

Evolutionary Stability of Behavioural Types in the Continuous Double Auction

Perukrishnen Vytelingum, Dave Cliff, and Nicholas R. Jennings

School of Electronics and Computer Science,
University of Southampton, Southampton, SO17 1BJ, UK
{pv03r,dc,nrj}@ecs.soton.ac.uk

Abstract. In this paper, we investigate the effectiveness of different types of bidding behaviour for trading agents in the Continuous Double Auction (CDA). Specifically, we consider behavioural types that are *neutral* (expected profit maximising), *passive* (targeting a higher profit than neutral) and *aggressive* (trading off profit for a better chance of transacting). For these types, we employ an evolutionary game-theoretic analysis to determine the population dynamics of agents that use them in different types of environments, including dynamic ones with market shocks. From this analysis, we find that given a symmetric demand and supply, agents are most likely to adopt neutral behaviour in static environments, while there tends to be more passive than neutral agents in dynamic ones. Furthermore, when we have asymmetric demand and supply, agents invariably adopt passive behaviour in both static and dynamic environments, though the gain in so doing is considerably smaller than in the symmetric case.

1 Introduction

The last decade has seen a significant change in the nature of electronic commerce with the emergence of trading agents [6], software that is capable of autonomous and flexible action to achieve its objectives and that is endowed with sophisticated strategies for maximising profit in different types of market mechanisms. Now, one of the most important such mechanisms is the Continuous Double Auction (CDA) [4], a symmetric auction with multiple buyers and sellers. CDAs are so important because they are the principal financial institution for trading securities and financial instruments (e.g. the NYSE and the NASDAQ both run variants of the CDA institution). However, developing agents that can participate in the CDA is difficult because it is not amenable to a game-theoretic analysis and there is no known optimal strategy [4]. Therefore, a number of heuristic strategies have been proposed [12,11,14], each of which has a particular behaviour in the market. Given this, in this paper, we are not concerned with developing yet another strategy, but rather we are interested in how a particular characteristic of such strategies impacts upon their behaviour and their effectiveness. The characteristic in question is the *aggressiveness* of the bidding behaviour, here defined as how eager an agent is to transact. We

focus on aggressiveness in particular because we believe it is a key determinant of success in the market. In particular, we consider three behavioural types¹:

1. The *neutral* agent always submits the quote (a bid or an ask) that maximises its expected profit. This is the most common type of behaviour and is one that is often hardwired into various strategies.
2. The *passive* buyer (seller) submits a lower (higher) quote than its neutral counterpart in order to try and obtain a more profitable transaction. Thus if it does transact, it makes more money because it pays less (if it is a buyer) or receives more (if it is a seller).
3. The *aggressive* buyer (seller) submits a higher (lower) quote than its neutral counterpart in order to try and improve its chances of transacting as it offers more (if it is a buyer) or asks for less (if it is a seller) with the aim of making sure that it can trade.

At this time, comparatively few researchers have considered the issue of varying aggressiveness in a market context. Abreu *et al.* [1] describe the evolutionary stability of behaviours in a reputational model of bargaining, and how players can be more profitable when adopting passive or aggressive behaviour. Thus, while they show that this attitude can have a significant effect on the outcomes experienced, they do not consider the CDA. Walsh *et al.* do consider the CDA and they use an evolutionary game-theoretic (EGT) analysis to examine the interaction of a number of common strategies [15,16]. Thus, their analysis provides an insight on the population proportion that will adopt each strategy. But, while they restrict their analysis to a particular set of well-known strategies, we are more interested in whether a particular strategy can perform better if it is passive, neutral or aggressive. Phelps *et al.* [9] also use an EGT analysis, but they compare two double auction mechanisms, the call and continuous, given that similar strategies are available for both mechanisms. Thus, they are interested in the performance of double auction mechanisms given particular strategies (including one that is evolved using a GA).

Against this background, we believe this study of behavioural types is important because it provides an insight into how this fundamental aspect of bidding behaviour affects the system's performance. Thus, these results apply to any CDA strategy that is capable of adjusting its behaviour along this dimension (e.g. GD [5], ZIP [3] or RB [14]). Such insights are important because if an agent can be more profitable by deviating to another behaviour, then it will do so. However, with every agent in the market doing this, the population distribution of types can change significantly. Now, an effective trading agent can use knowledge of such dynamics to decide on what behavioural type to adopt given the

¹ The nomenclature for the different behaviours varies over the literature, with some like Abreu *et al.* [1] considering passive and aggressive behaviours and others like Byde [2] considering the behaviour of different risk attitudes. However, they all refer to some similar behaviour. For example, a risk-averse agent adopts aggressive behaviour to improve its chance of winning, while a risk-seeking agent adopts passive behaviour to target higher profits. The risk neutral agent adopts a neutral behaviour.

particular population distribution of types. Furthermore, because behavioural aggressiveness is usually an endogenous aspect of a strategy, such an analysis can assist the strategy designer when assessing the effectiveness of various strategies. To perform our analysis, we adopt a similar EGT approach to that of Walsh *et al* and we choose EGT because it allows us to study the dynamics of the CDA when agents are allowed to evolve in terms of the behaviour they adopt. To do this, however, we need to describe the model in a form that abstracts the complex bidding that the CDA mechanism entails, into a simple normal-form game. By so doing, the CDA then becomes amenable to such an analysis. In particular, we develop a set of strategies that vary only in terms of their different behavioural types.

In Section 2, we formally define the CDA mechanism we use in this work and then in Section 3 we describe how we model the behavioural type of a strategy. In Section 4, we detail the model we use for our EGT analysis and in Section 5 we provide our actual analysis. Section 6 concludes.

2 The Continuous Double Auction

The CDA mechanism allows agents to submit quotes at any time during the auction, and a transaction can occur whenever a buyer's bid and a seller's ask can be matched. We conform to previous studies on the CDA in terms of our experimental setup [10,12]. In particular, a set of limit prices is endowed to buyers and sellers and these determine the market demand and supply. Thus, a desired demand and supply can be induced with the appropriate endowment to get the desired environment. Furthermore, the CDA consists of a number of fixed-duration trading periods (or trading days) at the beginning of which an agent is given an endowment to buy or sell. The CDA protocol includes the *NYSE spread-improvement rule* requiring that any bid (ask) submitted must be higher (lower) than the outstanding bid² (ask). Because of the non-deterministic nature of the CDA, payoffs to trading agents are averaged over a sufficient number of simulations to obtain statistically significant results, as computed by the Wilcoxon rank sum test (see Section 5).

3 Modelling the Behavioural Types

Software trading agents are usually designed to employ neutral bidding strategies. However, as discussed in Section 1, we are interested in the effect on the market if trading agents are also designed to be non-neutral and instead have passive or aggressive bidding behaviour. To this end, we model such behaviours by modifying the widely used GD strategy [5]. We choose this strategy out of all those available [12,3,14] because it is one of the most efficient [12] and it can readily be extended to incorporate different bidding behaviours. However, we

² The outstanding bid, o_{bid} , is the current highest bid and the outstanding ask, o_{ask} , is the lowest ask in the market.

could equally well have chosen any CDA strategy that can be suitably adjusted to model passive or aggressive behaviour, and our results generalise to these cases (not shown due to space restrictions).

Now, in the original GD strategy, while all GD agents have the same belief about the market, they form a neutral expected utility-maximising quote based on their own private preferences. In this section, therefore, we detail the GD strategy and how we modify it to be either passive or aggressive (in addition to its standard neutral perspective).

In more detail, the GD agent has a belief that its quote to buy or sell will be accepted in the market. The agent then submits the price that maximises its expected utility. The utility of a quote is the profit associated with such a quote given the agent’s private preferences (limit price, ℓ). The GD agent forms its belief on the basis of observed market information, H_M , that includes the frequencies of accepted quotes (transactions) and rejected quotes. Furthermore, the belief encompasses some form of recency so that only the most recent quotes during the past L transactions are considered. We now formally state the buyer’s belief, $\hat{q}(b)$, associated with a bid b , and seller’s belief, $\hat{p}(a)$, associated with an ask a :




Fig. 1. (a) Seller’s belief, $\hat{p}(a)$, that an ask a will be accepted in the market. (b) Buyer’s belief, $\hat{q}(b)$, that a bid b will be accepted in the market.

Definition 1. Bid Frequencies: D is the set of all permissible quotes in the market. $\forall d \in D$, $B(d)$ is the total number of bid quotes made at price d , $TB(d)$ is the frequency of accepted bids at d , and $RB(d)$ the frequency of rejected bids at d .

Definition 2. Ask Frequencies: D is the set of all permissible quotes in the market. $\forall d \in D$, $A(d)$ is the total number of ask quotes made at price d , $TA(d)$ is the frequency of accepted asks at d , and $RA(d)$ the frequency of rejected asks at d .

$$\hat{q}(b) = \frac{\sum_{d \leq b} TB(d) + \sum_{d \leq a} A(d)}{\sum_{d \leq a} TB(d) + \sum_{d \leq a} A(d) + \sum_{d \geq a} RB(d)} \tag{1}$$

$$\hat{p}(a) = \frac{\sum_{d \geq a} TA(d) + \sum_{d \geq a} B(d)}{\sum_{d \geq a} TA(d) + \sum_{d \geq a} B(d) + \sum_{d \leq a} RA(d)} \quad (2)$$

We use a cubic spline interpolation to calculate the belief at prices that are not registered in H_M (see [5] for further details). The seller's belief function is then modified to satisfy the *NYSE spread-improvement rule* (see Section 2). Thus, for any ask that is higher than the current outstanding ask, the belief function is set to 0. Similarly for the buyer, the belief that any bid below the outstanding bid is accepted is 0. An example of a buyer's and a seller's belief function is shown in Figure 1.

Given its belief function and neutral behaviour, the GD agent forms a quote that maximises its expected utility. In particular, the buyer's utility, $U_b(b)$, for a bid b is given in Equation 3 and the seller's utility, $U_s(a)$, for an ask a in Equation 4. In both cases, the expected utility is the product of the belief function and the utility function, and the bid b^* or ask a^* to submit is computed as the price that maximises the expected utility, as shown in Equation 5.

$$U_b(b) = \begin{cases} 0 & \text{if } b \leq \ell \\ (\ell - b) & \text{otherwise} \end{cases} \quad (3)$$

$$U_s(a) = \begin{cases} 0 & \text{if } a \leq \ell \\ (a - \ell) & \text{otherwise} \end{cases} \quad (4)$$

$$\begin{aligned} b^* &= \arg \max_{b \in (o_{ask}, o_{bid})} [U_b(b)\hat{q}(b)] \\ a^* &= \arg \max_{a \in (o_{ask}, o_{bid})} [U_s(a)\hat{p}(a)]. \end{aligned} \quad (5)$$

To modify the GD strategy so that it embodies the other behavioural types, we alter the agent's utility as a function of its limit price, ℓ , and some scalar parameter, θ , that represents its aggressiveness in the market. The new aggression-sensitive utility functions, $\tilde{U}_b(b)$ and $\tilde{U}_s(a)$ for buyers and sellers respectively, are given in Equations 6 and 7. Here, we vary θ from negative to positive to describe passive to aggressive behaviour, with $\theta = 0$ describing neutral behaviour. By way of an illustration, Figure 2 shows the different utility functions for passive ($\theta = -1$), neutral ($\theta = 0$) and aggressive behaviour ($\theta = 1$) for the buyer and the seller.

$$\tilde{U}_b(b) = \begin{cases} 0 & \text{if } b \leq \ell \\ (\ell - b)e^{\frac{\theta b}{\ell}} & \text{otherwise} \end{cases} \quad (6)$$

$$\tilde{U}_s(a) = \begin{cases} 0 & \text{if } a \leq \ell \\ 1 & \text{if } a > a_{max} \\ (a - \ell)e^{\frac{\theta(a_{max} - a)}{a_{max} - \ell}} & \text{otherwise} \end{cases} \quad (7)$$

Before we proceed to our EGT study of behaviours, we evaluate our new space of GD strategies to ensure that they remain reasonably efficient compared




Fig. 2. (a) Seller’s utility $U_s(a)$ for an ask a and $\ell = 2.0$. (b) Buyer’s utility $U_s(b)$ for a bid b and $\ell = 4.0$. The utility is given as a function of the agent’s limit price, ℓ , and the aggressiveness in bidding behaviour. $\theta = -1$ describes passive behaviour, $\theta = 0$ neutral behaviour, and $\theta = 1$ aggressive behaviour.

to their neutral counterpart. We do so in a homogeneous environment with a typical symmetric demand and supply. Specifically, we examine the market efficiency³ of the CDA with a 10-buyer and 10-seller homogeneous population with similar behaviour. To this end, Figure 3 gives the market efficiency over the aggressiveness parameter θ and different numbers of allocations to buy or sell. The latter gives an insight on how efficiency changes as the number of expected transactions in the market increases with larger allocations given to the traders. As can be seen, as the market size increases with the number of tradeable units, the market efficiency increases for all behaviours. However, *all* agents adopting neutral behaviour does indeed maximise market efficiency in all cases. Passive behaviour is slightly less profitable while aggressive behaviour performs worst. Thus, Gjerstad and Dickhaut’s decision to make GD neutral is validated in such a homogeneous environment (and indeed this is the result that they report in [5]). However, we cannot assume that all agents will remain neutral, as if it profits an agent to deviate to another behaviour, it will do so.

Given this, the question that naturally follows is whether an agent can indeed improve its performance by adopting a different behavioural type. Now, to answer that question, we use an evolutionary game-theoretic analysis, which we describe in the next section.

4 An Evolutionary Game-Theoretic Model

Evolutionary game-theory has traditionally been used to analyse simple games (such as the Prisoner’s Dilemma) in terms of the dynamics of the population of learning agents playing different strategies [17]. Here, however, we are interested in a much more complex game and so a slightly different model is needed. In particular, as explained in Section 1, we use Walsh *et al.*’s model which involves

³ The market efficiency is the ratio of actual total profit of *all* agents to the maximum profit that could be extracted in the market.




Fig. 3. Market efficiency given a homogeneous population of agents with similar behavioural type. We consider 3 different set of experiments with agents allocated 1, 5 and 10 limit prices.

considering the action of an agent as playing its bidding strategy during a game (lasting several trading days) and the payoff as being the total profit at the end of the game. Furthermore, they describe these payoffs as being *heuristic* because they are the output of a mapping of the strategies through a complex, non-deterministic interaction of the trading agents.

Given this, we first describe how we compute the heuristic payoff table that details the expected payoff to each agent (as a function of the S strategies that agents are allowed to play, and the combination of the A agents playing those strategies). We then describe how we use this table to compute the mixed Nash equilibrium of the game and the well-documented *replicator dynamics* model [17] (which is a standard way of representing the population distribution changes).

4.1 Computing the Heuristic Payoff Table

With the heuristic payoff table, we are interested in the expected payoff of a player playing a strategy, j , given the strategies adopted by the other $(A - 1)$ players. Now, because of the non-deterministic and complex nature of the CDA game, some simplifications are required:

1. The payoff of a strategy is the average payoff of an agent playing that strategy in the CDA game, given the different strategies all the A agents are playing in that game.
2. All agents have the same set of strategies to play, and have the same payoff when playing the same strategy. Thus, as described in [16], we can restrict our analysis to symmetric games [17], and significantly reduce the complexity of the problem. Rather than having a table of size S^A , we reduce it to $\binom{A+S-1}{A}$ entries.

Given these, we build our heuristic payoff table⁴ by considering the exhaustive set of strategies the A agents can play, and the number of agents playing each strategy (rather than considering which strategy each of the A agents is playing). Now, because payoff in the CDA game is non-deterministic, we require a significant number of independent simulations for each table entry to ensure that these are statistically significant. Thus, for each entry, we run a number of CDA games (typically 1000) with A agents, each assigned a strategy and a type (buyer or seller) to play, ensuring there is an equal number of buyers and sellers, with a probability of 0.5 that there will be an additional buyer or seller if A is odd.

Given our payoff table, we can now proceed with an EGT analysis as we would with a normal-form game.

4.2 Computing the Equilibrium

Here, we describe how to compute the mixed Nash equilibrium of the CDA game. An agent i chooses the strategy it plays according to its *mixed-strategy*, $\hat{p}_i = (\hat{p}_{i,1}, \dots, \hat{p}_{i,S})$ and $\sum_{j=1}^S \hat{p}_{i,j} = 1$, where $\hat{p}_{i,j}$ represents the probability that agent i plays strategy j . At the equilibrium, \hat{p}_i^* , an agent i cannot receive a higher payoff by unilaterally deviating to another mixed-strategy, assuming that the other agents do not change their strategies [17]. Now, because we assume a symmetric game, all agents have the same mixed-strategy and the same mixed Nash equilibrium. Furthermore, because we are considering a very large population, we can validate that p is equal to the mixed-strategy. Given this, we denote both the population distribution and the mixed-strategy as $p = (p_1, \dots, p_S)$, and the mixed Nash equilibrium as p^* hereafter.

In our EGT analysis, we denote the expected payoff of an agent playing a strategy j , given the mixed-strategy p , as $u(e^j, p)$. To compute $u(e^j, p)$, we consider the results from a large number of CDA games with an agent playing strategy j and $(A - 1)$ agents selected from the population, with a mixed-strategy p . For each game and every strategy, we average⁵ the individual payoffs (obtained from our heuristic-payoff table) of agents using strategy j . The mixed Nash equilibrium is then formulated as the argument to the minimisation problem given in Equations 8 and 9. Specifically, p^* is a mixed Nash equilibrium if and only if it is a global minimum of $v(p)$ [16,8], and we can validate that p is a global minimum if $v(p) = 0$.

$$v(p) = \sum_{j=1}^S \left(\max [u(e^j, p) - u(p, p), 0] \right)^2 \quad (8)$$

⁴ A table entry for a 20-player game with 3 strategies would be $(|S_1|, |S_2|, |S_3|, U_1, U_2, U_3)$ where S_j is the set of agents playing strategy j , $|S_j|$ is the number of agents playing strategy j , and U_j is the average payoff of an agent playing strategy j . Note $\sum_{j=1}^3 |S_j| = 20$ and we have 231 entries.

⁵ In effect, here, we are not running the CDA game for every simulation, but only selecting the appropriate payoff from our heuristic payoff table each time.

where $u(p, p) = \sum_{j=1}^S u(e^j, p)p_j$ is the average payoff of an agent in a population with distribution p .

$$p^* = \arg \min_{\bar{p} \in \Delta} [v(p)] \quad (9)$$

Solving such a non-linear minimisation problem is non-trivial and computationally demanding. Thus, we use a method provided by the Matlab optimization toolbox based on the Nelder-Mead method to find the zero-points of the function v . Because the algorithm used is a non-linear *local* minimiser, we restart the algorithm repeatedly at random points within the unit-simplex until no new equilibria are found for 20 runs.

Now, while the mixed Nash equilibrium gives a theoretical and static perspective of our simplified CDA game, the dynamics of the game and how the equilibria are reached often provide more insight. Given this, we turn to the replicator dynamics which have been shown to be a good model for a common kind of agent learning (namely Reinforcement Learning) [13] and describes how agents learn to reach the equilibrium.

4.3 Computing the Replicator Dynamics

The replicator dynamics, $\dot{p} = (\dot{p}_1, \dots, \dot{p}_S)$, describe how the population distribution p (where $p = (p_1, p_2, \dots, p_S), p \in \Delta$ is an element of a unit-simplex, Δ , and $\sum_{j=1}^S p_j = 1$) changes. This approach assumes that an agent deviates to another strategy that appears to be receiving a higher payoff. Specifically, \dot{p} is a vector given as follows:

$$\dot{p}_j = [u(e^j, p) - u(p, p)] p_j \quad (10)$$

The replicator dynamics show us the strategy trajectories and how they converge to an equilibrium, though they do not necessarily settle at a fixed point (see [17]). In this context, an equilibrium to which trajectories converge, and settle, is known as an *attractor*, while a *saddle point* is an unstable equilibrium at which trajectories do not settle. The region within which all trajectories converge to a particular equilibrium is known as the *basin of attraction* of that equilibrium. The basin is a very useful measure of the adoption of the attractor equilibrium and how likely the population is to converge to that equilibrium.

Here, we compute \dot{p} by starting at different population distributions p inside the unit-simplex and following the trajectory given by Equation 10.

5 An Empirical Analysis of the Behavioural Types

In these experiments, we compute the heuristic-payoff table for a 20-agent CDA game with 3 strategies and lasting 10 trading days. At the beginning of each day, buyers and sellers are each endowed with a single limit price drawn from uniform distributions U_b and U_s respectively. For the purposes of this paper, we consider different uniform distributions to model representative (symmetric and asymmetric) small markets (similar to those considered in previous studies on the CDA [5,3,14]):

- Market 1: $U_b = \mathcal{U}(1.5, 4.5)$ and $U_s = \mathcal{U}(1.5, 4.5)$. This is a symmetric market that has an expected equilibrium at 3.0.
- Market 2: $U_b = \mathcal{U}(2.5, 5.5)$ and $U_s = \mathcal{U}(2.5, 5.5)$. This is a symmetric market that has an expected equilibrium at 4.0.
- Market 3: $U_b = \mathcal{U}(1.5, 4.5)$ and $U_s = \mathcal{U}(2.8, 3.2)$. This is an asymmetric market in which the slope of the supply curve is greatly reduced (compared to market 1). The equilibrium is expected at 3.0.

We model a market shock by changing the demand and supply on trading day 6 (and the new demand and supply then remain the same for the ensuing trading days). In particular, we consider two market shocks, *MS12* where demand and supply changes from Market 1 to 2 and *MS13* where it goes from Market 1 to 3. The former is a market shock that simply results in an increase in equilibrium price, but produces no change in the symmetry of market. The latter is a more complex shock as the structure of the demand and supply changes, though the equilibrium price does not.

We split our analysis into static and dynamic environments to compare our interpretations of their dynamics. For the former, we consider one symmetric market (1) (the results are similar for Market 2) and the asymmetric market (3). For the latter, we consider both types of market shocks. We validate our result by running a Wilcoxon rank sum test [7] on randomly selected entries covering 10% of our heuristic payoff table in order to ensure statistical significance of our data. In some experiments where the change across the heuristic payoff table is particularly small (e.g. in Market 3), 2000 simulations were required for significance.

Now, because we are considering only 3 strategies (each corresponding to a behavioural type), the population distribution space is a unit-simplex in a three-dimensional space. Thus, we can visualise the replicator dynamics \dot{p} and the equilibria p^* by projecting the simplex onto a two-dimensional space [17]. The contour shading in our simplex is proportional to the magnitude of \dot{p} .

5.1 Static Environments

We first analyse behaviours in the CDA when there are no market shocks. In Market 1, we have 2 attractors at A and C and a saddle point at B (see Figure 4). As can be seen, the basin of attraction of C is considerably larger than that of A. This means the majority of the trajectories converge to neutral behaviour, with a non-significant proportion settling at passive. We can see that neutral behaviour is evolutionarily stable in symmetric markets and so most agents are likely to gravitate towards it. Thus, as a designer in such situations, the best thing is to do is make your agent neutral.

Now, when we consider the dynamics in asymmetric Market 3, we find a single attractor at A (see Figure 5). Thus, all trajectories converge towards passive behaviour. However, we note that the magnitude of the dynamics is considerably smaller compared to that in Market 1. This reflects the comparatively small gain that the agent achieves in Market 3 as it slowly adopts the evolutionary




Fig. 4. The replicator dynamics of a CDA game in Market 1. Here, we have 3 equilibria with attractors at $A=(1, 0, 0)$ and $C=(0, 1, 0)$ and a saddle point at $B=(0.95, 0.05, 0)$.

stable passive behaviour. Here, we can interpret the small gain as follows. At a particular population distribution, the agent might intend to deviate in opposing directions from a buyer's and a seller's perspective, particularly because of the asymmetry in the market. Because an agent can be either a buyer or a seller with the same behavioural type, the deviation is a composite of the buyer's and seller's intentions, which explains the small gain as the dynamics converge. Having said this, however, it is still more profitable for an agent to be passive in asymmetric markets, and it should be designed as such.

5.2 Dynamic Environments

The neutral GD strategy has been shown to react well to simple symmetric market shocks [5] in a homogeneous environment. Here, however, we are interested to see whether agents will still adopt neutral behaviour in the presence of market shocks, or whether there are behaviours that are more profitable in such dynamic environments.

With market shock *MS12*, we have 2 attractors at A and C, and a saddle point at B (see Figure 6), with a relatively larger basin of attraction for equilibrium A. Here, there are more agents that are likely to be passive than neutral, though they might all adopt neutral behaviour if there are insufficient passive agents in the market. Thus, it can be more profitable not to be neutral all the time when there are simple market shocks. The strategy designer can use the knowledge of how the dynamics reach the equilibrium to decide what strategy the agent should adapt given the population distribution of behaviours.




Fig. 5. The replicator dynamics of a CDA game in Market 3. Here, we have an attractor equilibria at $A=(1, 0, 0)$ and $C=(0, 0.29, 0.71)$ and a saddle point at $D=(0, 0.51, 0.49)$. Note that the magnitude of the replicator dynamics is very small compared to that of Market 1.




Fig. 6. The replicator dynamics of a CDA game and market shock *MS12*. Here, we have an attractor equilibrium at $A=(0.56, 0.44, 0)$ and $C=(0, 1, 0)$, and a saddle point at $B=(0.16, 0.84, 0)$.

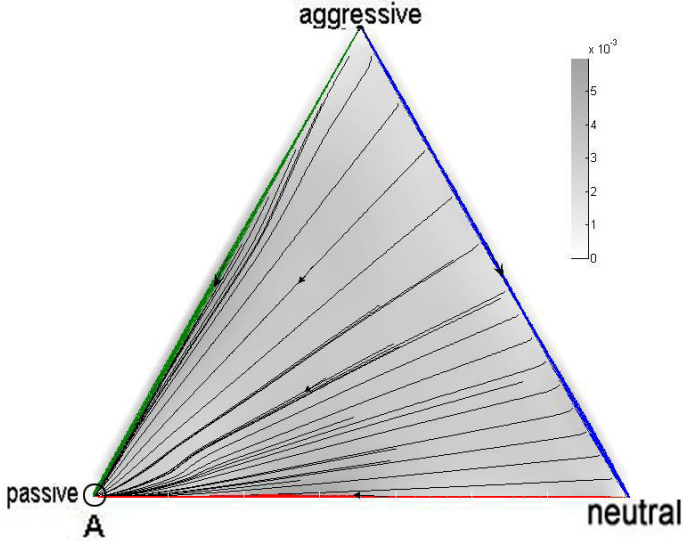


Fig. 7. The replicator dynamics of a CDA game and market shock *MS13*. Here, we have an attractor equilibrium at $A=(1, 0, 0)$.

With market shock *MS13*, we have a single attractor at *A*. Here, all the agents very slowly, but eventually, adopt passive behaviour. However, there is comparatively little to be gained in moving towards the evolutionary stable passive behaviour (for the same reasons we outlined in our discussion of Market 3).

6 Conclusions and Future Work

As electronic marketplaces become ever more common, we believe software agents will increasingly come to dominate the trading landscape. Their ability to quickly make informed decisions, based on the available data, make them ideal candidates for automated trading. To this end, analysing the impact of varying one of the fundamental characteristics of their bidding behaviour in a range of market situations is an important step. In particular, in this paper, we show that in a symmetric market, an agent is more likely to adopt an evolutionary stable neutral behaviour. However, when there are market shocks that increase the equilibrium price but maintain the symmetry of the market (meaning agents have to update their beliefs of the market) neutral is no longer the behaviour agents are most likely to adopt. In this case, more agents change to being passive. We also observe that changing behaviour is not particularly profitable for an agent in an asymmetric market.

For future work, we intend to look at other types of symmetric and asymmetric demand and supply, and other types of market shock in order to obtain further insights into how a trader's behaviour changes in yet other types of market.

For completeness, we also aim to address the limitation of our model where an agent has the same behavioural type when it is both a buyer and a seller. In particular, we believe that separately analysing these two roles can be more insightful, particular in asymmetric markets, where we do not expect the same behaviour from them.

Acknowledgments

We kindly thank anonymous reviewers for their helpful comments on the paper. The first author is funded by the DIF-DTC 8.6 project (www.difdte.com).

References

1. D. Abreu and R. Sethi. Evolutionary stability in reputational model of bargaining. *Games and Economic Behavior*, 44(2):195–216, 2003.
2. A. Byde. Applying evolutionary game theory to auction mechanism design. *ACM Conference on Electronic Commerce*, pages 192–198, 2003.
3. D. Cliff. ZIP60: Further explorations in the evolutionary design of online auction market mechanisms. Technical Report HPL-2005-85, 2005.
4. D. Friedman and J. Rust. *The Double Auction Market: Institutions, Theories and Evidence*. Addison-Wesley, New York, 1992.
5. S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.
6. M. He, N. R. Jennings, and H. Leung. On agent-mediated electronic commerce. *IEEE Trans on Knowledge and Data Engineering*, 15(4):985–1003, 2003.
7. M. Hollander and D. A. Wolfe. *Nonparametric Statistical Methods*. Wiley, 1973.
8. R. D. McKelvey and A. McLennan. Computation of equilibria in finite games. *Handbook of Computational Economics*, 1, 1996.
9. S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double auction market designs. *Proceedings of the 6th Workshop on Agent Mediated Electronic Commerce*, pages 192–198, 2004.
10. V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70:111–137, 1962.
11. G. Tesauro and J. L. Bredin. Strategic sequential bidding in auctions using dynamic programming. *Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 591–598, 2002.
12. G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. *Proceedings of the Third ACM Conference on Electronic Commerce*, pages 206–209, 2001.
13. K. Tuyls, T. Lenaerts, K. Verbeeck, S. Maes, and B. Manderick. Towards a relation between learning agents and evolutionary dynamics. *Proceedings of BNAIC 2002*, pages 21–22, 2002.
14. P. Vytelingum, R. K. Dash, E. David, and N. R. Jennings. A risk-based bidding strategy for continuous double auctions. *Proc. 16th European Conference on Artificial Intelligence*, pages 79–83, 2004.

15. W. Walsh, D. Parkes, and R. Das. Choosing samples to compute heuristic-strategy nash equilibrium. *Proceedings of the Fifth Workshop on Agent-Mediated Electronic Commerce*, 2003.
16. W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent games. *Proceedings of AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents (GTDT-02)*, 2002.
17. J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.