

Towards modeling securities markets as a society of heterogeneous trading agents

Paulo André Lima de Castro¹ and Simon Parsons²

¹ Technological Institute of Aeronautics - ITA,
São José dos Campos , SP, Brazil
`pauloac@ita.br`

² Brooklyn College, City University of New York
Brooklyn, NY 11210, USA
`parsons@sci.brooklyn.cuny.edu`

Abstract. In recent article, Farmer and Foley [3] claimed that the agent-based modeling may be a better way to help guide financial policies than traditional mathematical models. The authors argue that such models can accurately predict short periods ahead as long as the scenario remains almost the same, but fail in times of high volatility. Another real world problem that is rarely addressed in agent-based modeling is the fact that humans do not make decisions under risk strictly based on expected utility. This context inspired the goal of this work: modeling trading agents to populate an artificial market and use it to predict market price evolution in high and low volatility periods. We developed a set of simple trading agents and executed a set of simulated experiments to evaluate their performance. The simulated experiments showed that the artificial market prediction performance is better for low volatility periods than for higher volatility periods. This observation suggests that in high volatility period trading agent strategies are influenced by some other factor that is not present or is smaller in other period. These facts lead us to believe that in high volatility period human agents can be influenced by psychological biases. We also propose in this paper one simple trading agent model that includes prospect theory concepts in its decision making process. We intend to use such model in future work.

1 Introduction

Farmer and Foley [3] stated that agent based modeling could be a better way to help guide financial policies than traditional models. They grouped such traditional models in two big groups: (1) empirical statistical models that are fitted to previously collected data and (2) dynamic stochastic general equilibrium. They argue that the first group methods can successfully forecast short periods ahead:

“as long things stay more or less the same”

but they fail when there are great changes in the market scenario. The second group methods adopt convenient assumptions, such as:

“... assume a perfect world...” [3]

that simplify the problem. This way, they avoid too much complexity, that could make such problem cumbersome or intractable mathematically. However, the authors [3] claim these assumptions can make such models almost useless in high volatility periods, because these assumptions would be far from reality at the time. In fact, as stated by Phelps et al. in [9]:

“...in traditional mechanism design problem, analytical methods are used to prove that agents’ game-theoretically optimal strategies lead to socially desirable outcomes... however, there are many situations in which the underlying assumptions of the theory are violated due to the messiness of the real-world...”

This real-world messiness makes analytical methods hard to use or even impossible. However, the acceptance of suboptimal solutions and the use of iterative refinement methods can hopefully treat this complexity. In fact, significant research work has been carried out in automated mechanism design to overcome the complexity of creating mechanisms with some desirable features for situations inspired by real-world. As an example, Niu et al in [6] simulate agents able to trade in several possible markets.

However, several problems may be identified in agent-based modeling. For instance, it is hard to know how to specify the rules agents should use to make their decisions. Furthermore, it is possible that in volatile periods the rules are different or at least, slightly altered by components that are not present in normal periods. In order to address this question we developed a set of simple trading agents and simulated an artificial stock market in order to predict market price evolution. The rest of this paper is organized as follows. The next section, 2, describes our simple artificial market model and the trading agents that were used in the simulated experiments. These experiments are explained in section 3.1 and their results are presented in section 3.2 and analyzed in section 3.3. As a result of this analysis, we propose a new approach to modeling trading agents in section 4. It is interesting to note that the main motivation for such approach is reduce the market price prediction error by a better description of how human traders act rather than achieving better financial results in trading.

2 Our Simple Artificial Market Model

In this section we describe our simple market model.

2.1 Overview

Our approach for modeling markets is based on the following assumptions. The market price behavior is defined by the interactions among trader agents, i.e. their buy and sell orders. The trading agents’ strategies may be classified in two big groups: fundamentalist and technical strategies. The first group assumes that

the stock prices reflect the company's economic fundamentals, such as profit, market share and so on. The second group assumes that stock prices change according to some patterns and therefore it is possible to identify price trends analyzing past price behavior. Furthermore, the time is modeled as a discrete value that increases through the simulation session. The amount of resources traded by the agents and their orders define the stock price at each instant t as described in section 2.2.

We implemented three types of traders:

- **fundamentalist** traders, who have a fixed idea of the value of a good based on historical data;
- **technical** traders, who trade when the direction of price change alters (so, for example, they sell when the price stops raising); and
- **market makers** who provide liquidity to the market.

Market makers use the last market price to establish a buy order at that last price minus a certain spread and a sell order at last price plus the same spread. This way market making agent provides a lower and upper limits to the price with the result that there is always an trader who is prepared to trade near the current price. (This is exactly the role of market makers in real markets, ensuring that buy or sell orders do not go unmatched if they are made at sensible prices.)

In our model, we compare the price defined by our artificial society, a set of fundamentalist, technical and market making agents, with actual prices obtained in real stock exchange. The difference between the simulated and actual price is a prediction error (section 2.3). We use an algorithm based on hill climbing algorithm to adjust the artificial society parameters in order to reduce this prediction error, as detailed in section 2.4.

2.2 Market Price Formation

The price predicted by the artificial market at a given point in time, \overline{P}_t , is determined by the buy and sell orders made by the set of trader agents. The market acts as a continuous double auction, and the clearing process is performed by the Four heap algorithm described in [11]. In order to execute a deal, the sell price needs to be lower than the buy price and the transaction price is defined as the average of both prices³. The transaction volume is the smaller volume, but higher volume order remains in the book for posterior execution, see [11] for further details.

The market price for a given instant of time is defined as the average of all transaction prices weighted by the volume of each transaction. That way, one agent that makes a higher volume order is more relevant to the market price formation than another agent that submits small volume orders. One order is defined by its price, purpose (sell or buy) and volume. For simplicity, the volume is defined as an integer number of shares.




Fig. 1. An example simulation session showing price evolution and error. The blue line is the internal price, the price predicted by our market, the black line is the external price, the actual price, and the red line is the error for each day.

2.3 Prediction Error

The absolute difference between the price defined by the simulated transactions, that we call *internal price* and the price observed in the corresponding instant t at the real market, the so called *external price*, is the prediction error for a given instant of time t . The figure 1 presents a example of simulation session with historical prices from one real market. In this run, as in all our experiments, we use historical data from a real market. The price in that market on a given day $t - 1$ is the external price at $t - 1$. The traders use this price to determine what they will do on day t , and the buy and sell orders that they decide to place are used by the market to make a price prediction, the internal price, for day t . The difference between this price and the actual, external, price on day t is the error on day t .

This gives us an instantaneous, or daily, prediction error. However, the prediction error of a period of time is much more relevant than the daily error in order to compare one artificial market specification with another. Thus we look at the cumulative error over a period.

More formally, we define the prediction error at a given instant t , as:

$$E_t = |\overline{P}_t - P_t| \quad (1)$$

where \overline{P}_t refers to the price predicted by one artificial market at instant t , while P_t refers to the price observed in the real market at that time. For a given time period, we define the **session prediction error** (E), as the sum of the quadratic error at each round:

$$E = \sum_{t=1}^N (\overline{P}_t - P_t)^2 \quad (2)$$

If one artificial market specification M provides a smaller session error E , than another artificial market specification M' , then we may say that artificial market M is a better description or predictor than M'

It is worth noting that any change in the market specification does not alter traders' strategy, but their relevance to the market price definition as described in section 2.2. Given any trading strategy, it is possible to perform market adjustment, and such a process is described in section 2.4. We describe our trader model and trader agent optimization in section 2.5.

2.4 Artificial Market Adjustment

We use the fact that traders with higher volume have more relevance to the market price formation as described in section 2.2 to adjust the market population (i.e., the set of the agents) to fit data previously observed in real markets. For simplicity, each agent type has just one instance, and it trades one specific share quantity at each round. The **artificial market specification** is defined by three parameters: the share quantities of each one of the three kinds of agents:

³ The so-called $k = 0.5$ double auction market [4].

fundamentalist, technical and market making agents. The objective function is the session prediction error, defined in equation 2.

It is very hard to know a priori how a change in one of the specification parameters may affect the predicted price \bar{P}_t or the session error E . As a result, we used machine learning to discover a good specification. The specific approach we adopted was random-restart hill-climbing, a simple variant of the common hill climbing method [10], that uses a different random starting point at each time it finds a local minimum for the objective.

2.5 Trader Models

Trading agents are responsible for deciding what buy or sell orders to submit, doing this when their specific trading strategies indicate that this is the right thing to do, and making offers at a price determined by their strategy. For simplicity, each agent trades just one stock. In fact, our entire market models the trade in just a single stock. As mentioned above, our trader agents can be classified as **technical**, if they decide based on price and/or volume time series, or **fundamentalist**, if they decide according to information related to the company performance in its market, e.g., profits. In order to avoid unmatched orders, we also implemented **market makers** as described below.

Market Makers The market maker is responsible for making buy and sell orders in order to facilitate trading at every instant. The presence of market makers is important to guarantee an internal market price for each instant — without market makers it is possible that all the agents would decide to buy (or sell) leaving no sellers (buyers) and preventing trading from happening. With no trades, no price is defined since, as explained in section 2.2 the price is determined by a set of transactions weighted by the volume of each transaction. The price at which the market maker places an order is defined by the previous day's price (remember that we clear the market once a day, so this price is the price at the previous instant) plus a **spread**, a small percentage (in the case of a sell order) or the previous day's price minus the spread (in case of a buy order). Therefore, the market maker defines a lower and upper limit for the price. The spread was defined as 0.5% in our simulated experiments. However, the internal price is really defined by the technical and fundamentalist agent's orders and their respective volumes.

Technical Traders There are many technical strategies used in the stock market [2]. One of the simplest and most well-known strategies is the moving average (MA), and this is what we adopted. The moving average index tries to identify trends in stock prices. The average is defined by an observation period, usually defined between 14 and 60 days, and a calculation method that can be simple average (sum of all prices and divided by the number of values) or an exponential average that gives more relevance to newer prices rather than older prices. The moving average is interpreted using graphics with lines of moving average

and prices. The moving average line is a resistance for high trends and down trends. When prices are in high trend (or down trend) and the price line crosses the moving average line, it indicates a reversal of the trend. Therefore, when the moving average is crossed by the price line in a high trend, it means the price is temporarily rising, and this is a sell signal. Similarly when the moving average line is crossed by the price line in a down trend, it means that price is temporarily falling, and this is a buy signal. We used a version of MA that was adapted to provide an order price based on the market price on the previous day.

Fundamentalist Traders The modeling and implementation of fundamentalist traders may be much more complex than technical traders [1]. The data used by fundamentalist traders may be economic information about the company (such as profit, dividends policy and so on), about the economic sector (size and growth projections) and/or general economy (growth projections, volatility analysis, etc.). In this work, we used a very simple approach to fundamentalist trading based on the profit time series. We used this to predict a profit at a given future time t through simple linear regression. Then we assumed that the price/profit relation holds over this period, so it is possible to estimate the fundamental price at time t .

3 Simulated Experiments

In this section we describe the experiments that we carried out, and discuss the results we obtained.

3.1 Description

We performed a set of simulated experiments in order to test our simple model and evaluate the quality of predictions using real market data of several years. We implemented our trading agents using an adapted version of auction simulator called JASA [7]. JASA runs over an agent-based modeling toolkit called JABM [8]. The real market data includes nine years of Intel stock prices between 2003 to 2011 from the Nasdaq exchange. Figure 2 presents two graphs each one shows the simulated price, actual market price (or external price), the error at each round and the cumulative error for the whole simulation session. The left graph represent the results in a low volatility period and the other in high volatility period. Note that the cumulative error line is steeper in the graph for the high volatility period.

3.2 Results

We expected that our artificial market would achieve smaller errors in low volatility periods than in more volatile periods. We believe that this may happen because in volatile periods, human traders may let their emotions and feelings guide




Fig. 2. Examples of simulation sessions in low and high-volatility period. Simulated and actual prices, and daily error, is as in the previous figure. Cumulative error is given by the dotted line.

Table 1. Simulation Results of years 2003 to 2011.

Year	Variance	Volatility	Error	Performance
2003	542.3	high	344.3	good
2004	365.1	high	501.0	bad
2005	167.7	low	407.1	bad
2006	284.1	low	289.3	good
2007	269.9	low	349.0	good
2008	1116.6	high	263.4	good
2009	542.1	high	449.8	bad
2010	254.5	low	296.3	good
2011	292.6	low	234.5	good
High volatility average	641.5	high	389.6	bad
Low volatility average	253.7	low	315.2	good

their decisions, and so the prices from the real market would deviate from the prices predicted by our agents. The simulation results are presented in table 1. We simulated nine years of operation using the historical daily price of Intel Corporation stock on the Nasdaq Exchange. The *variance* of arithmetic returns were calculated for each year and used to classify the years from 2003 to 2011 as high or low *volatility* years. We simulated artificial markets as described in section 2.4 and calculated the smallest session error for each year after several execution of the artificial market adjustment process as explained in section 2.4. The smallest session *error* achieved for each year are presented in table 1. According to such errors, we classified artificial market *performance* as good or bad. We used the overall average error (348.3) as a guide to distinguish between good and bad performance. The low-volatility average error is 315.2, while the high-volatility average error is 389.6, as shown in table 1 (and this is heavily affected by the low outlier in 2008).

3.3 Analysis

The predictions made by the artificial market represented a good performance in four of five low volatility years, but only two of the four high volatility years. For the last two rows of table 1, we can see that the prediction performance (315.2) is better for low volatility periods than for high volatility periods (389.6). Therefore, we can conclude that the predictions made by our artificial market presented significantly better performance in low volatility periods than in high volatility periods, as we expected. One may argue that it is according to common sense, because more volatile periods are usually harder to predict. However, it is important to remark that as argued by Farmer and Foley [3], we believe that agent based models may bring more accurate predictions than traditional models specially when there are big changes in the market, but it will require better understanding about how agents reason in high volatility periods. This fact leads us to believe that in high volatility period trading is influenced by

some other factor that is not present or at least it is weaker in lower volatility periods. We believe that human agents can be influenced by psychological biases as described in Kahneman and Tversky’s work [5] in high volatility periods. We discuss this idea and how it can be used in trading agent modeling in section 4.

4 Prospect Theory and Trading Agent Modeling

One real-world problem that is not often addressed in artificial markets is the fact that human beings don’t make decisions under risk strictly based on expected utility. In fact, some alternative models are available, for example Prospect Theory. Here we describe the theory and how it may be applied in agent trading.

4.1 Prospect theory

Prospect theory was proposed by Kahneman and Tversky [5] and it can be seen as alternative to model and describe human decision making under risk. Kahneman and Tversky claim that several observed behaviors cannot be predicted or explained by expected utility theory [5]. For instance, people usually underweight outcomes, which are merely probable in comparison with outcomes that are obtained with certainty. This tendency is usually called the certainty effect, and contributes to risk aversion in choices involving sure gains and to risk seeking in choices involving sure losses. Another effect pointed by Kahneman and Tversky, describes the observed preference in their experimental studies with human beings for guaranteed small gains over uncertain large gains, and conversely for uncertain large losses over small certain losses, called reflection effect.

Auctions can be seen as decision making under risk, including continuous double auctions as observed in stock market. Prospect theory was developed for prospects with monetary outcomes and stated probabilities, but it can be extended to more complex options. The theory establishes one phase of editing and a subsequent phase of evaluation and selection. The **editing phase** consists of an analysis of the offered prospects, which may eliminate some possible outcomes to create simpler representation of the initial prospects. In the **evaluation phase**, the remaining prospects are evaluated through a value function proposed by the authors and the highest value prospect is chosen.

4.2 Trading agent modeling

$$\begin{aligned}
 Outcome &= (M_t - P_t * \theta_t) + (Q_t + \theta_t) * P_{t+1} \\
 &\quad - [M_t + P_t * Q_t] \\
 Outcome &= (P_{t+1} - P_t) * (Q_t + \theta_t)
 \end{aligned} \tag{3}$$

As described in section 2.5, our trading agents are able to define and submit orders to the market. Furthermore, each trading agent is able to make price prediction and use it to define one order among three possibilities: buy, sell or hold. As explained in section 2.4, the order volume is not defined by the agent

itself, but by the artificial market adjustment process. Furthermore, the trading agent selects the option that seems to him that it is going to bring the best outcome. Such an outcome is the difference between the position at time t and the next time, after an order is executed. This outcome may be calculated as stated in equation 3, where P_t refers to the price, M_t is the amount of money, Q_t is the number of shares at time t and θ_t is the number of shares, positive for buy orders or negative for sell orders, to be transacted by the order given at time t . Each order defines changes in Q_{t+1} and the market behavior defines the change in P_{t+1} . This price cannot be defined a priori, but it is estimated by our trading agents ($\overline{P_{t+1}}$), so we can calculate $\overline{P_{t+1}} - P_t$. Any order may bring different outcomes according to the market price in the next round P_{t+1} . In order to establish prospects of the possible orders, we would need to determine the probabilities given each possible outcome considering two possible decisions: buy or sell. The future price $\overline{P_{t+1}}$ is a continuous value and θ_t is a non-linear parameter, it is dependent of the trading strategy and the artificial market adjustment, so the outcome is itself a continuous non-linear function which would require a probability density function to represent the associated probabilities. The definition of such functions would be extremely complex or even impossible.

Therefore, we initially intend to use a simple approach based on discretization and the arbitrary reduction of the possible outcomes. The future price P_{t+1} may be approximately equal to the estimated price $\overline{P_{t+1}}$, i.e., P_{t+1} is in the interval $[\overline{P_{t+1}} - \delta, \overline{P_{t+1}} + \delta]$ (likely outcome). Furthermore the real price may be slightly higher or lower than the estimated price. It is slightly higher, if it is in the interval $(\overline{P_{t+1}} + \delta, \overline{P_{t+1}} + 2 * \delta]$. It is slightly lower, if it is in the interval $[\overline{P_{t+1}} - 2 * \delta, \overline{P_{t+1}} - * \delta)$. Assuming that the provided estimated $\overline{P_{t+1}}$ is usually close to the real price P_{t+1} and it is not biased to higher or lower values, we can assume a higher probability to the first scenario and two equal and smaller probabilities to the other two scenarios. The parameter δ may be defined according as percentage of the initial value of the stock and the probability that real price is outside the interval $(\overline{P_{t+1}} - 2 * \delta, \overline{P_{t+1}} + 2 * \delta)$ is assumed to be zero.

We believe that using these simplistic but reasonable assumptions, it is possible to construct one prospect for each possible action of the trading agent. Such a *prospect construction phase* takes place before the editing and evaluation phases and provides the information needed for them. The selected prospect in the evaluation phase is assigned to one action, which will be selected by the extended trading agent as his decision. We intend to use the proposed trading agent modeling based on prospect theory in future work.

5 Conclusions and Further Work

Traditional economic models include dynamic stochastic general equilibrium models and empirical statistical models that are fitted to previously collected data. These models may successfully forecast short periods ahead or

“as long things stay more or less the same” [3]

but they are not reliable for high volatility periods. Agent based modeling may become a better way to help guide financial policies, than traditional models according to some researchers [3]. However, several problems may be identified in agent based modeling. For instance, it is hard to know how to specify the rules agents should use to make their decisions. Furthermore, it is possible that in high volatility period the rules are different or at least, slightly altered by components that are not present in normal periods. In order to address this question we developed a set of simple trading agents and simulated an artificial stock market in order to predict market price evolution.

The simulated experiments showed that the artificial market prediction performance is better for low volatility periods. Furthermore, this observation suggests that in volatile period trading, agent strategies are influenced by some other factor that is not present in other periods. We believe that in volatile periods human agents can be influenced by psychological biases as described in Kahneman and Tversky's work [5]. Prospect theory may be seen as an alternative account of individual decision making under risk. The theory was developed for simple prospects with monetary outcomes and stated probabilities, but as the authors claims it can be extended to more involved choices [5].

We proposed a simple trading agent based on prospect theory that can be used to simulate artificial markets with this kind of agent. The model uses a prospect construction phase to be used within the trader agent reasoning process. Such phase happens before the two traditional prospect theory phases: editing and evaluation (section 4). We intend to use the proposed trading agent modeling based on prospect theory in future work to verify if artificial markets populated with this kind of agent may achieve better prediction performance.

References

1. Carlos H. Dejavite Araújo and Paulo Andre L. Castro. Towards automated trading based on fundamentalist and technical data. In *Proceedings of 20 th SBIA. LNAI*, pages 704–715, São Bernado do Campo, Brazil, 2010. Springer-Verlag.
2. Paulo Andre Castro and Jaime Simao Sichman. Towards cooperation among competitive trader agents. In *Proceedings of 9th ICEIS.*, pages 138–143, Funchal, Portugal, 2007.
3. J. Doyne Farmer and Duncan Foley. The economy needs agent-based modelling. *Nature*, 460:685–686, August 2009.
4. D. Friedman. The double auction institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 1, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.
5. Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, March 1979.
6. Jinzhong Niu, Kai Cai, Simon Parsons, and Elizabeth Sklar. Some preliminary results on the dynamic behavior of traders in multiple markets. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, Vancouver, British Columbia, 2007.

7. Steve Phelps. JASA: Java Auction Software API, 2007. <http://sourceforge.net/projects/jasa/>.
8. Steve Phelps. Applying dependency injection to agent-based modeling: the jabm toolkit, 2012.
9. Steve Phelps, Peter McBurney, and Simon Parsons. Evolutionary mechanism design: A review. *Journal of Autonomous Agents and Multi-Agent Systems*, 21:237–264, 2010.
10. Stuart Russell and Peter Norvig. *Artificial Intelligence A Modern Approach Second Edition*. Prentice Hall, Englewood Cliffs - NJ, 2003.
11. Peter R. Wurman, William E. Walsh, and Michael P. Wellman. Flexible double auctions for electronic commerce: theory and implementation. *International Journal of Decision Support Systems*, 24:17–27, 1998.