

Evolutionary Advantage of Foresight in Markets

Daniel Hennes, Daan Bloembergen, Michael Kaisers, Karl Tuyls and Simon Parsons*

Department of Knowledge Engineering, Maastricht University
P.O. Box 616, 6200MD Maastricht, The Netherlands
{daniel.hennes, daan.bloembergen, michael.kaisers, k.tuyls}@maastrichtuniversity.nl

*Department of Computer and Information Science, Brooklyn College
City University of New York, 2900 Bedford Avenue, Brooklyn, NY 11210, USA
parsons@sci.brooklyn.cuny.edu

ABSTRACT

We analyze the competitive advantage of price signal information for traders in simulated double auctions. Previous work has established that more information about the price development does not guarantee higher performance. In particular, traders with limited information perform below market average and are outperformed by random traders; only insiders beat the market. However, this result has only been shown in markets with a few traders and a uniform distribution over information levels. We present additional simulations of several more realistic information distributions, extending previous findings. In addition, we analyze the market dynamics with an evolutionary model of competing information levels. Results show that the highest information level will dominate if information comes for free. If information is costly, less-informed traders may prevail reflecting a more realistic distribution over information levels.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence — *Multiagent systems*

General Terms

Theory, Experimentation

Keywords

Stock markets, evolutionary game theory, value of information

1. INTRODUCTION

Markets play a central role in today's society, and range from stock markets to consumer-to-consumer e-commerce [1, 2]. Economic theory often starts from perfect competition as an idealized assumption about markets. It relies, among

other characteristics, strongly on a symmetric information structure. All traders have access to the same information about price and quality of goods. Many, if not all, of today's markets do not meet this utopian assumption and thus valorize the access to information. Undoubtedly, information is an important factor that has influence on trading success or losses. Insiders are clearly able to use their information to outperform the market. However, the relation between information level and success is not trivial.

Market forecasters and fund managers are generally assumed to be well informed, though for the most part they perform below market average. Cowles [4] has been the first to study this phenomenon and reports that a group of trained forecasters performed 4% below market average during a period of 4.5 years. These findings have since been confirmed by multiple studies; for an overview we refer to [9, 17], in particular we like to highlight that Malkiel [10] reports on returns of actively managed funds over a period of 30 years - less than 15% of the funds outperformed the market.

Toth et al. [17] study the relation between information and performance for traders with various information levels both in simulation and in human experiments. Average-informed traders perform below market level, while non-informed traders reach the market average; highly informed traders beat the market. These results suggest that if a trader has no inside information, trading based on current market prices (non-informed) is most sensible. Relying on outdated or average information has a negative impact on returns.

In [7, 9], the authors investigate whether this negative impact can be explained by behavioral patterns. In particular, the authors test the hypothesis that low performance of average-informed traders is the result of overconfidence, i.e. overestimating the value of (possibly outdated) information. Results show that traders do not exhibit overconfidence and low returns are caused by the asymmetric information structure itself. Huber [6] offers the following explanation: during trends, foresight is clearly advantageous. When the trend reverses, the averagely informed trader trusting its information performs worst due to outdated information. Non-informed (random) traders are safe from these systematic mistakes.

The vast body of previous work [6, 7, 9, 15, 16, 17] has evaluated the advantage of information in markets from various perspectives. However, only markets with a limited

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '12, July 7-11, 2012, Philadelphia, Pennsylvania, USA.
Copyright 2012 ACM 978-1-4503-1177-9/12/07 ...\$10.00.

number of traders and uniformly distributed information levels have been considered, and information was assumed to be free. Our work investigates several more realistic information distributions in larger markets. Furthermore, we extend the analysis by studying an infinite population of traders using an evolutionary model and demonstrate the influence of the price of information on market dynamics.

The remainder of this article is structured as follows: The following section provides background information on auctions, the value of information in auctions and existing evolutionary models of auctions. Section 3 introduces the market model and evaluation method used in the experiments. Evolutionary simulations and results are described and discussed in Section 4. Finally, Section 5 concludes the article.

2. BACKGROUND

This section contains an introduction to auctions and the value of information in auctions. Subsequently, the main concepts of evolutionary game theory are introduced, i.e. replicator dynamics and evolutionarily stable states.

2.1 Auctions

Auctions are highly efficient match making mechanisms for trading goods or services. As such, they are employed by a number of real markets, such as telecommunication spectrum rights auctions or the New York Stock Exchange (NYSE) [1, 11]. In practice, there is a variety of rules that may be used to conduct an auction. Each set of rules may result in different transaction volumes, transaction delays, or allocative market efficiency. One sided auctions, especially with one seller and many potential buyers, are popular in consumer-to-consumer e-commerce [2, 3]. Here, we focus on double auctions, which essentially provide a platform for buyers and sellers to meet and exchange a commodity against money. A taxonomy of double auctions especially tailored to automated mechanism design can be found in [12].

Double auctions maintain an open book of bids (offers to buy at a specified price) and asks (offers to sell at a specified price). Two principle forms are the clearing house auction and continuous operation auction. In a clearing house auction, orders are collected for a trading period (e.g., one day) and matched, or cleared, after the trading period is closed. This mode of operation allows for high allocative efficiency, but incurs delays in the transactions. In contrast, continuous operation immediately establishes a transaction as soon as some buyer is willing to pay more than a seller is asking for. This mode allows higher transaction rates at the cost of some allocative efficiency. Experiments in this article will use continuous operation mode, since it reflects the day-time operation mode of the NYSE [1].

2.2 Value of information

It is common sense that training and additional information should increase performance for any task. However, the value of information in markets is non-monotonic, i.e., having some information may be worse than having none.

In order to measure the value of information, experiments in auctions measure revenue. Since revenue is heavily dependent on market conditions dictated by the price signal, it is normalized to reflect the relative return. Assume trader i receives revenue r_i . The average profit $r_{avg} = \frac{1}{n} \sum_i r_i$ in




Figure 1: Relative market return over information level. 10 traders with information level 0 to 9 (1 trader for each level).

an n -trader market is used to compute the relative market return $u_i = \frac{r_i}{r_{avg}} - 1$ for each trader.

Figure 1 shows relative market return over information levels in a market with $n = 10$ traders, one trader for each of 10 information levels, where level 0 represents random traders (as formally defined in Section 3). The revenue follows a *J-curve*, which means that random traders perform at market average while weakly informed traders are exploited by insiders. This result holds both in abstract market models that can be simulated and in experiments with human participants [17].

2.3 Evolutionary game theory

Auctions provide a dynamic environment with a lot of traders (agents) that adapt to each other while competing for revenue. Learning in such multi-agent systems is generally complex and poses many challenges that inspire prescriptive, descriptive and normative research [14]. Evolutionary game theory provides a methodology to analyze multi-agent learning, replacing assumptions from game theory like rationality by evolutionary concepts such as pressure of natural selection [18].

The evolutionary perspective considers a population of individuals, where each individual belongs to one of several species. These species generally relate to atomic strategies, or to information levels within this article. Two core concepts are the *replicator dynamics*, describing how a population evolves, and *evolutionarily stable states*.

The replicator dynamics formally define the population change over time. The payoff function can be interpreted as the Darwinian fitness of each species.

$$\dot{x}_i = x_i \left[f_i(x) - \sum_j x_j f_j(x) \right] \quad (1)$$

Evolutionarily stable state are such population distributions x that are fixed points of the replicator dynamics, i.e., $\dot{x} = 0$, and where small perturbations $|\hat{x} - x| < \epsilon$ would be driven back to x by selection pressure, i.e., by following the replicator dynamics.

Previous research has demonstrated the viability of evolutionary game theory to analyze meta strategies in simulated auctions, and to compare clearing house against continuous double auctions [8, 13]. We will follow a similar analysis procedure but our data is generated by a different model described in the following section.

3. MARKET MODEL

In order to analyze the advantage of foresight, we simulate a stock market in which traders with different amounts of information on future prices, or *information levels*, trade a certain asset. We closely follow the market model as described by [7, 15] in order to be comparable. The market is based on a continuous double auction with open order book, in which all traders can place bids and asks for shares. The intrinsic value of the shares is determined by a dividend stream that follows Brownian motion:

$$D_t = D_{t-1} + \epsilon \quad (2)$$

where D_t denotes the dividend in period t with $D_0 = 0.2$, and ϵ is a normally distributed random term with $\mu = 0$ and $\sigma = 0.01$, i.e., $\epsilon \sim \mathcal{N}(0, 1)$. Figure 2 shows an example dividend stream.

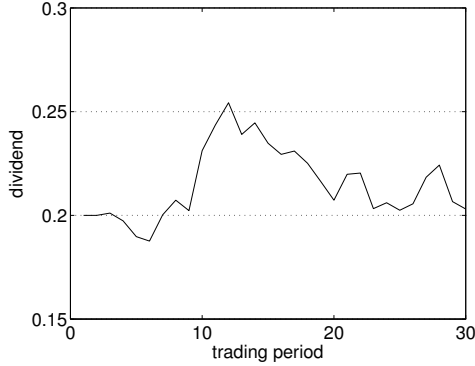


Figure 2: A Brownian motion dividend stream.

We simulate the market over 30 trading periods, each lasting $10 \cdot n$ time steps, where n is the number of traders present. All traders start with 1600 units cash and 40 shares, each worth 40 in the beginning. At the beginning of each period, all traders can put a bid or ask in the book (opening call). Hereafter, at every time step a trader is selected at random who can then place a bid or ask according to its trading strategy (see below). At the end of each period, dividend is paid based on the shares owned, and risk free interest rate (0.1%) is paid over cash. The performance of the traders is measured as their total wealth after the 30 periods, i.e., each share is valued according to the discounted future dividends (see below) and added to the cash reserves.

The different information levels are implemented by varying the amount of knowledge that traders have about the future dividends. In general, a trader with information level I_k knows the dividend of this and the next $(k - 1)$ periods. Traders with information level I_0 have no information about dividends and can only observe the current market price. This results in a cumulative information structure, where insiders know at least as much as averagely informed traders. The information that the traders receive each period is the conditional present value of the shares, conditioned on their information level. This value can be calculated using the dividend discount model (Gordon growth model) as

$$E(V|I_j, k) = \frac{D_{k+j-1}}{(1+r_e)^{j-2}r_e} + \sum_{i=k}^{k+j-2} \frac{D_i}{(1+r_e)^{i-k}} \quad (3)$$

where V denotes the value, I_j is the information level, k the period, and r_e the risk-adjusted interest rate (set to 0.5% in our experiments).

3.1 Trading strategies

We use two different trading strategies in our experiments. Traders that have at least some information about the dividend stream (I_1 and higher) use the *fundamentalist* strategy, that takes this information into account. Traders without any information (I_0) use the *random* strategy, in which their bids and asks are based purely on the current market price of the shares.

3.1.1 Fundamentalists

Fundamentalists completely rely on the information they receive. The fundamentalist strategy is explained in Algorithm 1 (see also [15, 16]). In essence, they compare their estimated present value $E(V|I_j, k)$ with the current best bid and ask in the book. If they find a bid (ask) with a higher (lower) value than their estimate, they accept the offer. Otherwise, they place a new order between the current best bid and ask prices. Naturally, the trader should own enough shares or cash to accept or place an order.

Algorithm 1 Fundamentalist trading strategy

```

pv ← E(V|Ij, k) {private value}
if pv < bestBid then
    acceptOrder(bestBid)
else if pv > bestAsk then
    acceptOrder(bestAsk)
else
    Δask = bestAsk - pv
    Δbid = pv - bestBid
    if Δask > Δbid then
        placeAsk(pv + 0.25 · Δask · N(0, 1))
    else
        placeBid(pv + 0.25 · Δbid · N(0, 1))
    end if
end if

```

3.1.2 Random traders

The random trading strategy only takes the current market price into account when deciding whether to accept or place an order. With equal probability the trader sells or buys shares. The random trading strategy is explained in Algorithm 2 (see also [15, 16]).

Algorithm 2 Random trading strategy

```

pv ← current market price {private value}
if U(0, 1) < 0.5 then
    ask = pv + 2 · N(0, 1)
    if ask < bestBid then
        acceptOrder(bestBid)
    else
        placeAsk(ask)
    end if
else
    bid = pv + 2 · N(0, 1)
    if bid > bestAsk then
        acceptOrder(bestAsk)
    else
        placeBid(bid)
    end if
end if

```

3.2 Simulations and results

We simulate the market with varying numbers of traders for each information level, in order to analyze the relative performance of traders with different amounts of foresight. To reduce the effect of randomness we run 100 sessions of 100 simulations each; the dividend stream is fixed for each session. Results are given as the relative performance with respect to the market average plotted against the information levels.

Figure 1 shows the results for a market of 10 traders in 10 information levels: one random trader, I_0 , and 9 fundamentalists, $I_1 \dots I_9$. As can be seen, performance does not necessarily increase with more information: the random trader performs at market average, whereas traders with limited amounts of information do significantly worse. Only highly informed traders are able to beat the market.

This result is in line with related work, where a similar shaped *J-curve* was reported [9, 15]. This relation between information level and performance, where more information is not always better, has also been observed in market experiments involving human traders [7]. A possible explanation is that random traders are by definition not predictable, and therefore hard to exploit by insiders. On the other hand, experts can more easily predict and exploit traders with limited or average information levels.

Previous work has mainly focussed on small scale markets, with uniform and static distributions of traders over information levels [6, 7, 9, 15, 16, 17]. We believe that this may lead to an overly simplified model of reality, which may in turn influence the reported findings. For example, a market will be more likely to contain only a small number of insiders, and a large group of averagely informed traders. Furthermore, having only one trader per information level rules out within-group trading, which could bias results.

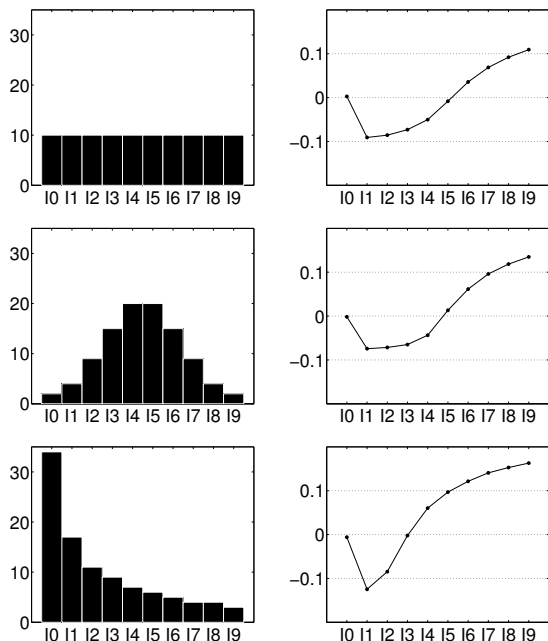


Figure 3: Relative market return over information level (right) for various information distributions (left) given a finite population of 100 traders.




Figure 4: Relative market return over information level; 1 trader for each of the information levels 0, 3 and 9.

We extend the experiments of [9, 15] by looking at markets with more traders and non-uniform distributions of traders over information levels. An overview of relative market returns for a selection of information distributions is given in Figure 3. Next to the uniform distribution used in previous work, we run simulations with a normal distribution and a power-law distribution over information levels. These distributions are chosen to reflect information distributions that are likely to be found in real markets. It is impossible to observe these distributions directly as the information level is private to the trader [7]. The two chosen distributions follow from natural assumptions: (1) normal distributions arise if access to information is cumulative based on independent and identically distributed bits of information; (2) the power-law distribution is motivated by an information flow in scale-free social networks where every trader has access to information of his social ties.

As can be seen from this figure, random traders perform at market average under all three distributions, and traders with limited information underperform the market. However, the shape of the curve does change considerably depending on the information distribution. Where in the uniform scenario only traders with information level I_6 or higher outperform the market, for the normal distribution this is the case for I_5 and for the power-law distribution for I_4 . However, we can conclude that the *J-curve* is relatively insensitive to changing information distributions and numbers of traders. Only in extreme cases (not shown here) does the curve change drastically.

Figure 4 shows that relative market returns follow the *J-curve* even in small markets with only three information levels: random traders, averagely informed traders, and insiders. Again, this is in line with previous work, where a similar setup was shown to reflect stylized facts such as autocorrelation observed within real markets [17]. Note that the obtained curve does not change qualitatively when varying the information level of average-informed traders: any choice between $\{I_0, I_1, I_9\}$ and $\{I_0, I_8, I_9\}$ results in a *J-curve*.

3.3 Discussion

We have reproduced the *J-curve* of relative market returns over information levels that has been observed in previous work. Furthermore, we have shown that the specific shape of

this curve prevails if the distribution over information levels changes. This indicates that the conclusions drawn from this may hold under more realistic settings as well.

This perspective still assumes that the distribution over information levels does not change over time. There are several ways this assumption may be violated in practice: First, traders may choose to acquire more information. For example, traders may or may not subscribe to financial news sources, which in turn determines their information level - possibly at a cost. The effect of having traders choose between trading strategies is investigated in [9, 15], with the conclusion that only highly informed traders will choose their fundamental strategy, taking their information into account. Second, traders may take over a larger market share due to their financial success while others are driven out of the market. This motivates the evolutionary analysis that accommodates evolving distributions over information levels, and elicits the market dynamics.

4. EVOLUTIONARY ANALYSIS

The previous section provides a method to compute expected relative market revenues for selected information distributions. This views information distributions as isolated and fixed in time. However, the market revenue can be interpreted as Darwinian fitness, such that traders performing below market average should be driven out of the market, while those with higher returns prevail. This section will first introduce the evolutionary analysis methodologically, list the results and discuss their implication.

4.1 Method

The evolutionary model assumes an infinite population. We cannot compute the payoff for such a population di-

rectly, but we can approximate it from evaluations of a finite population.

All possible distributions over k information levels can be enumerated for a finite population with n individuals. Let N be a matrix, where each row N_i contains one discrete distribution. The matrix will yield $\binom{n+k-1}{n}$ rows. Each distribution over information levels can be simulated with the market model, returning a vector of average expected relative market revenues $u(N_i)$. Let U be a matrix which captures the revenues corresponding to the rows in N , i.e., $U_i = u(N_i)$. A heuristic payoff table $H = (N, U)$ is proposed in [19] to capture the payoff information for all possible discrete distributions in a finite population.

In order to approximate the payoff for an arbitrary mix of strategies x in an infinite population distributed over the species according to x , n individuals are drawn randomly from the infinite distribution. The probability for selecting a specific row N_i can be computed from x and N_i :

$$P(N_i|x) = \binom{n}{N_{i,1}, N_{i,2}, \dots, N_{i,k}} \prod_{j=1}^k x_j^{N_{i,j}}$$

The expected payoff $f_i(x)$ is computed as the weighted combination of the payoffs given in all rows, compensating for payoff that cannot be measured. If a discrete distribution features zero traders of a certain information type, its payoffs cannot be measured and $U_{j,i} = 0$.

$$f_i(x) = \frac{\sum_j P(N_j|x)U_{j,i}}{1 - (1 - x_i)^k}$$

This expected payoff can be used in Equation 1 to compute the evolutionary change according to the replicator dynamics.




Figure 5: The central simplex shows the evolutionary dynamics of an infinite population mixing between the information levels 0, 3 and 9. Relative market revenue over information levels is given for four selected finite distributions: top-left (33, 33, 33) which reflects a uniform distribution, bottom-left (80, 10, 10), top-right (10, 10, 80), bottom-right (10, 80, 10).




Figure 6: Evolutionary market dynamics using a fixed cost (left) and quadratic cost function (right) for information levels 0, 3 and 9.

4.2 Experimental setup and results

The experiments of this section comprise two elements. An evolutionary analysis of an infinite population is performed to elicit the dependence of revenue on the presence of other information levels. In addition, selected population distributions are approximated with a finite population and illustrate revenue distributions for interesting points. The evolutionary analysis is based on the market model described in Section 3 and uses the method described in the previous section to compute payoffs and the replicator dynamics for an infinite population of traders with arbitrary and evolving information distributions. The heuristic payoff tables are computed for $n = 12$ traders distributed over the information levels I_0 , I_3 and I_9 , leading to 91 rows.

Figure 5 shows the evolutionary dynamics of the market model. Four representative population distributions are evaluated in more detail in a finite population of $n = 100$ traders to illustrate the revenue structure for the information levels. The evolutionarily stable state is a global attractor, where only insiders prevail. The relative market performance for the four selected finite distributions of traders highlights again the J-curve observed before. Even though non-informed traders perform close to market average, insiders take advantage of their knowledge and take over the market. However, their competitive advantage is vanishing as they are facing more and more competitors of the same information level (see top-right revenue graph of Figure 5).

Note that up till now, information was freely available to all traders. However, it is reasonable to assume that gathering more information is costly. In the most simple case, a fixed cost for information might be incurred leading to a possible advantage of non-informed traders as they do not have to pay this price. More realistically, costs could also increase with the amount of information gathered, for example using a quadratic cost function such that average-informed traders pay only a little whereas insiders pay the full price. This relates to a real-world scenario where average traders only subscribe to financial newspapers or magazines, whereas insiders may need to hire experts to gain advantage.

Figure 6 shows the market dynamics in both cost scenarios. The fixed cost is set to 5 units cash per trading period

for information levels I_3 and I_9 , non-informed traders pay nothing. The quadratic cost function used is

$$\frac{i^2}{9^2} \cdot 15,$$

where i is the information level, resulting in a maximum cost of 15 units cash for insiders per trading period. As can be observed, introducing cost leads to significantly different and more complex dynamics. In the constant cost scenario, the evolutionary advantage of insiders decreases in favor of non-informed traders, leading to an equilibrium state where insiders and non-informed traders co-exist. Using a quadratic cost function gives rise to an interior equilibrium, mixing between all information levels.

4.3 Discussion

Information does come at a cost in real markets, which has been neglected in much of the related work [6, 7, 9, 15, 16, 17]. Evolutionary analysis under different cost functions indicates that costs can significantly alter the market dynamics and allow less-informed traders to prevail.

Our results contribute to the ongoing debate about the strong-form efficient-market hypothesis, which has a large following and growing number of critics [5]. It states that prices in financial markets instantly reflect all information available to participating traders, including insider information. We found that evolutionary pressure drives a market toward an information distribution at which the market is strong-form efficient, possibly driving some information levels extinct in the process. However, the evolutionary process will only end in equilibrium for an isolated system; in real markets, traders that enter the market with information and money from other sources continuously perturb the system. As a result, real markets may be found off-equilibrium almost all the time. It is up to future experiments to quantify the influence of arriving traders on perturbation from the equilibrium.

5. CONCLUSIONS

Literature has established a link between human traders and a market model that can be rigorously analyzed in simulation. In this article, we have exploited this link and

contributed in the following ways: (1) The value of information in markets has been confirmed to follow a J-curve for several more realistic information distributions. (2) The evolutionary advantage of information makes insiders drive less-informed traders out of the market, with a diminishing competitive edge. (3) If information comes at a cost, less-informed traders may prevail in the market.

The experiments that have been carried out for this article were limited to an evolutionary analysis of three competing information levels. While this design choice is sufficient for demonstrating the arguments within this article, the evolutionary analysis naturally extends to four or more strategies. As such this article paves the way for larger scale comparisons. In addition, we suggest to test the hypothesis that a market's informational efficiency is perturbed by traders moving in or out of a market.

6. REFERENCES

- [1] J. J. Angel. Market mechanics: A guide to U.S. stock markets. Technical report, The Nasdaq Stock Market Educational Foundation, Inc., 2002.
- [2] P. Bajari and A. Hortacsu. The winner's curse, reserve prices, and endogenous entry: empirical insights from ebay auctions. *RAND Journal of Economics*, pages 329–355, 2003.
- [3] C. Barrot, S. Albers, B. Skiera, and B. Schäfers. Vickrey vs. ebay: Why second-price sealed-bid auctions lead to more realistic price-demand functions. *International Journal of Electronic Commerce*, 14(4):7–38, 2010.
- [4] A. Cowles. Can stock market forecasters forecast? *Econometrica*, 1(3):209–324, 1933.
- [5] J. Fox. *The myth of the rational market: a history of risk, reward, and delusion on Wall Street*. Harper Paperbacks, 2011.
- [6] J. Huber. 'J'-shaped returns to timing advantage in access to information - Experimental evidence and a tentative explanation. *Journal of Economic Dynamics and Control*, 31(8):2536 – 2572, 2007.
- [7] J. Huber, M. Kirchler, and M. Sutter. Is more information always better? Experimental financial markets with cumulative information. *Journal of Economic Behavior and Organization*, 65(1):86 – 104, 2008.
- [8] M. Kaisers, K. Tuyls, S. Parsons, and F. Thuijsman. An evolutionary model of multi-agent learning with a varying exploration rate. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems*, 2009.
- [9] M. Kirchler. Partial knowledge is a dangerous thing - On the value of asymmetric fundamental information in asset markets. *Journal of Economic Psychology*, 21:643–658, 2010.
- [10] B. Malkiel. The efficient market hypothesis and its critics. *Journal of Economic Perspectives*, 17(1):59–82, 2003.
- [11] J. McMillan. Selling spectrum rights. *The Journal of Economic Perspectives*, 8(3):145–162, 1994.
- [12] J. Niu, K. Cai, S. Parsons, M. Fasli, and X. Yao. A grey-box approach to automated mechanism design. *Electron. Comm. Res. Appl.*, 2012. To appear.
- [13] S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double-auction market designs. *Agent-Mediated Electronic Commerce VI. Theories for and Engineering of Distributed Mechanisms and Systems*, pages 101–114, 2005.
- [14] Y. Shoham, R. Powers, and T. Grenager. If multi-agent learning is the answer, what is the question? *Artificial Intelligence*, 171(7):365–377, 2007.
- [15] B. Tóth and E. Scalas. The value of information in financial markets: An agent-based simulation. In J. Huber and M. Hanke, editors, *Information, Interaction, and (In)Efficiency in Financial Markets*. Linde Verlag, 2007.
- [16] B. Tóth, E. Scalas, J. Huber, and M. Kirchler. Agent-based simulation of a double-auction market with heterogeneously informed agents. In *Potentials of Complexity Science for Business, Governments, and the Media*, 2006.
- [17] B. Tóth, E. Scalas, J. Huber, and M. Kirchler. The value of information in a multi-agent market model. *Eur. Phys. J. B*, 55:115–120, 2007.
- [18] K. Tuyls and S. Parsons. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence*, 171(7):406–416, 2007.
- [19] W. Walsh, R. Das, G. Tesauro, and J. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002.