

# Analysing Buyers' and Sellers' Strategic Interactions in Marketplaces: An Evolutionary Game Theoretic Approach

Perukrishnen Vytelingum, Dave Cliff and Nicholas R. Jennings

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

**Abstract.** We develop a new model to analyse the strategic behaviour of *buyers and sellers* in market mechanisms. In particular, we wish to understand how the different strategies they adopt affect their economic efficiency in the market and to understand the impact of these choices on the overall efficiency of the marketplace. To this end, we adopt a two-population evolutionary game theoretic approach, where we consider how the behaviours of both buyers and sellers evolve in marketplaces. In so doing, we address the shortcomings of the previous state-of-the-art analytical model that assumes that buyers and sellers have to adopt the same mixed strategy in the market. Finally, we apply our model in one of the most common market mechanisms, the Continuous Double Auction, and demonstrate how it allows us to provide new insights into the strategic interactions of such trading agents.

## 1 Introduction

Autonomous agents are increasingly being used in a variety of marketplace settings — be it in electronic market institutions such as eBay or the New York Stock Exchange (NYSE) or in market-based control applications [1]. Now, in most of these cases there is no known, analytically-determined dominant strategy that such trading agents can adopt. This means a multitude of heuristic strategies have been and continue to be developed, each trying to better exploit the observable market information. Unfortunately, this makes it difficult to determine *a priori* which strategies will be effective in which situations. This is a serious concern for agent designers because they have no principled way of selecting which strategy to adopt (important because the various strategies can perform very differently in different market settings). Moreover, it is also a concern for the market designer because she wants to deploy a mechanism that is efficient and stable, but this depends on the strategies that the various participants adopt. Thus, she would like a way of determining which strategies are more likely to be adopted in a given mechanism, so that she can determine which mechanism is likely to be best.

To address these needs, Walsh *et al.* [10] proposed a methodology for analysing strategic interactions in complex games, including markets. In particular, their analytical model is an evolutionary game theoretic (EGT) approach based on computing the mixed-Nash equilibrium of heuristic strategies and the dynamics analysis of equilibrium convergence [11]. Now, because an EGT analysis is infeasible for all but the

simplest games, Walsh *et al.* describe how complex games that involve repeated interactions with more elaborate actions and payoffs can be made amenable to such an analysis. Specifically, their model considers the high-level, heuristic strategies of the trading agents as simple actions, and the payoff to these strategies as the average profit extracted in the market (by so doing, they essentially abstract a complex iterated game to a simple normal-form one). To illustrate their approach, they apply it to two different scenarios, the Automated Dynamic Pricing (ADP) game and the Continuous Double Auction (CDA) game. In the former game, there is a set of sellers, each of which offers a good to a much larger set of buyers. Both types of trading agents try to maximise their profits, but sellers strategise over the price they ask for, while the buyers' strategies are fixed: half select a seller at random, and the rest use a shopbot to find the current lowest-price seller. Walsh *et al.* then analysed how the sellers, endowed with a set of heuristic strategies, interact in the ADP game, and what strategies these sellers are most likely to adopt. In the latter game, multiple buyers and sellers compete to transact at the same time (see Subsection 3.1 for more details). In this case, Walsh *et al.* analyse the strategic interaction of both buyers and sellers to observe how the market behaviour is evolving. Their methodology has now been widely adopted and, in particular, Phelps *et al.* [4] used it to compare two different auction mechanisms (the continuous and the call double auction mechanisms), given that similar strategies were available for both. Furthermore, Vytelingum *et al.* [8] used it to look at the evolutionary stability of various degrees of aggressiveness in bidding behaviour in CDAs.

However, a key assumption of all this previous work is that an agent will adopt the same heuristic strategy even when it has to perform different roles (such as being a buyer and a seller in the CDA). In games like the ADP, where agents have a single role (as a competing seller), such an assumption does not constrain the analysis and their methodology is appropriate. However, in double-sided games, like the CDA, with multiple buyers and sellers, such an assumption is both unrealistic and unnecessarily restrictive. In practice, buyers and sellers usually have different bidding behaviours whose efficiency depends on a number of factors (including what strategies other buyers and sellers adopt). The demand and supply of the market then determines the complex interactions of these strategies, which, in turn, determine their overall effectiveness. Given the aim of maximising its profit, an agent should be allowed to select whatever is the best strategy for it when acting as a buyer and whatever is the best for it as a seller. The present constraint of compromising on both and having to select the same strategy in each role can only have a negative effect on the agents' economic efficiency. For example, Vytelingum *et al.*'s analysis highlighted the loss in effectiveness of doing this when considering asymmetric demand and supply (with the magnitude of the dynamics decreasing by up to around 60% when buyers and sellers have different preferences compared to the symmetric and special case where the assumption is less problematic).

To address this shortcoming, we propose a two-population EGT model<sup>1</sup> to anal-

---

<sup>1</sup> It is also possible to separately analyse the strategic interactions of buyers with  $S_b$  strategies and sellers with  $S_s$  strategies by considering the one-population EGT model. An agent's strategy would then consist of a buying and a selling strategy from a strategy set of size  $(S_b \times S_s)$ . However, the analysis would then be more complex than with the two-population case. In particular, when we have 2 buyer and 2 seller strategies, we would

yse the complex interactions of buyers *and* sellers in the market. Specifically, we consider a game with two distinct populations. The two populations correspond to the two different types of trading agents, the buyers and the sellers, and they are endowed with distinct sets of heuristic strategies. Thus, our model makes no assumption that an agent must have the same strategy for buying and selling, although if this is the best thing to do, then our model will converge to it.

In developing this new model, we advance the state of the art as follows. First, we provide a novel analytical model that dissects the buyer and seller trading roles in the market and allows us to analyse the market efficiency and stability from both perspectives. This move from a one-population to a two-population EGT model is a non-trivial extension. As well as requiring a larger set of information about the buyers' and sellers' payoffs, we also need to reformulate the dynamics of the now co-evolving buyers' and sellers' strategies. Furthermore, we have to reformulate the problem to find the mixed-Nash equilibrium (which is now the mixed strategy of the buyers and of the sellers where it does not benefit both types to deviate to another mixed strategy). To this end, we had to devise a means of calculating the dynamics and the mixed-Nash equilibria of the market game with buyers and sellers. Second, to illustrate the power of our model, we analyse a typical market game, namely the CDA, and identify strategic interactions between buyers and sellers that have not been previously observed.

The remainder of the paper is structured as follows. In Section 2, we describe our model, including the equilibrium computation and the dynamics analysis. Section 3 describes its application to the CDA and compares our analysis with that of Walsh *et al.*'s. Section 4 concludes.

## 2 The Two-Population EGT Model

In a two-population game, a player  $i$  from population  $P$  selects its strategy from a set of  $S_P$  *strategies*, where a strategy is a policy that determines the agent's *action(s)* in the game. The payoff to each player is then a deterministic mapping of the strategies of the players from the two populations and is usually given by a corresponding payoff table. For generality, we assume that each player has a mixed strategy,  $x^i = (x_1^i, \dots, x_j^i, \dots, x_{S_P}^i)$  (where  $x_j$  is the probability that it plays pure strategy  $j$ ) that it plays in the game. Then, as rational behaviour dictates that each player will change its mixed strategy for a higher payoff, we have an evolution of behaviours as all the players in the game concurrently change their mixed strategy. EGT, then, allows us to analyse such an evolution. To date, however, it has typically been used to analyse simple games such as the two-population, two-player and two-strategy Prisoner's Dilemma (PD) [11]. Here, however, we are interested in more complex two-population market games with  $A_b$  buyers and  $A_s$  sellers that are endowed with a set of  $S_b$  and a set of  $S_s$  heuristic strategies respectively.

---

have a strategy space of 4 for the one-population case with the analysis in a 4-dimensional space. With the two-population model, the analysis would be in a two-dimensional space and so possible to visualise.

We simplify the complex market game to a simple normal-form game (as per Walsh *et al.*<sup>2</sup>). Moreover, as discussed previously, playing the complex buyer and seller strategies in the game can be considered as high-level actions similar to the *actions* in a normal-form game. The payoff to a buyer or a seller for such actions is then the total profit that they have extracted at the end of the game (which can last several trading days). In such cases, the payoffs are usually referred to as being *heuristic*, because they are the result of the complex, non-deterministic interaction of trading agents in the game. Thus, our analysis begins with the computation of the heuristic payoff table and, thereafter, we analyse our market games using a two-population EGT analytical approach. This analysis is the novel step of our model, and it covers the equilibrium computation and the dynamics analysis of the market game. In the former, we describe the ideal static properties of the population proportions using the different strategies in the system (i.e. the mixed-Nash equilibria of the game). In the latter, we detail how to calculate the dynamics  $\dot{p} = (\dot{p}_1, \dots, \dot{p}_{S_b})$  and  $\dot{q} = (\dot{q}_1, \dots, \dot{q}_{S_s})$  of the game, which describe how the buyer and seller population distributions,  $p = (p_1, \dots, p_{S_b})$  and  $q = (q_1, \dots, q_{S_s})$ , change<sup>3</sup>. Here, because we are considering very large populations, we can validate that the population distributions,  $p$  and  $q$ , are equal to the mixed strategies of the buyers and sellers. Hereafter, we will refer to these terms as the mixed strategies of the two populations.

## 2.1 Computing the Heuristic Payoff Table

The heuristic payoff table specifies the expected return to each agent as a function of the heuristic strategies played by all agents. Now, for a two-population, normal-form game with  $S_b$  buyer strategies,  $S_s$  seller strategies,  $A_b$  buyers and  $A_s$  sellers, we require  $(S_b^{A_b} \times S_s^{A_s})$  entries in the table. However, because the table size increases exponentially with the number of strategies, some simplifications are necessary to make the analysis tractable. In particular, Walsh *et al.* restrict their analysis to symmetric mixed-Nash equilibrium when they assume that each agent from the single population has the same mixed strategy and, hence, expects the same payoff when playing the same pure strategy. In our case, when considering two populations, the size of the payoff table then reduces considerably to  $\binom{A_b+S_b-1}{A_b} \times \binom{A_s+S_s-1}{A_s}$ . For example, for a market game of 10 buyers and 10 sellers, each endowed with 2

<sup>2</sup> Note that there are other techniques that have been used in the literature to simplify complex games such that they are more computationally tractable. Specifically, Wellman *et al.* [12] have worked on empirical game-theory where they demonstrated how a complex game can be reduced by iteratively selecting the best strategies in an evolutionary process and reducing the number of strategies. However, in this paper, we are not interested in how complex games can be simplified but rather on how we can analyse the complex interactions of buyers and sellers. Indeed, we do not reduce the CDA games but only assume symmetric payoff across buyers and across sellers. Such a technique of reducing games could be used in our analysis when we need to consider a large set of strategies. In such cases, the number of strategy profiles would explode as we increase the strategy set, and we would then need to reduce the CDA game such that the analysis is computationally tractable.

<sup>3</sup> subject to the constraints that  $\sum_{h=1}^{S_b} p_h = 1$  and  $\sum_{k=1}^{S_s} q_k = 1$ .

different strategies, the size of the payoff table reduces from  $1.05 \times 10^6$  to 121 (from asymmetric to symmetric payoffs in each population).

For the exhaustive set of strategy profiles<sup>4</sup>, we calculate the different payoffs (which is the most computationally intensive part of our analytical approach) to the different strategies adopted by the players from the different populations in the market. Because of the non-deterministic nature of the table, we require a statistically significant number of independent runs for each profile. Thus, we considered 2500 runs and validated our results at the 95%-confidence-interval by running the non-parametric Wilcoxon rank sum test [2] on the difference,  $[u_b(e^h, p, q) - u_b(p, p, q)]$ , between the actual and the expected payoff. We chose such a test because we cannot ensure the normality of our data set and because we want to ensure statistical significance of our dynamic analysis, particularly around mixed-Nash equilibria where that difference is significantly smaller.

Given the heuristic payoff table (see [7] for an example), we can now proceed to the equilibrium analysis.

## 2.2 Equilibrium Analysis

First, we describe how to calculate the mixed-Nash equilibrium given the heuristic payoff table. As discussed in [3], it is possible to formulate our solution as the global minimum of a real-valued function,  $v(p, q)$  (given in Equation 1), on a polytope, whose fixed points approximate the mixed-Nash equilibria ( $p_{nash}, q_{nash}$ ):

$$v(p, q) = \sum_{h=1}^{S_b} (\max [u_b(e^h, p, q) - u_b(p, p, q), 0])^2 + \sum_{k=1}^{S_s} (\max [u_s(e^k, q, p) - u_s(q, q, p), 0])^2 \quad (1)$$

In Equation 1,  $u_b(e^h, p, q)$  represents the expected payoff to a buyer adopting pure strategy  $h$  when the other buyers adopt a mixed strategy  $p$  and the sellers a mixed strategy  $q$ .  $u_b(p, p, q) = \sum_{i=1}^{S_b} u_b(e^i, p, q)p_i$  is the average payoff to a buyer in the market. Similarly,  $u_s(e^k, q, p)$  is the expected payoff to a seller adopting pure strategy  $k$  when all buyers adopt mixed strategy  $p$  and the rest of the sellers adopt mixed strategy  $q$ .  $u_s(q, q, p) = \sum_{j=1}^{S_s} u_s(e^j, q, p)q_j$  is the average payoff to a seller in the market. Now, when calculating the expected payoff of a buyer using a pure strategy  $h$ , we consider a significant number of games where one buyer adopts pure strategy  $h$ ,  $(N_b - 1)$  buyers adopt mixed strategy  $p$  and  $N_s$  sellers adopt mixed strategy  $q$ . The individual profit of all agents using the pure buyer strategy  $h$ , is then averaged over all the games as the required payoff. A similar procedure applies when calculating a seller's expected payoff.

The mixed-Nash equilibrium is calculated as the values that minimise function  $v(p, q)$ . Such a non-linear minimisation problem is non-trivial and, indeed, can be computationally demanding. Thus, we use the non-linear minimisation algorithm provided by the Matlab optimization toolbox to solve the problem of finding the

<sup>4</sup> A strategy profile  $[\rho^b, \rho^s]$  defines the number of buyers  $\rho^b = (\rho_1^b, \dots, \rho_{S_b}^b)$  and sellers  $\rho^s = (\rho_1^s, \dots, \rho_{S_s}^s)$  using the different buyer and seller strategies. An example of the 121 strategy profiles for the above 10 buyers and 10 sellers game would be  $[(1, 9), (2, 8)]$ , where we have 1 buyer using buyer strategy 1, 9 buyers using buyer strategy 2, 2 sellers using seller strategy 1 and 8 using seller strategy 2.

zero-points that minimise  $v(p, q)$ . Furthermore, because the toolbox can only find local minima, we repeatedly restart the algorithm at random points ( $p$  and  $q$ ) until no new minimum is found for 20 consecutive runs.

Having looked at the static perspective of our analysis of the market game, we now consider the dynamic perspective.

### 2.3 Dynamics Analysis

To analyse the dynamics of the market game, we use the well-researched *replicator dynamics* [11], which assumes that a trading agent will deviate to another strategy if it can be more profitable by so doing. We chose the replicator dynamics as it has been shown to be a good approximation to common agent learning models, such as reinforcement learning [6], which we would typically find in such markets. The following equations describe precisely how we calculate the dynamics,  $\dot{p}_h$  for pure buyer strategy  $h$  and  $\dot{q}_k$  for pure seller strategy  $k$ :

$$\dot{p}_h = [u_b(e^h, p, q) - u_b(p, p, q)] p_h \quad (2)$$

$$\dot{q}_k = [u_s(e^k, q, p) - u_s(q, q, p)] q_k \quad (3)$$

To observe the dynamics of the game, we calculate *trajectories* (i.e. how the mixed strategies change). In more detail, we start with any pair of mixed strategies  $(p, q)$ , and calculate the dynamics  $\dot{p}$  and  $\dot{q}$  given by Equations 2 and 3 respectively, as we progress along a certain trajectory. It is interesting to note that, at the mixed-Nash equilibrium, the dynamics of the two populations are zero, meaning that it does not pay for a buyer or a seller to change to another mixed strategy. In this case, a trajectory either converges or diverges from a mixed-Nash equilibrium. When we have convergence, we say that the equilibrium is an evolutionarily stable strategy and it is referred to as an *attractor*. Where we have divergence, the equilibrium is unstable to small fluctuations and it is known as a *saddle point*. The region within which all trajectories converge to a particular attractor is the *basin of attraction* of that equilibrium. More formally,

**Definition 1.** A *trajectory* is the change in mixed strategy, starting from a particular mixed strategy, and following the replicator dynamics.

**Definition 2.** An *attractor* is a mixed-Nash equilibrium towards which replicator dynamics (trajectories) converge.

**Definition 3.** A *saddle point* is a mixed-Nash equilibrium from which replicator dynamics (trajectories) diverge.

**Definition 4.** A *basin of attraction* of a mixed-Nash equilibrium is the space of mixed strategies from which trajectories will converge to that equilibrium. The area of the basin corresponds to the probability that this equilibrium will be reached if we assume that it is equally likely for an agent to start at any mixed strategy.

**Definition 5.** An *evolutionarily stable strategy (or ESS)* is a mixed-Nash equilibrium which, if adopted by buyers and sellers, cannot be invaded by any competing alternative strategy. An ESS is asymptotically stable in the replicator dynamics, such that trajectories do not necessarily have to settle at the equilibrium (which would be neutrally stable in that case) to be an ESS (see [11] for more details).

When we analyse a game, we identify all the attractors and saddle points, as well as the basins of attraction. The volume of the space covered by a basin is proportional to the probability that a mixed-Nash equilibrium will be adopted. Thus, basins of attraction provide insights into which equilibria are more likely to be adopted in the market.

### 3 Application to a Market Mechanism

Having detailed our analytical model, we now apply it to the CDA game, chosen for being one of the most popular and widely adopted market mechanisms in financial institutions and for its applications to market-based control, where decentralised allocation of scarce resource is required. Also, related work on the CDA has used Walsh *et al.*'s model [4, 8], and through the analysis given by our model, we intend to demonstrate the shortcomings of such previous applications of their model.

Our analysis is feasible for any number of buyers and sellers and for any number of strategies. However, a visual representation of the trajectories of the replicator dynamics and the equilibria is only possible when considering at most a set of two strategies in either population (because we can effectively plot the replicator dynamics in a two-dimensional space) and so this is the case to which we limit ourselves here. In this section, we first detail the CDA game and the agent strategies before undertaking the actual analysis of the strategic interactions of buyers and sellers.

#### 3.1 The CDA Game

To allow valid comparisons, we consider the same model of the CDA mechanism as in the papers mentioned above. Specifically, in such a model, buyers and sellers are allowed to submit bids and asks respectively. The market clears continuously whenever a transaction is possible, and at all times during the game we have an outstanding bid (highest uncleared bid) and an outstanding ask (lowest uncleared ask). The market mechanism is principally a set of market protocols that define the type of bids and asks allowed and the clearing rules that are in operation. For the model we consider here, single-unit bids or asks (offers to buy or sell a single unit of a good) are allowed, with the market clearing at the average of the accepted bid and ask. The CDA market protocol also includes the NYSE spread-improvement rule that states that a bid or an ask will be accepted in the market only if it improves upon the outstanding bid or ask previously submitted.

To run such a market game, we developed a discrete-time simulator implementing the buyers and sellers with their different strategies, and the market with its protocols and clearing rules. At the beginning of a trading day, the buyers and sellers are endowed with a number of limit prices corresponding to the units to buy or sell. Each trading day lasts 1000 time steps and, at each time step, an agent is randomly triggered to submit a bid or an ask. Furthermore, for the purpose of these simulations, we consider an arbitrary dynamic scenario with 10 buyers and 10 sellers, lasting 20 trading days with a market shock (a sudden change in the demand and the supply) on day 11. Specifically, by specifying the trading agents' limit prices, we can induce different market demand and supply. Here, we have a shock with an asymmetric demand and supply (with a flat supply and with buyers' and sellers' limit prices drawn

from two uniform distributions,  $U_b = \mathcal{U}(1.5, 4.5)$  and  $U_s = \mathcal{U}(2.8, 3.2)$ ) changing to a symmetric demand and supply (with  $U_b = \mathcal{U}(1.5, 4.5)$  and  $U_s = \mathcal{U}(1.5, 4.5)$ ). While the choice of the scenario is not fundamental to the analysis, we wish to demonstrate that an analysis is feasible even in such complex and dynamic environments (we observe similar trends for a wide variety of scenarios in [7]).

We now describe the strategies with which the buyers and sellers are endowed.

### 3.2 The CDA Strategies

We will consider the state of the art GDX [5] and AA strategies [7] (the latter is an extended version of the RB strategy [9]). In this context, the choice of strategies is not central to this analysis. Rather, our principal aim is to demonstrate how our EGT model can be used to evaluate the efficiency of strategies in marketplaces. Nevertheless, we do need to describe these strategies in more detail.

The GDX strategy is an extended version of the GD strategy, which adopts an expected-profit maximising approach. Specifically, it calculates a belief that a bid or an ask will be accepted in the market, based on the past market information. This belief is then extrapolated over the space of feasible bids or asks, and the bid or the ask that maximises the expected profit is submitted in the market. The GDX extension adds a new dimension to the strategy, namely time. Specifically, it calculates the number of bidding opportunities that the agent still has left and further strategises with the knowledge that the agent can be more profitable if it waits longer. The expected profit is now calculated to factor in this number of bidding opportunities and a discount factor associated with the bid or ask. Based on the simulations in [5], we use a discount factor of 0.9, which was empirically demonstrated to be the most efficient.

The AA strategy is an improved version of the RB strategy. In particular, it has a short and a long-term learning mechanism that adapt its bidding behaviour to the changing market conditions. The short-term mechanism updates the agent's bidding aggressiveness to immediately respond to short-term market fluctuations. The agent's long-term mechanism, on the other hand, learns its aggressiveness model (which determines how the degree of aggressiveness translates to a bid or an ask to submit in the market) to adapt to long-term changes in market conditions and to enable the agent to perform efficiently in dynamic environments, such as the one under consideration here.

Given the market mechanism and strategies, we now analyse the buyers' and the sellers' strategic interactions using our two-population EGT model, before we compare this with the corresponding one-population analysis.

### 3.3 The Analysis


We analyse how buyer and seller behaviours change assuming they can choose between the two strategies. First, we generate our heuristic payoff table (which requires several hours of simulations for each of the exhaustive set of the 121 strategy profiles given the 2500 runs of 20 trading days each for each profile), and we then go on to perform the actual EGT analysis for our dynamic scenario with the GDX and AA strategies.



As we are considering a set of only 2 strategies in either population, we can plot the replicator dynamics,  $(\dot{p}_1, \dot{q}_1)$ , as vectors at different  $(p_1, q_1)$  in a two-dimensional figure. In such cases, the horizontal axis represents the buyer population proportion,  $p_1$ , and the vertical axis represents the seller population proportion  $q_1$ . Then, the different vertices correspond to different pure buyer and seller strategies. An example of such an EGT analysis is given in figures 1 and 2 (for 10 buyers and 10 sellers). The former plot gives the replicator dynamics of the analysis, with its shading denoting the magnitude of the dynamics,  $(|\dot{p}_1| + |\dot{q}_1|)$ , given the mixed strategies of the buyers and sellers. As the magnitude of the dynamics decreases (and the shading is darker), there is less and less incentive to deviate to another strategy, until the magnitude is 0 at a mixed-Nash equilibrium, at which point it does not pay to deviate to another buyer or seller strategy. The latter plot gives the magnitude of the buyers' and sellers' dynamics, with a mixed-Nash equilibrium occurring when the magnitude of both dynamics is 0. We consider these magnitudes to compare the buyers' and sellers' payoff difference when deviating to the more efficient strategy. Furthermore, we also consider the market efficiency at the different mixed strategies and, specifically, at the mixed-Nash equilibrium (see Figure 3 for an example). This replicator dynamics is rational from a selfish agent's perspective as it deviates to the most profitable strategy when it attempts to maximise its individual profit. However, from the market designer's perspective, maximising social welfare (i.e. market efficiency) is the objective. But, because the market designer cannot specify the behaviour of the market (i.e the behaviour of every buyer and seller), it can only influence the market efficiency by designing the structure of the market mechanism, which, in turn, influences the market behaviour. Now, the ESS is not reached instantaneously but rather as the result of an online learning process and the market designer is interested in how the market efficiency changes during that process to finally settle at the ESS. Thus, the market designer can use such an analysis of efficiency to identify how market efficiency evolves as well as the *evolutionarily stable efficiency* (i.e. the efficiency at the ESS) which would be insightful in modifying the structure of the mechanism to maximise its evolutionary stable efficiency. Formally:

**Definition 6.** *The evolutionarily stable efficiency is the market efficiency (analogous to the average population fitness in the evolutionary computing literature) at the evolutionarily stable strategy.*


In the CDA game we consider, we have three attractors A (at pure strategy AA for the buyer and seller), B (at pure strategy GDX for the buyer and seller) and C (at pure strategy GDX for the buyer and pure strategy AA for the seller), and two saddle points D and E. Specifically, the area of the basin of attraction of A is 0.78, that of B is 0.12 and that of C is 0.10. Based on this, we can infer that there is a 78% chance that AA will eventually be adopted by all buyers and sellers in the market, a 12% chance that GDX will eventually be adopted by all buyers and sellers and, finally, a 10% chance that all buyers will adopt GDX and all sellers will adopt AA. Furthermore, if we aggregate these results, there is a 12% chance that GDX will eventually be adopted by all buyers and sellers, and a 10% chance that, eventually, all buyers will adopt GDX and all sellers will adopt AA. We can further infer that there is a 88% chance the sellers will adopt AA, a 12% chance they will adopt GDX, and there is a 78% chance the buyers will adopt AA and a 22% chance they will adopt GDX.



**Fig. 1.** The replicator dynamics for the market with 10 buyers and 10 sellers. Here, we have three attractors: A at (1,1), B at (0,0) and C at (0,1) and two saddle points: D at (0.19,1) and E at (0,0.40). The dotted line denotes the boundary between the basins of attraction.

When we consider the trajectories when GDX buyers are in the minority (right-hand part of the dynamics plot), we observe that the sellers nearly always deviate to AA. As GDX becomes more popular among buyers (with trajectories flowing towards the left, i.e. GDX), sellers now deviate to either GDX or AA (shown by the trajectories flowing either to the top at AA or to the bottom at GDX), depending in which basin of attraction they are. When we consider the magnitude of the buyers' and the sellers' dynamics (Figure 2), we observe that the latter is larger than the former, specifically when AA buyers are in the majority (when  $p_1$  is close to 1). This implies that there is, then, a fast convergence of the seller's strategy to AA, suggesting that AA sellers are most profitable when competing against AA buyers and the strategies quickly evolve to AA. Furthermore, we observe that there is significantly less economic incentive to deviate to another strategy when GDX is in the majority (shown by the lower magnitude of the buyers' and sellers' dynamics). Thus, it takes longer for the strategy to evolve to either the mixed-Nash equilibrium B or C.

Furthermore, as discussed in Section 1, such an analysis allows the market designer to observe how the behaviours of the buyers and sellers evolve. Now, because the market efficiency changes as the behaviours change, the market designer is interested in how efficient the market will eventually be. To this end, we also calculate the efficiency of the market given the different buyers' and sellers' behaviours (see Figure 3). In particular, we observe that stable outcomes (attractors A, B and C) of the game are not the most efficient (which is at mixed strategy X — see Figure 3). Such an analysis therefore enables the market designer to observe how the market



**Fig. 2.** Magnitude of the buyers' and sellers' dynamics for the market with 10 buyers and 10 sellers.

efficiency changes with the co-evolving behaviours of the buyers and sellers, as well as the evolutionarily stable efficiency of the market at the ESS. The designer can use such insights to decide how effective is the market protocol (which defines how buyers and sellers interact and which is the only aspect of the marketplace that he can change), and how to modify it to maximise the evolutionarily stable efficiency of the marketplace.


### 3.4 Comparison with Walsh *et al.*'s Model

We have previously argued that our two-population model offers better insights into the behaviour of the CDA (and other mechanisms where agents have more than one role). Now, in order to see this directly, we compare the two in a given scenario. Specifically, the EGT plot using the two-population model is the scenario we considered above (for 10 buyers and 10 sellers) and the corresponding one-population model is given in Figure 4. Note that such a comparison is possible as we use the same set of strategies for buyers and sellers in our two-population model. Here, as we are considering the same two strategies that an agent can choose from, the dynamics are given by a one-dimensional plot, where GDX is the pure strategy at 0 and AA is the pure strategy at 1. We begin by reviewing the one-population model. Here, the replicator dynamics,  $\dot{x} = (\dot{x}_1, \dots, \dot{x}_S)$ , describe how the population distribution<sup>5</sup>  $x$  changes (where  $x = (x_1, x_2, \dots, x_S)$ ,  $x \in \Delta$  is an element of a unit-simplex  $\Delta$ ,  $\sum_{j=1}^S x_j = 1$  and  $S$  is the number of strategies an agent can select from). In this case,  $\dot{x}$  and the mixed-Nash equilibrium  $x_{nash}$  are calculated as follows:

$$\dot{x}_j = [u(e^j, x) - u(x, x)] x_j \quad (4)$$


$$x_{nash} = \arg \min_{x \in \Delta} \sum_{j=1}^S (\max [u(e^j, x) - u(x, x), 0])^2 \quad (5)$$

<sup>5</sup> Note that because buyers and sellers are assumed to adopt the same behaviour,  $p = x$  and  $q = x$ .



**Fig. 3.** Efficiency of the market with 10 buyers and 10 sellers. A, B and C are the attractors. Mixed strategy X at (0.92,1) is the behaviour that gives the highest market efficiency.

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



**Fig. 4.** The one-population EGT analysis for the market with 10 buyers and 10 sellers. We have two attractors: A' at 1 and B' at 0 and a saddle point: C' at 0.30.

From Figure 4, we observe two attractors A' at 1 and B' at 0, and a saddle point C' at 0.30. By considering the space of trajectories that converge to either attractors, we infer that there is a 30% chance that all agents (buyers and sellers) will eventually adopt the pure strategy GDX, while there is a 70% chance that all agents will eventually adopt the pure strategy AA. We also observe that AA is more efficient than GDX (as more and more agents adopt AA) when in the majority (shown by the higher magnitude when close to 1, i.e. AA), while GDX is more

efficient than AA when in the majority (shown by the higher magnitude when close to 0, i.e. GDX). At the saddle point  $C'$ , the strategies are equally efficient and, thus, it does not pay for an agent to deviate to another strategy. When we consider market efficiency, we also observe that the agent's behaviour that maximises the social welfare is not necessarily evolutionarily stable (shown by the fact that the efficiency at ESS  $B'$  is lower than at  $C'$ ). Thus, while the behaviour at  $C'$  is better than that at  $B'$ , from the market designer's perspective,  $C'$  is not a stable outcome, and the behaviour will change (even if the market efficiency decreases whether the agents select GDX or AA) and converge to an ESS.

As we would expect, this shows certain similarities with the dynamics of the two-population model. Specifically, we also have two attractors at the pure strategies, AA and GDX, but with a probability of 0.78 and 0.12 that they will be adopted respectively. Thus, the one-population model incorrectly predicts the probabilities of the outcomes and the eventual state of the market behaviour. Moreover, because the one-population model cannot predict whether buyers and sellers will adopt different strategies, it can miss different mixed-Nash equilibria and, in this case, it missed the ESS at C. Now, the impact of using the wrong ESSs and probabilities is that the market designer incorrectly predicts the evolutionary stable efficiency and its likelihood. While it is possible that the one-population model predicts a high evolutionary stable efficiency, it is possible that the market behaviour settles at an unpredicted and very poor market efficiency (where buyers and sellers have different behaviours). This is not desirable, and is exactly what the market designer must ensure never occurs. Furthermore, the analysis of our two-population model explicitly shows that the assumption that buyers and sellers adopt the same behaviour is wrong. If we consider an example at  $x_1 = 0.25$  in the one-population model, we can see that in this case, the agent deviates to GDX. In the two-population model, at  $p_1 = 0.25$  and  $q_1 = 0.25$ , the buyer deviates to GDX, while the seller deviates to AA. Thus, buyers and sellers do not necessarily adopt the same strategy and such a wrong assumption simply cannot be made. This is a fundamental and intrinsic shortcoming of Walsh *et al.*'s approach which this two-population model addresses.

In summary, the one-population model only considers the case where  $p = q$ . The one-population analysis is really an approximation of the dynamics of this model as the cross-section along the  $p_1 = q_1$  axis in the two-population model. However, because of this approximation along the one dimensional space (where buyers and sellers deviate to the same strategy), it effectively misses all the possible phenomena that occur in the two-dimensional space (specifically how buyers and sellers generally deviate to different strategies). Furthermore, the observations that buyers and sellers do indeed adopt different strategies when they are given the freedom to do so is proof that our method is a more accurate predictor of behaviour. Indeed, we showed that the assumption that buyers and sellers behave in a similar manner is flawed and there is a whole space of dynamics that is overlooked in the original model.

## 4 Conclusions

With software agents becoming increasingly prevalent, a better understanding of their interactions in marketplaces is highly desirable. To this end, we advance the state of the art by developing a novel analytical model of marketplaces with mul-

tiple buyers and sellers. Our model removes the key restriction of previous work that agents must adopt the same strategy as both a buyer and a seller. In so doing, our model provides a more insightful analysis in terms of buyers' and sellers' strategic behaviours. To illustrate this, we showed how our model identified important interactions in the CDA that Walsh *et al.*'s model could not.

For future work, we want to look at other double-sided market mechanisms with multiple buyers and sellers (such as the call double auction). We also intend to look at more general market games where we have players with two or more different roles that are not necessarily buyer and seller. First, we intend to look at games with multiple buyers and an auctioneer in which the buyers are endowed with a set of buyer strategies and the auctioneer with a set of clearing strategies. Then, with each type of player being rational and trying to maximise their economic efficiency, our model will provide us with insight into how the market structure (determined by the market clearing) and bidding behaviours would co-evolve. Second, we intend to look at markets with buyers and sellers *and* additional players such as arbitrageurs and observe how their different behaviours co-evolve in the marketplace.

## References

1. S. H. Clearwater, editor. *Market-Based Control- A paradigm for distributed resource allocation*. World Scientific, 1996.
2. M. Hollander and D. A. Wolfe. *Nonparametric Statistical Methods*. Wiley, 1973.
3. R. D. McKelvey and A. McLennan. Computation of equilibria in finite games. *Handbook of Computational Economics*, pages 18–20, 1996.
4. S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double auction market designs. *Proc. 6th AAMAS Workshop on Agent Mediated Electronic Commerce*, pages 192–198, 2004.
5. G. Tesauro and J. L. Bredin. Strategic sequential bidding in auctions using dynamic programming. *Proc. First international joint conference on Autonomous agents and multiagent systems*, pages 591–598, 2002.
6. K. Tuyls, T. Lenaerts, K. Verbeeck, S. Maes, and B. Manderick. Towards a relation between learning agents and evolutionary dynamics. *Proc. 14th Belgian-Dutch Conference on Artificial Intelligence*, pages 21–22, 2002.
7. P. Vytelingum. *The structure and behaviour of the Continuous Double Auction*. PhD thesis, School of Electronics and Computer Science, University of Southampton, 2006.
8. P. Vytelingum, D. Cliff, and N. R. Jennings. Evolutionary stability of behavioural types in the continuous double auction. *Proc. AAMAS Joint Workshop on Trading Agent Design and Analysis and Agent Mediated Electronic Commerce VIII*, pages 153–166, 2006.
9. 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.
10. W. E. Walsh, R. Das, G. Tesauro, and J. O. Kephart. Analyzing complex strategic interactions in multi-agent games. *Proc. AAAI Workshop on Game-Theoretic and Decision-Theoretic Agents*, pages 109–118, 2002.
11. J. W. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1995.
12. M. P. Wellman, D. M. Reeves, K. M. Lochner, S. Cheng, and R. Suri. Approximate strategic reasoning through hierarchical reduction of large symmetric games. *Proc. 20th AAAI Conference*, pages 502–508, 2005.