Using evolutionary game-theory to analyse the performance of trading strategies in a continuous double auction market

Kai Cai[1], Jinzhong Niu[1], and Simon Parsons[1,2]

[1] Department of Computer Science, Graduate Center, City University of New York, 365 5th Avenue, New York, NY 10016, USA
{kcai,jniu}@gc.cuny.edu

[2] 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. In agent-based computational economics, many different trading strategies have been proposed. Given the kinds of market that such trading strategies are employed in, it is clear that the performance of the strategies depends heavily on the behavior of other traders. However, most trading strategies are studied in homogeneous populations, and those tests that have been carried out on heterogeneous populations are limited to a small number of strategies. In this paper we extend the range of strategies that have been exposed to a more extensive analysis, measuring the performance of eight trading strategies using an approach based on evolutionary game theory.

1 Introduction

An *auction*, according to [2], is a market mechanism in which messages from traders include some price information — this information may be an offer to buy at a given price, a *bid*, or an offer to sell at a given price, an *ask* — and which gives priority to higher bids and lower asks. The rules of the auction determine, on the basis of the offers that have been made, the allocation of goods and money between traders. Auctions have been widely used in solving real-world resource allocation problems [9], and in structuring stock or futures exchanges [2]. Auctions are used for three reasons: (i) to increase the speed of sale by providing a public forum where buyers and sellers can look for trading partners (ii) to reveal information about traders' valuations allowing efficient transactions to take place, and (iii) to prevent dishonest dealing between the representatives of the seller and the buyer.

There are many different kinds of auction. One of the most widely used is the *double auction* (DA), in which both buyers and sellers are allowed to exchange offers simultaneously. The flexibility of double auctions means that their study is of great importance, both to theoretical economists and those seeking to implement real-world market places. The *continuous double auction* (CDA) is a DA in which traders make deals continuously throughout the auction (rather than, for example, at the end of the auction). The CDA is one of the most common exchange institutions, and is in fact the primary institution for trading of equities, commodities and derivatives in markets such as New York Stock Exchange (NYSE) and Chicago Mercantile Exchange.

Models of CDAs have been extensively studied using both human traders and computerized agents. Starting in 1955, Smith carried out numerous experiments investigating the behavior of such markets, documented in papers such as [22, 23]. The experiments in [22], for example, involved human traders and showed that even with limited information available, and only a few participants, the CDA can achieve very high efficiency, comes close to the theoretical equilibrium, and responds rapidly to changing market conditions. This result was in contrast to classical theory, which suggested that high efficiency would require a very large number of traders. Smith's results led to the suggestion that double auction markets are bound to lead to efficiency irrespective of the way that traders behave. Gode and Sunder [6] tested this hypothesis, introducing two automated trading strategies which they dubbed "zero-intelligence". The two strategies Gode and Sunder studied were *zero intelligence without constraint* (ZI-U) and *zero intelligence with constraint* (ZI-C). ZI-U traders make offers at random, while ZI-C traders make offers at random, but are constrained so as to ensure that traders do not make a loss (it is clear that ZI-U traders can make a loss, and so can easily lead to low efficiency markets). In the experiments reported in [6], the ZI-C traders gained high efficiency and came close enough to the performance of human traders that Gode and Sunder claimed that trader intelligence is not necessary for the market to achieve high efficiency and that only the constraint on not making a loss is important[3].

This position was attacked by Cliff and Bruten [1], who showed that if supply and demand are asymmetric, the average transaction prices of ZI-C traders can vary significantly from the theoretical equilibrium[4]. They then introduced the *zero intelligence plus* (ZIP) trader, which uses a simple machine learning technique to decide what offers to make based on previous offers and the trades that have taken place. ZIP traders outperform ZI-C traders, achieving both higher efficiency and approaching equilibrium more closely across a wider range of market conditions (though [1][page 60] suggests conditions under which ZIP will fail to attain equilibrium), prompting Cliff and Bruten to suggest that ZIP traders embodied the minimal intelligence required.

A range of other trading algorithms have been proposed — including those that took part in the Santa Fe double auction tournament [18, 19], the reinforcement learning *Roth-Erev* approach (RE) [17] and the expected-profit maximizing *Gjerstad-Dickhaut* approach (GD) [5] — and the performance of these algorithms have been evaluated under various market conditions. However, many of the studies of trader behavior leave something to be desired. In particular, those described above, with the honorable exception of the Santa Fe tournament [18], concentrated on the efficiency of markets as a whole and on markets in which the population of traders was homogeneous (in other words they all used the same strategy for deciding what to bid).

---

[3] In fact, for the markets tested in [6], even the ZI-U traders achieved pretty high efficiency, they were just outperformed by ZI-C traders in this regard.

[4] The experiments in [6], while reflecting typical market conditions, might be considered to represent easy conditions from which to attain equilibrium. In contrast, the experiments in [22] show convergence to equilibrium from a much wider range of initial conditions.

Both of these aspects are unsatisfactory from the perspective of someone who is interested in deciding whether to use a specific automated trader in a given market. If you want to adopt a trading agent to bid on your behalf, you don't much care about the efficiency of the market. What you care about is the profit you will make, and you'll quite happily use a ZI-C trader if it makes you more profit than a ZIP trader. Furthermore, even if we look at profit, it is not enough to know what a given type of trader will do in a homogeneous population. You're only going to want to use that ZI-C trader if you know that it will get you a good profit across all possible combinations of traders that you will encounter (in a game-theoretic sense you'd like adopting the ZI-C trader to be a dominant strategy). Tesauro and Das addressed both these problems [24]. In their paper, they examined the profit generated by a modified version of GD (MGD), ZI-C, ZIP, and the Kaplan strategy [19] from the Santa Fe double auction tournament in both homogeneous populations and mixed populations. The mixed populations studied in [24] were made up of two different kinds of trader, with one trader of one type, and the remainder of the traders being of the second type[5].

One way to consider the results of the kind of study carried out in in [24] is as an analysis of the stability of a homogeneous population. If the analysis shows that a single trader using strategy A in a population of B traders gets a higher profit than a homogeneous population of traders using strategy B, then there is an incentive to introduce a single A trader into a homogeneous B population, and that population is not stable. However, this kind of analysis does not say whether introducing a second A trader, or a third, or a fourth will necessarily be appropriate. As a result, these "one-to-many" experiments, while they will tell us something about the relative merits of A and B, will not give us any idea of the optimal mixture of traders (or, alternatively, what is the best strategy to adopt given the existing mix). To get closer to identifying the optimal mix, Walsh *et al.* [26] adopted techniques from evolutionary game theory, and applied them to more complex mixtures of trading strategies than were used in [24], an approach that has become known as *heuristic strategy analysis*. In particular, one can compute plausible equilibria for heterogeneous populations, and thus identify combinations of trading strategies that are likely to be adopted (assuming that traders are picked from a limited pool of possible strategies).

This paper extends the work of [24] and [26] exploring a larger set of trading strategies, thus expanding our understanding of the interaction between trading strategies, and giving us a more complete understanding of the possible equilibria that may arise in a continuous double auction. Such an analysis can also provide the groundwork for learning new kinds of trading strategy, as illustrated in [13], as well as for evaluating new varieties of auction such as those in [11].

2 Preliminaries

In this section we describe the precise scenario that we analyse in the rest of the paper.

[5] The same kind of analysis was later carried used by Vytelingum *et al.* to evaluate their risk-based bidding strategy [25].

2.1 The market

We are concerned with a specific kind of continuous double auction market (CDA). We have a population of traders, each of which is either a buyer or seller. Buyers have a supply of money which they seek to exchange for a certain kind of good, and sellers have a supply of that good which they seek to exchange for money. Each trader has a *private value* that specifies the value that they place on each unit of the good. Once the market opens, buyers place *bids*, specifying to all other traders in the market the amount of money that they are willing to exchange for a unit of the good (though we deal with traders that wish to trade multiple units of the good, they do this sequentially). Sellers make *asks*, specifying the amount of money they require in exchange for a unit of the good. We use the terms *offer* and *shout* to mean either a bid or an ask.

The market is controlled by an auctioneer, who notes all the offers, and, as each offer is made (offers are made sequentially in the implementation we use) compares the highest bid with the lowest ask. If the highest bid is higher, or equal to, the lowest ask, the offers are *matched*, and the auctioneer establishes a *trade price* or *sale price*. The trade price is constrained to be no greater than the bid price and no less than the ask price — the auctioneer chooses the trade price to fall in this bid/ask spread[6]. A trader with an offer that is matched is obligated to make the exchange at the trade price. (The existence of the auctioneer, and the obligation to trade once offers have matched distinguish our setup from, for example, that in the Santa Fe tournament where traders identified matches for themselves, and could choose whether or not to exchange when matches occurred [18].) If a bid is higher than two or more asks, the auctioneer gives priority to the lower ask, and if an ask is made that is lower than two or more bids, the auctioneer gives priority to the higher bid.

2.2 The traders

The traders we consider in this paper are all automated — what economists would call *program traders*. Each trader uses a specific strategy to choose what offers to make. The trading strategies we study in this paper are a mixture of established strategies from the literature, and some that we came up with ourselves. Those from the literature are:

- Zero Intelligence with Constraint (ZI-C), as introduced by Gode and Sunder [6]. Traders employing this strategy submit offers that are generated randomly subject to a simple constraint. This constraint states that bids are drawn from a uniform distribution between the buyer's private value and a specified lower bound (typically 0) while asks are restricted to the range between seller's private value and a specified higher bound (a value higher than any trader thinks the good in question is worth).
- Zero Intelligence Plus (ZIP), as introduced in [1]. ZIP traders use a simple heuristic to adjust their offers. Broadly speaking, traders increase their profit margin[7] if recent market activity suggests that doing so will still allow them to trade, and reduce

[6] Typical rules for choosing where to set the trade price are to set it in the middle of the bid/ask spread, or to set it to the value of the earlier of the two offers to be made.

[7] The profit margin for a trader is the difference between their private value and their offer price.

their profit margin if recent market activity suggests they are making offers too far from where the market is trading. The traders employ a simple form of machine learning to adjust their offers, smoothing out fluctuations in the market.

- Truth-Telling (TT). Traders using this strategy submit shouts equal to their private value for the resource being traded. TT is an interesting strategy to experiment with since in strategy-proof markets[8] TT will be a dominant strategy. The failure of TT to dominate is thus an indication of the degree to which traders in a particular market can benefit by clever strategic behavior.
- Pure Simple (PS), is an inadvertent copy of the strategy "Gamer" which was an entrant in the Santa Fe tournament [19][page 90][9], and traders using PS bid a constant 10% below their private value. This is not a strategy that one would expect to perform well — Gamer placed 24th out of 30 entries in the Santa Fe tournament — but, like TT is a useful control, and one that comfortably out-performs TT. Indeed, as shown in [27], with the right choice of margin, PS can be very efficient.
- Roth-Erev (RE), introduced in [17], is a strategy that considers the problem of what offer to make as being a reinforcement learning problem. RE experiments, making offers and recording how many times they are successful, and makes choices based on the expected value of each possible offer, computed using the past probability of success. We set the free parameters of RE as described in [10].
- Gjerstad-Dickhaut (GD) as introduced in [5]. A GD trader makes its decision on what to offer based on previous offers, but unlike RE, GD uses offers made by all other traders. A GD trader uses this list of past offers to estimate the likelihood of any sensible bid (that is one in the gap between the highest bid and the lowest ask at the time the offer is made) being accepted, and uses this probability distribution to compute the offer with the highest expected profit.

Those we came up with are:

- Linear Gjerstad-Dickhaut (GDL). GD runs more slowly than other trading strategies that we have been using, and it spends most of its time computing the probability of offers being accepted — it computes this by fitting recent offers to a cubic equation, and then uses the cubic to define the cumulative probability of a given offer being accepted. Frustrated by the running time of experiments that used GD, we replaced the cubic with a piecewise linear approximation to create GDL, which runs considerably faster, hoping that the performance drop would not be too great.
- Estimated Equilibrium Price (EEP). If all traders are rational (in other words make profitable offers) and make offers around the theoretical equilibrium, then the market will be efficient. Thus bidding at the theoretical equilibrium is good for the market as a whole. We were interested to test whether bidding at the theoretical equilibrium is also good for individual agents and EEP is an attempt to evaluate this. EEP seeks to make offers at the theoretical equilibrium, estimating this as the mid-point of the highest accepted ask and the lowest accepted bid so far, and so our estimate of the equilibrium is similar to of [25].

[8] A strategy proof market, such as that discussed in [8], is one in which traders cannot manipulate results in their favor by misrepresenting the extent to which they value resources.

[9] The copy was inadvertent since we devised PS in ignorance of the existence of Gamer.

This is, clearly, not an exhaustive selection — we could, for instance, have included the RB strategy from [25] — but is a large enough set of strategies to be going on with.

Note that though the many of the strategies we use are *adaptive*, in the sense that the offers they make change over time in response to other offers, a given trader uses the same strategy throughout a given auction. This contrasts with the work of Posada [15, 16] which studies agents that are allowed to switch bidding strategy during an auction.

2.3 The simulation environment

All of the experiments reported here are based on the open-source JASA auction simulator [12], devised by Steve Phelps of the University of Liverpool. The current version of JASA implements a CDA marketplace much as described in [24] as well as all the trading strategies described above. In JASA the auction runs for a number of *days*, and each day is broken up into discrete *rounds*. In each round, every trader is selected to make an offer, and this selection takes place in a random order. At the end of every day, every trader has its initial allocation of goods and money replenished, so that trading on every day in a given experiment takes place under the same conditions, but trading strategies that record information will remember what took place in previous days.

We ran every experiment described here for five trading days, and each day consisted of 300 rounds. The private values of traders are drawn at the start of the first trading day of each experiment from a uniform distribution between 100 and 200. Every experiment was repeated 100 times.

3 Heterogeneous trading populations

In this section we describe the first series of experiments we carried out with mixed populations of traders. The methodology used for this series of experiments is that of [24], outlined above. For the first group of experiments we used 20 traders, 10 buyers and 10 sellers. For each of the eight trading strategies, we ran an experiment in which all but one agent used that strategy and the remaining agent used another strategy, carrying out one such "one-in-many" experiment for each of the other strategies. In other words, we tested every "one-in-many" combination. For all these experiments, we measured the average profit of traders using both the trading strategies under test.

Tables 1 and 2 show the results of "one-in-many" tests for the first group of experiments, those involving 20 agents. Note that the standard deviations of the payoffs are usually high, as a result of the fact that we are picking the private value of the "one" agent at random. As a result it is inevitable that there will be times when the "one" agent is an extra-marginal trader[10] because it has a low private value (the "one" agent is always a buyer) and in a market of savvy traders will not make any profit. Such occurrences will increase the standard deviation. Since the high standard deviations make direct comparisons of the profits difficult, we carried out hypothesis tests (in particular t-tests) to find out the confidence level for the "one" to "many" pairs of payoffs.

[10] An extra-marginal trader is one with a private value to the right of the intersection of the supply and demand curves for the market, and so should not trade if the market operates at its theoretical equilibrium.

	Many EEP	Many GD	Many GDL	Many PS	Many RE	Many TT	Many ZIC	Many ZIP
1-EEP	8.027	8.529: 10.038	8.33: 10.036	10.348: 9.121	11.837: 9.302	12.431: 7.584	12.383: 9.483	10.216: 9.695
stdev	(1.958)	(12.943): (0.696)	(12.514): (0.679)	(11.661): (0.852)	(13.955): (0.776)	(12.695): (0.78)	(13.957): (0.722)	(11.75): (0.636)
rel	-	$<$	$<$	$>$	$>$	$>$	$>$	$>$
conf	-	85.00%	90.00%	85.00%	95.00%	99.95%	97.50%	$< 75\%$
1-GD	9.29: 8.198	9.972	8.89: 10.018	12.249: 9.194	13.271: 9.308	13.492: 7.632	13.487: 9.475	11.6: 9.637
stdev	(13.189): (1.93)	(0.024)	(13.229): (0.699)	(15.651): (0.845)	(15.749): (0.806)	(15.864): (0.861)	(14.33): (0.755)	(13.313): (0.685)
rel	$>$	-	$<$	$>$	$>$	$>$	$>$	$>$
conf	75.00%	-	80.00%	95.00%	99.00%	99.95%	99.50%	90.00%
1-GDL	9.598: 8.289	9.028: 10.02	9.96	12.501: 9.177	13.449: 9.285	13.87: 7.605	13.806: 9.459	11.693: 9.646
stdev	(13.773): (1.801)	(13.316): (0.704)	(0.038)	(15.571): (0.837)	(15.824): (0.829)	(15.805): (0.859)	(14.753): (0.775)	(15.1): (0.768)
rel	$>$	$<$	-	$>$	$>$	$>$	$>$	$>$
conf	80.00%	75.00%	-	97.50%	99.50%	99.95%	99.75%	90.00%
1-PS	4.956: 8.108	5.405: 10.189	5.313: 10.191	9.281	8.926: 9.5	10.566: 7.719	9.281: 9.708	7.129: 9.838
stdev	(7.262): (2.005)	(7.821): (0.437)	(7.617): (0.421)	(0.349)	(9.333): (0.587)	(9.039): (0.7)	(8.884): (0.522)	(8.088): (0.509)
rel	$<$	$<$	$<$	-	$<$	$>$	$<$	$<$
conf	99.95%	99.95%	99.95%	-	$< 75\%$	99.75%	$< 75\%$	99.90%

Table 1. Profits for agents using different trading strategies in a 20 agent CDA market. The top line of each cell gives the average value of the profits — the value to the left of colon is the average profit of the one agent, the value to the right is the average profit of the majority populations. The second line gives the standard deviation of the profit. The third line indicates whether the "one" performs better ($>$) or worse ($<$) than the "many" on average. The fourth line gives the confidence in this relationship.

	Many EEP	Many GD	Many GDL	Many PS	Many RE	Many TT	Many ZIC	Many ZIP
1-RE	6.88: 8.32	7.046: 10.119	6.956: 10.118	9.814: 9.319	9.422	11.339: 7.69	10.236: 9.664	10.454: 9.652
stdev	(9.081): (1.542)	(9.555): (0.519)	(9.462): (0.508)	(10.278): (0.552)	(0.237)	(9.879): (0.758)	(10.914): (0.585)	(11.549): (0.651)
rel	$<$	$<$	$<$	$>$	-	$>$	$>$	$>$
conf	90.00%	99.90%	99.90%	$< 75\%$	-	99.95%	$< 75\%$	75.00%
1-TT	2.951: 8.291	3.514: 10.29	3.322: 10.294	6.175: 9.364	6.011: 9.565	7.755	5.95: 9.802	4.737: 9.928
stdev	(4.353): (1.929)	(5.189): (0.307)	(4.909): (0.29)	(5.827): (0.513)	(6.072): (0.463)	(0.595)	(5.373): (0.376)	(6.301): (0.398)
rel	$<$	$<$	$<$	$<$	$<$	-	$<$	$<$
conf	99.95%	99.95%	99.95%	99.95%	99.95%	-	99.95%	99.95%
1-ZIC	6.864: 8.485	7.657: 10.059	7.585: 10.065	9.282: 9.386	9.795: 9.437	10.503: 7.727	9.715	8.021: 9.743
stdev	(10.387): (1.451)	(11.024): (0.593)	(10.893): (0.583)	(11.969): (0.584)	(12.907): (0.711)	(12.04): (0.84)	(0.125)	(9.485): (0.563)
rel	$<$	$<$	$<$	$<$	$>$	$>$	-	$<$
conf	90.00%	97.50%	97.50%	$< 75\%$	$< 75\%$	97.50%	-	95.00%
1-ZIP	8.592: 8.325	9.397: 9.984	9.704: 9.977	10.091: 9.327	11.128: 9.385	12.736: 7.678	12.098: 9.554	9.712
stdev	(12.087): (1.558)	(14.36): (0.753)	(14.507): (0.771)	(11.059): (0.637)	(11.449): (0.624)	(12.204): (0.758)	(10.166): (0.524)	(0.132)
rel	$>$	$<$	$<$	$>$	$>$	$>$	$>$	-
conf	$< 75\%$	$< 75\%$	$< 75\%$	75.00%	90.00%	99.95%	99.00%	-

Table 2. Profits for agents using different trading strategies in a 20 agent CDA market. The top line of each cell gives the average value of the profits — the value to the left of colon is the average profit of the one agent, the value to the right is the average profit of the majority populations. The second line gives the standard deviation of the profit. The third line indicates whether the "one" performs better ($>$) or worse ($<$) than the "many" on average. The fourth line gives the confidence in this relationship.

These results give some suggestion of the complexities of bidding in continuous double auctions. If we think of Tables 1 and 2 as payoff matrices for the game where one player picks the strategy for the "one", and the other picks the strategy for the "many", we can immediately rule out TT as a choice — it is dominated. This is the same kind of analysis that is used in [25] to argue for the success of RB traders. However, we also found more complex relationships than in [24, 25]. Thus, once we have eliminated TT from consideration, PS can be eliminated as a strategy for the "one", as it performs worse than any of the "many" against which it might be played, but it works as a "many" strategy against ZIC. In a similar way, ZIC is not a great performer, but as a "many" strategy will outperform PS, and as "one" strategy will outperform RE. RE performs poorly as a majority strategy, but can generate higher profits than ZIP, a strong performer, when it is the "one" (though the low confidence we have for this results suggests that this performance is not consistent).

Looking at the high performing strategies, if an agent with a strategy other than GD or GDL is in an otherwise homogeneous GD or GDL populations, that agent will do better by switching to GD or GDL. In other words, GD and GDL come close to being dominant strategies for the "one". However each prevents the other from dominating. The performance of GDL is rather impressive — it even performs slightly better than GD does when it's the lone strategy amongst a population of PS, RE, TT, EEP or ZIC strategies (in most cases both in terms of the raw average payoff and confidence that it outperforms the general population). Thus, it seems that the switch from cubic to linear approximation might not only not hurt the strategy, but might even improve it.

When we look at slightly less well-performing strategies than GD and GDL, the situation is less clear. Indeed from Tables 1 and 2 it is hard to get a good feel for the relative merits of RE, ZIP and EEP. A lone RE trader will outperform a set of ZIP traders, a lone ZIP trader will outperform sets of EEP traders and RE traders, while a lone EEP trader will outperform sets of RE and ZIP traders.

4 Evolutionary game-theoretic analysis

Since we can't easily see how some combinations of strategies stack up against one another using the analysis in the previous section, we turn to a more sophisticated approach, *heuristic strategy analysis*. Heuristic strategy analysis was first proposed by Walsh et al. [26] precisely for the analysis of double auctions, and we have used it for this purpose in several papers [13, 14] though on a rather smaller scale than here.

4.1 Heuristic strategy analysis

The idea behind the heuristic strategy analysis is as follows. If we wanted to obtain a game theoretic solution to the continuous double auction, we would need to compute a payoff matrix that gives the expected outcome for an agent that bids in a particular way. Indeed, since there is no dominant strategy[11] we would need to compute such a payoff matrix for *all* possible offers or combinations thereof (since the CDA offers

[11] Unlike, for example, the case of the buyer's bid double auction [7].

multiple opportunities for making offers we would need to consider all possible offers that might be made at all opportunities). Clearly such a matrix would be extremely large, and that is why there is no analytical solution to the auction [21]. However, we can get around the need to consider all possible combinations of offers. Since there are a number of powerful strategies for computing the best offer to make — exactly the ones we have been studying so far — we can reasonably assume that each trader in the auction picks one of these *heuristic strategies* and lets that strategy pick offers. Under such an assumption, not only does the game we are trying to analyse become a single step game, but the number of possible strategies reduces to those that we know work well.

Now, for small numbers of players and heuristic strategies, we can construct a relatively small normal-form payoff matrix which we can analyse using game theory. This *heuristic payoff matrix* is calibrated by running many simulations of the auction. If we restrict the analysis to symmetric games in which each agent has the same set of strategies and the same distribution of private values (or *types* in the usual terminology of game theory), we can reduce the size of the payoff matrix, 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 present entries in the heuristic payoff matrix as vectors of the form:

$$p = (p_1, ...p_k) \tag{1}$$

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

$$q = (q_1, ...q_k) \tag{2}$$

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

$$s = \frac{(n + k - 1)!}{n!(k - 1)!} \tag{3}$$

For small n and small k this results in payoff matrices of manageable size. For $n = 20$, $k = 3$, as in the experiments we consider here, the symmetric payoff matrix contains just 231 entries.

Given the payoff matrix, we have a full description of a game in which traders pick between the heuristic strategies, and we can carry out an equilibrium analysis on that game. Any equilibria that we find are only equilibria for the game of choosing between heuristic strategies, not for the game of choosing a sequence of bids in a double auction — it is possible, for example, for traders to use different heuristic strategies than the ones we have analysed, in which case the equilibrium analysis will not help. However, as argued in [13], the equilibria of the heuristic strategy game are useful precisely because they only consider strategies that are commonly known and widely used. If we consider an exhaustive set of widely used strategies, we can be confident that no commonly known strategy will generate different equilibria from the ones we find, and thus the equilibria stand some chance of persisting until new trading strategies become established.

4.2 Evolutionary game theory

Now, even given the heuristic payoff matrix, standard game theory does not tell us which of the many possible Nash equilibrium strategies will result. *Evolutionary game theory* [3, 20] and its variants attack this problem by positing that, rather than computing the Nash strategies for a game using brute-force and then selecting one of these to play, traders are more likely to gradually adjust their strategy over time in response to to repeated observations of their own and others' payoffs. One approach to evolutionary game-theory uses the *replicator dynamics* equation to specify the frequency with which different pure strategies should be played depending on the payoffs of different strategies:

$$\dot{m}_j = [u(e_j, \boldsymbol{m}) - u(\boldsymbol{m}, \boldsymbol{m})]\, m_j \tag{4}$$

where $\boldsymbol{m}$ is a mixed-strategy vector, $u(\boldsymbol{m}, \boldsymbol{m})$ is the mean payoff when all players play $\boldsymbol{m}$, and $u(e_j, \boldsymbol{m})$ is the average payoff to pure strategy j when all players play $\boldsymbol{m}$, 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 process of learning by copying, in which agents switch to strategies that appear to be more successful[12]. For any initial mix of strategies we can find the eventual outcome from this *co-evolutionary* process by solving $\dot{m}_j = 0$ for all j to find the final mixed-strategy of the converged population. This model has the attractive properties that: (i) all Nash equilibria of the game are stationary points under the replicator dynamics; and (ii) all focal points of the replicator dynamics are Nash equilibria of the evolutionary game.

What this means is that the Nash equilibrium solutions are a subset of the stationary points of the direction field of the dynamics specified by equation 4. Although not all stationary points are Nash equilibria, we can use the direction field to see which solutions are more likely to be discovered by *boundedly-rational* agents. The Nash equilibria at which a larger number of initial states will end up, are equilibria that are more likely to be reached (assuming an initial distribution that is uniform, and that the replicator dynamics is an accurate reflection of the way that traders adjust their strategy[13]).

4.3 Results

We applied the analysis as described so far to sets of strategies we used in the "one-to-many" experiments, concentrating on the strategies which we felt had the most interesting interactions. Since the computational complexity of establishing the payoff matrix depends on not only the number of traders, but also on the number of strategies, we restricted our analysis to sets of three strategies (which also makes the results easier to visualize), and for every strategy vector p, allocated the given set of strategies randomly between all traders (so that a given strategy has equal probability of being used by a buyer or a seller). Some of the results we obtained may be found in Figures 1 and 2.

[12] Though they switch *between* auctions rather than in the middle as in [15, 16].

[13] Though the Nash equilibria cannot be disputed, the route by which they are reached is dependent upon the precise assumptions encoded in the replicator dynamics, and those, like all assumptions, are open to argument.

Figure 1 analyses the performance of GDL and GD. While the "one-to-many" experiments suggested that neither of these strategies dominates the other, the upper replicator dynamics plot in Figure 1 suggests that, at least in the presence of ZIP — which as [1] and our own analysis suggest is a pretty good strategy and thus a likely choice in trading scenarios — there is one equilibrium in which all traders adopt GDL, and there is another in which about half of the traders use GDL, the rest adopting GD. If we switch GD for a lesser strategy, such as EEP, as in the lower part of Figure 1, then the only equilibrium is when all traders adopt GDL.

The results in Tables 1 and 2 suggest that the relationship between ZIP and EEP deserves a little more attention since one EEP trader out performs the average ZIP trader when the latter are in a majority, while one ZIP trader will outperform the majority EEP traders. In other words, neither dominates the other. The upper part of Figure 2 shows us how this relationship plays out when the other possible strategy is TT. Here ZIP is powerful enough that it is a pure strategy equilibrium, but there is a second equilibrium in which roughly half of the traders use EEP and half use ZIP. The lower part of Figure 2 shows us that switching ZIP for RE allows EEP to become a pure strategy equilibrium and that RE is also a pure strategy equilibrium. Overall this suggests that, when faced with EEP, RE is a less powerful strategy than ZIP.

5 Conclusions

The main point of this paper is to report on work that has extended the analysis of the continuous double auction, and, in particular, the relative performance of trading strategies for making offers in the continuous double auction. As things stand, it is not clear whether there is a dominant strategy for the auction. However, if there is, then we will only discover it empirically, and the best way that we currently have for making this discovery is to continue to analyse the performance of different strategies against one another. The approach taken in this paper is one, we believe promising, way to do this. The "one-to-many" experiments that we started with allow us to identify pairs of strategies where one strategy does not dominate the other (that is when "one" of both strategies outperforms the "many" of the other). The heuristic strategy analysis experiments then home in on the relative merits of these strategies, giving us a way to compute equilibrium solutions for the continuous double auction under the assumption that traders are restricted to pick from a fixed set of trading strategies. The results we get are, as one would expect from a heuristic analysis, approximate, and not as exhaustive as the analysis of the double auction in [27]. However, unlike those of [27], our results are not restricted to a single trading strategy.

From this perspective, we can conclude three things. First, we can conclude that our analysis has shown, once again, the value of evolutionary game theory in analysing complex games. Second, we can conclude that the analysis has highlighted the powerful performance of our GDL variant of GD. Third, we can conclude that EEP, while not a winning strategy is also not a losing strategy in every situation. All of these results, though, should be taken with a pinch of salt — all performances in the CDA, as we have stressed above, are conditional on the mix of strategies present, and as [13] shows, it is perfectly possible to find (indeed, automatically generate) a strategy that beats GD.

However, given the dependence of results on the mix, the only course open to us is to keep expanding the set of strategies that are analysed in competition to each other, and with that aim our work is a straightforward extension of that of [24] and [26]. Of course, we can go further in this direction, and a natural way to do this is to extend the set of strategies with RB from [25] and the meta-strategy studied in [15, 16].

Finally, we should note that the research described here, like that of [24] and [26], only views matters from the perspective of the traders. The analysis is all couched in terms of the profits generated by different strategies — as described above, this is an analysis that is appropriate from the perspective of selecting a trader to operate on one's behalf. This research will not, in contrast, tell one much about the effect of the different trading strategies on the market as a whole. For that, one must turn to work like that of [4].

Acknowledgments

This work was supported by the National Science Foundation under grant IIS-0329037. We are grateful to Steve Phelps and Peter McBurney for generously sharing code and ideas with us, to Marek Marcinkiewicz for his work on JASA, to Jeff Mackie-Mason and Michael Wellman for helpful comments, and to the anonymous reviewers for helping us to improve the paper.

References

1. D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviours in market-based environments. Technical report, Hewlett-Packard Research Laboratories, 1997.
2. D. Friedman. The double auction institution: A survey. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 3–25. Perseus Publishing, Cambridge, MA, 1993.
3. D. Friedman. Evolutionary economics goes mainstream: A review of "The Theory of Learning in Games". *Journal of Evolutionary Economics*, 8(4):423–432, December 1998.
4. J. Gil-Bazo, D. Moreno, and M. Tapia. Price dynamics, informational efficiency and wealth distribution in continuous double auction markets. *Computational Intelligence*, 23(2):176–196, 2007.
5. S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
6. D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.
7. J. H. Kagel and W. Vogt. Buyer's bid double auctions: Preliminary experimental results. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity, chapter 10, pages 285–305. Perseus Publishing, Cambridge, MA, 1993.
8. R. P. McAfee. A dominant strategy double auction. *Journal of Economic Theory*, 56:343–450, 1992.
9. J. McMillan. *Reinventing the Bazaar: A Natural History of Markets*. W. W. Norton & Company, 2003.

10. 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, 2001.
11. J. Niu, K. Cai, S. Parsons, and E. Sklar. Reducing price fluctuation in continuous double auctions through pricing policy and shout improvement. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, 2006.
12. S. Phelps. Java auction simulator API. `http://sourceforge.net/projects/jasa/`.
13. S. Phelps, M. Marcinkiewicz, S. Parsons, and P. McBurney. A novel method for automated strategy acquisition in n-player non-zero-sum games. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*, Hakodate, Japan, 2006.
14. S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparision of two double auction market designs. In *Agent-Mediated Electronic Commerce VI*. Springer-Verlag, 2004.
15. M. Posada. Srategic software agents in a continuous double auction under dynamic environments. In *Intelligent Data Engineering and Automated Learning*, volume 4224 of *Lecture Notes in Computer Science*, pages 1223–1233. Springer Verlag, Berlin, 2006.
16. M. Posada, C. Hernández, and A. López-Paredes. Strategic behaviour in continuous double auction. In C. Bruun, editor, *Advances in Artificial Economies: The Economy as a Complex Dynamic System*, volume 584 of *Lecture Notes in Economics and Mathematical Systems*, pages 31–43. Springer Verlag, Berlin, 2006.
17. A. E. Roth and I. Erev. Learning in extensive-form games: Experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior*, 8:164–212, 1995.
18. J. Rust, J. H. Miller, and R. Palmer. Behaviour of trading automata in a computerized double auction market. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 155–199. Perseus Publishing, Cambridge, MA, 1993.
19. J. Rust, J. H. Miller, and R. Palmer. Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.
20. L. Samuelson. *Evolutionary Games and Equilibrium Selection*. MIT Press, Cambridge, MA, 1997.
21. M. A. Satterthwaite and S. R. Williams. The Bayesian theory of the k-double auction. In D. Friedman and J. Rust, editors, *The Double Auction Market: Institutions, Theories and Evidence*, pages 99–123. Perseus Publishing, Cambridge, MA, 1993.
22. V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70(2):111–137, April 1962.
23. V. L. Smith. Experimental auction markets and the Walrasian hypothesis. *The Journal of Political Economy*, 73(4):387–393, August 1965.
24. G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. In *Proceedings of the 3rd ACM Conference on Electronic Commerce*, 2001.
25. P. Vytelingum, R. K. Dash, E. David, and N. R. Jennings. A risk-based bidding strategy for continuous double auctions. In *Proceedings of the 16th European Conference on Artificial Intelligence*, pages 79–83, Valencia, Spain, 2004.
26. W. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *Proceedings of Workshop on Game-Theoretic and Decision-Theoretic Agents*, 2002.
27. W. Zhan and D. Friedman. Markups in double auction markets. Technical report, LEEPS, Department of Economics, University of Santa Cruz, 2005.

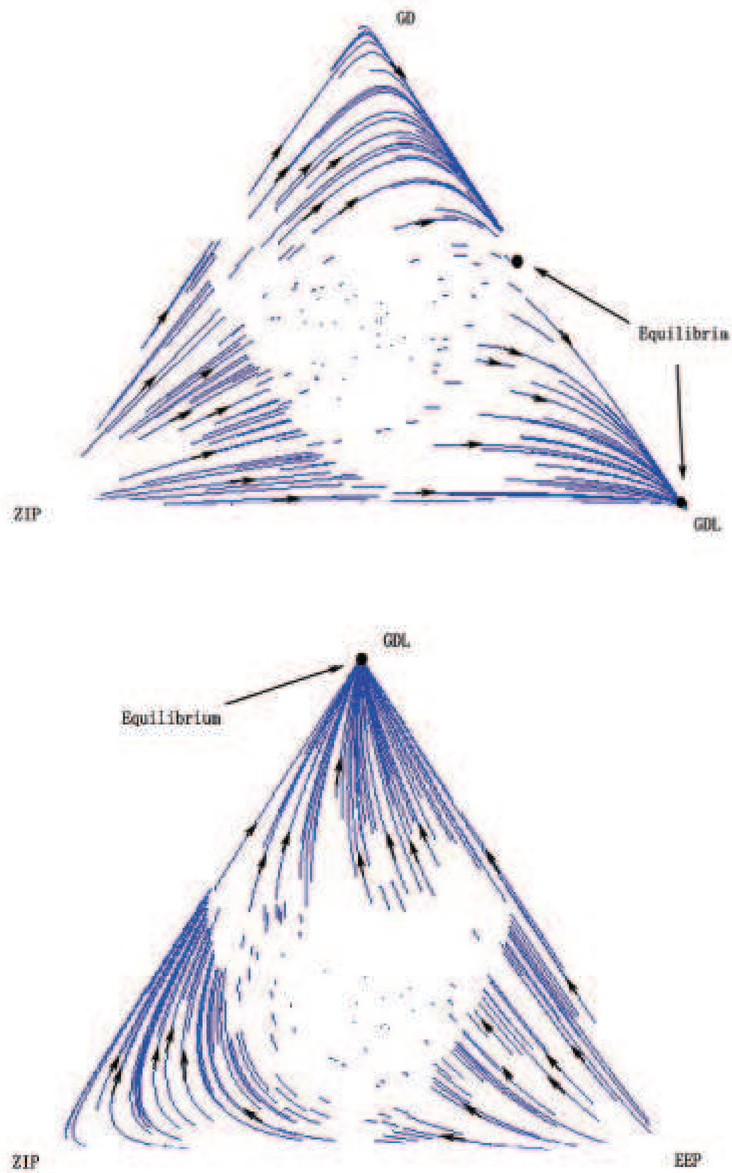

Fig. 1. Replicator dynamics direction field for 20 traders in a CDA where (top) the traders choose between the ZIP, GD and GDL strategies, and (bottom) the traders choose between the ZIP, GDL and EEP strategies

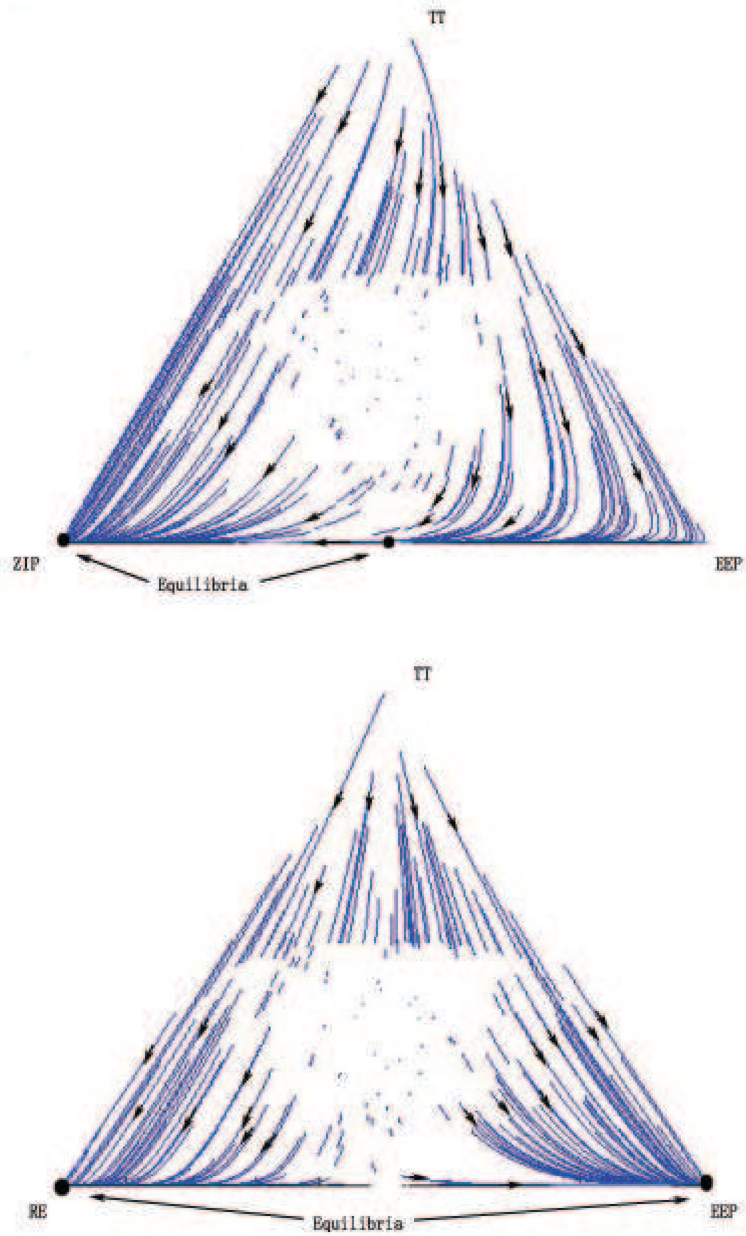

Fig. 2. Replicator dynamics direction field for 20 traders in a CDA where (top) the traders choose between the ZIP, GD and GDL strategies, and (bottom) the traders choose between the ZIP, GDL and EEP strategies