

An Evolutionary Game-Theoretic Comparison of Two Double-Auction Market Designs

Steve Phelps¹, Simon Parsons², and Peter McBurney¹

¹ Department of Computer Science, University of Liverpool,
Liverpool L69 7ZF UK

{sphelps, p.j.mcburney}@csc.liv.ac.uk

² Department of Computer and Information Science, Brooklyn College,
City University of New York, Brooklyn NY 11210 USA
parsons@sci.brooklyn.cuny.edu

Abstract. In this paper we describe an analysis of two double auction markets—the clearing house auction and the continuous double auction. The complexity of these institutions is such that they defy analysis using traditional game-theoretic techniques, and so we use heuristic-strategy approximation to provide an approximated game-theoretic analysis. As well as finding heuristic-strategy equilibria for these mechanisms, we subject them to an evolutionary game-theoretic analysis which allows us to quantify which equilibria are more likely to occur. We then weight the design objectives for each mechanism according to the probability distribution over equilibria, which allows us to provide more realistic estimates for the efficiency of each mechanism.

1 Introduction

A double-auction mechanism is a generalization of an auction in which both buyers and sellers are allowed to exchange offers simultaneously. Since double-auctions allow dynamic pricing on both the supply side and the demand side of the marketplace, their study is of great importance, both to theoretical economists, and those seeking to implement real-world market places. On the one hand, economists who are interested in theories of price formation in idealized models of general markets have often turned to exchange-like models such as Walrasian tâtonnement, to describe and understand the price-formation process [2], and on the other, variants of the double-auction are used in large real-world exchanges to trade commodities in marketplaces where supply and demand fluctuate rapidly, such as markets for stocks, futures, and their derivatives [7].

However, the models of exchanges traditionally used by economists in general equilibrium theory are often simplified for the purposes of analytical tractability to such an extent that they are of scant relevance to the designers of real-world exchanges, and even, it is sometimes argued, of scant relevance to the theoretical modelling of markets. For example, one important simplification often made is that the number of agents participating in a market is very large; this simplification allows relative market power and consequent *strategic effects* to be ignored. However, in many real-world marketplaces, such as deregulated wholesale electricity markets, there may be relatively few competitors on one or both sides of the market. With small numbers of participants, general

equilibrium models break down because they fail to allow for market power, and the potential gains of strategic behavior, of participants.

An alternative approach is a sophisticated micro-theory of marketplaces called *auction theory*, in which the rational behavior of individual agents faced with different pricing institutions is analyzed using game theoretic techniques. Whereas neoclassical equilibrium theory often abstracts away from the details of individual agents, game-theoretic models allow economists to build sophisticated micro-models of individual agents' reasoning and preferences. In many scenarios, especially in analyzing single-sided monopoly markets, these models have been spectacularly successful to the extent where they have been directly applied to the design of real-world auctions for high-value government and corporate assets [9]. However, in other practical scenarios, especially when it comes to analyzing and designing double-sided markets, such as exchanges, there are still a number of problems with the theory, which we shall briefly review.

Auction-theorists typically analyze a proposed market institution by defining a set of design objectives, and then proceed to show that these design objectives are brought about when rational agents follow their best strategies according to a game-theoretic analysis. The typical design objectives considered by auction-theorists are:

Allocative efficiency: The outcome of using the mechanism should be optimal in some defined sense, for example, the total surplus generated should equal the available surplus in competitive equilibrium.

Budget balance: No outside subsidy inwards or transfers outwards are required for a deal to be reached.

Individual rationality: The expected net benefit to each participant from using the mechanism should be no less than the net benefit of any alternative.

Strategy-proofness: Participants should not be able to gain an advantage from non-truth-telling behavior.

In many applications, auction-theory demonstrates the existence of market mechanisms that satisfy all of these properties when agents follow rationally prescribed bidding strategies. However, the impossibility result of [13] demonstrates that no *double-sided* auction mechanism can be simultaneously efficient, budget-balanced and individually-rational. Moreover, many of the underpinnings of the theory assume that designers' interests are restricted to only the aforementioned properties. For example, the revelation principle states that, without loss of generality, we may safely restrict attention to mechanisms in which agents reveal their types truthfully. However, this result does not take into account the potential cost or other impracticalities of polling agents for their type information. Once minimizing the cost of revelation is introduced as a design objective, the revelation principle ceases to hold, because there may exist partial-revelation mechanisms with non-truthful equilibria which sacrifice strategy-proofness for expedience of revelation. This is of more than academic interest, since in real-world electronic exchanges it is rarely possible to poll *all* agents for their valuations before clearing the market; hence the *continuous* double-auction, in which we execute the clearing operation as new offers arrive, thus increasing transaction throughput at the expense of strategy-proofness.

In designing market places, as with any other engineering problem, we often need to make such tradeoffs between different objectives depending on the exact requirements

and scenario at hand. We can often satisfactorily solve such multi-objective optimisation problems, provided that we have some kind of quantitative assessment of each objective, yet classical auction-theory provides only a binary yes or no indication of whether each of its limited design objectives is achievable, making it extremely difficult to compare the different trade-offs.

Further complications arise when we attempt to use auction-theory to analyze existing (“legacy”) market institutions. Exchanges such as the London Stock Exchange have been in existence far longer than game-theory and auction-theory, thus, unsurprisingly, the original rules of the institution were not necessarily based on sound game-theoretic or auction-theoretic principles. Moreover, it is unrealistic to expect that core financial institutions such as these radically alter their rules overnight in response to the latest fashionable developments in auction-theory or game-theory. Rather, it may be more salient to view financial institutions *evolving* gradually and incrementally in response to a changing environment [12]. Similarly, agents participating in these institutions do not necessarily instantaneously and simultaneously adjust their trading behavior to the theoretical optimum strategy; for example, adoption of a new trading strategy may spread through a population of traders as word of its efficacy diffuses in a manner akin to mimetic evolution.¹ Thus, we may think of the institutions we see today as the outcome of a *co-evolutionary* adaptation between financial institutions on the one hand, and trading strategies on the other.

The issue of legacy institutions has ramifications for auction-design; in these contexts the choice of adjustments to the auction rules may be tightly constrained by existing infrastructure, both physical and social, thus it may be necessary to examine the *attainability* of equilibria under the new design given existing strategic behavior in the legacy design. Classical auction theory relies on classical game-theory which in turn says nothing about the dynamics of adjustment to equilibrium.

For such applications, we need to turn to models of evolution and learning in strategic environments; models that we collectively categorize under the banner of *evolutionary game theory*. Models of learning and evolution as applied to agents’ strategies are not new. Where our approach differs, however, is in the application of models of learning and evolution to the market mechanism itself, a new field we call *evolutionary mechanism design* [16,3].

In this paper, we extend our previous work on evolutionary mechanism design by describing a more sophisticated means of analyzing the performance of a mechanism. Previously we have either evolved trading strategies along with the mechanism [16], or used a single heuristic bidding strategy [17]. Here we use a mix of heuristic strategies, and describe a rigorous and fully automated way of evaluating a mechanism using this mixture. We start in Section 2 with a description of the mechanisms we are studying here, and, in Section 3 with use of several heuristic strategies. Then, in Section 4, we describe our experimental set up, and in Section 5 how we use these results to establish the evolutionary behavior of the markets. Section 6 gives results, and Section 7 analyses them before Section 8 describes the work that we will pursue next.

¹ The adoption by derivatives traders of the Black-Scholes equation for option pricing provides an example [10].

2 The Continuous Double Auction Verses the Clearing-House

In a typical exchange, the market institution attempts to match offers to buy with offers to sell in such a way that the overall surplus extracted from the market is maximized. If offers are considered as signals of agents' valuations for a resource, and assuming agents signal truthfully, then an auctioneer can maximize allocative efficiency by matching the highest buy offers with the lowest sell offers. In this paper we compare two types of exchange:

- a market in which trades are executed as new offers arrive, and
- a market in which we wait for all traders to place offers before clearing the market.

Following the terminology of [6], we refer to the former as the continuous double-auction (CDA) and the to the latter as the clearing-house (CH).

On casual inspection of the CDA, we might expect that it is designed according to the revelation principle, and so should maximize allocative efficiency when agents signal truthfully. Surprisingly, however, it turns out that surplus extraction in a CDA is extremely *poor* under direct revelation—typical values are approximately 80 percent, which is extremely low compared with outcomes of almost 100 percent which are observed with the non-truthful strategies that are actually adopted by real traders.

The reason for this poor efficiency is easy to spot; the continuous clearing rule results in myopic matching; when the clearing operation is performed the auctioneer has only a partial view of the aggregate supply and demand in the market place. In order to maintain a high throughput of actual transactions, the auctioneer impatiently clears the market before every trader has the opportunity to place their bid. The extremely surprising thing about this institution, however, is that rational agents acting locally to maximize their own profit are able to compensate for this efficiency loss by placing extra-marginal, non-truthful bids, which collectively result in high-efficiency outcomes.

Much analysis of the CDA has focused on showing that although the CDA is not a direct revelation mechanism (DRM), it can be considered an almost-DRM by virtue of the fact that trading strategies with only a minimal amount of intelligence are able to extract high surpluses from the market [4]. However, such approaches are unsatisfactory because they fail to demonstrate that such minimalist strategies are *dominant* against more sophisticated strategies.

Ideally, we would like to find the game-theoretic solution for the CDA, and show that although truth-telling or other minimalist strategies are not dominant, we can still find the theoretical mix of strategies that are best-responses to each other, and demonstrate that the institution performs well in game-theoretic equilibria. However, even at this point, the CDA along with other variants of the double-auction market, confounds auction theorists by admitting of no clear equilibrium solution [19].² Hence in the absence of robust analytical tools, much analysis of this institution has used an ad-hoc mixture of computer simulation and laboratory experiments [7]. These techniques are invaluable, since they are able to faithfully incorporate many of the complex details of

² Though this reference is dated, to the best of our knowledge it is still the case that the CDA has no such solution.

the market institution which lead to intractability under conventional analysis. However, the results thus obtained are often criticised for being difficult to generalize in the absence of compelling models that explain the observed outcomes.

Recently, however, techniques have been developed that combine simulation-based approaches with an approximated game-theoretic analysis. In the following sections, we describe and then adopt the technique proposed by Walsh and colleagues [23]. However, our work extends the scope of Walsh *et al.*'s use of their technique. Whereas the original work focuses on designing *strategies* for a given institution, specifically the CDA institution, we build on this work by applying the same technique for *mechanism* design issues, using it to compare the CDA and CH institutions.

3 Heuristic-Strategy Approximation

Walsh *et al.* introduce an approximation technique for analyzing games such as the CDA where the sheer size of the strategy and player-type spaces makes an exhaustive game-theoretic solution impractical [23].

3.1 Basic Approach

The central idea is simple. Rather than considering every possible pure strategy and type in the multi-stage game, Walsh *et al.* simplify the analysis by considering a limited number of high-level *heuristic* strategies, such as Cliff's *Zero-Intelligence Plus* (ZIP) strategy [4], and treat these high-level strategies as if they were simple pure strategies in a normal form game. For small numbers of players and high-level strategies, this gives rise to a relatively small normal form game payoff matrix which is amenable to game-theoretic solution. This *heuristic* payoff matrix is calibrated by running many simulations of the market game; variations in payoffs due to different player-types are averaged over many samples of type information resulting in a single mean payoff to each player for each entry in the payoff matrix. Players' types are assumed to be drawn independently from the same distribution, and an agent's choice of strategy is assumed to be independent of its type, which allows the payoff matrix to be further compressed, since we simply need to specify the number of agents playing each strategy to determine the expected payoff to each agent. Thus for a game with k strategies, we represent entries in the heuristic payoff matrix as vectors of the form

$$\mathbf{p} = (p_1, \dots, p_k)$$

where p_i specifies the number of agents who are playing the i th strategy. Each entry $p \in P$ is mapped onto an outcome vector $q \in Q$ of the form

$$\mathbf{q} = (q_1, \dots, q_k)$$

where q_i specifies the expected payoff to the i th strategy. For a game with n agents, the number of entries in the payoff matrix is given by

$$s = \frac{n^k - 1}{(k - 1)!} \quad (1)$$

For small n and small k this results in payoff matrices of manageable size; for $k = 3$ and $n = 6, 8,$ and 10 we have $s = 28, 45,$ and 66 respectively. For very large n the game becomes intractable, but this is not of major concern since our interest is specifically in markets with small numbers of traders where strategic effects are likely to be prominent.

3.2 Choice of Heuristic Strategies

For moderately large values of k , combinatorial explosion rapidly leads to intractability. This constraint is of more concern than the constraint on n , since in any realistic trading environment we might, a priori, expect agents to be confronted with a vast number of high-level strategies from which to choose. There are, for example, many automated strategies that have been proposed in the literature [4,5,8,18,22]. However, there is evidence to show that in many real-life market scenarios traders choose from a limited number of heuristic strategies. For example, [15] discusses the observed strategic interaction between human agents and two predominant automated bidding strategies commonly used on two real-world auction institutions with different auction designs (Amazon and eBay).

Following these results, we base our work in this paper on the premise that we are modeling the effect of the adoption of automated trading agents in the CDA and CH markets. Thus we compare the behavior of traders using a well known automated strategy for the double-auction [18], and one that has been developed to emulate human strategic behavior in market settings [5]. By comparing these representative heuristic strategies we hope to gain insight into whether non-homogeneous populations of human and agent-based traders are strategically stable, and the likely market outcomes when human and agent-based traders interact. In addition, because we are interested in the strategy-proofness of the mechanisms themselves under different conditions, we also introduce the truth-telling strategy. If a mechanism is strategy proof, it should not be possible to do better than when truthfully report one's limit price. Thus we have $k = 3$ heuristic strategies, which is well within the limits of tractability for the Walsh approximation technique.

4 Experimental Setup

In order to compare the CDA and CH, we must first generate a heuristic payoff matrix for each institution by sampling many simulations of the market game. We made use of the JASA auction simulator³ which implements a CDA marketplace as described in [22], as well as a CH marketplace where the market is not cleared until offers from all agents have been received.

In order to model human-like trading behavior, we adopt a trading strategy based on a modified version of the Roth-Erev learning algorithm [5] as described in [14], which we abbreviate *RE*. This is the same version of the Roth-Erev algorithm that we have used in our previous work [16,17]—basically a reinforcement learning approach that builds up a probability distribution over the space of possible bids. We pit this against a strategy based on ZIP, but modified for persistent-shout markets as described by Preist

³ <http://www.csc.liv.ac.uk/~sphelps/jasa>

and Van Tol [18], which we abbreviate PvT , and the truth-telling strategy which simply bids at the agent's limit price, which we abbreviate TT .

As in [23], at the start of each game half the agents are randomly assigned as buyers and the remainder are chosen as sellers. For each run of the game, we choose limit prices from the same uniform distribution as [23], but limit prices remain fixed across periods in order to allow agents to attempt to learn to exploit any market-power advantage in the supply and demand curves defined by the limit prices for that game (This is common in much experimental work in this area [4,21], and makes it possible for both artificial traders and humans to exploit memory to quickly converge on trade prices). Additionally, although we discard limit-prices which do not yield an equilibrium price, we do not ensure that a minimum quantity exists in competitive equilibrium as this introduces a floor effect which fails to expose the inferior efficiency of a CDA. We use the Mersenne Twister random number generator [11] to draw all random values used in the simulation. Each entry in the heuristic payoff matrix is computed by averaging the payoff to each strategy across 2000 simulations.

5 Dynamic Analysis

Once the heuristic payoff matrix has been computed, we can subject it to a game-theoretic analysis. In conventional mechanism design, we solve the game by finding either a dominant strategy or the Nash equilibria: the sets of strategies that are best-responses to each other. However, because classical game-theory is a static analysis, it is not able to make any predictions about which equilibria are more likely to occur in practice. Such predictions are of vital importance in mechanism design problems. Since our design objectives depend on outcomes, we should give more consideration to outcomes that are more likely than low probability outcomes. For example, if there is a Nash equilibrium for our mechanism which yields very low allocative efficiency, we should not worry too much if this equilibria is extremely unlikely to occur in practice. On the other hand, we should give more weight to equilibria with high probability.

As in [23], we use *evolutionary* game-theory [20] to model how agents might gradually adjust their strategies over time as they learn to improve their behavior in response to their payoffs. We use the replicator dynamics equation

$$\dot{m}_j = [u(e_j, \mathbf{m}) - u(\mathbf{m}, \mathbf{m})] m_j \quad (2)$$

where \mathbf{m} is a mixed-strategy vector, $u(\mathbf{m}, \mathbf{m})$ is the mean payoff when all players play \mathbf{m} , and $u(e_j, \mathbf{m})$ is the average payoff to pure strategy j when all players play \mathbf{m} , and \dot{m}_j is the first derivative of m_j with respect to time. Strategies that gain above-average payoff become more likely to be played, and this equation models a simple *co-evolutionary* process of mimicry learning, in which agents switch to strategies that appear to be more successful. For the three heuristic strategies that we have chosen to analyze, we can interpret this process as modeling the potential uptake of ZIP-like automated trading agent technology; for example, managers bidding using human-like trading strategies may switch to a ZIP-like strategy if they observe a rival firm obtaining better than average profits by using automated trading agents.

For any initial mixed-strategy we can find the eventual outcome from this coevolutionary process by solving $\dot{m}_j = 0$ for all j to find the final mixed-strategy of the converged population. Unlike co-evolutionary approaches that use evolutionary computing to do the search, for instance [16,1], this model has the attractive properties that:

- all Nash equilibria of the game are stationary points under the replicator dynamics; and
- all focal points of the replicator dynamics are Nash equilibria of the evolutionary game.

Thus the Nash equilibrium solutions are embedded in the stationary points of the direction field of the dynamics specified by equation 2. Although not all stationary points are Nash equilibria, by overlaying a dynamic model of learning on the equilibria we can see which solutions are more likely to be discovered by *boundedly-rational* agents. Those Nash equilibria that are stationary points at which a larger range of initial states will end up, are equilibria that are more likely to be reached (assuming an initial distribution that is uniform).

We capture this idea of “range of initial states” with the notion of a *basin of attraction*. The basin of attraction for a stationary point is the range of mixed strategies within which all strategies will, under the replicator dynamics, lead to the stationary point. The bigger the basin, the bigger the region of strategy-space which leads to the attractor, and hence the stronger the attractor.

6 Results

Since $\sum m_i = 1$, each vector \mathbf{m} lies in the unit-simplex. For $k = 3$ strategies we can project the unit-simplex onto a two dimensional space and then identify the switching between strategies. We plot this switching in Figures 1—4 which show plots of the *direction field* defined by equation 2 for each institution. The direction field gives us a map which shows the trajectories of strategies of learning agents engaged in repeated interactions, from a random starting position. Thus, for Figure 1, each agent participant has a starting choice of 3 pure strategies (*TT*, *RE* and *PvT*) and any mixed (probabilistic) combination of these three. The pure strategies are indicated by the 3 vertices of the simplex (triangle), while mixed strategies are indicated by points on the boundaries or in the middle of the simplex.

An agent is assigned a random (mixed or pure) strategy to start, and then progressively adjusts this strategy over time in repeated interactions as a result of the learning mechanism described by Equation 2. The paths shown in Figure 1 trace this sequence of adjustments. In order not to overload the display, we have not placed arrows on these paths, but the overwhelming majority of paths start inside the simplex and head outwards, towards the edges and the three vertices. This indicates that the three pure strategies act as attractors for randomly-selected mixed starting strategies. The set of oriented paths leading to each vertex indicates the basin of attraction of the corresponding pure strategy. We can assess the relative likelihood of one strategy relative to another by comparing the size of their respective basins of attraction.

Each plot shows trajectories generated from 250 randomly sampled initial \mathbf{m} vectors. For now, we assume that every initial mixed-strategy is equally likely to be adopted

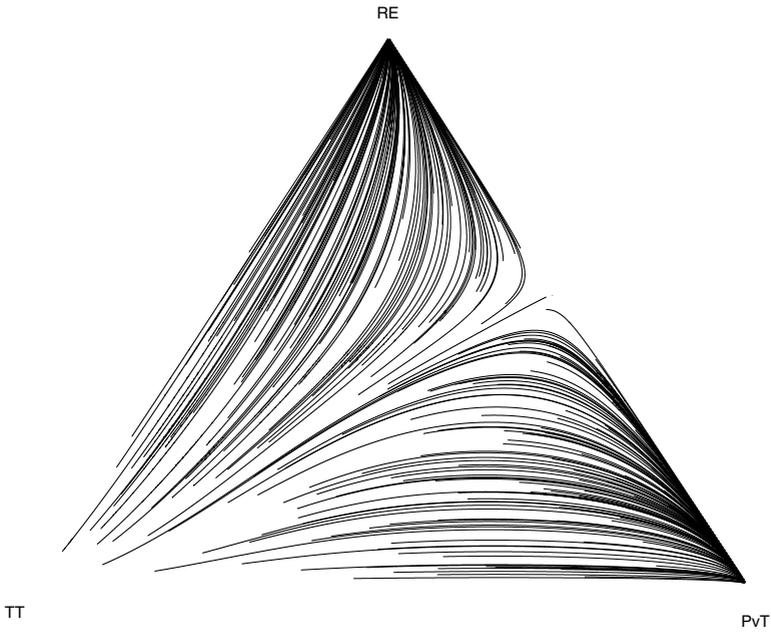


Fig. 1. Replicator dynamics direction field for CH with 6 agents

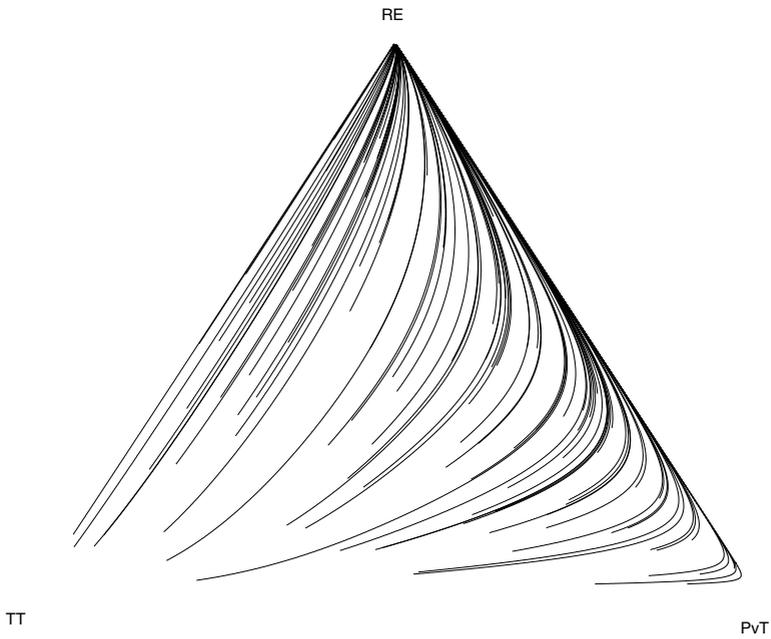


Fig. 2. Replicator dynamics direction field for CDA with 6 agents

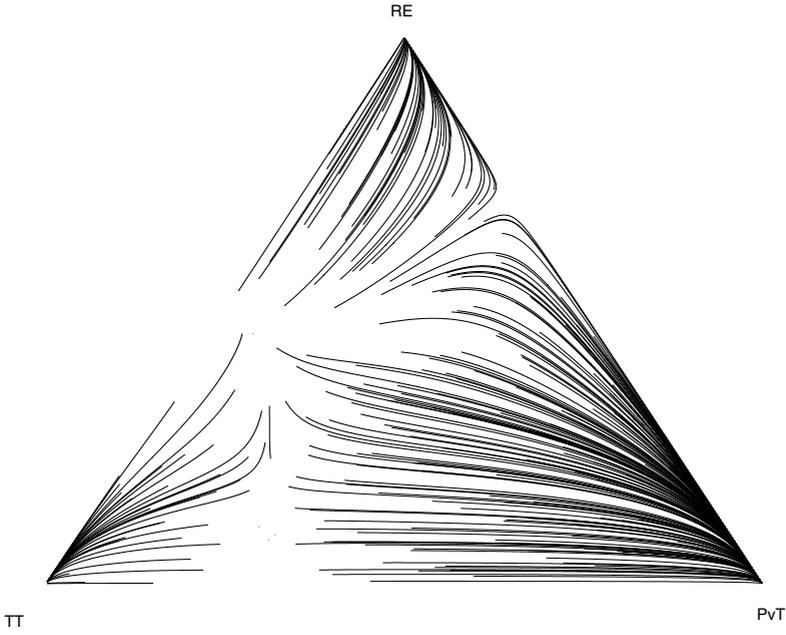


Fig. 3. Replicator dynamics direction field for CH with 8 agents

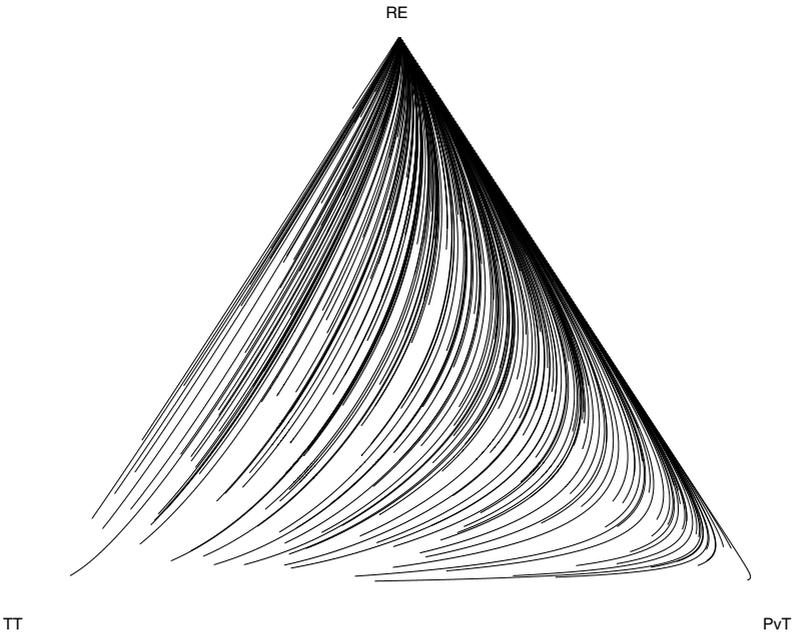


Fig. 4. Replicator dynamics direction field for CDA with 8 agents

Table 1. Pure-strategy equilibria probability distribution for 6 agents

Equilibrium	CH probability	payoff	CDA probability	payoff
TT	0.00	1.00	0.00	0.87
RE	0.50	1.00	1.00	0.98
PvT	0.50	1.00	0.00	0.93

Table 2. Pure-strategy equilibria probability distribution for 8 agents

Equilibrium	CH probability	payoff	CDA probability	payoff
TT	0.17	1.00	0.00	0.85
RE	0.18	1.00	1.00	0.98
PvT	0.65	1.00	0.00	0.92

as a starting-point for the co-evolutionary process, and so we randomly sample the initial values of m from a uniform distribution and plot their trajectories as they evolve according to equation 2.

To automate the analysis of institutions, we need to be able to provide some metric that allows us to quantify their performance in this kind of analysis. In other words, we would like to measure the size of the basin of attraction of each equilibria, in order to arrive at a probability of the equilibria actually occurring. Different equilibria will yield different outcomes and different values of our design objectives, such as market efficiency, and we would like to weight these according to their likelihood.

Tables 2 and 1 show the stationary points of 1000 randomly sampled trajectories together with the proportion of trajectories that terminate at that point. Given the random start, this probability is an estimate of the probability of each equilibrium. In the absence of a static analysis, we discount the stationary-points that occur with less than 1% probability. Since we know the payoffs of the various points from the heuristic payoff matrix, we can then compute expected payoffs, which are also shown in the table. With probabilities over outcomes, we are in a position to assess the design of each mechanism.

7 Discussion

First of all it is clear that TT is not dominant, and hence neither the CH or CDA mechanism is strategy-proof. However, it is interesting to note that when we increase the number of agents from 6 to 8 in a CH, we see the emergence of a truth-telling equilibrium. This agrees with the approximate analysis presented in [19], and suggests that truth-telling may become a strategy adopted by more traders as the market grows larger.

In a CH market, we see that the most likely strategy to be played is the ZIP-like trading agent strategy, whereas in a CDA, the human-like RE strategy is dominant.

As expected from our discussion above, we see that payoffs under truthful bidding in a CDA are relatively low; 86% in this case. This might suggest that the CDA itself has a rather low efficiency. However, in order to assess the efficiency of the CDA we must take into account the fact that truth-telling is not an equilibrium. Since the RE strategy

is dominant we should assume that all agents will eventually use this strategy, and so our efficiency will be equal to the pure-strategy payoff for RE, which in this case is 98%. In order to calculate the efficiency for the 8-agent CH market, we need to take into account that there are three possible pure-strategy equilibria. In this particular scenario, each equilibria results in the same efficiency of 100%, so we can conclude that the CH market will yield 100% efficiency in all eventualities.

Although the CDA yields lower surplus, it is not as inefficient as we might expect had we assumed that it was designed according to the revelation principle. As [6] points out, the main reason for choosing a CDA rather than a CH is to handle larger volumes of trade, and our results here suggest that this is a reasonable trade-off. Switching to a CDA from a CH as the New York Stock Exchange did in the 1860s, does not seem likely to entail a large loss of efficiency.

The above analysis assumes that all initial points in the mixed-strategy phase-space are equally likely to be selected. However, if we are in a situation where we are proposing to make changes to an existing “legacy” exchange with existing traders, our observations of current trading behavior in the legacy mechanism may influence our beliefs about likely behavior in any proposed altered version of the mechanism. For example, we may be tasked with assessing the likely impact in switching from a CH clearing rule to an exchange with continuous clearing. If we observe that traders bid truthfully in the existing mechanism, then when we come to perform the dynamic analysis for the new design, we may decide to weight our distribution of initial mixed-strategies in favor of truth-telling to reflect current observations.

8 Further Work

What we have demonstrated in this paper is an approach that provides an approximate game-theoretic analysis, involving equilibria over multiple heuristic strategies, for mechanisms that do not admit an analytical solution. This is fully automated, and gives us a means of *analyzing* and hence comparing auction mechanisms. Our previous work has demonstrated proof-of-concept for the idea of *evolving* auction mechanisms, for example using genetic programming to evolve parts of the pricing mechanism for a double auction market [17], establishing the quality of the market using a single heuristic strategy.

Since all parts of the approach we have detailed here are fully automated, it is possible to combine these two lines of work. This will enable us to create new auction mechanisms and then use the kind of analysis described here to rate them, thus searching the space of possible mechanisms while rigorously analyzing them. With our current implementation running on a 1.4Ghz Athlon AMD processor, it takes approximately 24-hours to generate the heuristic payoff matrix and perform the dynamic analysis for a single 10-agent mechanism. We hope to significantly reduce this evaluation cost by

- using a more selective sampling, as in [24] for example;
- further optimizing our code, and
- reducing the number of samples at the expense of accuracy whilst using an optimization algorithm that will be robust to the additional noise.

Finally, we recognize that the existence of dominant strategies, and the lack of mixed-strategy attractors are probably as a result of us not taking a representative set of heuristic strategies for these early experiments. Future work will extend the number of heuristic strategies that are used in our analysis.

With these techniques we will move closer to our overall goal of completely automated mechanism design.

Acknowledgements

This work was partially supported by NSF IIS-0329037. The authors would like to thank William Walsh, Gerry Tesauro and Karl Tuyls for their assistance, the workshop attendees for pointing out inaccuracies in a previous draft, and Bob Brisco for his helpful comments during a talk given at BTEExact.

References

1. P. J. Angeline and J. B. Pollack. Competitive environments evolve better solutions for complex tasks. In S. Forrest, editor, *Genetic Algorithms: Proceedings of the Fifth International Conference (GA93)*, 1993.
2. W.D.A. Bryant. Information, adjustment and the stability of equilibrium. Technical Report 6/2000, Department of Economics, Macquarie University, Sydney, Australia, 2000.
3. D. Cliff. Evolution of market mechanism through a continuous space of auction-types. Technical Report HPL-2001-326, HP Labs, 2001.
4. D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report, HP Labs, 1997. <http://www.hpl.hp.com/techreports/97/HPL-97-91.html>.
5. I. Erev and A. E. Roth. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–881, 1998.
6. D. Friedman. The double auction market institution: A survey. In Friedman and Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Studies in the Sciences of Complexity, chapter 1, pages 3–26. Westview, 1993.
7. D. Friedman and J. Rust, editors. *The Double Auction Market: Institutions, Theories and Evidence*, volume XIV of *Santa Fe Institute Studies in the Sciences of Complexity*. Westview, 1993.
8. S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behaviour*, 22:1–29, 1998.
9. P. Klemperer. How (not) to run auctions: the European 3G telecom auctions. *European Economic Review*, 46(4–5):829–845, 2002.
10. D. A. MacKenzie. An equation and its worlds: *Bricolage*, exemplars, disunity and performativity in financial economics. *Social Studies of Science*, 33:831–868, 2003.
11. M. Matsumoto and T. Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation*, 8(1):3–30, 1998.
12. P. Mirowski. *Machine Dreams: Economics Becomes a Cyborg Science*. Cambridge University Press, Cambridge, UK, 2002.
13. R. B. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

14. J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, October 2001.
15. A. Ockenfels and A.E. Roth. The timing of bids in internet auctions: market design, bidder behaviour, and artificial agents. *AI Magazine*, 2002.
16. S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Co-evolutionary mechanism design: a preliminary report. In J. Padget, D. C. Parkes, N. M. Sadeh, O. Shehory, and W. E. Walsh, editors, *Agent Mediated Electronic Commerce IV*, Bologna, Italy, 2002. AAMAS 2002.
17. S. Phelps, S. Parsons, E. Sklar, and P. McBurney. Using genetic programming to optimise pricing rules for a double auction market. In *Proceedings of the Workshop on Agents for Electronic Commerce*, Pittsburgh, 2003.
18. C. Preist and M. van Tol. Adaptive agents in a persistent shout double auction. Technical Report HPL-2003-242, Hewlett Packard Labs, 2003.
19. M. A. Satterthwaite and S. R. Williams. The Bayesian theory of the k-double auction. In Friedman and Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Studies in the Sciences of Complexity, chapter 4, pages 99–123. Westview, 1993.
20. J. M. Smith. *Evolution and the Theory of Games*. Cambridge University Press, 1982.
21. V. L. Smith. An experimental study of competitive market behaviour. *The Journal of Political Economy*, 70(2):111–137, April 1962.
22. G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Third ACM Conference on Electronic Commerce*, pages 206–209, 2001.
23. W. E. Walsh, R. Das, G. Tesauro, and J.O. Kephart. Analyzing complex strategic interactions in multi-agent games. In *AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents*, 2002.
24. W. E. Walsh, D. C. Parkes, and R. Das. Choosing samples to compute heuristic-strategy Nash equilibrium. In *Proceedings of the Fifth Workshop on Agent-Mediated Electronic Commerce*, 2003.