually news provided by the information-source component, for example concerning the fundamental value, is available. The limit-order-book is theoretically closed, however, depending on the market-maker's bid-ask quoting strategy the best ask and offer might be visible. In the future we might consider experimenting with open book trading as well.

  Post-trade information, regarding the transactions, is made public to all the participants. The identity of the traders participating in transactions is registered for statistical reasons, but it can be hidden from other participants if wished so.

4 Price formation and behavioural aspects

While in the previous section we focused on the organizational aspects of the markets that can be implemented within the framework, in this section we focus on the "black-box" feature of the markets and describe the design bed of price formation mechanisms and traders. We base the description again on the behavioural factors identified by Boer et al. (2005). In order to allow for the flexible representation of price formation mechanisms and experiments with different trading strategies, the framework incorporates only the skeletons for the three trader types, that is market makers, brokers and investors. Skeletons represent the basic structure and generic behaviour of the traders. On top of the skeletons traders' specific behaviours (strategies) can be implemented. Since we have chosen for an artificial agent-based implementation of the traders, we illustrate them by focusing on their environment-sensing, decision making and acting behaviour (see Russell and Norvig (1995)).

4.1 Order-placing behaviours

Orders are primarily initiated by investors as a result of some portfolio management process. From the point of view of an outsider, an investor performs the following generic behaviour:

- it continually senses its environment;

- makes trading decisions that result in orders;

- contacts financial agents to ask them to execute the generated orders.

What we do not know, however, is the process and reasoning behind trading decisions, that is for us a sort of partially black box. The framework is designed in a way to reflect the generic behaviour of investors (Figure 2). Accordingly, we represent the trading decision of investors as a "black box".

As suggested by the behavioural finance literature, illustrated by the various ASM implementations, and discussed by Boer et al. (2005), arbitrary many strategies might exist behind trading decisions. In general, as theory describes, trading decisions are driven by portfolio management processes (Reilly and Brown, 2003). Accordingly, trading decisions are based on the individual investment strategies that try to meet the particular policy statement guidelines, and reflect the portfolio construction and maintenance decision of the investors resulting in orders.

Figure 2: The generic behaviour of investors

Given that the strategy behind the trading decision is a kind of varying "black box", what we need to take care of when designing the generic behaviour of the investor is:

- how and when to trigger the trading decision, and

- to process the resulted orders.

In the framework, a trading decision can be triggered, whenever:

- confirmation to the investor arrives, that one, or part of its placed orders is executed;

- news is published or event happens on the market that is sensed by the investor, e.g., new bid and ask quotes are published by the market maker, transaction occurs or news related to the fundamental value of a stock is diffused;

- spontaneously, the investor decides to review and manage its portfolio.

The trading decision can result in a set of orders, but investors might decide as well not to place any order for the moment. Focusing on one type of stock for the time being, in the experiments conducted, investors decide to place either a single order or no order at all.

9

The framework contains only the skeleton, that provides the generic investor-type of behaviour, it does not include however a specific trading strategy, but an empty placeholder on top of which different strategies can be implemented. In this way experiments with arbitrary types of investors can be conducted in a flexible way. Various investors can be set to behave, for example, like the traders at SantaFe (LeBaron, 2002), KapSyn (Loistl and Vetter, 2000) or at a sort of "business school" (Chen and Yeh, 2001) trying to maximize a specific utility function, or use arbitrage opportunities like traders in other markets (e.g.; Das (2003), Chen et al. (2001)). Further, if required they could generate orders with random values like traders in (Shatner et al., 2000; Chan and Shelton, 2001) or (Das, 2003), trade a predefined fraction of their portfolio (Raberto et al., 2001; Shatner et al., 2000) and use a variety of fundamental and/or technical strategies to forecast future values. An investor component in the framework is characterized in this way, besides its portfolio, by the particular strategy applied to make trading decisions.

4.2   Order execution behaviours

4.2.1   Order execution by brokers

Brokers on financial markets are primarily entitled to execute orders on behalf of investors. In accordance to this specific role, the decisions they are faced with are discussed by Boer et al. (2005). Similarly to the investors' skeleton the design of the brokers within the framework involves just their generic behaviour, the chosen strategies behind the decision problems being at the users' freedom.

As illustrated by Figure 3, a broker's common behaviour consists of the following main tasks:

- receive orders from investors, and

- eventually, listen to other information;

- decide which received orders to execute first, and how to execute them;

- make transactions if decided to execute orders.

Figure 3: The generic behaviour of brokers

There are several selection scenarios possible for deciding which of the orders to handle next. A broker might select, for example:

- the order with the earliest arrival time (FIFO mechanism); or

- with the best execution probability (considering current market conditions); or

- aggregate and try to execute more orders at once with similar parameters.

An order is selected in view of possible trading mechanisms. In a general case the broker has three choices:

- match orders internally: if, for example, there are other received orders in the order book that clear at a price close to the current market price;

- try to negotiate with other brokers within the market makers' quoted spread, for instance, through double-auction, like at NYSE;

- define a new (improved) limit price for the order based on market conditions in order to submit it to a third party for execution, such as market maker or central matching system.

11

Brokers are allowed to negotiate among each-other only within the specified bid/ask spread. It does not make sense to negotiate outside it, as then it would be more profitable for one of the parties to trade with the market maker. Brokers might apply a variety of negotiation strategies. The decision they take during the negotiation strategy results either in a new quote or the acceptance of a quote made by another market participant. Several decisions has to be taken if brokers choose to negotiate, such as: the offered negotiation price, the values with which the negotiation price should change and the time-length for which they are trying to negotiate. These values might depend on the actual quote of the market maker, the offers that other brokers make for negotiation purposes, the initial limit price of the selected order, etc.. If the broker decides to accept an order, or his placed order is accepted, a transaction price is defined, a deal is made and the transaction price of this deal is published as the new market price.

The way brokers analyze information, and interpret it to select and execute orders, or to define a negotiation strategy can take arbitrary many forms. Therefore we build the architecture of the brokers in such a way to allow for the implementation and experiments with multiple strategies. Since, unfortunately, it is not clear how in reality brokers solve all the decision problems that they face, we have to experiment with a number of possible solutions. Allowing for multiple strategies enables us to study how the brokers' success and the market dynamics depend on the strategy applied. Hence, the framework again, contains the skeleton that provides the implementation of the generic behaviour of broker-agents without a concrete strategy-implementation.

4.2.2 Market makers

We define market makers as financial agents on quote-driven markets with the specific role to provide liquidity for the stocks they are responsible for. Accordingly, when designing the behaviour of market makers we focus on the representation of this particular feature. According to the generic behaviour of market makers described by (Boer et al., 2005), and illustrated by Figure 4, the tasks that market makers need to repeatedly carry out, is to:

- determine bid and ask quotes at which they are willing to trade;

- receive orders and execute them against the quoted bid-ask values if possible;

- enter the received orders into the limit order book if they could not have been (completely) executed.

The question and the main decision problem market makers are faced with, is:

- *when* and *how* to *set the bid-ask quotes* in order to reflect market conditions and to provide liquidity?

The strategy that governs their decision depends on the market organization where they interact, but also on individual and situational characteristics. In principle, they reconsider the bid and ask quotes in two main situations:

- whenever new orders arrive; or

- they perceive the stock they are responsible for is not liquid enough, that is reflected for example, by the fact that no orders arrive after a while.

Additionally, other considerations, like belief that current quotes do not reflect market conditions, can lead market makers to revise the bid and ask.
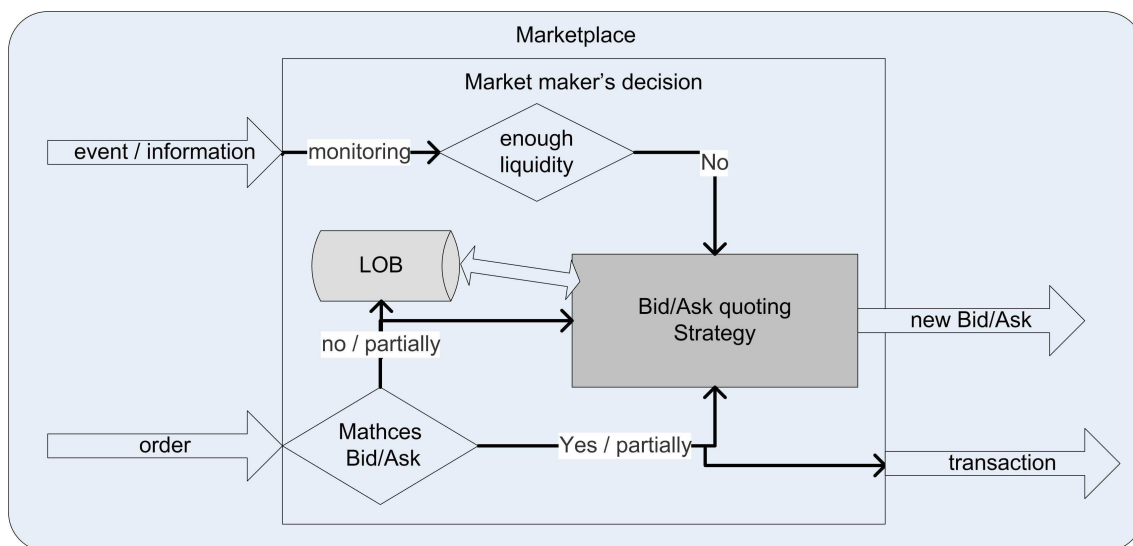

Figure 4: The behaviour of market makers

Market makers apply various strategies to set their bid and ask quotes. In order to allow for experiments with multiple market maker types, the skeleton of the market makers again

13

does not provide a concrete solution to this decision problem, but an empty placeholder to be extended with user-defined strategies. Experiments with several market making algorithms can be conducted and compared in this way: such as bid-ask determinations based on Bayesian learning as implemented by (Das, 2003), or based on position imbalance and a threshold like in (Chan and Shelton, 2001), or simply based on the content of the limit order book, like in KapSyn (Loistl and Vetter, 2000) or in continuous automated auctions.

4.2.3 Automated and equilibrium-based execution of orders

Order execution on automated continuous markets can be in fact considered a special bid-ask quoting strategy of the market maker. Therefore, in this last short section, we would like to focus only on the way equilibrium prices are and/or can be determined during call auctions. In the framework we let this task as well to be carried out by market makers as a specific auctioneer behaviour. Based on the study in (Boer et al., 2005) we define the generic auctioneer behaviour of market maker as consisting of the following steps (Figure 5):

- collect orders during a call;

- determine new equilibrium price based on the received orders;

- execute the orders that fit this price.

Figure 5: Price formation at equilibrium

Again, the strategy used for determining the equilibrium price varies from market to market and, in fact, often is not clear how it is solved in real markets. In the framework therefore, we just provide an empty placeholder for it, and implementations plugged in by users determine the final strategy. Consequently, experiments with several different strategies can be conducted, that can be, for example, the supply and demand intersection as implemented by Raberto et al. (2001) or Brock and Hommes (1998), based on excess demand like in (Chen and Yeh, 2001), or at highest volume, as determined in (Loistl and Vetter, 2000).

5 Verification and validation

In the previous sections we have described the framework and discussed how can be strategies and market types configured on top of it in a flexible way. In this section we focus on benchmark cases in order to test the accuracy of the framework, in various situations. It should be noted, that, since the framework is just an environment to support the design of several market organizations and trading strategies, it has to be always validated together with an implementation. Validation of market models with respect to real markets is, however, extremely difficult if not impossible, given the "diffuse priors" as discussed by LeBaron (2001). Therefore, in this section we deal with the problem of "building the model right", that is verification, and, for the time being, we set aside the problem of validation, that deals with "building the right model" (see Balci (1998)). In this section we aim thus, to test whether the focused implementation functions correctly, as expected. Consequently, within the scope of this article we implement simple analytically tractable models, the functioning of which can be separately tested as well.

5.1 Experimental settings

In the experiments that we describe here there is one stock traded, one market maker, and more investors are represented. Trading takes place in continuous quote-driven trading sessions. Investor-agent(s), representing one or more investors, continuously place orders based on their strategy and the information that they observe (e.g. fundamental value, market prices, bid-ask quotes of the market maker). The market price is formed at the bid or ask quote of the market maker, depending on whether a sell or a buy order was matched against it. Market makers in different experiments define the bid and ask quotes based on the orders they receive and

particular strategies implemented and plugged in on top of the skeletons.

At the beginning of the experiments, there is an initial P_0 market price and s-spread value given, and the bid and ask price quotes of the market maker are set to $s/2$ distance from the market price. In case of markets where only market orders are placed with the same size, the size is set to 1, in the rest of the cases there is no size quoted, so transactions will not occur at this initial quote. In the experiments described below: P_0 is set to 100, and s is set to 1.

5.1.1 The Roll model

The first market type we have conducted experiments with is based on the model introduced by Roll, as described in (Campbell et al., 1997). The reason why we have chosen for this model, is that it is analytical and as such, the properties of the generated price series can be analytically deduced and compared thus to the properties of the simulated results.

In the Roll model, the market maker sets the bid and ask quotes around the fundamental value (FV) he perceives $(FV - s/2$, and $FV + s/2)$. Investors are represented by a single investor-agent, who continuously places market orders with the same probability regarding the trading side. Market prices result from matching the orders to the quotes of the market maker.

We have conducted experiments with two different settings regarding the fundamental value. In the first one, fundamental value is fixed and does not change during the experiments, in the second one, the fundamental value changes randomly following a normal distribution with mean 100 and deviation 1. The market maker senses the changes regarding the fundamental value and adapts the bid-ask quotes accordingly.

5.1.2 The continuous automated model

In the second type of the experiments conducted the strategy applied by the market maker to determine the bid and ask quotes reflects the automated order matching mechanism. The bid is defined as the buy-order from the limit order book with the highest quoted price, and respectively, the ask is defined as the buy-order from the limit order book with the lowest quoted price. Investors continuously place limit orders. The orders that match the quoted bid-ask are executed, the rest is entered into the limit order book. The ask price is always higher than the bid price because orders in the order book cannot be matched, as otherwise a transaction would have been occurred beforehand. If there are no orders on one side of the order book, the quoted

price is of one initial spread difference from the quote of the other side, and there is no volume quoted.

Investors are again represented by a singe investor-agent, who continuously generates random orders. The side of an order is generated with an equal probability to be buy or sell, the number of shares offered or asked is multiple of 100 and below 10000. The price quoted is pseudo-randomly generated, from a normal distribution around a certain mean, with a given standard deviation.

We conducted a number of experiments within this model, where the type of the investors changes regarding the value of the mean and the standard deviation. The mean depends on the "memory" of the market maker, that is the average of the past k market prices, k taking the values of 1, 10 and 100. Three round of experiments are conducted with these mean settings and a predefined standard deviation, that is set to 1.

5.2   Experimental results

For each type of market setting we have conducted five rounds of experiments for time intervals of: 5, 10, 15, 30 and 60 minutes. The size of the time series generated varies from around 1500 (in 5-minute experiments) to around 30000 (in 1-hour experiments). In general 300-500 transactions have been made during a minute. The number of transactions depends of course on the market structure with respect to the time it takes the investor to place new orders and with respect to the time it takes the market maker to handle orders and to set new bid and ask quotes. Given that the behaviour of market maker is computationally more intricate than the investors' behaviour, we delayed investors' decision with 0.1 seconds. By this we tried to avoid to overload the market makers with orders. In the continuous automated model based on best bid and ask there are in general slightly less transactions conducted than in the Roll model. The reason behind this phenomenon is the fact that there is not always a match when limit orders are placed, while in the Roll model a transaction is possible with any received order.

5.2.1   The Roll model

As proven in the analytical model introduced by Roll, described also in Campbell et al. (1997), the bid-ask spread has an impact on the time series properties of the returns, namely negative serial correlation. The explanation in the simplest case is, that if the fundamental price does

not change, the bid and ask quotes will not change either, and as a consequence the measure of change between two consecutive market prices is either 0, or the spread, or the negative spread. There are, thus, never two consecutive increases or decreases in the price, and accordingly, the value of correlation is independent of the spread and equals -0.5. Furthermore, even if the fundamental price changes, it is shown that the serial correlation of returns is non positive, given that changes in the fundamental are serially uncorrelated and independent of the probability of generated order side.

Based on the theoretical findings with respect to the Roll model, we analyzed the auto-correlation of returns generated by the experiments we have conducted with the implemented version of the Roll model. We have found that, indeed, in all the experiments the autocorrelation is -0.5 in case of the simple Roll model, and it is negative as well in experiments conducted with changing fundamental value. This finding proves the correctness of the functioning of the framework with the implementation of the Roll model as a form of bid-ask quoting strategy.

It is interesting to observe that the kurtosis of the returns in the general model is close to 3, and the skewness is close to 0 in all the experiments, indicating that returns are almost normally distributed. The reason behind this phenomenon is the fact that fundamental values based on which bid and ask quotes are determined are randomly generated, taken from a set of normally distributed values.

Further, the autocorrelation of the squired returns is close to 0.2 in all the experiments, suggesting that there might be "volatility clusters" present. Similar results are found by Alexander (2001, pp67) for the empirical data studied. Strong autocorrelation is however, not a proof in itself for volatility clusters, it can just be an implication of it!

Regarding the Roll measure, developed by Roll to estimate the spread when quote date is not available, we found that in the simple model, the estimated spread equals indeed the quoted spread, that is 1. However, in the general model, this is not the case, the estimated spread being more than twice higher than the real spread, which is still 1. The reason behind this phenomenon is that: as fundamental value changes every time, bid and ask quotes change continuously, and thus the difference between two consecutive transactions is very probably high. And, since fundamental value always changes, there are never two transactions conducted at the same quotes.

5.2.2 The continuous automated model

The characteristics of the time series generated by the continuous automated model slightly differ from the results of the Roll model, fact that proves the impact of price formation mechanism on price dynamics. In contrast to the first model, we found that prices in the continuous automated market are strongly autocorrelated, around 0.8-0.9. This phenomenon can be implied by the way limit orders are generated, that is around the last market price. Further, sometimes strong autocorrelation of squared returns can be noticed, feature that indicates the possible presence of volatility clusters.

Such as in the Roll model, negative autocorrelation of return values can be observed again. The value in this case is around -0.2. This findings strengthens the observation that the bid-ask has a negative impact on returns.

We have also looked at the estimated spread using the Roll measure and found that it is lower with around 0.2 than quoted spread. Similar findings are presented in the experiments described by Theissen (2000), pointing out that the Roll measure is not an accurate measure of the transaction costs. The possible explanation behind this phenomenon is that not all the bid-ask quotes placed end immediately in transactions, based on which the estimated spread is calculated. Further, transactions can be fractions of quoted volumes. Interestingly, the quoted spread is around 0.5 in all the experiments.

We have additionally conducted experiments, in which we have adapted the standard deviation to the deviation of the historical time series of the market prices in similar way as the mean. As expected, the limit order prices and as a consequence, the market prices quickly converged to a single price in that case.

6 Discussion and further research

Based on the experiments we can say that the framework functions correctly with the implementations considered, as it behaves in similar way as the analytical form of the implemented models. Real markets are, however, much more intricate than analytically tractable models and, given that their dynamics is not understood, validating the framework against them is extremely difficult, if not impossible. On the other hand, the complex feature of market dynamics is the reason why models are needed to understand them. Although, we have presented and tested

simple implementations in this paper for verification purposes, our aim with the framework, as pointed out by the way it is designed is to conduct experiments in more realistic trading environments.
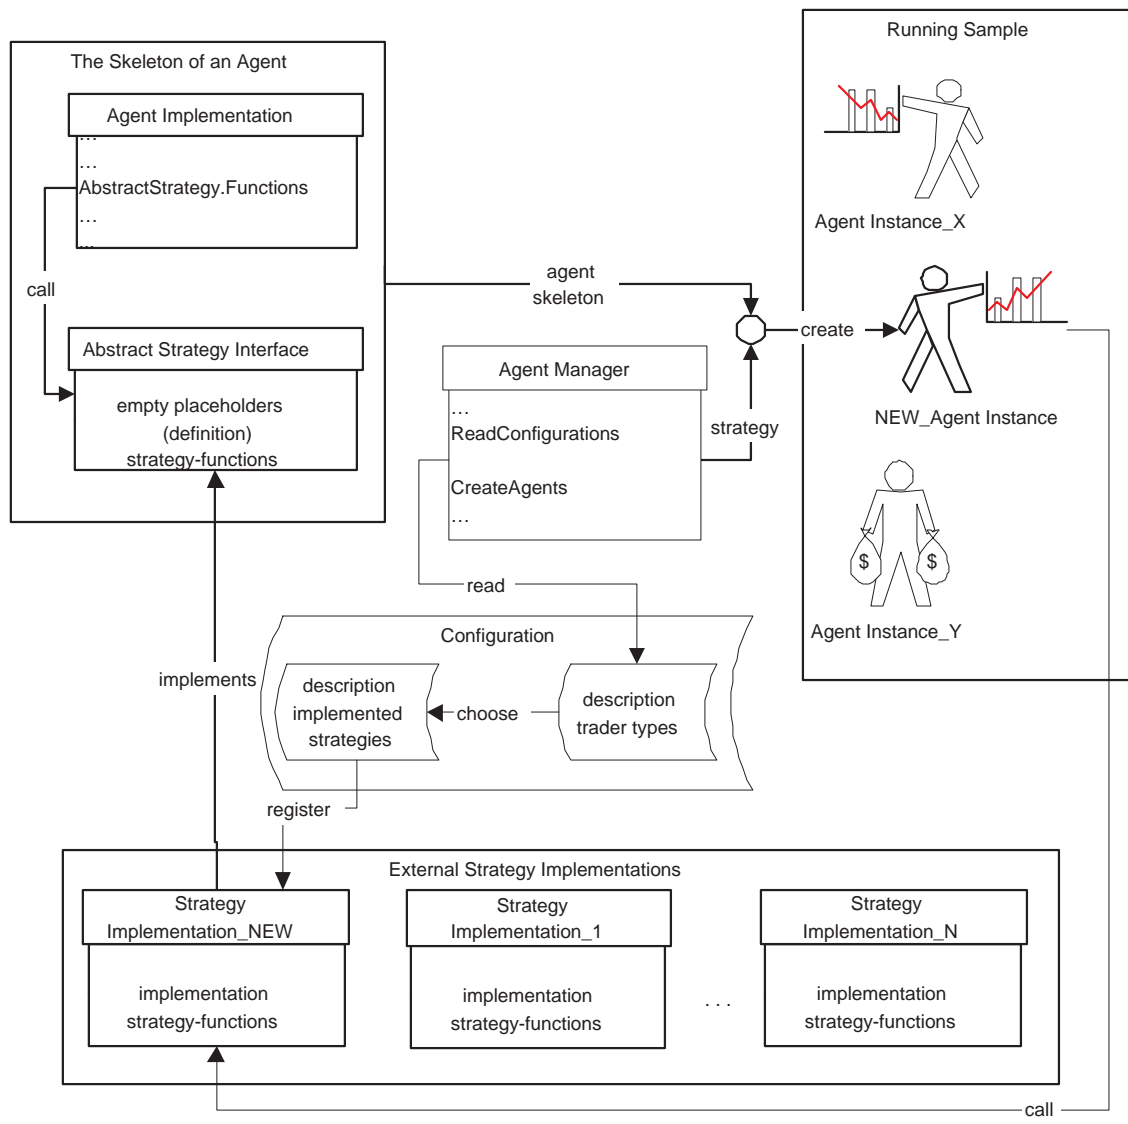

Figure 6: Sample skeleton with implemented strategies

We believe that in order to be able to provide acceptably accurate explanation and analysis of the workings of the financial markets several features of them should be represented. This is why we have chosen for an agent-based micro-simulation approach based on market microstructure literature when designing the artificial stock market framework. In order to allow for the representation of "diffuse priors" regarding the workings of financial markets we do not choose

for a certain market mechanism, but for a framework with the help of which experiments with arbitrary many types of trading strategies in various trading environments can be conducted and compared (Figure 6).

During the design of the framework aimed for studying stock markets that we propose in this paper we have striven to represent the perceived shortcomings of ASM's in the literature, discussed in a previous article (see Boer et al. (2005)). In accordance, the framework allows for continuous trading, asynchronous and autonomous decision making and considers the different roles traders have to fulfill in the market. In contrast to most of the ASM's in the literature, the framework is not centralized. Traders are not centrally directed, but are individual, autonomous elements. They decide when to place an order in function of the type of the market where they interact. Autonomy results from the agent-based implementation that we have chosen.

The framework we propose in this article allows for testing for previous findings in the literature; for studying how different market structures with the joint influence of various types of agents affect market characteristics; and for analyzing whether some behaviours can be more successful than others in certain environments. Further, it allows for studying which features of the market prices are due to learning, adaptation and which are coming from the structure of the market itself. All these studies serve to test the quality of markets having different structures and might indicate the changes that need to be made in order to improve market quality.

References

Alexander, C., 2001. Market Models: A Guide to Financial Data Analysis. John Wiley and Sons, LTD.

Balci, O., 1998. Handbook of Simulation. Principles, Methodology, Advances, Applications, and Practice. John Wiley and Sons, Inc., Ch. Verification, Validation and Testing, pp. 335–393.

Boer, K., de Bruin, A., Kaymak, U., 2005. On the design of artificial stock markets. ERIM Report Series Research in Management (ERS-2005-001-LIS).

Boer, K., Kaymak, U., 15-16 September, Porto, Portugal 2003. Microsimulation of artificial

stock markets based on trader roles. In: International Workshop on Data Mining and Adaptive Modelling Methods for Economics and Management (IWAMEM-03). pp. 61–72.

Brock, W. A., Hommes, C. H., Sep. 1997. A rational route to randomness. Econometrica 65 (5), 1059–1095, aRED-cobweb.

Brock, W. A., Hommes, C. H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics and Control (22), 1235–1274.

Campbell, J. Y., Lo, A. W., MacKinlay, A. C., 1997. The Econometrics of Financial Markets. Princeton University Press.

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

Chen, S.-H., Lux, T., Marchesi, M., 2001. Testing for non-linear structure in an artificial financial market. Journal of Economic Behaviour and Organization 46, 327–342.

Chen, S.-H., Yeh, C.-H., 2001. Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market. Journal of Economic Dynamics and Control (25), 363–393.

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

Demarchi, M., Foucault, T., October-December 2000. Equity trading systems in europe: A survey of recent changes. Annales D'Economie et de Statistique (60), 73–115.

Gaunersdorfer, A., Hommes, C. H., Wagener, F. O., 2001. Bifurcation routes to volatility clustering. Discussion Paper TI 2001-015/1, Tinbergen Institute.

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

Hommes, C. H., 2001. Financial markets as nonlinear adaptive evolutionary systems. Quantitative Finance 1, 149–167.

LeBaron, B., 2001. A builder's guide to agent based financial markets. Quantitative Finance 1 (2), 254–261, draft.

22

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

Loistl, O., Vetter, O., 2000. KapSyn Computer Modelled Stock Exchanges, User Manual. Ver. 3.02.

Raberto, M., Cincotti, S., Focardi, S., Marchesi, M., 2001. Agent-based simulation of a financial market. Physica A (299), 319–327.

Reilly, F. K., Brown, K. C., 2003. Investment Analysis and Portfolio Management. South-Western.

Russell, S., Norvig, P., 1995. Artificial Intelligence. Prentice Hall, Upper Saddle River, New Jersey.

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

Theissen, E., 2000. Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets. Journal of Financial Markets 3, 333–363.

Publications in the Report Series Research* in Management

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

2005

On The Design Of Artificial Stock Markets
Katalin Boer, Arie De Bruin And Uzay Kaymak
ERS-2005-001-LIS
http://hdl.handle.net/1765/1882

Knowledge sharing in an Emerging Network of Practice: The Role of a Knowledge Portal
Peter van Baalen, Jacqueline Bloemhof-Ruwaard, Eric van Heck
ERS-2005-003-LIS
http://hdl.handle.net/1765/1906

A note on the paper Fractional Programming with convex quadratic forms and functions by H.P.Benson
J.B.G.Frenk
ERS-2005-004-LIS

A note on the dual of an unconstrained (generalized) geometric programming problem
J.B.G.Frenk and G.J.Still
ERS-2005-006-LIS

A Modular Agent-Based Environment for Studying Stock Markets
Katalin Boer, Uzay Kaymak and Arie de Bruin
ERS-2005-17-LIS

Keywords Lagrangian duality, cone convexlike functions
J.B.G. Frenk and G. Kassay
ERS-2005-019-LIS