Strategic Bidding in Continuous Double Auctions

P. Vytelingum [a] D. Cliff [b] N. R. Jennings [a]

[a] *School of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, United Kingdom.*
{pv, nrj}@ecs.soton.ac.uk

[b] *Department of Computer Science, University of Bristol, Bristol BS8 1UB, United Kingdom.*
dc@cs.bris.ac.uk

Abstract

In this paper, we describe a novel bidding strategy that autonomous trading agents can use to participate in Continuous Double Auctions (CDAs). Our strategy is based on both short and long-term learning that allows such agents to adapt their bidding behaviour to be efficient in a wide variety of environments. For the short-term learning, the agent updates the *aggressiveness* of its bidding behaviour (more aggressive means it will trade off profit to improve its chance of transacting, less aggressive that it targets more profitable transactions and is willing to trade off its chance of transacting to achieve them) based on market information observed after any bid or ask appears in the market. The long-term learning then determines how this aggressiveness factor influences an agent's choice of which bids or asks to submit in the market, and is based on market information observed after every transaction (successfully matched bid and ask). The principal motivation for the short-term learning is to enable the agent to immediately respond to market fluctuations, while for the long-term learning it is to adapt to broader trends in the way in which the market demand and supply changes over time. We benchmark our strategy against the current state of the art (ZIP and GDX) and show that it outperforms these benchmarks in both static and dynamic environments. This is true both when the population is homogeneous (where the increase in efficiency is up to 5.2%) and heterogeneous (in which case there is a 0.85 probability of our strategy being adopted in a two-population evolutionary game theoretic analysis).

Key words: Continuous Double Auction, Bidding Strategy, Evolutionary Game theory.

Preprint submitted to Elsevier Science *21 May 2008*

1 Introduction

The emergence of software agents that are capable of flexible, autonomous actions and interactions is changing the online trading landscape [12]. In particular, one of the most significant applications of such agents is in the Continuous Double Auction (CDA) [9], where multiple buyers and sellers compete with one another to buy and sell goods and services. Such CDAs are one of the most common forms of marketplaces and have emerged as the dominant financial institution for trading securities and financial instruments. Indeed, today, the major exchanges like the NASDAQ and the New York Stock Exchange (NYSE) and the major foreign exchanges (FX) use variants of the CDA institution (with the total value of trades on the NYSE standing at around a yearly 12.4 trillion dollars [1] and foreign exchanges worth in excess of 1.9 trillion dollars [2] of daily transactions). Other significant applications are in market-based control [2], where CDAs provide a dynamic and efficient approach to the decentralised allocation of scarce resources. Examples of such market-oriented applications range from the allocation of air pollution permits [15], to air-conditioning systems [3], to complex resource allocation problems where suppliers have limited capacity and consumers have inelastic demand [8]. In all of these cases, the CDA is used because it exploits the dynamics of the free market to balance demand and supply efficiently in a highly responsive and decentralised system. This is in contrast to most auctions that have a single, centralised auctioneer responsible for the matching of participants and resources in the market. In the CDA, on the other hand, the resource allocation is an emergent behaviour of the complex interactions of the individual trading agents, with transactions corresponding to allocations. Thus, there is no single agent responsible for the allocation. This decentralisation allows the system to be very robust because its performance degrades gracefully if agents fail, rather than having a single point of failure in the case of the centralised auctioneer.

Given its prominence and importance, considerable research endeavour has been invested into devising strategies for agents that participate in CDAs. However, there is no known dominant strategy [9]. Thus, many strategies have been developed as heuristic-based, decision-making algorithms that attempt to best exploit the observable market information available to the agents in order to maximise their profits (see Section 3). Indeed, several of these strategies have been shown to outperform human traders in laboratory experiments [7]. However, we believe that more efficient strategies can still be developed and in this paper we develop just such a strategy.

---

[1] `http://www.nyse.com/pdfs/movolume0505.pdf`

[2] `http://www.bis.org/publ/rpfx05t.pdf`

In particular, to date, the extant CDA strategies have typically been developed assuming that the market is *static*, meaning there is no change in demand and supply at the beginning of each trading day. However, real markets such as NASDAQ and the NYSE are typically very *dynamic*, with changes frequently occurring (they are called *market shocks*). Thus, we believe the efficiency of strategies in dynamic environments is central to their application in practice. Now, although some of the designers have made initial attempts to show their strategies will still do well in dynamic markets, these strategies were not developed explicitly for such environments. This is a shortcoming because we believe that there are fundamental differences between static and dynamic environments. Specifically, these are primarily to do with the sporadically changing competitive market equilibrium price which is where demand meets supply in the market (see Section 2 for more details). When we have a market shock, the micro-economics of the free market dictates that the market dynamics, through the interaction of the traders, will force the transaction prices to a new competitive equilibrium price [16]. Given this, the efficiency of a strategy depends on how effective it is at adapting its bidding behaviour to the new market conditions, and thus to the new competitive equilibrium price. From this, our intuition is that different behaviours are needed when the market is relatively stable and when it is changing. In particular, in the static case, the agent can be effective by assuming that the competitive equilibrium does not change significantly, whereas in the dynamic case, it can make no such assumption and must learn, assuming that this competitive equilibrium may change. Furthermore, for maximum generality, we want to ensure that our strategy performs well in *homogeneous* populations in which all the agents use it (as would typically be the case in market-based control applications) and in *heterogeneous* populations in which agents can adopt a range of alternate strategies (as would be the case in financial institutions). Given this, we simply assume that an agent is selfish and tries to maximise its individual return and that it is unaware of whether it is trading in a homogeneous or a heterogeneous environment.

Against this background, we have developed a novel bidding strategy for CDAs. In particular, we employ a short-term and a long-term learning mechanism to update the agent's bidding aggressiveness [3] to remain competitive in the market. We focus on aggressiveness, in particular, as we believe it is the key determinant of success in the market. It is central because it describes how the agent manages the trade off between profit and probability of transaction. Thus, an *aggressive* trading agent tries to increase its chance of transacting by placing bids, that are not necessarily highly profitable. In contrast, its *passive*

[3] In some work, the trader's risk attitude has been used to describe broadly the same behaviour [1]. However, we believe that such a property is intrinsic to the trader, and thus, is not an appropriate term to describe our changing behaviour in this case.

counterpart tries to transact at more profitable prices, but has to trade off its chance of actually transacting. When the agent is not able to transact, it could choose to become more aggressive, such that it increases its chances of being able to transact, and, conversely, when it can transact, it could choose to become more passive in order to try to increase its profits. In other words, the agent could react to the market information by being more or less aggressive based on how it is performing in the market. Given this, we employ a short-term learning mechanism to fine-tune the agent's aggressiveness whenever it submits a bid or an ask, or a transaction occurs (if the bid and the ask match) in the market. The actual way in which the degree of aggressiveness translates to a bid or an ask to submit in the market can be fixed or can be linear, in which case the aggressiveness would be similar to the agent's profit margin. However, we believe that this mapping should be updated depending on the prevailing market conditions. Thus, we employ a long-term learning strategy that adapts this mapping, in a non-linear fashion, to the changing conditions and, in particular, to the volatility of transaction prices. We refer to this learning as long-term because the occurrence of bids and asks is a fraction of the number of transactions [4] that occur and because the benefit of learning this mapping is only really observable over a number of trading days. The purpose of the long-term learning is especially evident in dynamic markets where market conditions can change drastically and a different mapping should then clearly be adopted. Hereafter, we refer to our strategy as the *Adaptive-Aggressiveness* (or AA) strategy.

This work advances the state of the art in the following ways. First and most importantly, we develop a novel bidding strategy, AA, with a short-term and a long-term learning mechanism that adapts to market conditions by being more or less aggressive. We then empirically show that our strategy outperforms (by up to 5.2% [5]) the state of the art in homogeneous populations and

[4] During a trading round, multiple bids and asks are submitted until a bid and an ask match and a transaction occurs at the end of the round. There are typically several trading rounds in a trading day.

[5] While at first glance, the improvements may seem relatively small, there are a number of contextual factors that should be taken into account. First, the CDA has already been shown to be an efficient mechanism in many settings regardless of the trading strategy [11], such that the scope for improvement is limited. Second, over the past 15 years, several strategies have been designed and improved upon, to be ever more efficient. The evaluation of these strategies in a common, but static, market setting [22] shows how they have systematically improved starting with the ZI strategy [11] (with a baseline market efficiency of 98.3%), through the ZIP strategy [6] (with a market efficiency of 99.7%) and finally with the family of GD strategies [10] (with a market efficiency of up to 99.7%). The work on the GDX strategy [21] showed that GD [10] was outperformed by the improved GDX which extracted 1.0% more profit, though the paper fails to provide the market efficiency using the GDX strategy. In our work, we were able to show that the AA strategy improves

heterogeneous populations (by being dominant or with probabilities of over 85% of being adopted in all the different environments we consider). Second, we advance the state of the art in the methodology for analysing the CDA within both homogeneous and heterogeneous populations. For the former, we look at the daily market efficiency and price volatility, rather than simply the overall efficiency as is commonly done in the literature. This is an advance because our methodology provides more insights into how the efficiency of a strategy changes as the strategies learn over the different trading days and because it identifies the drastic decrease in efficiency after a market shock. For the latter, we develop a methodology with a novel two-population game theoretic model that analyses the buyers' and sellers' strategic behaviours in the CDA. Here, we advance the state of the art by analysing the evolution of buyer and seller strategies in the market which contrasts with the current state of the art that assumes an agent adopts the same strategy as both a buyer and a seller. Finally, our work is the first to systematically compare state of the art CDA strategies in a wide variety of (static and dynamic) market settings, based on different types of demand and supply. It is also the first work that considers a market setting whereby the equilibrium price changes constantly over the trading days, which better reflects real markets, and, indeed, we considered data from a real financial market when designing such a setting.

The remainder of this paper is structured as follows. We begin, in Section 2, with a detailed description of the CDA mechanism. We provide a detailed discussion of previous strategies that have been developed for the CDA and methodologies that have been used to analyse their effectiveness in Section 3. Section 4 details the AA strategy. We provide an empirical evaluation of AA in both static and dynamic environments for both homogeneous and heterogeneous populations in Section 5. Section 6 concludes.

2 The Continuous Double Auction Mechanism

Market trading is governed by a market mechanism; defined by a *market protocol* that determines the nature of bids and asks allowed in the market, the clearing rule that indicates when a transaction occurs, the pricing rule that indicates the price at which a transaction occurs and the information published to the buyers and sellers in the market. The CDA is one such mechanism. However, there exist many variants, based on different market protocols. For example, in financial institutions like the NYSE, some traders have different

the market efficiency of the CDA significantly (with a considerably bigger leap than moving from ZIP to GD or to GDX), demonstrating a genuine advancement of the state of the art.

levels of privilege with better access to other traders' messages than is available to unprivileged traders (usually to improve the overall efficiency of the system) or Dash *et al.* describe a variant of the CDA for market-based control applications [8] (with modified clearing rules for inelastic demand).

These examples of the CDA are highly domain specific and difficult to generalise from. Thus, most research in this area (e.g. [22, 6, 11]) has generally been structured around the market protocol initially proposed by Smith [20]. In this, multiple buyers and sellers are allowed to submit bids and asks in a market for homogeneous, single-attribute goods, and the market clears (with a single trade) whenever a bid and an ask match (hence, the continuous nature of the CDA), and clears at the average of the bid and the ask. Furthermore, the protocol includes the NYSE *spread-improvement* and the *no-order queuing* rules. The former requires that a submitted bid or ask improves on the outstanding bid (the highest unmatched bid) or the outstanding ask (the lowest unmatched ask) respectively, while the latter specifies that offers are single-unit, are not queued in the system, and are simply erased when a better offer is submitted. The CDA lasts several trading days, with a trading day itself lasting several trading rounds which is the period during which bids and asks are submitted (with the bid-ask spread decreasing) until the market clears.

To more formally analyse the CDA, we now explore some of these basic notions in more detail:

Definition 1 *A **trading day** is the period (with a deadline) during which traders are allowed to submit bids and asks (resulting in transactions whenever these match), at the end of which the market closes. At the beginning of a trading day, traders are endowed with a set of goods to buy or sell (that determine the market demand and supply).*

Definition 2 *The **outstanding bid**, o_{bid}, is the current maximum (uncleared) bid submitted in the market.*

Definition 3 *The **outstanding ask**, o_{ask}, is the current minimum (uncleared) ask submitted in the market.*

Definition 4 *The **bid-ask spread** is the difference between o_{bid} and o_{ask}.*

Definition 5 $\boldsymbol{\Delta}$ *is the minimum bid or ask increment in the market.*

Definition 6 ***MAX*** *is the maximum bid or ask allowed in the market (to prevent unreasonably high asks and speed up the trading process).*

Definition 7 *A **trading round** is the period during which bids and asks are submitted until there is a match and a transaction occurs. There are typically several trading rounds in a trading day. At the beginning of the trading round,*

$o_{bid} = 0$ *and* $o_{ask} = MAX$.

Furthermore, as defined by the model proposed by Smith and adopted in the literature, at the beginning of each trading day, each agent is endowed with a set of limit prices corresponding to the goods it would like to buy or sell:

Definition 8 *The* ***limit price*** *is the maximum bid a buyer is currently willing to offer, and the minimum ask a seller is willing to offer.*

Definition 9 $\boldsymbol{\ell_i}$ *is the limit price of buyer i; that is the highest bid price it is willing to submit.*

Definition 10 $\boldsymbol{c_j}$ *is the limit price of seller j; that is the lowest ask price it is willing to submit.*

Finally, we define the following notions that we employ to analyse and evaluate the CDA mechanism:

Definition 11 *The* ***competitive market equilibrium*** *is when demand meets supply in a free market populated by rational and selfish agents. According to the classical micro-economic theory, the transaction prices in the CDA are then expected to converge towards that competitive equilibrium price* p^*. *As* p^* *can only be calculated if the demand and supply are available, which is not the case here because of the decentralised nature of the CDA,* p^* *cannot be known a priori.*

Definition 12 *The* ***market efficiency*** *is the ratio of all agents' surpluses in the market to the maximum possible surplus that would be obtained in an allocation where the profits of all buyers and sellers are maximised.*

Definition 13 *The* ***efficiency of a bidding strategy*** *is the ratio of the profits of the agents adopting that strategy during a trading day to the maximum profit these agents could extract in an efficient, centralised allocation. In the homogeneous scenario, this is identical to the market efficiency, while in a heterogeneous scenario, the mean efficiency of all the strategies is equal to the market efficiency.*

In our work, we consider a discrete-time simulator of such a CDA model, and at each time step, an agent is randomly triggered to submit a bid or an ask in the market. In line with previous work, we impose a deadline on the duration of a trading day with the auction closing after 1000 time steps. Now, for controlled experiments, we specify single-unit or multi-unit allocations [6]

[6] A single-unit allocation is when a trader is given a single unit (with a corresponding limit price) to buy or sell. A multi-unit allocation is when a trader is given multiple units (with a corresponding set of limit prices) to buy or sell.

endowed to buyers and sellers at the beginning of a trading day, to induce a desired demand and supply for the market. In particular, limit prices are drawn from uniform distributions U_b and U_s for buyers and sellers respectively. We chose a uniform distribution in order to obtain an expected linearly decreasing demand curve and an expected linearly increasing supply curve, commonly found in the micro-economics literature [16]. For the purposes of this paper, we consider the following different uniform distributions to model representative (symmetric [7] and asymmetric) markets similar to those considered in previous studies [5, 6, 22] (see Figure 1), and we further describe how we use such markets to induce market shocks.

- Market 1 (M1): $U_b = \mathcal{U}(1.5, 4.5)$ [8] and $U_s = \mathcal{U}(1.5, 4.5)$. This is a symmetric market that has an expected competitive equilibrium price, p^*, $= 3.0$.
- Market 2 (M2): $U_b = \mathcal{U}(1.5, 4.5)$ and $U_s = \mathcal{U}(2.8, 3.2)$. An asymmetric market with a flat supply curve. $p^* = 3.0$.
- Market 3 (M3): $U_b = \mathcal{U}(2.8, 3.2)$ and $U_s = \mathcal{U}(1.5, 4.5)$. An asymmetric market with a flat demand curve. $p^* = 3.0$.
- Market 4 (M4): $U_b = \mathcal{U}(2.5, 5.5)$ and $U_s = \mathcal{U}(2.5, 5.5)$. A symmetric market with $p^* = 4.0$.

Definition 14 *A* ***market shock*** *is a sudden change in agents' preferences (their limit prices) and, hence, in the market demand and supply, but not necessarily in the competitive equilibrium price. In our case, the shock occurs at the beginning of a trading day with the new set of endowment of limit prices to the trading agents. There are different types of dynamic changes in real markets that are not referred to as market shocks, e.g. rallies (sustained upward movement of the competitive equilibrium price), sell-offs (sustained downward movement of the competitive equilibrium price), movements/trends (less sustained upward or downwards shifts). However, because it is not a central aspect of this work, we generalise the meaning of market shocks to cover all of these in this paper.*

In more detail, M1 and M4 have symmetric demand and supply (they differ only in their competitive equilibrium price). By considering cases with a flat demand (M3) or a flat supply (M2), we want to see how such extreme asymmetry will affect the efficiency of buyer and seller strategies in the CDA. Specifically, we did not want to constrain our study to symmetric demand and supply, which others have done [22, 21], because we want to contrast the efficiencies of strategies in markets with different demand and supply, and observe what effect this has on the strategic interactions of agents. Finally, the

[7] In a symmetric market, the ratio of the gradient of the demand curve and that of the supply curve is -1. M1 is an example of a symmetric market, while M2 is that of an asymmetric market (see Figure 1).

[8] $\mathcal{U}(u, v)$ is a uniform distribution between u and v.

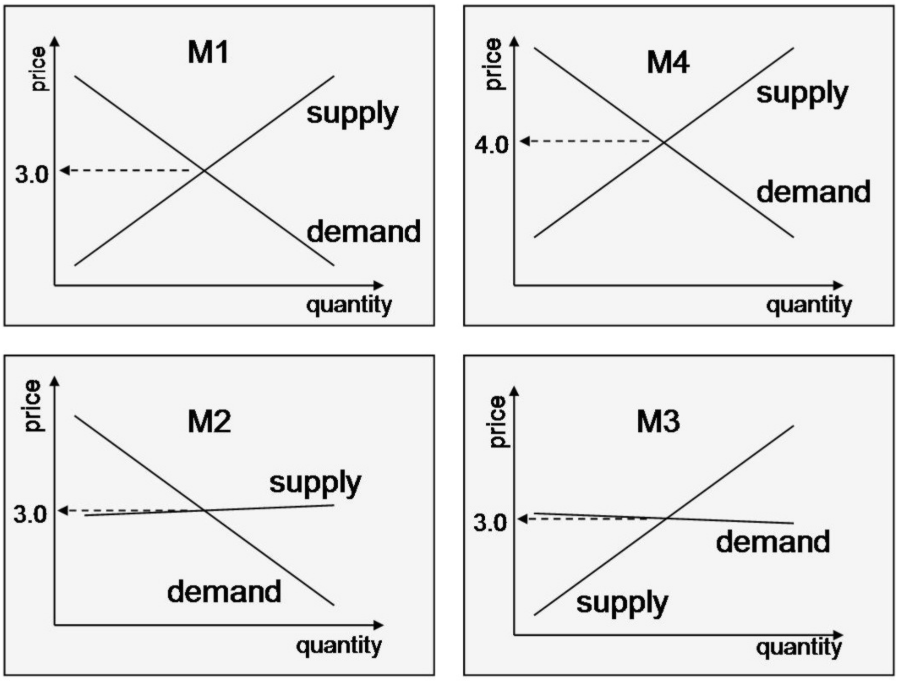

Fig. 1. Expected demand and supply of markets M1, M2, M3 and M4. The dashed arrow points to the competitive equilibrium price, p^*.

purpose of M4 is to observe how the strategies perform when the competitive equilibrium price changes during a market shock. This is important to see how fast the strategies can adapt to the new equilibrium, given that the slower they adapt, the longer they remain inefficient.

In a static environment, a CDA typically lasts several trading days, with trading agents receiving the same set of limit prices at the beginning of each trading day. However, because dynamic environments are commonplace (see Section 1) we also need to investigate the efficiency of our strategy in such situations. To this end, as per previous work including Smith's, we induce a market shock by changing the market demand and supply. This is effected by changing the endowment of limit prices at the beginning of a trading day. For example, for a CDA lasting 20 trading days, we could use M1 for the first 10 and M4 for the last 10 trading days, effectively inducing a shock on Day 11 (see Figure 2). Hereafter, we identify a market shock (MS) by the different markets it involves, and in our example, such a shock would then be identified as *MS14* to signify the fact that it is moving from M1 to M4.

Furthermore, and in order to be more realistic, we consider a scenario based on real market data from NASDAQ. In particular, we consider the history of Google shares (NASDAQ:GOOG) to analyse the performance of CDA strategies in a market where the demand and supply changes constantly every day (which could potentially affect learning strategies). The purpose of using real market data is to analyse how the changing demand and supply affects the

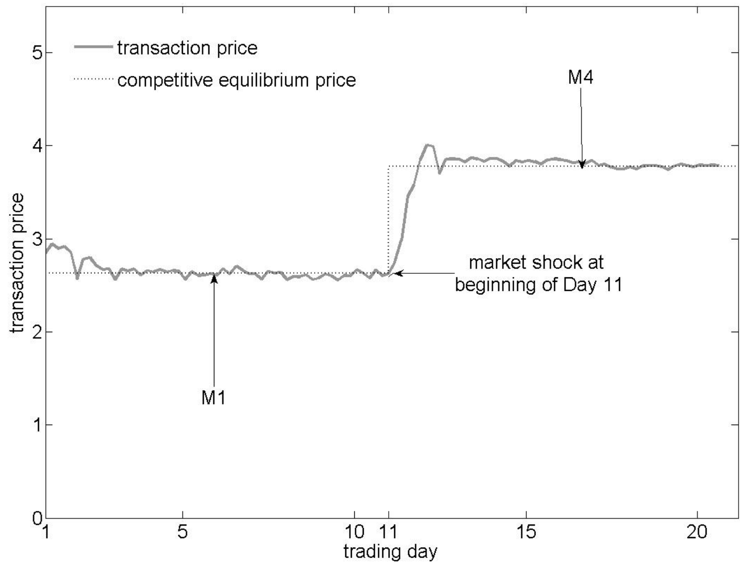

Fig. 2. Example of a transaction history with scenario MS14.

efficiency over trading days, rather than during a trading day. Given this, such an estimation of the equilibrium price is satisfactory for the CDA model where we assume that the demand and supply does not change during a trading day. Now, because in a real market, the market demand and supply changes during the trading day (while it does not in the standard CDA model), we estimate the equilibrium price (see Figure 3) as an average of the highest and the lowest transaction prices during each trading day. Within this context, in our experiments, we consider two specific sub-markets. The first one is based on the segment A to B (see Figure 3), which we will refer as the GOOG market hereafter. This is typical of the market fluctuations that are most common in real markets. The other one is a more dynamic market, based on segment C to D (see Figure 3), which we will refer to as GOOGshock, where the market experiences a sharp and sustained drop in the transaction prices.

3 Related Work

Research on CDAs was significantly advanced by Smith's seminal work [20] on competitive market behaviour. He showed that in CDAs populated by a relatively small number of selfish human traders, the market efficiency achieved in such a decentralised environment, where no single agent has complete and perfect information about the system, was close to one (the maximum possible value — see Definition 12). This result was ground-breaking as it showed that markets governed by a decentralised mechanism, such as the CDA, do not have to be large to be efficient, as had previously been assumed [16]. As we

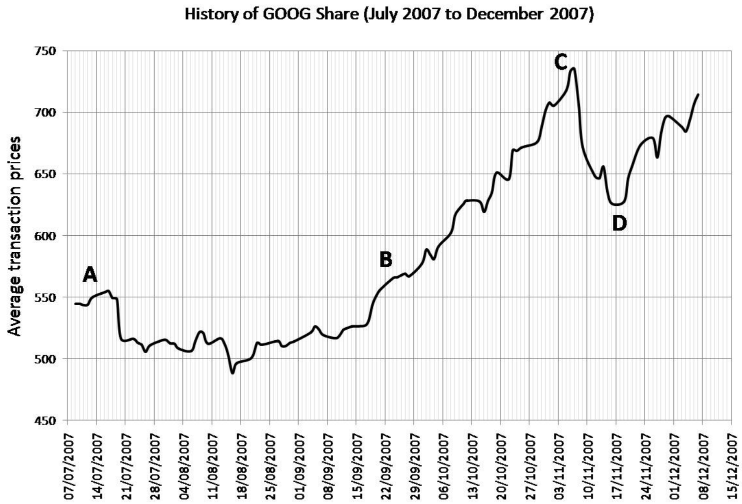

Fig. 3. History of GOOG share prices (July 2007 to December 2007). Segment A to B corresponds to the GOOG market while segment C to D correspond to the GOOGshock market.

will see, many subsequent research endeavours in this area have been heavily influenced by this work and we have also adopted the same broad methodology for the market and agent setup. In this section, we review, in turn, the work on strategies for the CDA and on analysing strategic interactions in such an institution since both of these are necessary to understand the contributions of this work.

3.1 Strategies for CDAs

In more detail, many strategies for the CDA have subsequently been developed, and over the last decade, there has been considerable emphasis on strategies for software trading agents with the emergence of electronic markets. In 1992, for example, the Santa Fe Trading Agent Competition [9] was organised to find the most efficient CDA strategy (see Definition 13) at the time. The competition winner was Kaplan [9] which is categorised as a sniping strategy that reacts to current market conditions and submits an offer to buy or sell only when one of the following conditions is met:

(1) For the buyer (seller), the best ask (bid) is less (higher) than the minimum transaction price in the previous trading day.
(2) For the buyer (seller), the best ask (bid) is less (higher) than the maximum transaction price in the previous trading day and the ratio of the bid-ask spread and the best ask (bid) is less than some spread factor, while the expected profit is more than some minimum profit factor.

(3) The fraction of time remaining in the trading day is less than some time factor (hence the reference to being a sniping strategy).

However, it was also shown in an evolutionary tournament [9] that the Kaplan strategy does not perform well when a majority of the trading agent population adopt it [9]. This is because Kaplan has a simple reactive behaviour and does not learn to be more efficient in the market given the changing market conditions. Indeed, it works primarily by exploiting the more complex (adaptive) behaviours of other strategies that are learning to be more efficient in the market, and as it becomes a bigger proportion of the market, there are fewer agents to exploit. Hence, its performance decreases.

Another important piece of work in the literature is the Zero-Intelligence (ZI) strategy by Gode and Sunder [11]. A ZI agent makes an uninformed, but profitable decision that is not based on observed market information. In particular, a ZI buyer submits an offer drawn from a uniform distribution between 0 and its limit price, and a ZI seller between its limit price and the maximum offer, MAX, allowed in the market. Gode and Sunder showed that CDAs populated by these non-intelligent trading agents were still highly efficient. They then conjectured that the high market efficiency was principally due to the structure of the market mechanism (i.e. its protocol), rather than how intelligent the agents were. Their work was subsequently critiqued by Cliff and Bruten who argued that the high efficiency observed was in fact an artefact of the symmetric demand and supply considered, and that at least a minimal intelligence is necessary to achieve efficiency that is comparable to that of CDAs with human traders [20]. They justified their critique by developing the Zero-Intelligence Plus (ZIP) strategy [6] which, indeed, was shown to achieve market efficiency close to that of human traders, and to considerably outperform the ZI strategy. The ZIP strategy's minimal intelligence is that it learns its best profit margin (to maximise its profits while attempting to trade) given a set of six different rules that decide whether to increase or decrease the profit margin based on market information. In particular, the profit margin is updated after every market event (a bid or an ask being submitted, or a transaction occurring), and there are 8 parameters covering things such as the learning mechanism (including the learning rate and the momentum coefficient [10]) that determines how this margin is updated, the initial profit margin and the relative and absolute increase and decrease in profit margin.

---

[9] In an evolutionary tournament, at the beginning of each run of the CDA, the proportion of strategies being adopted corresponds to their relative efficiencies [9]. Thus, the more efficient strategies will be more common in the market.

[10] The momentum-based update of the ZIP strategy considers the movement of the bids and asks submitted in the market to minimise the effect of high frequency change in these bids and asks. The momentum coefficient is then the weight given to previous bids and asks [6].

An alternate approach is adopted by the GD strategy [10] and its subsequent extension GDX [21]. GD is an expected profit-maximising and belief-based strategy. Specifically, it calculates its belief that a bid or an ask will be accepted in the market based on a set of the most recent transactions and submitted, but unaccepted, bids and asks, and the expected profit associated with such bids and asks. Then, the bid or ask that maximises the expected profit is submitted in the market. GDX calculates the belief in a similar manner, but uses dynamic programming (coupled with the expected profit maximisation process) to decide on the best price and when to submit a bid or an ask. This improved version adds another dimension to the decision-making process, namely time as the expected number of bidding opportunities before the auction closes. This means the GDX agent has the opportunity to trade later on during the trading day and can thus wait for more profitable transactions. The GD family contrasts with ZIP particularly in how the latter learns its profit margin. In fact, ZIP uses a scalar parameter, based on the latest bid or ask or transaction price, while GD builds up its belief using a set of bids, asks and transaction prices over a number of the most recent trading rounds. The GD approach is significantly more computationally intensive than the ZIP approach, particularly when coupled with the dynamic programming, but has been proven to be more effective in symmetric markets given homogeneous and heterogeneous populations [21].

Other strategies have also been developed, including the FL strategy [13] which uses fuzzy logic to form a bid or an ask to submit in the market and the modified Roth-Erev strategy based on a myopic reinforcement-learning algorithm [17, 18]. However, in this paper, we will consider only ZIP and GDX when benchmarking our AA strategy. We chose these because they are the two most widely used benchmarking strategies in the literature (being used in [7, 28]). Moreover, they have also been shown to be the most efficient [22, 21].

3.2 Evaluating Strategic Interactions in CDAs

Given the benchmarks, we now require a methodology for analysing the strategic interactions of agents adopting our strategy or these benchmarks in both homogeneous and heterogeneous populations. For the homogeneous case, the emergent behaviour of the market is most interesting from a system designer's perspective. For the heterogeneous case, the emergent behaviour of the market is also of interest, though, from an agent's perspective, the efficiency of the different strategies in the market is also important. Thus, these two types of population require a different type of evaluation methodology and, so, we deal with each in turn.

In homogeneous populations, previous work has typically looked at the average efficiency of a strategy over several trading days, in static markets with symmetric demand and supply (e.g. [22, 10]). However, we believe this has a number of shortcomings.

First, analysing the *daily efficiency* of a strategy provides more insight into how effective it is in learning from market interactions. This view is partly supported by [6], though they focus on the daily price volatility. Specifically, we believe that as the agent learns to be more competitive in a static market and the transaction prices converge towards the competitive equilibrium price, we expect its efficiency to improve. Thus, it is important to measure efficiency on a daily basis because it gives us insights into how the behaviour of the market is changing and, in particular, how it is improving. Moreover, such observations would not be possible if we just focused on average efficiency because we end up with a scalar value that does not say anything about the trend of the daily efficiencies.

Second, we believe that daily price volatility, α, should also be looked at. To date, however, only Cliff and Bruten consider such a metric:

$$\alpha = \frac{1}{p^*}\sqrt{\frac{\sum_{i=1}^{N}(p_i - p^*)^2}{N}} \tag{1}$$

where p_i is the price of transaction i, and N is the number of transactions over which we are investigating the convergence. Because the competitive market equilibrium is usually central in a strategy, the price volatility, calculated as Smith's parameter (see Equation 1), is important because it gives insights into how the agents adjust their behaviours such that the transaction prices converge to that equilibrium. The rate of convergence usually determines how fast the market reaches a high efficiency and, thus, is useful in analysing the effectiveness of a strategy in a homogeneous population.

Third, only Cliff and Bruten have looked at dynamic environments with different market demand and supply. However, they only describe how transaction prices change, and not how daily efficiency and price volatility change in such environments. This is important because, as discussed in Section 1, we are considering decentralised resource allocation in both static and dynamic environments. Thus, it is interesting to analyse how daily efficiency and price volatility change in both types of environments. Furthermore, because demand and supply cannot be known *a priori*, we must ensure that the strategies are evaluated in markets with different types of representative demand and supply, and not simply the standard cases (with symmetric demand and supply) to ensure the significance of our analysis.

In sum, given these shortcoming, we will use an analytical method that considers both market efficiency and price volatility, on a daily basis, to highlight this learning in both static and dynamic markets with different market demand and supply (see Section 2 for more details).

When we consider methodologies for evaluating strategies in heterogeneous populations, we come across two principal approaches. The first one (adopted in [22, 21, 27]) consists of comparing the efficiency of strategies in balanced populations (where strategies are adopted in equal proportions). However, this approach fails to consider unbalanced populations where strategies are present to different degrees. The second one, proposed by Walsh *et al.* (2002) and adopted in [26, 18], does allow unbalanced populations. This approach is important because a strategy might perform better or worse based on the number of buyers and sellers that adopt it, an insight which would allow us to better evaluate a strategy and, thus, we consider this approach in this work.

In particular, Walsh *et al.* propose an evolutionary game-theoretic (EGT) approach based on computing the mixed-Nash equilibrium of heuristic strategies and the dynamics analysis of equilibrium convergence [29]. Such an EGT analysis is insightful because it has been shown to approximate the learning of agents (using a standard learning technique such as reinforcement learning) in a multi-agent system [23, 24], which, in our case, translates to traders learning to adopt the better strategies in the market. Now, because an EGT analysis is infeasible for all but the simplest games (such as the Prisoner's Dilemma [29]), Walsh *et al.* describe how complex games that involve repeated interactions with more elaborate actions and payoffs, can be made amenable to such analysis. Specifically, their model considers the high-level, heuristic strategies of the trading agents as simple actions, and the payoff to these strategies is 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 games, namely the Automated Dynamic Pricing (ADP) game and the CDA game. In the former, they analyse how sellers endowed with a set of heuristic strategies interact in the market, and what strategies these sellers are most likely to adopt. In the latter, they consider the strategic interaction of agents that use the same strategy as a buyer and a seller. Their methodology has now been widely adopted and, in particular, [18] used it to compare two different auction mechanisms (the continuous and the call double auction mechanism) where similar strategies were available for both.

The EGT approach to evaluating strategies in heterogeneous populations (in the more general case of unbalanced populations) is indeed more insightful than simply comparing the efficiencies of strategies in balanced populations. However, a key assumption of this approach 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 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, 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, and the demand and supply of the market then determines the complex interactions of these strategies which in turn, determines their overall effectiveness. To maximise its profit, we believe 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 for both roles can only have a negative effect on the agent's economic efficiency. We believe that such an assumption should not be made because this approach may well miss some important phenomena. Thus, an EGT approach to analyse how the buyer *and* seller strategies separately evolve in the market is needed. Now, just such an approach is developed in [25] and it is shown to offer new insights into the strategic behaviour of agents in the CDA when compared with Walsh *et al.*'s model. Given this, we also deploy this method in this work (see Subsection 5.1.4 for more details).

4 The AA Bidding Strategy

In this section, we describe our AA strategy in detail. As argued in Section 1, given that an agent can make a more informed decision based on every additional piece of information observed in the market, being adaptive to new market conditions can allow it to be more profitable. Here, the market conditions describe all observed information and how it leads to the history of transaction prices and the current outstanding bid and ask. For this reason, in the AA strategy, we adapt the degree of aggressiveness in the agent's bidding strategy to reflect its beliefs about the prevailing market situation. Specifically:

Definition 15 ***Aggressiveness*** *is defined as the inclination to interact more actively in the market. The* ***aggressive*** *trader submits better offers than what it believes the competitive equilibrium price to be, to try and improve its chance of transacting, and trades off profit for that purpose. The* ***passive*** *trader is less inclined to transact and more inclined to try and win a profitable transaction and thus submits offers that are worse than what it believes the competitive equilibrium price to be. The* ***active*** *trader submits offers at what it believes is the competitive equilibrium price, which is the expected transaction price.*

Thus, the agent can adopt behaviours that have different levels of aggressiveness, $r \in [-1, 1]$, ranging from aggressive ($r < 0$), through active at $r = 0$, to passive ($r > 0$), coupled with a learning mechanism to decide upon this level. Specifically, an agent that adopts a *passive* strategy waits for more profitable transactions than its estimate of p^* (hereafter the estimate is denoted by $\hat{p}^*$) and is willing to trade off its chance of transacting for higher expected profit. In contrast, an *aggressive* strategy trades off profit to improve its probability of transacting in the market. The *active* agent attempts to transact at $\hat{p}^*$ which is the expected transaction price. Now, because market conditions keep changing, different levels of aggressiveness are likely to be best at different times and so the agent needs a means for updating r. Thus, an AA agent has two principal decision-making components: (i) a bidding one that specifies what bid or ask should be submitted based on its current degree of aggressiveness, and (ii) a learning mechanism to update its behaviour according to the prevailing market conditions. In more detail, these two components can be represented by two distinct layers, the bidding layer and the adaptive layer (see Figure 4).

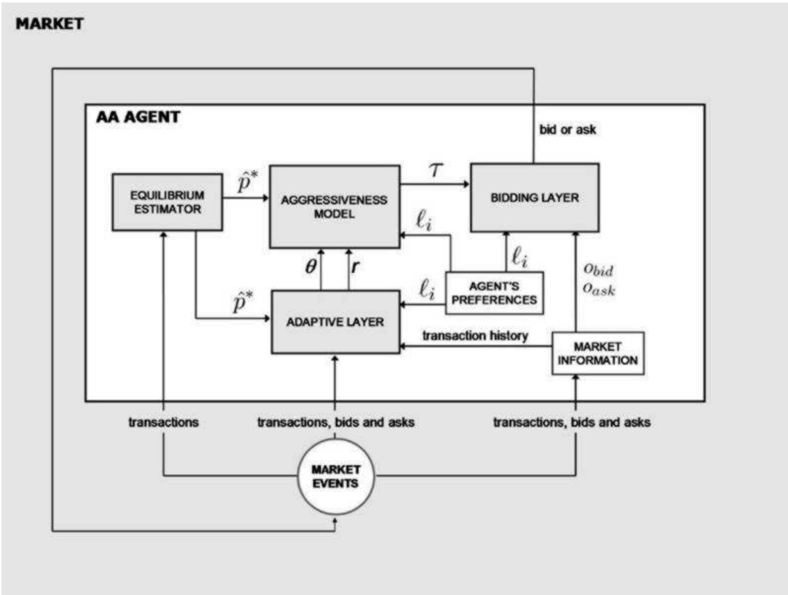

Fig. 4. The AA bidding strategy.

The first layer determines which bids or asks to submit given a set of bidding rules (see Subsection 4.4). These rules specify how to react to the current market conditions given the target price τ which represents the agent's most competitive price in the market. A bid (or ask) is competitive if it is the agent's most profitable bid (or ask) that it believes would be accepted in the market. Note that a bid is always less than or equal to, and the ask always more than or equal to, its limit price. This is similar to the price given by the ZIP agent's profit margin and the price that maximises the GDX agent's expected profit. Now, the *aggressiveness model*, as described in detail in Subsection 4.2, gives a mapping function to τ of the agent's current degree of aggressiveness, its limit price, $\hat{p}^*$ (which is provided by the *equilibrium estimator* described in Subsection 4.1), and an intrinsic parameter θ. In particular, θ determines the shape of that mapping function (see Subsection 4.2 for more details).

The second layer represents the adaptive part of the strategy where the agent updates its bidding behaviour, when triggered by a market event (when a transaction occurs or a new bid or ask is submitted). This update causes the agent to adopt a more passive behaviour if it believes it can transact at a higher profit or a more aggressive one if it believes it is targeting too high a profit to transact. In particular, we have short-term and long-term learning mechanisms that update the agent's bidding behaviour. The former updates the degree of aggressiveness, r, whenever a bid or ask is submitted and is described in more detail in Subsection 4.3.1. The latter updates θ in the aggressiveness model after every transaction and is described in more detail in Subsection 4.3.2.

Each component of our strategy is now described in turn, in the following subsections, before the efficiency of the strategy is empirically benchmarked (in Section 5).

4.1 The Equilibrium Estimator

Because p^* cannot be known *a priori*, we use the *moving average* method for calculating it, based on the history of transaction prices (see Equation 2). We make this choice because the moving average is an objective analytical tool that gives the average value over a time frame spanning from the last transaction. Moreover, it is sensitive to price changes over a short time frame, but over a longer time span, is less sensitive and filters out the high-frequency components of the signal within the frame. Moving average thus allows us to emphasise the direction of a trend and smooth out large price fluctuations and, thus, we believe this is a reasonable choice. Based on our assumption that the transaction prices converge to the competitive equilibrium price, we introduce the notion of *recency* in the moving average by giving more weight to the more recent transaction prices. Specifically, Equation 2 describes how we calculate the estimate of the competitive equilibrium price, denoted by $\hat{p}^*$, given a set of the N most recent transaction prices:

$$\hat{p}^* = \frac{\sum_{i=T-N}^{T}(w_i \times p_i)}{N} \; where \sum_{i=T-N+1}^{T} w_i = 1, \; w_{i-1} = \rho w_i \tag{2}$$

where $(w_{T-N+1}, \ldots, w_T)$ is the weight given to the N most recent transaction prices, $(p_{T-N+1}, \ldots, p_T)$, and T is the latest transaction. Based on simulations, we set ρ to a value of 0.9 to emphasise any converging pattern in the history (see Figure 9 for an example).

The role of the aggressiveness model is to generate the current target price, τ, given the agent's current degree of aggressiveness r. In this context, an agent can be of two types; namely, intra-marginal and extra-marginal. A buyer (seller) is intra-marginal if its limit price is higher (lower) than the competitive equilibrium price. In contrast, the extra-marginal buyer's (seller's) limit price is lower (higher) than the competitive equilibrium price. Now, in a centralised mechanism with an efficient allocation (market efficiency is 1), only intra-marginal agents transact, while extra-marginal ones do not. However, in a decentralised mechanism, while intra-marginal agents are expected to transact, and extra-marginal counterparts are not, the latter do sometimes succeed in transacting. This is because transaction prices are never exactly at p^*, and thus, extra-marginal buyers can exploit asks below p^*, and extra-marginal sellers bids above p^*. In such cases, when the extra-marginal traders do transact, the allocation is no longer efficient and the market efficiency dips below 1.

Our aggressiveness model differs for these two type of traders fundamentally because of their limit prices and whether they are expected to transact given p^*. Given this, we consider them each in turn.

First, we consider the intra-marginal trader. In its aggressiveness model, a target price equal to $\hat{p}^*$ implies that the trader is active. When an intra-marginal agent adopts a passive behaviour, it considers a target price that is below (for the buyer) or an ask that is above (for the seller) p^*, in order to obtain a higher (than expected at $\hat{p}^*$) profit margin. Conversely, an aggressive attitude implies that the intra-marginal trader targets bids above (asks below) the competitive equilibrium price, which improves the probability of its bids (asks) being accepted (but decreases its profit margin).

For the intra-marginal aggressiveness model, we identified the following constraints that it should satisfy over the different degrees of aggressiveness. In particular, when the buyer is completely aggressive ($r = -1$), it targets a bid at its limit price and when it is completely passive ($r = 1$), it targets a bid at 0 (for maximum profit but no chance of actually transacting). The active buyer ($r = 0$) targets a bid at $\hat{p}^*$. Therefore, the aggressiveness function is defined at these three specific aggressiveness levels. Similar intuitions apply for the seller's aggressiveness function. However, when ($r = -1$), the seller submits the maximum ask, MAX (see Definition 6), allowed in the market. Given these constraints, there is an infinite solution space for such a function and so we choose a parameterised function [11] (see Figure 5) within the solu-

[11] While a range of functions could have been used, our choice was a basic function family where $f_\theta(x) = (e^{\theta x} - 1)/(e^\theta - 1)$, which, for $x \in [0, 1]$, takes values between 0 and 1, with $f_\theta(0) = 0$; $f_\theta(1) = 1$, and second derivative proportional to $1/(e^\theta - 1)$,

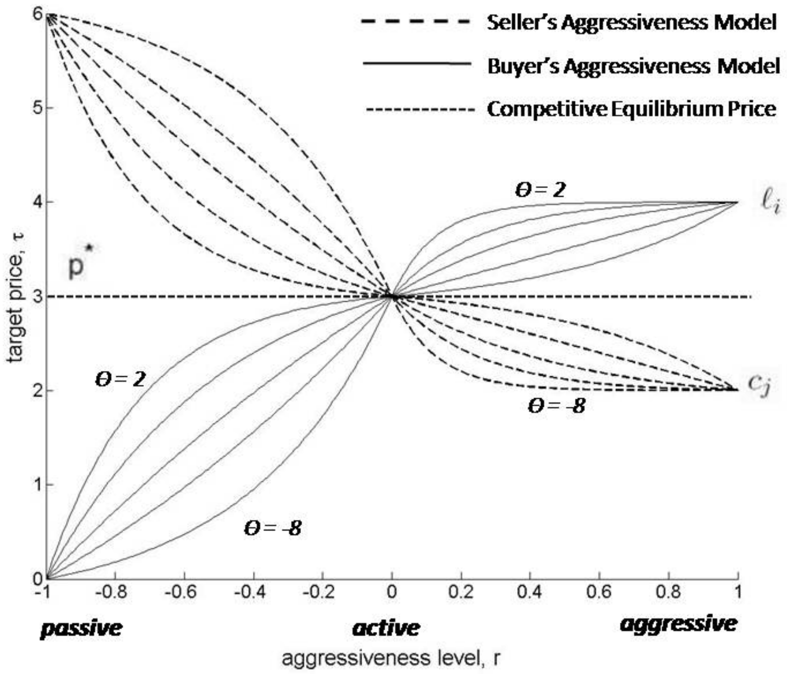

Fig. 5. Aggressiveness for the intra-marginal trader for different θ. Solid lines represent the buyer's function, and the dashed lines the seller's function.

tion space with θ determining the behaviour of the function (that is its rate of change with respect to the degree of aggressiveness r).

Specifically, equations 3 and 4 detail the intra-marginal buyer's and seller's aggressiveness model (and its relationship between r and τ). We adopt these particular functions because they are continuous (and thus, we do not have sudden jumps of τ as r changes) and θ allows the agent to explicitly specify the properties of the function. When θ is high, the magnitude of the gradient tends to 0 at $r = 0$ and increases as θ tends to -1. Conversely, when θ is low, the magnitude of the gradient is high at $r = 0$ and thus allows faster update of the target price as r changes. A slow update is required when the transaction prices are converging to $\hat{p}^*$, while a fast update is required at the beginning of the auction or after a market shock, when market conditions are changing considerably. Indeed, experimental results described in [27] suggest that the effectiveness of our bidding strategy depends on the value θ. In particular, we observed from market simulations that a high θ is more beneficial when the prices are converging towards $\hat{p}^*$ and it is not profitable to deviate too much from $\hat{p}^*$. When faced with a high price volatility (with all agents still exploring the market), an agent is then better off with a low θ to also explore the market by allowing a faster update of its degree of aggressiveness. In Subsection 4.3.2, we describe how updating θ, and thus the aggressiveness model, after every

so that f_θ is concave when θ is negative, convex when θ is positive.

transaction can be beneficial in the long term.

For an intra-marginal buyer,

$$\tau = \begin{cases} \hat{p}^*(1 - \frac{e^{-r\underline{\theta}}-1}{e^{\underline{\theta}}-1}) & \text{if } r \in (-1, 0) \\ \hat{p}^* + (\ell_i - \hat{p}^*)(\frac{e^{r\theta}-1}{e^{\theta}-1}) & \text{if } r \in (0, 1) \end{cases} \tag{3}$$

where $\underline{\theta}$ is calculated such that the function is continuous as $r = 0$, that is there is no jump in the first derivative of τ.

For an intra-marginal seller,

$$\tau = \begin{cases} \hat{p}^* + (MAX - \hat{p}^*)(\frac{e^{-r\underline{\theta}}-1}{e^{\underline{\theta}}-1}) & \text{if } r \in (-1, 0) \\ c_j + (\hat{p}^* - c_j)(1 - (\frac{e^{r\theta}-1}{e^{\theta}-1})) & \text{if } r \in (0, 1) \end{cases} \tag{4}$$

where $\underline{\theta}$ is calculated such that the function is continuous as $r = 0$, that is there is no jump in the first derivative of τ.

Now, for the above equations, the marginal trader is a limiting case, where $\ell_i = \hat{p}^*$ and $c_j = \hat{p}^*$. However, these equations are not valid in the extra-marginal case where the seller cannot ask below $\hat{p}^*$ and the buyer cannot bid above $\hat{p}^*$. In such situations, the extra-marginal buyer and seller modify their aggressiveness functions to that of Figure 6. This reflects the fact that the extra-marginal trader cannot be aggressive and its degree of aggressiveness, r, is clipped at 0 such that it will submit its limit price to maximise its chance of transacting. For this case, Equations 5 and 6 describe that aggressiveness function precisely:

For an extra-marginal buyer,

$$\tau = \begin{cases} \ell_i(1 - \frac{e^{-r\theta}-1}{e^{\theta}-1}) & \text{if } r \in (-1, 0) \\ \ell_i & \text{if } r \in (0, 1) \end{cases} \tag{5}$$

For an extra-marginal seller,

$$\tau = \begin{cases} c_j + (MAX - c_j)(\frac{e^{-r\theta}-1}{e^{\theta}-1}) & \text{if } r \in (-1, 0) \\ c_j & \text{if } r \in (0, 1) \end{cases} \tag{6}$$

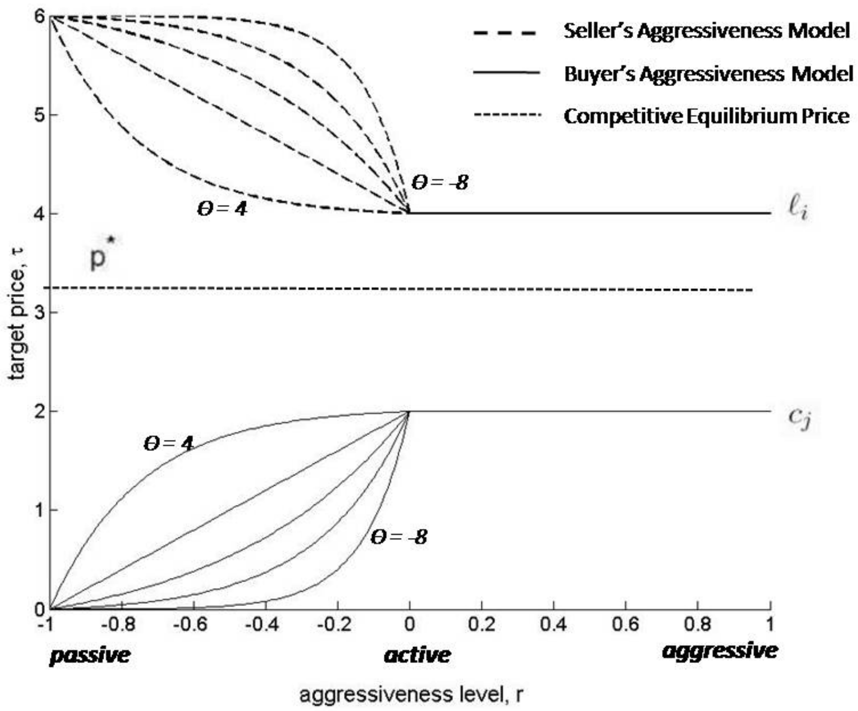

Fig. 6. Aggressiveness for the extra-marginal traders for different θ. Solid lines represent the buyer's function, and the dashed lines the seller's function.

We next look at the adaptive layer of the AA strategy, where the agent learns its degree of aggressiveness and its aggressiveness model.

4.3 The Adaptive Layer

The adaptive layer consists of the short and long-term learning mechanisms that update r and θ respectively. In the following subsections, we describe each of these in more detail.

4.3.1 Short-Term Learning

In the short-term mechanism, the agent uses a set of learning rules (summarised in Figure 7) to update its aggressiveness, every time a bid or an ask is submitted or a transaction occurs in the market. It performs this in order to better fit the prevailing market conditions. Specifically, a simple continuous learning algorithm, the Widrow-Hoff algorithm [30] (initially adopted in the ZIP strategy), is used to increase or decrease the aggressiveness, $r(t)$, at time step t (see Equation 7).

In more detail, the aim here is to adapt the agent's aggressiveness to the current *desired aggressiveness*, $\delta(t)$, which represents the degree of aggressiveness that would allow the buyer to bid the minimum of its limit price and a price slightly higher than the outstanding bid or the seller to ask the maximum of its limit price and a price slightly lower than the outstanding ask. Here, $\delta(t)$ is a factor of r_{shout}, the degree of aggressiveness that would form a price equal to the bid b, if the agent is a buyer and the last event was a bid, or to the ask, a, if the agent is a seller and the last event was an ask, or to the transaction q, whether the agent is a buyer or a seller and if the last event was a transaction (see bidding rules in Figure 7). When the agent is decreasing its degree of aggressiveness (to be more profitable), it sets $\delta(t)$ to slightly lower than r_{shout} (negative λ_r and λ_a) so that the target price is higher than the outstanding bid or lower than the outstanding ask. When it is increasing its degree of aggressiveness (to improve its chance of transacting), it sets $\delta(t)$ to slightly higher than r_{shout} (positive λ_r and λ_a). The algorithm then enacts a continuous-space learning process that backprojects a fraction of the error between the desired degree of aggressiveness, $\delta(t)$, and the degree of aggressiveness, $r(t)$, onto the same degree of aggressiveness $r(t)$. As $r(t)$ updates, it gradually follows the changing $\delta(t)$ at a rate dependent on the learning parameter β_1. A reasonable value of β_1 is chosen [12]. Specifically,

$$\begin{aligned} r(t+1) &= r(t) + \beta_1(\delta(t) - r(t)) \\ \delta(t) &= (1 \pm \lambda_r) r_{shout} \pm \lambda_a \end{aligned} \tag{7}$$

where $\beta_1 \in (0,1)$ is the learning rate of the algorithm which influences the rate of change of $r(t)$ and, hence, of the target price, τ. λ_r and λ_a are the relative and absolute increase or decrease in r_{shout} respectively.

The learning rules employed here are broadly similar to those of the ZIP strategy. We employ its learning mechanism because it has been shown to effectively exploit market information. However, rather than updating a profit margin, we employ the mechanism to update the agent's degree of aggressiveness. We also simplify the adaptive mechanism by not considering a momentum-based update, since the manner in which the aggressiveness is updated with respect to the competitive equilibrium price minimises any high-frequency change in the bid or ask prices. In more detail, when the buyer's target price is greater than the transaction price, this implies that the buyer can transact and so it should try to be more profitable in the next round by being less aggressive. If its target price is less than the transaction price, this suggests that the buyer cannot transact at its target price, and thus should increase it by being more aggressive. Similar intuitions apply for the seller's learning rules. An example of how the level of aggressiveness changes in a specific scenario is given in Figure 9. Here, we observe how the AA trader is generally passive (and r

[12] As we will see in Subsection 5.1.6, λ and β_1 are not sensitive to the performance of the strategy and, thus, to the results we report here.

Learning Rules for Buyer i:

if (transaction occurs at price q)

  if ($\tau \geq q$) *buyer must be less aggressive*

  else *buyer must be more aggressive*

else if (bid, b, submitted)

  if ($\tau \leq b$) *buyer must be more aggressive*

Learning Rules for Seller j:

if (transaction occurs at price q)

  if ($\tau \leq q$) *seller must be less aggressive*

  else *seller must be more aggressive*

else if (ask, a, submitted)

  if ($\tau \geq a$) *seller must be more aggressive*

Fig. 7. Short-Term Learning Rules.

is negative), waiting for a profitable transaction (just below the equilibrium price for the buyer and just above for the seller). Sometimes, the AA traders becomes aggressive (and r is positive) when it is not transacting, which is typically at the beginning of a trading day or following a market shock.

4.3.2 Long-Term Learning

As described in Subsection 4.2, θ influences the bidding behaviour. Given this, we now describe how we can learn such a parameter on a long-term basis, after every transaction, to improve the efficiency of AA. The underlying intuition here is that different values of θ are best within different market conditions and, in particular, the best values of θ depend on the price volatility. Given this, we update θ (after every transaction) through a learning process based on the price volatility, which we measure as an approximation of Smith's α-parameter (see Section 3), given that the agent only has an estimate of the competitive equilibrium price. Equation 8 describes the learning mechanism:

$$\theta(t+1) = \theta(t) + \beta_2(\theta^*(\alpha) - \theta_t)$$

$$\alpha = \frac{\sqrt{\frac{1}{N}\sum_{i=T-N+1}^{T}(p_i - \hat{p}^*)^2}}{\hat{p}^*} \tag{8}$$

where $\beta_2 \in (0,1)$ is the learning rate of the algorithm that determines how θ adapts. In particular, $\theta^*(\alpha)$ is a function (see Figure 8) that determines the desired θ parameter given the current price volatility calculated as Smith's coefficient of convergence α over a window of the N latest prices. p_i is the price of transaction i, and T is the most recent transaction. $\theta^*(\alpha)$ is given by Equation 9 and shown in Figure 8. Based on simulation results for different environments, we chose this particular function as an approximation to the optimal θ parameter that maximises performance given the price volatility. Our function is only an approximation since it is averaged over the optimal θ for a number of different market environments. The exact environment and, thus, the exact optimal θ, are unknown *a priori*. Specifically, :

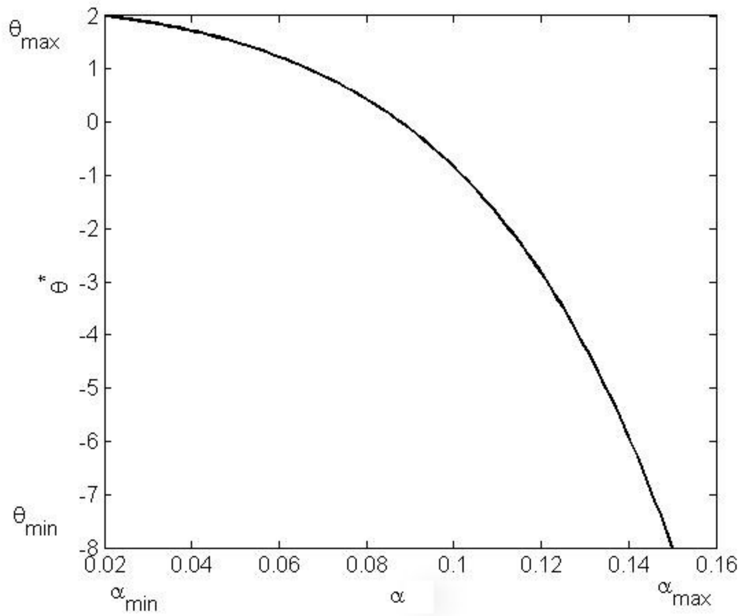

Fig. 8. Function $\theta^*(\alpha)$ gives the desired θ^*.

$$\theta^*(\alpha) = (\theta_{max} - \theta_{min})(1 - (\alpha - \alpha_{min})/(\alpha_{max} - \alpha_{min}) \\ e^{\gamma((\alpha-\alpha_{min})/(\alpha_{max}-\alpha_{min})-1)} + \theta_{min} \tag{9}$$

where $[\theta_{min}, \theta_{max}]$ is the range over which we update θ, α_{max} is the maximum α that occurs in the market, and α_{min} is the minimum α. γ determines the shape of the function. We will show that the choice of γ is not central to this work through a sensitivity analysis (see Subsection 5.1.6).

Given the mechanism, we now consider an example of how θ changes in a specific scenario in Figure 9. In particular, θ updates after every transaction, as specified by Equation 8, it is fixed at a reasonable value of -4 for the first few transaction prices (until a reasonable estimate of the competitive equilibrium price is obtained) and then updated to settle at θ_{max} (around 2) as the transaction prices converge to the competitive equilibrium price (2.65). When a market shock is identified by the sudden increase in α (at Day 11), θ gradually decreases towards θ_{min} (until α starts decreasing) to give a faster update of the target price τ (see Subsection 4.2). As the agents' behaviours gradually adapt to the new market demand and supply and the transaction prices converge towards the new competitive equilibrium price (3.82), θ gradually increases back to a high value (around 2) that is more suitable for a low price volatility.

Furthermore, in Subsection 5.4, and more specifically in Figure 25, we report on the efficiency of our strategy without the learning mechanism to demonstrate the importance of adapting the aggressiveness. In detail, we fixed the aggressiveness level at 0 (that is the AA trader is always active). We observed that the AA strategy performed considerably worse, with the larger drop in performance when there were large market shocks (see Days 8, 19 and 27 in Figure 25). With the aggressiveness fixed at 0, the AA strategy can only adapt through its running estimate of the equilibrium price, which is significantly slower (especially when the equilibrium price changes considerably during large market shocks) than with the learning mechanism. By so doing, we show that adapting the aggressiveness is central to the AA strategy.

Having looked at the aggressiveness model (that outputs τ given r and θ), and the adaptive layer (that updates r and θ), we now need to describe the bidding layer where the agent forms a bid or an ask to submit in the market, based on the current market conditions, its limit price and τ.

4.4 The Bidding Layer

In the bidding layer, the agent employs a set of bidding rules to decide whether or not to submit a bid or an ask, and at what price if it decides to do so. If the buyer's (seller's) limit price is lower (higher) than the current o_{bid} (o_{ask}), it cannot submit any bid (ask), and waits for the beginning of the next round. On the other hand, if the agent can submit a bid or ask in the market, it considers its set of bidding rules to form a price. In this, we identify two cases when an agent bids: during the first trading round, where it cannot estimate the competitive equilibrium price, and the subsequent rounds where it can. In particular, Figure 10 gives the bidding rules, and equations 10 and 11 detail the price formation process in the two cases:

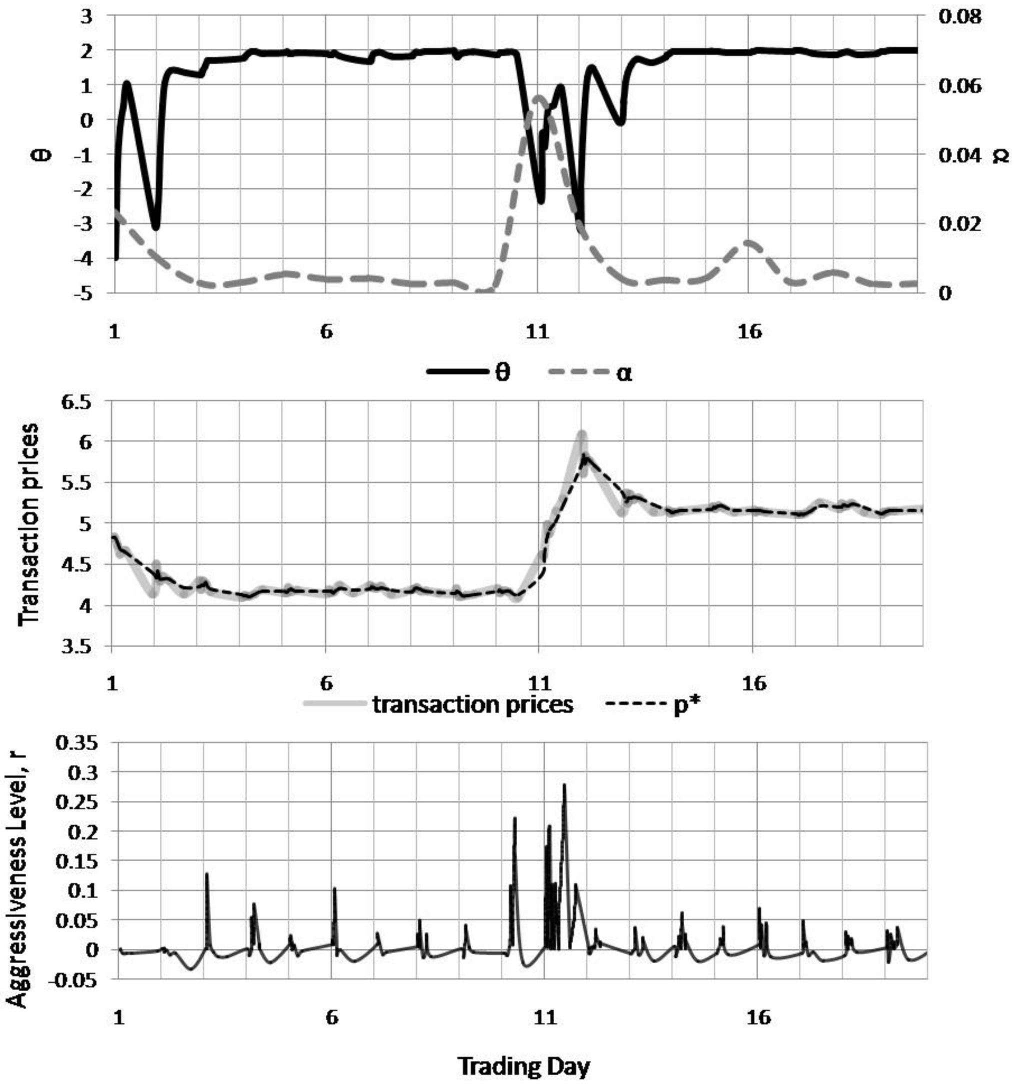

Fig. 9. The history of transaction prices and $\hat{p^*}$ (top plot), the short-term learning of r (middle plot) and the long-term learning of θ (bottom plot). Note that we have a market shock on Day 11, with θ updated to match the change in price volatility.

$$bid_i = \begin{cases} o_{bid} + (\min\{\ell_i, o^+_{ask}\} - o_{bid})/\eta & \text{if first round} \\ o_{bid} + (\tau - o_{bid})/\eta & \text{otherwise} \end{cases} \tag{10}$$

$$ask_j = \begin{cases} o_{ask} - (o_{ask} - \max\{c_j, o^-_{bid}\})/\eta & \text{if first round} \\ o_{ask} - (o_{ask} - \tau)/\eta & \text{otherwise} \end{cases} \tag{11}$$

where $o^+_{ask} = (1 + \lambda_r)o_{ask} + \lambda_a$, $o^-_{bid} = (1 - \lambda_r)o_{bid} - \lambda_a$ and $\eta \in [1, \infty)$ is a

Bidding Rules for Buyer i**:**

if ($\ell_i \leq o_{bid}$) *submit no bid*

else

  if (first trading round) *submit bid given by* Equation 10

  else

    if ($o_{ask} \leq \tau$) *accept* o_{ask}

    else *submit bid given by* Equation 10

Bidding Rules for Seller j**:**

if ($c_j \geq o_{ask}$) *submit no ask*

else

  if (first trading round) *submit ask given by* Equation 11

  else

    if ($o_{bid} \geq \tau$) *accept* o_{bid}

    else *submit ask given by* Equation 11

Fig. 10. Bidding Rules.

constant that determines the rate of increase (decrease) of the bids (asks).

At the beginning of the first trading round, the agent has *no information* other than its limit price. Now, because if the buyer submits too high a bid, it can transact at a not very profitable price (with respect to p^*), it starts with low bids that progressively approach the minimum of its limit price, ℓ_i, and the outstanding ask, o_{ask}, (see Equation 11) to explore the market. Similarly, the seller, j, submits an ask towards the maximum of its cost price, c_j, and the outstanding bid o_{bid} (see Equation 10). Thus, the agent effectively reduces the bid-ask spread with an exponentially decreasing trend (since the bid increase should be decreasing to reflect the decreasing bid-ask spread) determined by η and its limit price. Here, a low η implies a faster rate of convergence of bids or asks until they are matched at a transaction price and, conversely, a high η implies a more conservative bidding approach and a slower convergence. With the latter, while being more profitable if it transacts, the agent risks missing out on a transaction if other agents adopt a more conservative strategy (similar to that of an AA agent with a lower η). However, with a lower η, the agent might be too hasty, and adopting a more conservative approach might be more profitable. In our simulations, we choose a value of 3 for η, which was observed to be a good compromise over a multitude of environments. Furthermore, the buyer can only submit a bid if its limit price is higher than o_{bid}, or otherwise, it remains idle until the beginning of the next round. We use similar intuitions to design the behaviour of the seller.

After the first trading round, the agent has an initial estimate of the competitive equilibrium price, which it subsequently updates after each transaction. Initially, we set the agent's aggressiveness factor, r, to 0 (meaning it adopts an active attitude) because of the lack of market information. Based on the target price, τ, and the set of bidding rules that dictate how the agent should react to the current market conditions, the trader then forms a bid or ask to submit in the market. In more detail, if the target price is higher than the outstanding ask at any time during the bidding process, the buyer accepts the outstanding ask (which is a better offer than it was targeting). Otherwise, it submits a bid, given by Equation 10, that approaches the (changing) target price in a similar manner as in the first trading round. We use similar intuitions to design the seller's bidding rules. Here, if the target price is lower than the outstanding bid, the seller accepts the outstanding bid. Otherwise, it submits an ask given by Equation 11. Furthermore, as in the first trading round, η affects the bidding process in a similar manner and is set to 3 throughout the trading day.

Having described all the components of the AA strategy in detail, we now evaluate it in a number of different environments in the following section.

5 Empirical Evaluation

In this section, we first detail the methodology for analysing the strategic interaction of the AA agent in two different types of populations; homogeneous in which all the agents adopt the same strategy and heterogeneous in which they adopt different strategies. We then proceed to the actual empirical study of the strategies in these two cases. Finally, we apply our strategy in a market based on real market data, where the demand and supply constantly changes. The purpose of this final exercise is to evaluate our strategy within a systematically dynamic environment, rather than with a single market shock as per the first part of this evaluation.

5.1 The Methodology

There are two main parts to the methodology for benchmarking a strategy for the CDA: the market setup and the empirical evaluation of the strategy or strategies adopted in the market. First, we describe the market setup, invariant of the strategies adopted. Second, we describe how we analyse the strategic interactions of these agents (given the adopted strategies in the market) and give the metrics we use to analyse the performance of a strategy in the market, in both homogeneous and heterogeneous populations. Next, we discuss how we ensure statistical significance in our empirical study. Finally, we analyse the sensitivity of the parameters used in AA strategy using a standard sensitivity analysis.

5.1.1 The Market Setup

In all of our simulations, the market is populated by a set of 10 buyers and 10 sellers. In particular, we look at different experiments with markets M1, M2, and M3, and market shocks MS14, MS21, MS31 and MS23 (see Section 2 for terminology) and, finally, the real dynamic GOOG market (see Section 2). In the static environment, we only look at these three markets as we would have the same behaviour in M4 as we would in M1 since there is only an upward shift in the demand and supply and the absolute differences between the agents' preferences remain the same. In the dynamic environment, on the other hand, we are mostly interested in how the strategies adapt from their best behaviour in one market to their best behaviour in the new market. Now, if we have more than one market shock, as in scenario MS214, we would observe how the strategy adapts from M2 to M1 and finally to M4. Now, we would identify the same behaviour in MS21 and MS14 as the agent's behaviour at the end of MS21 would be the same as just before the shock in MS14. While we have a transitive property, we do not have a reflective one. Thus, the observations from MS21 and MS12 would be different. Indeed, in the former, we would observe how the agent adapts from the flat supply of M2 to the normal supply of M1, while in the latter, we would observe how it adapts from a normal to a flat supply. However, due to space limitations, we only analyse in detail a subset of the single market shocks, and generally look at how the agent adapts from an extreme to a normal demand or supply, or to a change in the competitive equilibrium price. A brief analysis of the remaining cases is given in Appendix A. Finally, for the reasons outlined earlier, we also consider real market data. Specifically, we look at the real market (GOOG) where the equilibrium price changes every trading day (see Page 14).

In most of our experiments, each buyer and seller is endowed with a single unit to buy or sell to induce the market demand and supply found in M1, M2, M3, M4 or the GOOG market. Now, given markets with multi-unit allocation are already intrinsically very efficient, because of the size of the market [13], the challenge in this area is to improve the efficiency of the smaller, less efficient single-unit allocation market. Thus, we focus on single-unit allocation experiments in this paper (as do almost all the previous work on the design of CDA strategies). Nevertheless, we do provide a multi-unit allocation experiment in our GOOG market to ensure that AA still performs efficiently in such a dynamic setting.

Now, in all our experiments, each agent is endowed with a limit price corresponding to a unit of good to buy or sell. For the static scenario, the CDA lasts 10 days. For the dynamic scenario, it lasts 20 days with the market demand and supply kept constant during the first 10 days and changed thereafter, effectively inducing a market shock on Day 11 (see Figure 2 for an example). For the homogeneous scenario, we can avoid redundancy in our experiments when evaluating the strategies within a static environment for markets M1, M2 and M3, by looking at the performance of the strategies before the market shocks for the dynamic cases. Given the market setup, we consider a statistically significant number of runs of the CDA (see Subsection 5.1.5 for more details), each lasting 10 or 20 trading days.

5.1.2 The Agent Setup

We now look at the agent setup. For the setup of the AA agents, based on simulations for a wide range of demand and supply, we set the size of window of transactions (N) over which we calculate $\hat{p^*}$ to 5 (see Equation 2), the parameter η in the bidding layer (see equations 10 and 11) to 3, and λ_a and λ_r (see Equation 7) to 0.01 and 0.02 respectively. The learning rates β_1 and β_2 (see equations 7 and 8) are drawn from a uniform distribution $\mathcal{U}(0.2, 0.6)$ while γ is set to 2. Thus, by considering the performance of AA for many different demand and supply situations, the fixing of these parameters means they are not fine-tuned for any single market. Thus, in any given situation, superior performance could be obtained by optimising these parameter choices to the prevailing situation. To emphasise this point still further, Subsection 5.1.6 investigates the sensitivity of the performance to the choice of these parameters.

For GDX agents, the discount factor is set to 0.9 based on Tesauro and Bredin's simulations (see [21] for more details) and, finally, the ZIP agents are initialised with the set of parameters evolved in [4].

[13] The efficiency of the CDA using ZI agents is only attributed to the efficiency of the market structure, rather than its behaviour. Given ZI agents, we evaluated the efficiency of a ten-unit allocation at 99.3%, while that of a single-unit allocation market was 95.7%.

Given the market and agent setup, we now study how the choice of strategy determines an agent's efficiency in different environments. We do so for both homogeneous and heterogeneous populations.

5.1.3 Analysing Efficiency in Homogeneous Populations

As discussed in Section 3, we believe that analysing the daily efficiency of the strategy provides more insights into how effective a strategy is in learning from market interactions (as we can observe how the efficiency of the strategy is changing as it is learning its best behaviour). This means our evaluation methodology does not favour the AA strategy, which was designed to perform differently and better in the 'first round' than the benchmarks (because we do not aggregate this advantage in a single average efficiency measure). Now, as the agent learns to be more competitive in a static market and the transaction prices converge towards the competitive equilibrium price, we expect its efficiency to improve. Thus, we calculate the efficiency of a strategy at the end of each trading day, as well as the market volatility, α, calculated as Smith's α-parameter over all the transactions during each day (see Equation 1), which describes how the transaction prices converge at the competitive equilibrium price.

5.1.4 Analysing Efficiency in Heterogeneous Populations

Next, we describe our methodology for benchmarking strategies in heterogeneous environments. As in the homogeneous case, we still need to determine performance of a strategy by its efficiency. However, unlike the homogeneous case, agents do not all adopt the same strategy, and different numbers of buyers and sellers can adopt different strategies. Therefore, for the reasons outlined in Subsection 3.2, the one-population EGT model is insufficient and, thus, we adopt the *two-population* EGT model we have previously developed to analyse buyers' and sellers' strategic interactions [25]. In more detail, this model assumes that buyers and sellers will adopt the buyer and seller strategies that are the most efficient for them in the market. Thus, we calculate how the proportions of buyers and sellers adopting the different strategies change. In our two-population EGT analysis, the first step is to calculate the payoff table [14] with the payoffs to each buyer and seller strategy given the exhaustive set [15] (of size 121) of the 10 buyers and 10 sellers adopting the different

[14] A more detailed description on computing the heuristic payoff table is given in [25] (see Chapter 6), with an example of the heuristic payoff table (see Appendix A).

[15] The size of the payoff table is given by $\binom{A_b+S_b-1}{A_b} \times \binom{A_s+S_s-1}{A_s}$ where S_b is the number of buyer strategies, S_s is the number of seller strategies, A_b is the number of buyers and A_s is the number of sellers (see [25] for more details).

buyer and seller strategies. Given the payoff table, we can then calculate the *replicator dynamics* that describe how the population mix of agents adopting the different buyer and seller strategies (if these strategies are more efficient) changes, and the *mixed-Nash equilibrium* of the CDA game that describes the particular proportion mix where it does not pay for any agent to adopt another strategy. Together, the replicator dynamics and the mixed-Nash equilibrium allow us to analyse the performance of the different strategies over *all* population mixes (and not only in a balanced population) and they also indicate how the population mixes will change as buyers and sellers adopt the more efficient of the strategies at the different mixes. Specifically, the following equations describe how we calculate the buyer's dynamics $\dot{p_h}$, and the seller's dynamics $\dot{q_k}$:

$$\begin{aligned} \dot{p_h} &= \left[u_b(e^h, p, q) - u_b(p, p, q)\right] p_h \\ \dot{q_k} &= \left[u_s(e^k, q, p) - u_s(q, q, p)\right] q_k \end{aligned} \tag{12}$$

where p denotes the buyer's mixed strategy and p_h is the probability that the buyer adopts pure buyer strategy h. q denotes the seller's mixed strategy and q_k is the probability that the buyer adopts pure buyer strategy k. $\dot{p} = \{\dot{p}_1, \dot{p}_2, \ldots\}$ is the replicator dynamics along the buyer axis and $\dot{q} = \{\dot{q}_1, \dot{q}_2, \ldots\}$ the dynamics along the seller axis.

Furthermore, the mixed-Nash equilibrium of the analysis is given as the (p, q) that minimises Equation 13:

$$\begin{aligned} v(p,q) = &\sum_{h=1}^{S_b}(\max\left[u_b(e^h, p, q) - u_b(p, p, q), 0\right] p_h)^2 + \\ &\sum_{k=1}^{S_s}(\max\left[u_s(e^k, q, p) - u_s(q, q, p), 0\right] q_k)^2 \end{aligned} \tag{13}$$

Other definitions that are relevant to the EGT analysis include:

Definition 16 *A* ***trajectory*** *is the change in mixed strategy, starting from a particular mixed strategy, and following the replicator dynamics.*

Definition 17 *An* ***attractor*** *is a mixed-Nash equilibrium towards which the replicator dynamics (trajectories) converge.*

Definition 18 *A* ***saddle point*** *is a mixed-Nash equilibrium from which replicator dynamics (trajectories) diverge.*

Definition 19 *A* ***basin of attraction*** *of a mixed-Nash equilibrium is the* *space of mixed strategies from which trajectories will converge to that equilibrium.*

Now, while the two-population EGT model can be used to analyse the interaction of any number of buyer and seller strategies, a visual representation of the analysis is possible for up to two different buyer and two different seller strategies since an analysis of three buyer and three seller strategies would be in four dimensions. Though Walsh *et al.*'s model would allow a visual analysis with up to three strategies (assuming that these strategies can be used for both buying and selling), it is inaccurate as outlined in Subsection 3.2. For these reason and, as is common practice in related work on the bidding strategies for the CDA [22, 21, 27], we compare the performance of two strategies at a time.

5.1.5 Statistical Significance

In all cases, we consider a statistically significant number of runs, 2500, of the CDA. We validate our results at the 95%-confidence-interval by running the non-parametric Wilcoxon rank sum test [14] on the daily efficiency of the strategies and on the difference between the actual and the expected payoff ($\left[u_b(e^h, p, q) - u_b(p, p, q)\right]$) in heterogeneous populations. 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. Finally, we provide error bars at the 95%-confidence interval in the daily efficiency of strategies within homogeneous populations (as shown in Figure 12).

5.1.6 Sensitivity Analysis

We also investigate how sensitive our parameters are using a standard one-at-a-time sensitivity analysis [19] (where we separately vary the different parameters). In particular, we do a sensitivity analysis in two different markets (see Figure 11), namely the dynamic GOOG market (with a changing equilibrium price) and the static M1 market (with a static expected equilibrium price). In the former case, we observe that N and η are the only sensitive parameters, while in the latter case only η is sensitive. Thus, depending on the type of market we are considering, η and N have to be carefully chosen to maximise the efficiency. Having said that, we also observed that suitable values for η and N in both static and dynamic markets can be identified (which we set to 3.0 and 4 respectively here).

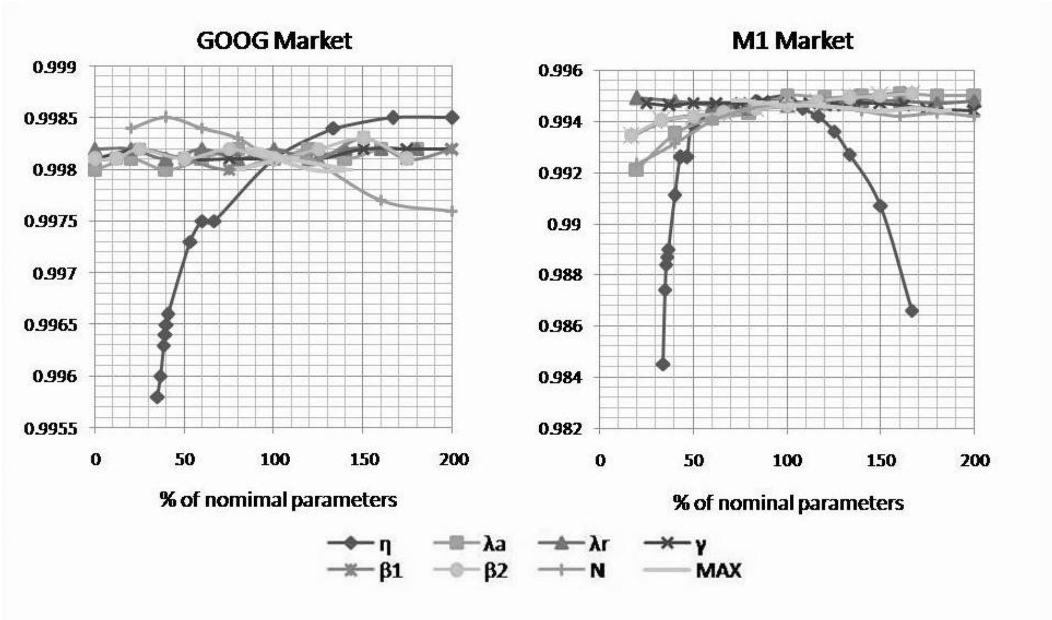

Fig. 11. Sensitivity analysis of parameters in the static M1 market on the left and in the dynamic GOOG market on the right.

5.2 Performance in Homogeneous Environments

First, we consider the homogeneous scenario. In Figures 12 to 15, we look at the performance of the AA strategy and the benchmarks GDX and ZIP in the different markets highlighted earlier, and Table 1 details the efficiency of the buyers and the sellers, and the efficiency of all the agents in these markets. Note that apart from the symmetric Market M1, buyers do not expect the same profit as sellers due to the asymmetric nature of the demand and supply, and, thus, the efficiency of the strategy is not the average of the buyers' and sellers' efficiency. By dissecting the efficiency of the buyers and the sellers separately, we can observe whether the buyers or the sellers are performing better given the particular demand and supply.

5.2.1 The Static Scenario

We first analyse the efficiency of the strategies within a static environment, with markets M1, M2 and M3. In M1 (see Day 1 to 10 in Figure 12), we can see that our strategy outperforms both benchmarks on every trading day, with an average efficiency of 0.997. We also note that with AA agents, the transaction prices converge faster (with a lower α) and, on average, remain closer to p^* than with GDX or ZIP agents (AA has the smallest α on Day 10). On the first day, we observe that AA has the highest efficiency because the AA agents assume that there is no information on the first round and adopt a conservative approach (submit bid and ask with a slowing increasing

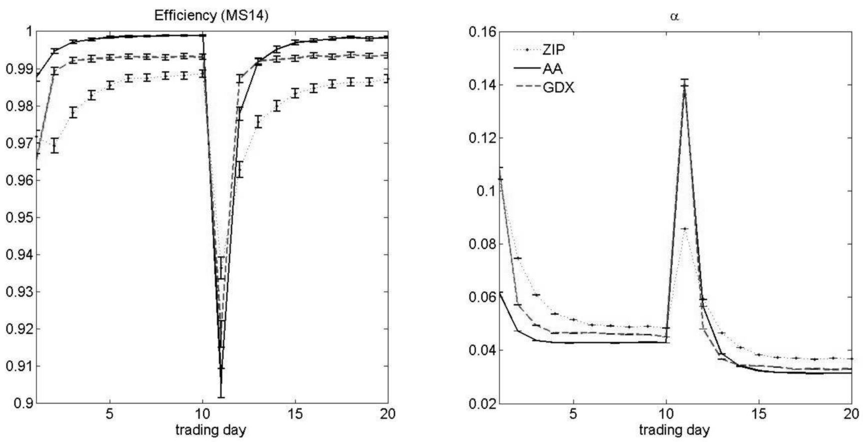

Fig. 12. Scenario MS14. The market efficiency of AA is 0.992, of ZIP, 0.979 and of GDX, 0.988. If we consider the static scenario for Market M1, the market efficiency of AA is 0.997, of ZIP 0.982, and of GDX 0.990.

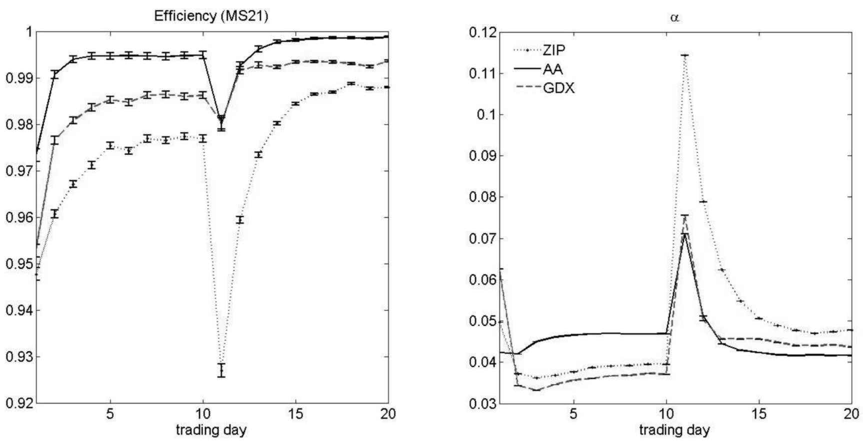

Fig. 13. Scenario MS21. The market efficiency of AA is 0.993, of ZIP 0.974, and of GDX 0.987. If we consider the static scenario for Market M2, the market efficiency of AA is 0.992, of ZIP 0.971, and of GDX, 0.981.

trend) and they have a faster update of their target price. ZIP makes no such assumption and starts with a random profit margin, while GDX suffers from the lack of information (bids, asks and transaction prices). After a few days, the efficiency of all three strategies converges to some value, which is highest with AA agents, and lowest with ZIP. This validates our market setup of 10 days for each market, since we can observe that even if we would consider a larger number of days, the efficiency of the subsequent days would not change. Furthermore, it also validates our analytical method to look at daily efficiency, since we can observe that the efficiency is different on different days for different strategies. Moreover, the daily efficiency converges to different maxima for each strategy, suggesting that the AA strategy is best at learning to be more efficient in the market. With the traditional analytical method (as detailed in Subsection 3), we would only calculate the average efficiency over all the

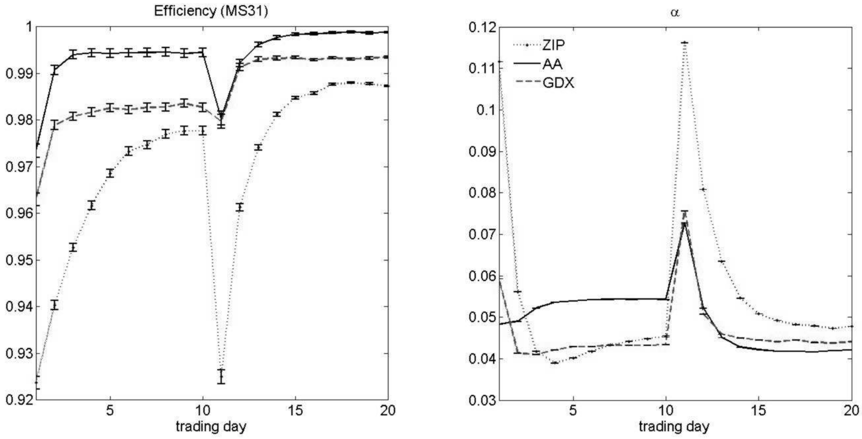

Fig. 14. Scenario MS31. The market efficiency of AA is 0.996, of ZIP 0.968, and of GDX 0.987. If we consider the static scenario for Market M3, the market efficiency of AA is 0.996, of ZIP 0.960, and of GDX 0.981.

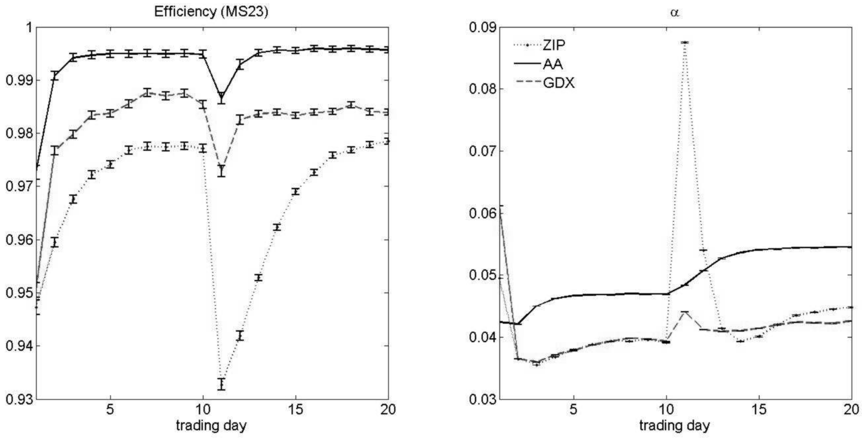

Fig. 15. Scenario MS23. The market efficiency of AA is 0.994, of ZIP 0.968, and of GDX 0.980. As in MS21, when we consider the static scenario in the first 10 days, the market efficiency of AA is 0.992, of ZIP 0.971, and of GDX, 0.981.

trading days and would not observe the fact that efficiency is capped after a few trading days to a maximum, while strategies like AA and GDX learn to be efficient at a much faster rate than ZIP.

In markets M2 and M3 (see days 1 to 10 in Figures 13 and 14), we also observe that AA is the most efficient (99.7% in M1, 99.1% in M2 and 99.6% in M3). In particular, it does much better than the other strategies in asymmetric markets than it does in the symmetric Market M1 (around 2.1% better in asymmetric cases, compared to 1.1% better in the symmetric case – see Table 1). This is because the competitive equilibrium price does not change significantly and, thus, the target price remains close to p^* on Day 11. The competitive equilibrium price still changes, as we are dealing with uniformly distributed limit prices, and the non-deterministic demand and supply is ex-

Table 1
Efficiency of strategies in homogeneous environments (over all trading days).

	AA			**GDX**			**ZIP**		
Scenario	buyer	seller	all	buyer	seller	all	buyer	seller	all
M1	0.969	1.025	0.997	0.981	0.998	0.990	1.010	0.960	0.982
M2	1.212	0.459	0.992	1.145	0.708	0.981	1.143	0.660	0.971
M3	0.389	1.247	0.996	0.595	1.161	0.981	0.896	1.069	0.960
MS14	1.044	0.981	0.992	1.022	0.989	0.988	1.033	0.980	0.979
MS21	1.088	0.754	0.993	1.054	0.876	0.987	1.085	0.817	0.974
MS31	0.667	1.159	0.996	0.778	1.103	0.987	0.932	1.049	0.968
MS23	0.968	1.250	0.994	0.894	1.228	0.980	1.015	0.863	0.968

pected to be as in M1 to M4. However, in these asymmetric markets, the α-parameter of AA is the highest, while being lowest in the symmetric market. We explain this difference by separately considering the buyers' and the sellers' efficiencies (see Table 1). In Market M2 (with a flat supply curve), the fact that the buyers' efficiency is higher than the sellers' means that the transaction prices are, on average, less than p^*, with the buyers having more profitable transactions. This, in turn, indicates that the buyers are more successful at driving the market price (i.e. forcing transaction prices to be lower and be more profitable from their perspective) when the supply is flat. We make similar observations with Market M3 which has a flat demand curve. While α is still highest for AA, the AA sellers' efficiency is higher than the AA buyers', indicating that sellers are driving the market price to be higher than p^*, and are being more profitable from their perspective. As with M1, the daily efficiency converges with all three strategies, with AA still having the highest efficiency on Day 10.

5.2.2 The Dynamic Scenario

We now analyse the daily efficiency of strategies when faced with market shocks. At the beginning of Day 11, the strategies are all tailored to perform best in Day 10. Now, with a market shock, the conditions to which those strategies have adapted are different, forcing those agents to relearn the best strategic behaviour in the market. Essentially, a robust strategy should be able to rapidly adapt to the new market conditions, since the longer it takes to do this, the more inefficient it is.

In scenario MS14 (see Figure 12), the market demand and supply structure remains the same, with an increase in the competitive equilibrium price p^*. In this case, on average, AA still outperforms the benchmarks with an efficiency of 0.992. Because there is a significant shift of p^*, we observe a significant decrease in the efficiency of the AA strategy, as p^* has to be re-estimated gradually (as transaction prices diverge from the old equilibrium and converge to the new one). However, with the higher α, and thus a lower θ, the AA target price changes at a faster rate than it would with a fixed θ (see Subsection 4.2), forcing transaction prices to converge at a faster rate to the new p^*. Here, we also observe that the efficiency of the benchmarks, GDX and ZIP on Day 11, is only slightly better than AA, though the latter's efficiency improves after a few trading days to be better than the benchmarks. This can be explained by the fact that $\hat{p}^*$ is a fundamental parameter of the AA strategy, such that a significant change in p^* affects its performance. Furthermore, AA and GDX have the highest α because p^* is central to the AA's aggressiveness model, and because GDX's belief function approximates a step function at p^*. On the other hand, ZIP does not consider p^* explicitly when it forms a bid or an ask. In fact, it only considers its latest profit margin on Day 10 when starting to bid (with a new limit price given the market shock) on Day 11.

In scenario MS21 (see Figure 13) where p^* does not change significantly, we initially have a flat supply followed by a symmetric demand and supply. Again, AA performs best with the highest average efficiency and it is the most efficient strategy with the fastest adaptivity to the new market conditions (with the lowest α). Indeed, GDX and AA have the lowest α, which is considerably smaller than in scenario MS14 where p^* changes significantly. ZIP suffers the most from a market shock, with a significant drop in efficiency and slow adaptability. This is because ZIP reuses the same profit margin at the beginning of the following day, and given the significant change in preferences (limit prices) after a market shock, its profit margin is no longer tailored for the new market, and the decrease in efficiency then depends on how different the preferences in the two consecutive markets are. Thus, the decrease is considerable as we are looking at an extreme change for the sellers' preferences. We observe similar behaviour for the three strategies in scenario MS31, with AA outperforming the other strategies.

Furthermore, AA outperformed the benchmarks with the best margin in scenario MS23 (see Figure 15), where the market goes from a flat supply to a flat demand. ZIP suffers considerably here as the profit margin, which had been tailored to Market M2, is now used in Market M3 at the beginning of Day 11. With the supply curve now ranging from 1.5 to 4.5 (rather than between 2.8 and 3.2), the same set of sellers' profit margins gives a wider range of asks that are no longer profitable in the market. On the other hand, GDX and AA do not suffer such a drastic change in α as ZIP does. Indeed, we observe that the magnitude of the peak in α on Day 11 for GDX and AA in MS14 is about twice that in MS21 and MS31 where we change either the demand or the supply curve, while there is no peak when we change both the demand and supply curves in MS23. We explain this by considering the limit prices of buyers and sellers. Indeed, in MS14, even though the market demand and supply remain the same, the buyers' and sellers' individual preferences change drastically, with the extreme case where extra-marginal traders become intra-marginal and intra-marginal traders become extra-marginal. In MS21 with a flat supply (MS31 with a flat demand) the change in sellers' (buyers') preferences is not as significant as buyers' (sellers'). Since market behaviour is affected by both buyers' and sellers' behaviours, the change in preferences is then reflected in the change of market efficiency and α. Thus, in MS23 with no extreme changes in demand and supply observed in MS14, in demand in MS21 and in supply in MS31, the drop in efficiency for GDX and AA is even smaller, with no peak in α on Day 11.

5.3 Performance in Heterogeneous Environments

Here, we look at the EGT plots for different scenarios in static and dynamics environments, given the methodology outlined in Subsection 5.1.4. Now, because we consider only pairwise comparisons, a visual representation of the above equations is possible, and an example is given in Figure 16, where we have two sub-plots. The left sub-plot gives the replicator dynamics of the analysis, with the vertices corresponding to different pure strategies, and its shading denotes the magnitude of the dynamics 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 and then, it does not pay off to deviate to another buyer or seller strategy. Note that trajectories (see Definition 16) can converge towards or diverge from a darker region, depending on whether that region contains an attractor or a saddle point. Finally, the right sub-plot gives the magnitude of the buyer's and seller's dynamics, with a mixed-Nash equilibrium occurring when the magnitude of both dynamics is 0. In particular, we consider these magnitudes to compare the buyer's and seller's payoffs when deviating to the more efficient strategy, which is particularly insightful when we have an asymmetric demand and supply, as we will now see.

5.3.1 The Static Scenario

First, we evaluate the strategies in a static scenario with no market shock, and in turn consider populations with AA against ZIP, and AA against GDX. In Market M1 with AA and ZIP agents, we have a single mixed-Nash equilibrium at AA (see Figure 16), implying that all buyers and sellers adopt the *dominant*

AA strategy. We also observe that the dynamics have comparable magnitudes and that the magnitude of buyer's dynamics is higher when AA sellers are in the majority, and that the seller's dynamics are higher when AA buyers are in the majority (with higher magnitude here implying faster convergence to AA). Thus, here, AA agents are most efficient when they are in the majority. Similarly, with AA and GDX agents in M1, we have a single *attractor* (mixed-Nash equilibrium towards which trajectories converge) at A and *saddle points* (mixed-Nash equilibrium that trajectories diverge away from) as can be seen in Figure 17, with the majority of buyers and sellers eventually adopting the AA strategy (and only 4% of buyers and 21% of sellers adopting GDX). Here, the magnitude of convergence to A is highest when there are a majority of GDX buyers and sellers, implying that AA buyers and sellers are most efficient when they are in the minority. We also observe that buyers and sellers do not necessarily select the same buyer and seller strategy respectively (e.g. when AA buyers and ZIP sellers are in the majority, buyers tend to deviate to ZIP, and sellers to AA), which would not have been identified with the traditional one-population model.

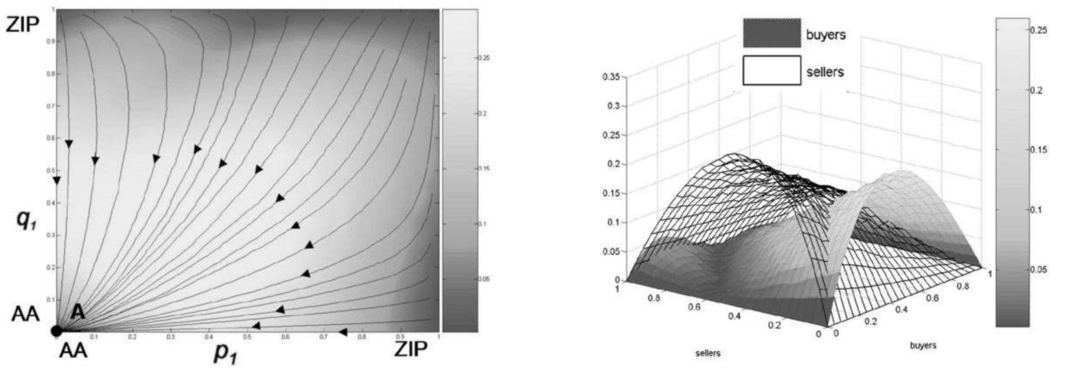

Fig. 16. Scenario M1 with AA and ZIP agents. Here, we have a single dominant strategy at (0,0). The magnitudes of the buyer's and seller's dynamics are of comparable magnitude.

Next, we look at Market M2 with a flat supply. With a population of AA and ZIP (see Figure 18), we have a single dominant strategy, A, with all buyers and sellers adopting AA. The obvious observation here is that the magnitude of the seller's dynamics is considerably smaller than that of the buyer's. This suggests that there is more economic incentive for buyers to adopt AA than for sellers to do so. This happens because of the market's flat supply, meaning the sellers' have considerably lower expected profits than buyers, and, thus, gain less in profit when deviating to another seller strategy (in contrast with the buyer case). Furthermore, we observe that when the majority of buyers adopt ZIP, the sellers tend to adopt ZIP, and when the majority of buyers adopt AA, the sellers tend to adopt AA.

Now, with a population of AA and GDX in M2 (see Figure 19), we have two

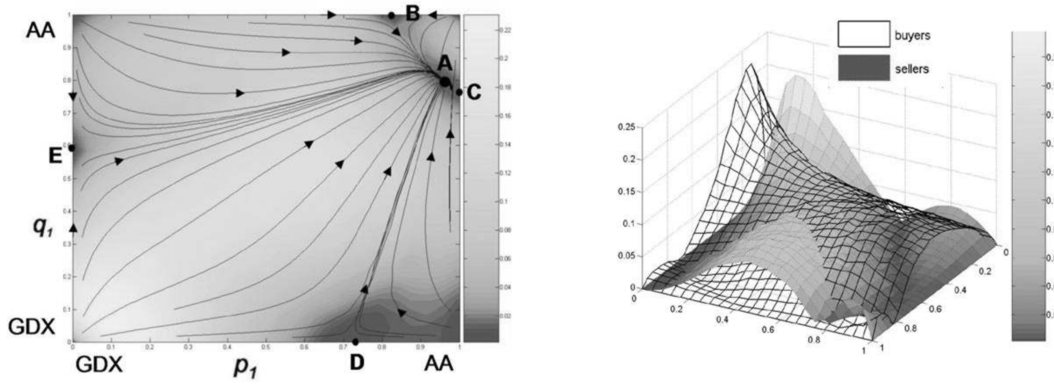

Fig. 17. Scenario M1 with AA and GDX agents. The replicators converge towards the single mixed-Nash equilibrium A at (0.96,0.79). Thus, buyers and sellers are more likely to adopt the AA strategy, with a relatively small proportion adopting the GDX strategy. The magnitudes of the buyer's and seller's dynamics are comparable.

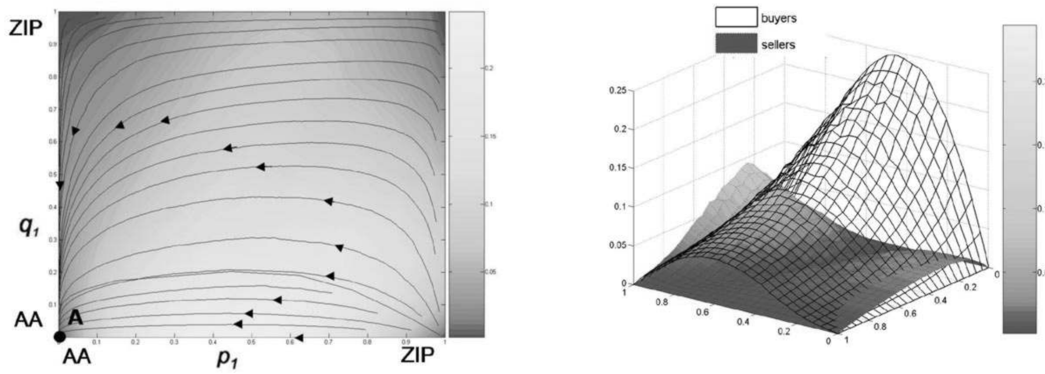

Fig. 18. Scenario M2 with AA and ZIP agents. Here, we have a dominant strategy at (0,0). All buyers and sellers eventually adopt the AA strategy. The magnitude of the seller's dynamics is considerable smaller than that of the buyer's.

equilibria, A at $(0.92, 0)$ and B at $(1, 1)$. Because the basin of attraction of A is considerably larger than that of B, there is an equally larger probability (0.884 compared to 0.116) that the mixed-Nash strategy A will be adopted (and all agents will eventually select AA). Thus, there is still a small probability of 0.116 that 8.0% of buyers will adopt GDX and all sellers will adopt GDX, such that AA is not dominant. When we consider the magnitude of the dynamics, we observe that the sellers' magnitude is considerably smaller than the buyers', and we explain this with the same intuition as with AA and ZIP in M2. Furthermore, when GDX buyers are in the majority, the sellers are more inclined to adopt GDX, and when AA buyers are in the majority, sellers tend to adopt AA, though if GDX sellers are in the majority, then sellers are likely to adopt GDX.

Now, because of the reflective nature of M2 and M3, we only report on our analysis of the strategic performance in M2. However, we observe reflective be-

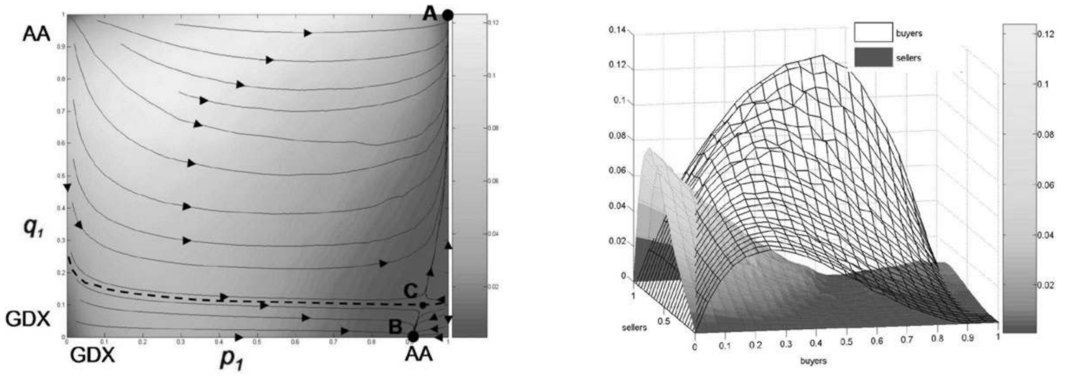

Fig. 19. Scenario M2 with AA and GDX agents. Here, we have two attractors: A at (1,1) and B at (0.92,0), and a saddle point C at (0.93,0.10). The area of the basin of attraction for A is 0.884, and for B is 0.116. The magnitude of the seller's dynamics is higher when GDX buyers are in majority, and considerably lower as AA buyers are represented more.

haviours in M3 (see Appendix A), with the magnitude of the buyers' dynamics being considerably smaller than the sellers' in this case.

5.3.2 The Dynamic Scenario

We now turn to the performance of the strategies in dynamic environments with market shocks. In particular, we look at scenarios MS14 and MS21, and provide further results for MS31 and MS23 in the appendix (which further validate our claim that AA is better than both ZIP and GDX).

In scenario MS14 with AA and ZIP strategies (see Figure 20), we have two attractors at A and B, and a saddle point at C. The basin of attraction of A is considerably larger than that of B, with the higher probability of 0.978 that all the buyers and sellers will eventually adopt the AA strategy. As in M1, the magnitude of the buyer's and seller's dynamics is highest when AA agents are in the majority, which again suggests that AA is most efficient when it is in the majority. Furthermore, as in M1, we observe that the magnitude of the buyer's and the seller's dynamics are comparable, and this is because we are still dealing with symmetric markets where buyers and sellers expect similar payoffs. However, unlike in M1, AA is no longer dominant, and there is now a small probability of 0.022 that ZIP will be eventually adopted in the market. Similarly, with AA against GDX in MS14 (see Figure 21), we have two attractors at A and B, and a saddle point at C, and the basin of attraction is much larger for attractor A. As with AA against ZIP, the market shock causes AA to no longer be dominant, and there is now a small probability of 0.065 that GDX will be eventually adopted in the market.

In scenario MS21 (see figures 22 and 23), where the supply changes, we observe

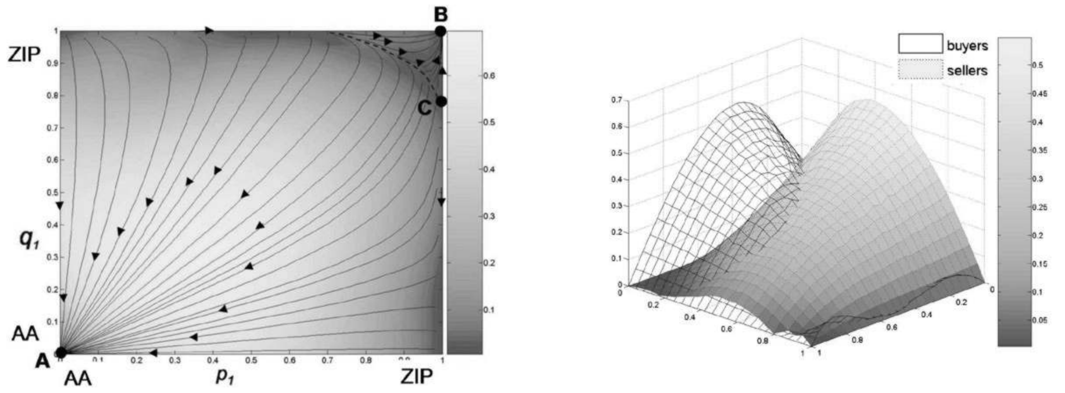

Fig. 20. Scenario MS14 with AA and ZIP agents. Here, we have two attractors: A at (0,0), B at (1,1) and a saddle point, C at (1,0.78). The area of the basin of attraction of A is 0.978, and that of B is 0.022.

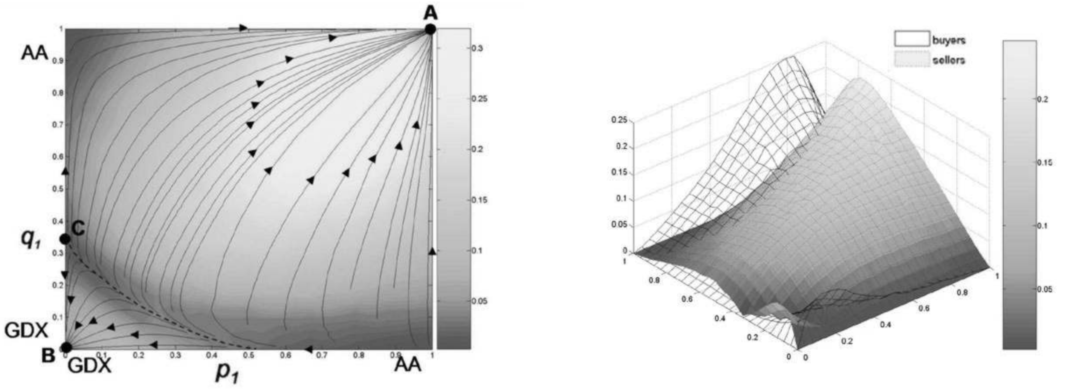

Fig. 21. Scenario MS14 with AA and GDX agents. Here, we have two attractors: A at (1,1), B at (0,0) and a saddle point, C at (0.0.34). The area of the basin of attraction for A is 0.935 and that of B is 0.065.

a similar set of attractors as in MS14 (but with a probability of 0.961 that AA will be adopted against ZIP, and a probability of 0.869 that it will be adopted against GDX). However, the dynamics of how these equilibria are reached differ, with sellers having a slight tendency to adopt more ZIP or GDX than in MS14 when AA buyers are in the minority. In that case, the magnitude of the seller's replicator dynamics is higher than that of the buyer's (because of the asymmetric demand and supply, and sellers expect higher profits than buyers) and thus influence more the dynamics of the CDA. As the AA buyer strategy becomes increasingly popular, the buyer's dynamics have increasingly more weight and increasingly influence the dynamics of the market. In some cases (in the basin of attraction of equilibrium B), the change in dynamic is not sufficiently in favour of AA buyers, and the GDX and ZIP buyers then take over.

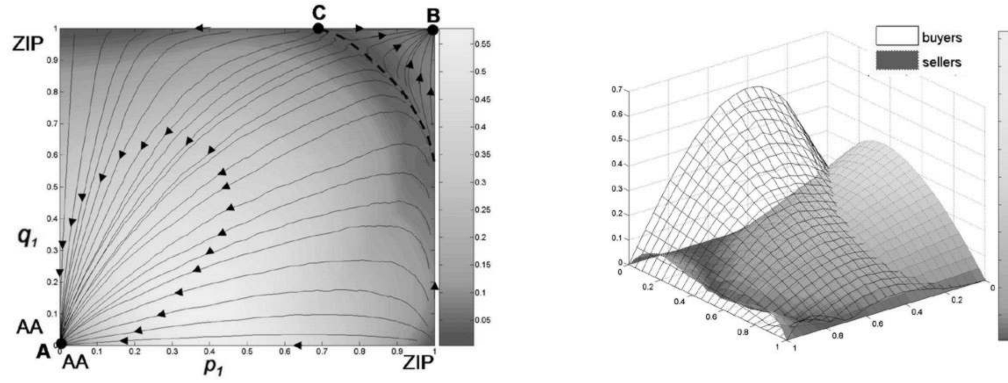

Fig. 22. Scenario MS21 with AA and ZIP agents. Here, we have two attractors: A at (0,0), B at (1,1) and a saddle point, C at (0.68,1). The area of the basin of attraction of A is 0.961, and that of B is 0.039.

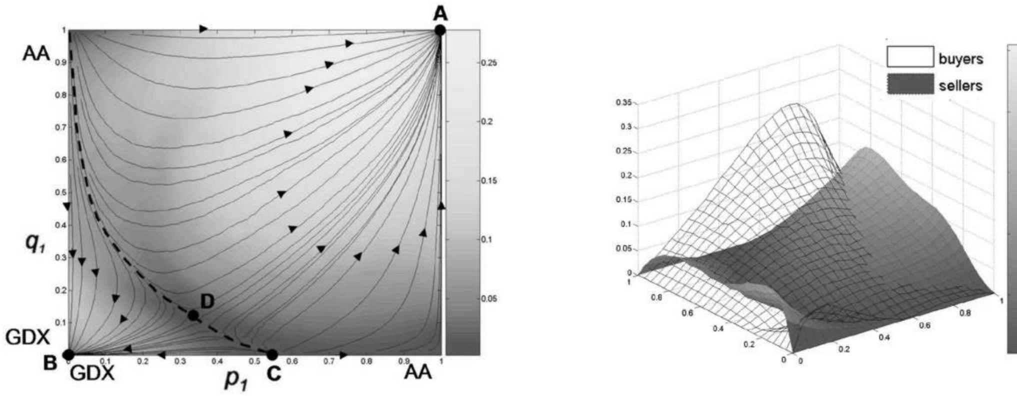

Fig. 23. Scenario MS21 with AA and GDX agents. Here, we have two attractors: A at (1,1) and B (0,0), and two saddle points: C at (0.55,0) and D at (0.33,0.11). The area of the basin of attraction of A is 0.869, and that of B is 0.136.

5.4 Performance in Real Markets

For the reasons outlined at the end of Section 2, we now analyse the strategies in the GOOG and the GOOGshock market. The aim of this exercise is to observe how the efficiency of the different strategies changes given a demand and supply that changes on a daily basis for single-allocation and multi-allocation scenarios.

As can be seen in Figure 24, AA outperforms ZIP and GDX in both single-allocation and multi-allocation scenarios. In a market populated solely with AA traders, the efficiency is typically higher than 99.9% even in the presence of market shocks in both scenarios (with the efficiency of the multi-allocation scenario being, as expected, considerably higher). Futhermore, when we consider the more dynamic GOOGshock market with considerably larger changes in the equilibrium price (as a result of a market shock lasting several trading days) and asymmetric demand and supply (see Figure 25), AA performs

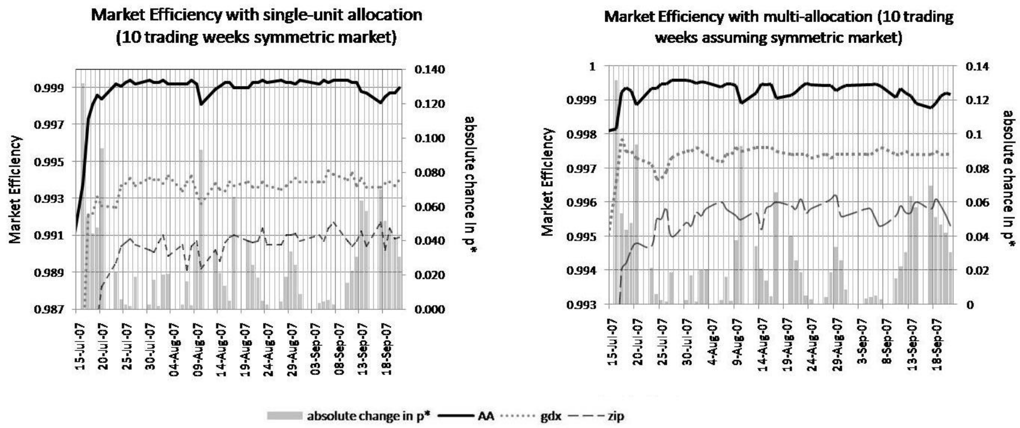

Fig. 24. Market efficiency with single-unit allocation on the left and with multi-unit allocation on the right, in a GOOG market. Note that the market shocks are given as the change in equilibrium prices.

considerably better than GDX and ZIP (by a margin of 5.2% and 3.3% respectively), demonstrating its superiority even in highly dynamic markets. Furthermore, we also notice that ZIP now outperforms GDX, suggesting that it performs better given more variable environments. It further suggests that adaptive strategies, such as AA and ZIP, are more efficient in very changing markets than belief-based strategies, such as GDX, which have to build a belief of a market that is constantly changing and, therefore, can possibly never catch up with the actual situation.

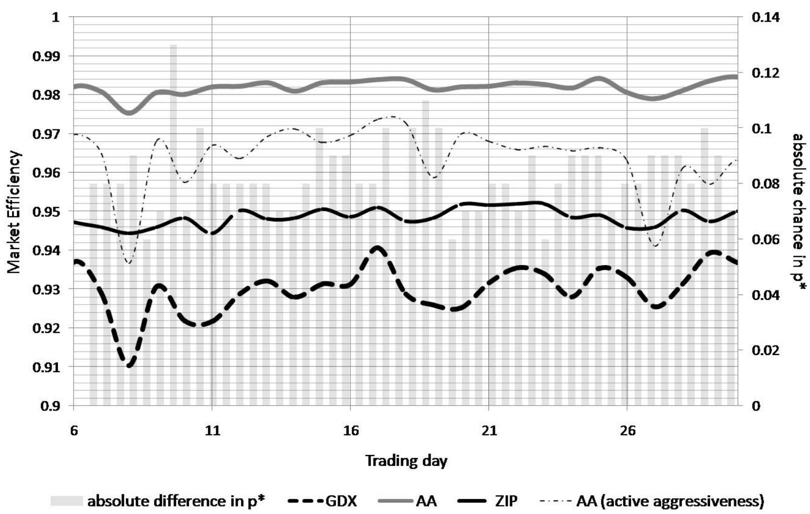

Fig. 25. Efficiency in the GOOGshock market (we ignore the first few days when agents are still very inefficient). Note that the changes in equilibrium price are considerably higher than in the GOOG market (see Figure 24). The thin dotted line represents the efficiency of the AA strategy when it is always active (aggressiveness level fixed at 0).

In summary, we observe that AA is the most efficient strategy and, therefore, the one most likely to be eventually adopted by the population of buyers and sellers in the market. This is true for all the market environments we have tested and not just the examples shown here.

6 Conclusions

In this paper, we presented a novel adaptive-aggressiveness strategy, AA, that software agents can use to bid in Continuous Double Auctions. The AA strategy is principally based on a short-term and a long-term learning of the agent's bidding behaviour. For the short-term learning, the motivation was to immediately respond to fluctuations in the market conditions, and the agent updates the *aggressiveness* of its bidding behaviour based on market information observed after every bid or ask appears in the market. The motivation for the long-term learning mechanism, on the other hand, was to respond to more systematic changes in the market conditions and, in particular, to market shocks. To achieve this, our strategy updates an aggressiveness model that determines how the agent's degree of aggressiveness influences its choice of bids or asks to submit in the market, based on market information observed after every successful transaction.

We then went on to describe a novel methodology for benchmarking the efficiency of strategies in homogeneous populations by considering daily efficiency and price volatility, rather than simply the overall efficiency as in previous work. Our approach provided more insights into how the learning mechanisms of the strategies affect their efficiencies in the market. In particular, we showed how the efficiencies of strategies change as they learn to be more efficient in the market or when there is a market shock. These would not have been observable using the standard approach because it calculates only the average market efficiency over all trading days and compares the (scalar) efficiencies of different strategies. We also provided a novel methodology for heterogeneous populations where we adopt a two-population evolutionary game theoretic model for a comprehensive analysis of buyer's and seller's choices of strategy in the market. In particular, we looked at the separate evolution of buyer's and seller's strategies in the market. Again this would not be observable in the traditional one-population approach because it assumes that buyers and sellers adopt the same strategy in the market and, thus, the evolution of buyers' and sellers' strategies is the same. Furthermore, this work also represents the first attempt to analyse the strategic interactions of agents in dynamic environments with different symmetric and asymmetric demand and supply. This is an advance over the analysis that is typically performed in heterogeneous populations with a static setting (i.e. no market shock) and only a symmetric demand and supply.

Using these methods, we benchmarked our AA strategy against the state of the art ZIP and GDX strategies. In so doing, we empirically demonstrated how it outperforms these benchmarks in different static and dynamic environments, in both homogeneous and heterogeneous populations. Specifically, within homogeneous populations, the AA strategy outperformed the benchmarks, in terms of market efficiency, by up to 3.6% in the static case and 2.8% in the dynamic case. Furthermore, we empirically showed the AA outperforms the state of the art GDX by 5.2% and ZIP by 3.3% in a market based on real market data. It is interesting to point out that learning strategies such as AA and ZIP are relatively more efficient in the dynamic situations such as the GOOGshock market than the belief-based GDX strategy. This is so because GDX has to update its whole (non-scalar) belief of the market every trading day and therefore it is slower in adapting to these variable market conditions than AA or ZIP which simply learn their (scalar) aggressiveness and profit margin respectively and are better able to adapt to drastically changing market conditions. Finally, within heterogeneous populations, based on our evolutionary game theoretic analysis, we showed out that there was a probability above 85% that the AA strategy will eventually be adopted by buyers and sellers in the market.

For future work, we first will study the use of other selection mechanisms than the Replicator Dynamics, including mutation and cross-over. Second, and more importantly, we intend to consider the use of the AA strategy in more complex variants of the CDA, and in particular financial exchanges such as the NYSE or NASDAQ. This would entail considering the additional (or lack of) information in these variants and modifying AA accordingly. The particular complication we foresee with the real financial markets would be in terms of being robust to market shocks, which rather than occurring over trading days as in our model of the CDA, occur over trading hours, such that we would require a more complex equilibrium estimator with some form of trend analysis. Thereon, we intend to show how parameters including λ_a, λ_r, γ, β_1 and β_2, η and N can be fine-tuned for particular markets by adopting a similar evolutionary approach as Cliff's [5] using a GA-search. In real markets, such fine-tuning is today's norm in a very competitive environment where improvements of the order of 0.01% are highly desirable (and translate into profits of hundreds of thousands of pounds in the investment banking industry with annual trading profits worth billions of dollars).

Acknowledgement

The first author was funded by the DIF-DTC 8.6 project (`www.difdtc.com`) We wish to thank the reviewers for their valuable comments that have enabled us to improve the paper. Specifically, this paper is a significantly extended

version of our ECAI paper [27]; we have revised the terminology used, provided the novel long-term learning mechanism to update the θ-parameter based on price volatility, and devised the methodology for evaluating CDA strategies within both homogeneous and heterogeneous populations.

A Empirical Study within a Heterogeneous Population

In Section 5, we benchmarked the AA strategy against the state of the art ZIP and GDX strategies for different scenarios. Here, we provide an analysis of AA against ZIP and GDX in the remaining cases that have not been considered in the main body of the paper. For each case, we give the different attractors and saddle points and the probability that each of these attractors will be adopted. We observe that AA always outperforms ZIP and GDX in line with our observations in the main body of the paper.

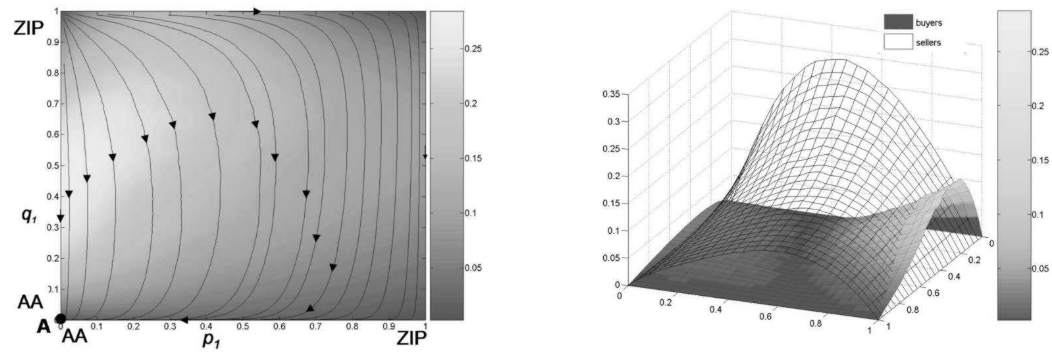

Fig. A.1. Scenario M3 with AA and ZIP agents. Here, we have one attractor: A at (0,0). AA is a dominant strategy that will eventually be adopted in the market.

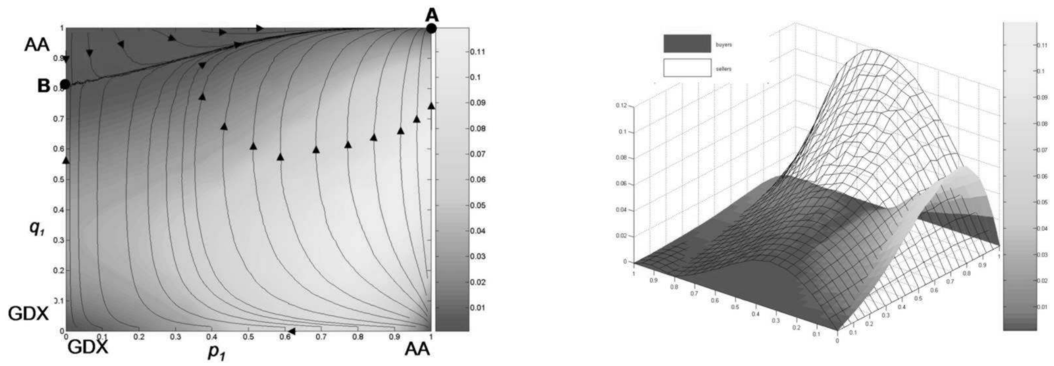

Fig. A.2. Scenario M3 with AA and GDX agents. Here, we have one attractor: A at (1,1) and one saddle point: B at (0,0.81). AA is a dominant strategy that will eventually be adopted in the market.

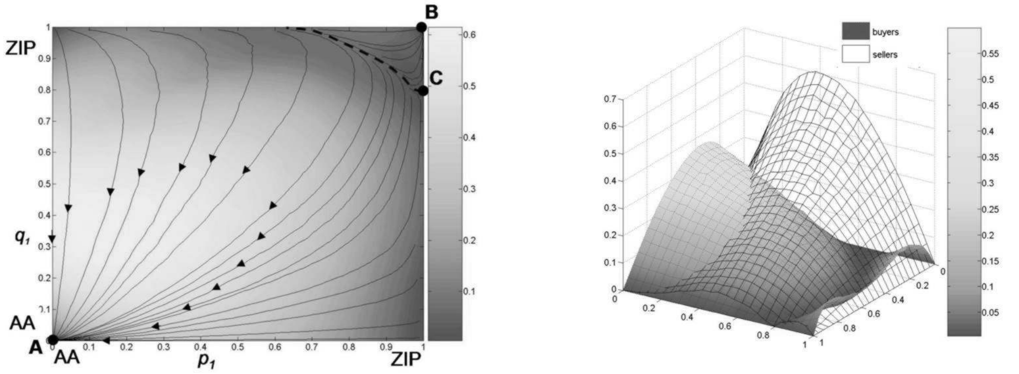

Fig. A.3. Scenario MS31 with AA and ZIP agents. Here, we have two attractors: A at (0,0) and B at (1,1) and a saddle point: C at (1,0.80). The probability that A will be adopted is 0.952 and that B will be adopted is 0.048.

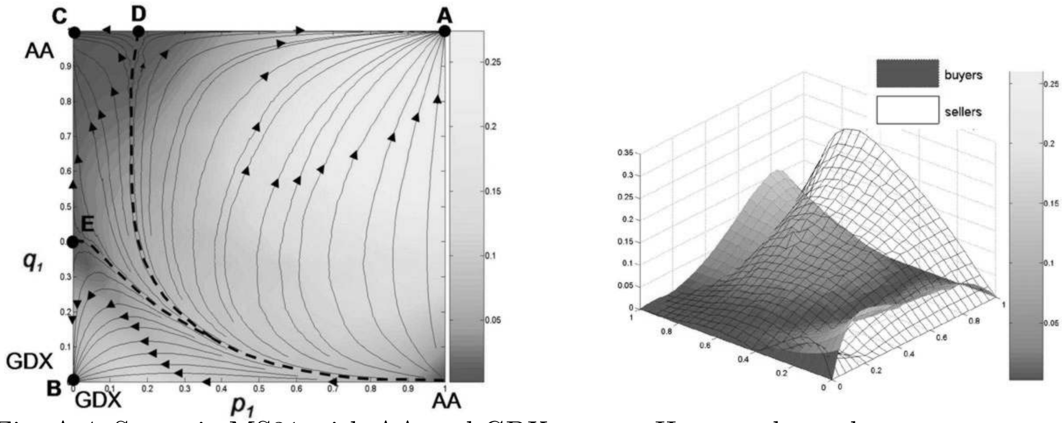

Fig. A.4. Scenario MS31 with AA and GDX agents. Here, we have three attractors: A at (1,1), B at (0,0) and C at (0,1) and two saddle points: D at (0.19,1) and E at (0,0.40). The probability that A will be adopted is 0.776, that B will be adopted is 0.125 and that C will be adopted is 0.099.

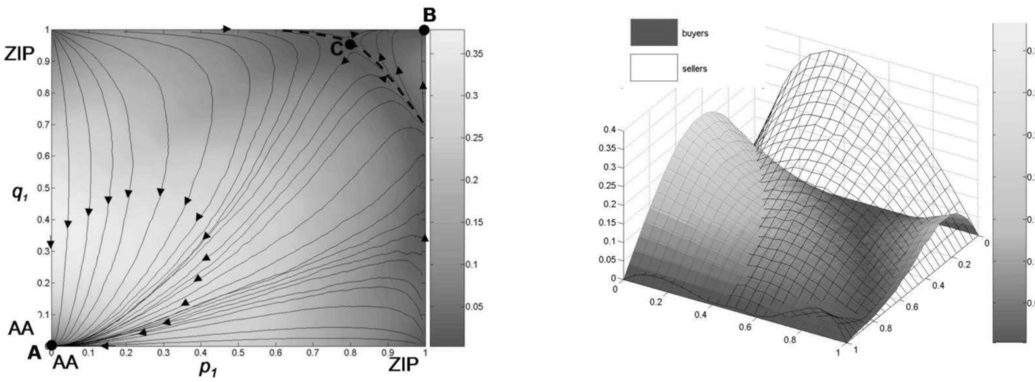

Fig. A.5. Scenario MS23 with AA and ZIP agents. Here, we have two attractors: A at (0,0) and B at (1,1) and a saddle point: C at (0.80, 0.95). The probability that A will be adopted is 0.965 and that B will be adopted is 0.035.

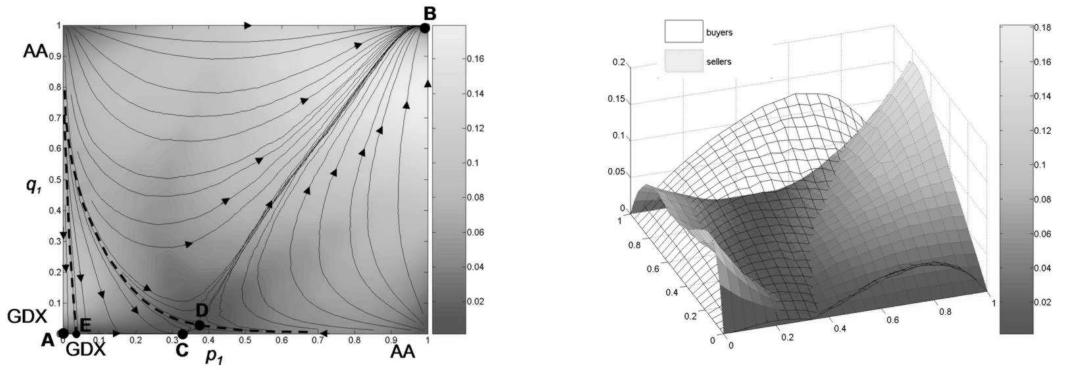

Fig. A.6. Scenario MS23 with AA and GDX agents. Here, we have three attractors: A at (0,0), B at (1,1) and C at (0.32.0) and two saddle points: D at (0.37,0.04) and E at (0.03,0). The probability that A will be adopted is 0.010, that B will be adopted is 0.902 and that C will be adopted is 0.088.

References

[1] A. Byde. Applying evolutionary game theory to auction mechanism design. *ACM Conference on Electronic Commerce*, pages 192–198, 2003.

[2] S. H. Clearwater, editor. *Market-Based Control - A paradigm for distributed resource allocation*. World Scientific, 1996.

[3] S. H. Clearwater, R. Costanza, M. Dixon, and B. Schroeder. Saving energy using market-based control. *S. H. Clearwater (Ed.), Market-Based Control: A Paradigm for Distributed Resource Allocation*, pages 253–273, 1996.

[4] D. Cliff. Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets. Technical Report HPL-2001-99, 2001.

[5] D. Cliff. ZIP60: Further explorations in the evolutionary design of trader agents and online auction-market mechanisms. IEEE Transactions on Evolutionary Computation (to appear), 2008.

[6] D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, 1997.

[7] R. Das, J. E. Hanson, J. O. Kephart, and G. Tesauro. Agent-human interactions in the continuous double auction. *Proc. Seventeenth International Joint Conference on Artificial Intelligence*, pages 1169–1176, 2001.

[8] R. K. Dash, P. Vytelingum, A. Rogers, E. David, and N. R. Jennings. A market-based task allocation mechanism for limited capacity suppliers. *IEEE Trans. on System, Man and Cybernetics (Part A)*, 37(3):391–405, 2007.

[9] D. Friedman and J. Rust. *The Double Auction Market: Institutions, Theories and Evidence*. Addison-Wesley, New York, 1992.

[10] S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.

[11] 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.

[12] 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.

[13] M. He, H. F. Leung, and N. R. Jennings. A fuzzy logic based bidding strat-

egy for autonomous agents in continuous double auctions. *IEEE Trans on Knowledge and Data Engineering*, 15(6):1345–1363, 2003.

[14] M. Hollander and D. A. Wolfe. *Nonparametric Statistical Methods.* Wiley, 1973.

[15] D. B. Marron and C. W. Bartels. The use of computer-assisted auctions for allocating tradeable pollution permits. *S. H. Clearwater (Ed.), Market-Based Control: A Paradigm for Distributed Resource Allocation*, pages 274–299, 1996.

[16] A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic Theory.* Oxford University Press, 1995.

[17] 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.

[18] S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparision of two double auction market designs. *Proc. Sixth Workshop on Agent Mediated Electronic Commerce*, pages 192–198, 2004.

[19] A. Saltelli, K. Chan, and E. M. Scott. *Sensitivity Analysis.* Wiley, 2001.

[20] V. L. Smith. An experimental study of competitive market behaviour. *Journal of Political Economy*, 70:111–137, 1962.

[21] 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.

[22] G. Tesauro and R. Das. High-performance bidding agents for the continuous double auction. *Proc. Third ACM Conference on Electronic Commerce*, pages 206–209, 2001.

[23] K. Tuyls and A. Nowe. Evolutionary game theory and multi-agent reinforcement learning. *The Knowledge Engineering Review*, 20(1):63–90, 2005.

[24] K. Tuyls and S. Parsons. What evolutionary game theory tells us about multiagent learning. *Artificial Intelligence Journal*, 171(7):406–416, 2007.

[25] P. Vytelingum. *The structure and behaviour of the Continuous Double Auction.* PhD thesis, School of Electronics and Computer Science, University of Southampton, http://eprints.ecs.soton.ac.uk/13234/, 2006.

[26] 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.

[27] 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.

[28] 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*, 2002.

[29] J. W. Weibull. *Evolutionary Game Theory.* MIT Press, Cambridge, MA, 1995.

[30] B. Widrow and M. E. Hoff. Adaptive switching circuits. *IRE WESCON Convention Record*, 4:96–104, 1960.