Expected Utility or Prospect Theory: which better fits agent-based modeling of markets?

Paulo André Lima de Castro, Anderson Rodrigo Barreto Teodoro

Autonomous Computational Systems Lab Aeronautics Institute of Technology (ITA) São José dos Campos , SP, Brazil

pauloac@ita.br, andersonrbteodoro@gmail.com

Luciano Irineu de Castro

Department of Economics the University of Iowa Iowa City, USA luciano-decastro@uiowa.edu

Simon Parsons

Department of Informatics King's College London Strand, London, WC2R 2LS simon.parsons@kcl.ac.uk

Abstract

Agent-based simulations may be a way to model human society behavior in decisions under risk. However, it is well known in economics that Expected Utility Theory (EUT) is flawed as a descriptive model. In fact, there are some models based on Prospect Theory (PT), that try to provide a better description. If people behave according to PT in finance environments, it is arguable that PT based agents may be a better choice for such environments. We investigate this idea in a specific risky environment, a financial market. We propose an architecture for PT-based agents. Due to some limitations of the original PT, we use an extension of PT called Smooth Prospect Theory (SPT). We simulate artificial markets with PT and traditional (TRA) agents using historical data of many different assets over a period of twenty years. The results showed that SPT-based agents provided behavior that is closer to real market data than TRA agents, and that the improvement when using SPT rather than TRA agents is statiscally significant. It supports the idea that PT based agents may be a better pick to model the behaviour of agents in risky environments.

Key words: Multiagent systems, Agent-based modeling

Preprint submitted to Elsevier

1. Introduction

The price dynamics, in a financial market, are defined by the sum of the actions of all agents in such market. Thus, an alternative approach to modeling prices as an exogenous stochastic process is to model the agents' behavior and verify if we can reproduce the market outcome with such agents. This paper contributes to this artificial economics research program, by comparing how two different sets of artificial agents resemble the outcome of real markets. The two sets of agents that we use differ in how they make decisions under risk.

The most traditional and widely used model for decision under risk is Expected Utility Theory (EUT), in particular the use of the principle of expected utility maximization. However, there is evidence that people do not make decisions under risk strictly based on expected utility [1]. In fact, some experiments [2] show that financial professionals (who likely are aware of Expected Utility Theory) also may behave according to Prospect Theory and violate expected utility maximization. Prospect theory was proposed [1] and later improved by Kahneman and Tversky [3]. It is an alternative model of human decision making under risk. PT may describe some behaviors that cannot be explained by Expected Utility Theory. For instance, there is a clear preference for guaranteed small gains over uncertain large gains, and conversely for uncertain large losses over small certain losses even when EUT would point to the reverse option. This is usually called the reflection effect [1].

If financial professionals behave according to Prospect Theory in financial markets, the agent-based modeling of such markets could benefit from prospect theory based agents. Our idea is to create trading agents based on Prospect Theory and simulate an artificial market populated with such agents. If investors' behavior is consistent with PT, such simulations should provide results closer to historical data from real markets than those provided by traditional agents based on EUT.

We modeled traditional trading agents and created a new class of trading agent based on Prospect Theory. The class of traditional trading agents are briefly described in Section 2. The proposed Prospect Theory-based agent is fully described in Section 3. These agents were instantiated to populate an Artificial Financial Market. This market, and its population with those PT and traditional agents is explained in Section 4. We performed many simulated experiments using these markets. These experiments are described and the results we obtained are presented and discussed in Section 5. Finally, we state some conclusions and some open questions in Section 6.

2. Trading Agent Modeling and Prospect Theory

We modeled three classes of trading agents: agents based on traditional techniquies (TRA), which are broadly based on EUT, agents that play the role of market makers (MM) and agents that use Smooth Prospect Theory SPT. The first two classes are based on well known strategies from the literature and

are briefly described in section 2.1 and section 2.2, respectively. The SPT class is based on Prospect Theory and it is fully described in section 3.

2.1. TRA Agents

Automated trading strategies are not new and a significant number of papers have been published proposing such strategies. Most of them are based on analysis of time series (usually called technical strategies), while some others are based on the analysis of economic and financial fundamentals of companies and/or economic sectors (usually called fundamentalist strategies). They all try to maximize some expected value function. Some common functions are financial return, return variance (as proxy of risk), or a trade-off between risk and return (see. pp. 212-213, [4]). They can be seen as Expected Utility Theory (EUT) strategies, because they are not concerned in reproducing human behavior but maximizing some value (or utility). The TRA agents used in this study is very simple, given that we are not focused on maximizing trader's performance, but reproducing price behavior observed in real market using automated traders.

We therefore picked one of the simplest and most well-known technical strategies: the moving average (MA). The moving average index tries to identify trends in stock prices. The average is defined by an observation period and a calculation method that can be simple average (sum of all prices and divide it by the number of values), front-weighted triangular method or exponential average to give more relevance to newer prices rather than older prices [5]. We used MA with simple average and adapted it to provide order price based on the last market price.

2.2. Market Maker Agents

Any market used to uncover the value of an asset may benefit from an agent that stands ready to interact with traders, providing liquidity [6]. Such an agent, usually called a market maker agent, enables other trading agents to trade at every round. Its presence is very important to guarantee that a price is defined at each round.

This price is determined by the sum of all business transactions, as explained in section 4. In case of a sell order, the price of the order is defined by yesterday's price plus a small percentage, the *spread*. In the case of a buy order, the price of the order is yesterday's price or minus the spread. Therefore, the price offered by the market maker defines a lower and upper limit for the transaction price. However, the exact treansaction price is really defined by the other agents' orders (i.e., the orders placed by the TRA and SPT agents). In our study, the spread was defined as a fixed 0.5%.

3. Our Agent Model based on Prospect Theory

The modeling of agents based on Prospect Theory is not straight-forward and it has to deal with some strong difficulties. The original Prospect Theory, as proposed by Kahneman and Tversky [1], establishes one phase of editing and a subsequent phase of evaluation and selection. It deals only with prospects that have at most two non-zero outcomes (simple prospects). However, many real world problems present prospects with more than two non-zero outcomes and even prospects with continuous distributions (complex prospects). Furthermore, the editing phase as originally proposed is not well defined [7]. These issues make it really hard to use the original version of Prospect Theory in an agent model.

Therefore, we chose an alternative extension of PT, called Smooth Prospect Theory (SPT) [7]. This extension can be used for complex prospects and even continuous prospects. We are aware of the criticism about it [8] and the existence of other proposals to extend PT to complex prospects, such as Cumulative Prospect Theory [3]. The violation of first-order stochastic dominance is usually pointed as a major problem. However, as pointed out by Rieger et al [7], it is not necessarily a weakness for a descriptive model because it is known that individuals may frequently choose dominated lotteries especially when stochastic dominance is unclear to them. Therefore, we chose SPT due to the fact it incorporates the editing phase into the calculation and avoids the unclear part of the original form of Prospect Theory.

Smooth Prospect Theory is explained in Section 3.1. Our model requires that for each possible action of the agent, there is one representative prospect. these prospects are created in the Prospect Construction Phase, which is described in Section 3.2. We present the SPT Agent model in Section 3.3,

3.1. Smooth Prospect Theory for Agents

The Smooth Prospect Theory as proposed by Rieger and Wang [7] computes a SPT value for each prospect and selects the highest value prospect. The SPT value (SPTV) of a discrete prospect with an arbitrary number of outcomes x_i and respective probabilities p_i is given by Equation 1.

$$SPTV = \frac{\sum_{i=1}^{n} w(p_i)v(x_i)}{\sum_{i=1}^{n} w(p_i)}$$
(1)

where the value function v(x) is chosen as:

$$v(x) = \begin{cases} x^{\alpha} & x \ge 0\\ -\lambda(-x^{\beta}) & x < 0 \end{cases}$$
(2)

and $\lambda \approx 2.25$ is a *loss-aversion* coefficient and α, β are the risk-attitudes parameters for gains and losses. Furthermore, the weighting function is defined as:

$$w(p) := \frac{p^{\gamma}}{(p^{\gamma} + (1-p)^{\gamma})^{1/\gamma}}$$
(3)

The parameter γ reflects the amount of over or underweighting in the weighting function.

SPTV could also be calculated for continuous distributions [7], but we deal only with discrete prospects in this study.

3.2. Prospect Construction Phase

The product of the Prospect Construction Phase is a set of prospects (one for each possible action of the agent). We assume that each trading agent deals with only one asset and it has an estimate of the fair value for such an asset. Furthermore, the agent's decision-making process has to place an order at each moment of time t. An order is defined by its volume and price. An order's volume θ_t is defined as an integer number $\in [-M, M]$ at a given moment t. The value M is the maximum number of shares that can be bought or sold by the agent in one cycle. Positive values of θ mean a buy order, while negative θ means a sell order and $\theta = 0$ means to keep the current position.

The agent's order and market price dynamics will define the agent's outcome. We assume that trading agents are concerned about their return and orders will always be executed at the market price P_{t+1} . Thus, the outcome is the difference between the value of an agent's assets at time t and its value at the next time (t + 1), right after the order θ_t is executed at price P_t . This outcome (x) may be calculated as stated in Equation 5, where P_t refers to the asset price, M_t is the amount of money, Q_t is the number of shares at time t:

$$x = [(M_t - P_t * \theta_t) + (Q_t + \theta_t) * P_{t+1}] -[M_t + P_t * Q_t] x = (P_{t+1} - P_t) * (Q_t + \theta_t),$$
(4)

Which may be simplified to:

$$x = (P_{t+1} - P_t) * (Q_t + \theta_t).$$
(5)

The Section 4 describes how the market price is calculated. We also assume all orders will be executed, so each order defines changes in Q_{t+1} and the market behavior defines the market price P_{t+1} . The market price P_{t+1} cannot be defined a priori, but it can be estimated by the trading agents. Let $\overline{P_{t+1}}$ be this estimate. So, an agent can calculate $\overline{P_{t+1}} - P_t$ and then estimate the outcome (x) for any $\overline{P_{t+1}}$.

Naturally, any order may bring different outcomes according to the real market price in the next round P_{t+1} . In order to establish prospects, given an order θ_t , we would need to determine the probabilities for each possible outcome. The estimate of market price $\overline{P_{t+1}}$ is a continuous value and θ_t is dependent on the trading strategy (certainly non-linear) and market state. So, the outcome is itself a continuous non-linear function. It would require a probability density function, p(x), to represent the associated probabilities for each possible outcome (x).

We adopted Markowitz's assumption that returns present a Gaussian distribution [9], so the price P_{t+1} is a also a random variable with a Gaussian probability distribution. Thus, the density probability function p(x) of the outcome x can be given by equation 6, where σ is the standard deviation and μ the expected value or $\overline{P_{t+1}}$:

$$p(x) = \frac{\exp^{-(P_{t+1}-\mu)^2/2\sigma^2}}{\sigma * \sqrt{2 * \pi}}$$
(6)

It is easy to see that the outcome, x, (Eq. 5) is a linear function of P_{t+1} . Let $a_t = Q_t + \theta_t$ and $b_t = P_t(Q_t + \theta_t)$. If so, we can rewrite Equation 5 to determine a new expression for P_{t+1} and use it in Equation 6 to find an expression for the distribution probability function p(x) for the outcome x. Such an expression is given by Equation 7, where a_t and b_t are known at time t. It can be used for the calculation of SPT as stated in Equation 1.

$$p(x) = \frac{\exp^{-(x-b_t-a_t\mu)^2/2a_t^2\sigma^2}}{\sigma * \sqrt{2*\pi}}$$
(7)

Each prospect would have infinite possible outcomes and could be calculated by SPT for continuous distributions. In order to avoid such complexity and since prices are limited to cents, we decided to limit the possible outcomes to a finite set. We adopt a step ϵ in (0,1) for prices. Thus, As each order θ_t is limited to [-M, M] and if we assume that P_{t+1} is limited to $[0, 2P_t]$, it is easy to verify using equation 5 that the outcome x is limited to interval $[-P_t(Q_t + M), P_t(Q_t + M)]$. Therefore, the number of possible outcomes is limited to $2P_t(Q_t + M)/\epsilon$ for each prospect. We limit the orders of agent to three: sell (-M), hold (0) and buy (M).

3.3. The SPT Agent Model

The SPT agent model is a simple extension of the classic Utility based Agent Model [10], where the SPTV value (Eq. 1) is used instead of utility to select an agent's action from the set of possible actions, shown in Figure 1. There is a oneto-one relationship between prospects and possible actions. This relationship is created by the **Prospect Construction Phase**, which is fully explained in Section 3.2.

In the **prospect evaluation phase**, the agent computes SPTV for each prospect and selects the action related to the highest value SPT prospect. The selected action is then executed by the agent through its actuators (Figure 1).

It isn't clear to me from the figure what the acutators are – are they buy/sell decisions, or something else?

It is a straight forward approach, but it requires that the agent have information about all possible outcomes and respective probabilities for each possible action. Such information is used to construct one prospect for each possible action.

4. Artificial Financial Markets populated with SPT and TRA agents

In our Artificial Financial Market, each trading agent gives orders that are stored in a buy or sell book as in a real stock market. The clearing process is performed by the Four Heap algorithm described in [11]. The market price



Figure 1: SPT Agent Model

for a given instant of time P_t is defined as the average of all transaction prices weighted by the volume of each transaction. A **market specification** defines a set of agents. Each agent is described by its class (SPT, TRA or MM) and its volume order.

My understanding is that the work compares two kinds of market, one with MM and TRA agents, and the other with MM and SPT agents. If this is the case, I would say so explicitly.

That way, an agent that gives a higher volume order is more relevant to the market price formation than other agent that submits small volume orders [11]. This price, called **market price**, is compared with actual prices obtained from actual stock exchange data, called the **external price**. The difference between the market $\text{price}(\overline{P_t})$ and external $\text{price}(P_t)$ is the prediction error in a given instant of time t.

It is relevant to observe that the prediction error of a period of time is much more relevant than just one moment to state that one set of trading agents is better adapted than the other one. Therefore, we define the **session error** (E) as the sum of squared errors (eq. 8). If a market specification A provides a smaller session error (E), than market specification B, then we may say that artificial market A is a better description of the real market than B.

$$E = \sum_{t=1}^{N} (\overline{P_t} - P_t)^2 \tag{8}$$

4.1. Parameter Calibration

As pointed out by LeBaron, a common criticism of agent-based markets is that they usually have too many parameters and the impact of these parameters is not well understood ([12],pp. 1222). In this study, we are making a direct comparison between agent-based markets with SPT and TRA agents. Furthermore, we set the parameters using the same algorithm for all markets. So, we believe that it is a fair comparison.

We used a search algorithm to adjust the volume orders of a market specification in order to reduce the Session Prediction Error (Eq. 8). Given the fact that trading agents with higher volume order have more relevance to market price formation, we adjusted the market specification (i.e., the volume order of the trading agents) to fit data previously observed in real markets. For simplicity, each agent type had just one instance, and it traded one specific share quantity at each round. The **market specification** was defined by three parameters: the share quantities of each one of the three kinds of agents: SPT, TRA, and market maker agents.

It is very hard to know a priori how a change in one of the specification parameters may affect the market price $\overline{P_t}$ or the session error (E). Therefore as search algorithm, we used the random-start gradient descendent method to find minimum points of the objective function E, which is a variant of the common hill climbing methods [10]. The method uses a new random starting point each time it finds a local minimum for the objective function. It is worth noting that any change in the market specification does not alter trader strategy, but their relevance to the market price definition.

5. Simulated Experiments, Results and Discussion

In this section, we describe the simulated experiments performed (Section 5.1) and their results (Section 5.2). Then, we discuss such results and answer some questions (Section 5.3).

5.1. Simulated Experiment Setup

We performed simulated experiments in many different scenarios. Each scenario was defined by a year from 1994 to 2013 (twenty years) and one company picked from the twenty biggest companies in the NYSE or Nasdaq during the period. Table 1 lists these companies. This gave us four hundred scenarios (20×20) . The number of data points in each scenario changed according to the number of business days in the respective year. Typically, each scenario had about 247 data points. Each one was defined by the asset's closing price or **external price** and the respective business day, providing about 98,000 historical prices $(247 \times 20 \times 20)$. These prices were downloaded from the Yahoo Finance service.

The simulations were performed considering periods with high price volatility (crisis) and low price volatility (non-crisis periods). We implemented our trading agents using an adapted version [13] of an auction simulator, called JASA [14]. JASA runs over an agent modeling toolkit, called JABM [15].

Stock	Stock
AMD	JPMorgan Chase
Apple	JC Penney
AT&T	Microsoft
BarrickGold	Nike
Citigroup	Pfizer
Ford Motor	Rite Aid
General Eletric	Sprint
HP	Verizon
IBM	Wells
Intel	Xerox

Table 1: Companies whose stock was used in the simulated experiments

5.2. Results

We compared SPT agents (that is the market with SPT and MM agents) against TRA agents (the market using TRA and MM agents) in 400 scenarios. A smaller Session Error in a scenario was considered better performance. We observed which agent type (SPT or TRA) presented better performance in each scenario. The simulation results are presented in Table 2. The SPT agents achieved better performance in 288 scenarios against only 112 in which TRA agents achieved higher performance. We also analyzed the performance by segregating the scenarios into two categories: crisis (high volatility) and noncrisis (low volatility). We arbitrarily drew the boundary between crisis scenarios and non-crisis scenarios as being where the level of volatility was 80. Crisis scenarios (CRs) were defined as those scenarios where the level of volatility was greater than or equal to 80 and non-crisis scenarios (NCRs) were defined as those scenarios in which volatility was smaller than 80. Table 2 presents the number of scenarios where SPT agents achieved better performance (smaller session error) in the first row. The second row shows the number of scenarios where TRA agents performed better. For each element in the table we show the absolute number of scnearios and the respective percentages of the total number of scenarios. It is arguable that SPT agents performed slightly better in CR scenarios than in NCR scenarios.

Table 2: Number of scenarios where SPT or TRA achieved better performance in crisis (CR), non-crisis (NCR) and all scenarios. Results are given both as the total number of scenarios and the percentage of scenarios.

	CR		NCR		All	
SPT	60	(75%)	228	(71%)	288	(72%)
TRA	20	(25%)	92	(29%)	112	(28%)
Total	80	(100%)	320	(100%)	400	(100%)

5.3. Discussion

Based on the results, we may state that SPT performed better in crisis (CR) or non-crisis (NCR) situations. We tested the hypothesis that the performance is the same for TRA agents and SPT agents. The data presented in Table 2 allows us to reject such hypothesis with 99.9% confidence using the χ^2 test for all situations: CR (20.0), NCR (57.8) and general (77.4), since $\chi^2_{0.999}$ is 10.8 for one degree of freedom. Therefore, we may state that the data supports the hypothesis that SPT agents are a better description than TRA agents.

We also analyzed the results with several different levels of volatility as crisis limit in Section 5.3.1. In Section 5.3.2, we verify how far the TRA's and SPT's performances are from each other.

5.3.1. What happens if the crisis definition changes?

We used the volatility (variance) of the prices as proxy of risk. If a scenario presented volatility equal or higher than the specified limit, it was classified as crisis and it was classified as non-crisis in the negative case. We used limit levels from 0 to 1600. We considered a tie when the difference between performances were equal or below 0.1%. In limit level 0, all 400 scenarios were classified as crisis. In limit level 1600, only six were so classified. SPT performed better in all six. The results achieved by SPT and TRA are presented in Figure 2.

I think that the figure would be easier to understand if different line styles (dotted, dashed, etc) were used for the lines rather than different colours. Also, it is confusing to label the TRA agent performance as EUT. Why do some values exceed 100%?

SPT presented better performance in about 70% of the scenarios for limits from 0 to 400. It seems that SPT performed even better when the limit went from 800 to 1600 (very high volatility or very bad crises). However, the small numbers of crisis scenarios in such conditions (16 to 6) are not statiscally significant to reject the hypothesis that SPT performance would have the same performance in such cases.

5.3.2. How far apart are the agents' performances?

We also analyzed the results with several levels of tie for agent's performance. The tie scale went from 0% to 30%. The results are presented in Figure 3.

The same comment about line type applies to this figure. It would be good if the y-axes and the measurement of scenario number on the x-axis were the same in both figures — you want people to be able to compare between the two.

For instance, when the difference between SPT and TRA performance was equal or less then 1%, there were 119 ties, SPT was superior to TRA in 214 scenarios and TRA was superior to SPT in only 67 scenarios. All performance differences were equal or smaller than 30%. Therefore, all scenarios were classified as tie in the last column of Figure 3. It is relevant to note than in all circumstances with different performances, SPT out performed TRA.



Figure 2: SPT vs TRA performances for multiple crisis definitions. Different volatilities are plotted on the x-axis (recall that 80 is the value used above), and for each value we plot the percentage of scenarios in which SPT agents have better performance, the percentage of scenarios in which TRA agents have better performance, and the percentage of scenarios in which the two types of agent tie.



Figure 3: SPT vs TRA performances for multiple tie definitions. Different values of the tie parameter (the percentage difference in performance under which different results for SPT and TRA are considered the same) are plotted on the x-axis, and for each value we plot the number of scenarios in which SPT agents have better performance, the number of scenarios in which the two types of agent tie.

5.3.3. Result Analysis

The predictions made by SPT agents presented better performance (smaller errors) than TRA agents in most cases. In fact, we can reject the hypothesis that SPT and TRA agents present the same performance with 99.9% confidence for all scenarios. Furthermore, we can state that SPT was better in 70% of the scenarios with 75% confidence. We expected SPT would perform better in crisis, but it would present poor performance in low volatility scenarios. We expected that because we believe that psychological biases can present higher influence on investment decisions in crisis periods. Surprisingly, the results support the idea that SPT agents are better fit to real data than TRA agents regardless of volatility. Nevertheless, we still believe that SPT superiority would be more evident in high volatility scenarios and the results seem to point in such direction (see Figure 2). However, the current data does not allow us to conclude that the hypothesis is true with a suitable degree of statiscal significance. Furthermore, the fact that SPT agents presented smaller errors than TRA agents, supports the idea that the strategies of real traders may be influenced by psychological biases as described in PT [1].

6. Conclusions and Future Work

Agent based modeling (ABM) may become a better way to help guide financial policies than traditional models according to some researchers [16]. However, several problems may be identified in this approach. For instance, human beings do not make risky decisions strictly based on expected utility theory (EUT) as usually assumed in ABM and perhaps an alternative descriptive models as Prospect Theory (PT) may be a better model for agent's decision process. In fact, a recent study that uses prospect theory based agents and fits the model to experimental data in the context of Behavioral Mechanism Design, have shown different equilibria when agents are based on expected utility theory than those observed when they are based on prospect theory [17].

We addressed the modeling and simulation of PT-based agents. We used an extension of PT called Smooth Prospect theory (SPT) to develop an agent model. Such a model uses a prospect construction phase that creates a oneto-one relationship between a prospect and an agent's action. This model was used to build a trading agent. We populated simulated artificial markets with this kind of agent (SPT) and with traditional (TRA) agents. Those agents were used in a significant number of simulated experiments.

The results showed that the artificial market populated with SPT agents performed significantly better than EUT agents. In fact, we were able to reject the hypothesis that SPT and TRA agents present the same performance with 99.9% confidence for all scenarios. Furthermore, we can state that SPT is better in 70% of the scenarios with 75% confidence. These results support the idea that real human trading agents are influenced by psychological biases, such as described in PT [1]. It may be pointed out that the agents are relatively simple and do not fully represent human behavior in financial market. For instance, the

prospect construction phase could deal with continuous distributions rather than discrete ones. In future work, we intend to address these issues. Furthermore, we believe that two other open questions that are really worth to study: (1) the use of an alternative version of Prospect Theory that does not violate the first-order sthocastic dominance (such as CPT) and (2) try PT based agents on different risk environments, such as games, to observe if they are still a better choice over expected utility theory based agents as descriptive model.

Acknowledgments

Paulo A.L. Castro was partially funded by FAPESP (11/18325-8), Brazil.

References

- D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk, Econometrica 47 (2) (1979) 263–291.
- [2] M. Abdellaoui, H. Bleichrodt, H. Kammoun, Do financial professionals behave according to prospect theory. an experimental study, Theory and Decision 74 (3) (2013) 411–429.
- [3] A. Tversky, D. Kahneman, Advances in prospect theory: Cumulative representation of uncertainty, Journal of Risk and Uncertainty 5 (1992) 297–323.
- P. A. L. Castro, J. S. Sichman, Automated asset management based on partially cooperative agents for a world of risks, Applied Intelligence 38 (2013) 210-225. doi:10.1007/s10489-012-0366-8. URL http://dx.doi.org/10.1007/s10489-012-0366-8
- [5] J. Katz, D. McCormick, The Encyclopedia of trading strategies, McGraw-Hill, New York, 2000.
- [6] D. M. R. Abraham Othman, David M. Pennock, T. Sandholm, A practical liquidity-sensitive automated market maker, ACM Transactions on Economics and Computation 1 (3) (2013) 14:1–14:25.
- [7] M. O. Rieger, M. Wang, Prospect theory for continuous distributions, Journal of Risk and Uncertainty 36 (2008) 83–102.
- [8] V. S. Amit Kothiyal, P. P. Wakker, Prospect theory for continuous distributions: A preference foundation, Journal of Risk and Uncertainty 42 (2011) 195–210.
- [9] H. M. Markowitz, Portfolio selection, Journal of Finance 7 (1) (1952) 77–91.
- [10] S. Russell, P. Norvig, Artificial Intelligence A Modern Approach Second Edition, Prentice Hall, Englewood Cliffs - NJ, 2003.

- [11] P. R. Wurman, W. E. Walsh, M. P. Wellman, Flexible double auctions for electronic commerce: theory and implementation, International Journal of Decision Support Systems 24 (1998) 17–27.
- [12] B. LeBaron, Agent-based computational finance, North-Holland, Amsterdam, The Netherlands, 2006.
- [13] P. A. L. Castro, S. Parsons, Towards modeling securities markets as a society of heterogeneous trading agents, in: Proceedings of the 15th International Workshop on Agent-Mediated Electronic Commerce (AMEC-13), AAMAS, Saint Paul, USA, 2013, pp. 17–29.
- [14] S. Phelps, Jasa java auction software simulator api (2007). URL http://sourceforge.net/projects/jasa/
- [15] S. Phelps, Applying dependency injection to agent-based modeling: the jabm toolkit (2012). URL http://sourceforge.net/projects/jabm/
- [16] J. D. Farmer, D. Foley, The economy needs agent-based modelling, Nature 460 (2009) 685–686.
- [17] D. Easley, A. Ghosh, Behavioral mechanism design: Optimal crowdsourcing contracts and prospect theory, in: The Sixteenth ACM Conference on Economics and Computation (EC'15), Association for Computing Machinery (ACM), Portland, OR, 2015, pp. 679–701.