Department of Applied Mathematics, University of Venice

WORKING PAPER SERIES

Shira Fano, Marco Li Calzi, Paolo Pellizzari

Convergence of outcomes and evolution of strategic behavior in double auctions

Working Paper n. 196/2010
February 2010

ISSN: 1828-6887

Convergence of outcomes and evolution of strategic behavior in double auctions*

Shira Fano
[shira.fano@gmail.com]
Dept. Applied Mathematics
University of Venice

Marco LiCalzi
[licalzi@unive.it]
Dept. Applied Mathematics and
Advanced School of Economics,
University of Venice

Paolo Pellizzari
[paolop@unive.it]
Dept. Applied Mathematics and
Advanced School of Economics,
University of Venice

(February 2010)

Abstract. We study the emergence of strategic behavior in double auctions with an equal number n of buyers and sellers, under the distinct assumptions that orders are cleared simultaneously or asynchronously. The evolution of strategic behavior is modeled as a learning process driven by a genetic algorithm. We find that, as the size n of the market grows, allocative inefficiency tends to zero and performance converges to the competitive outcome, regardless of the order-clearing rule.

The main result concerns the evolution of strategic behavior. Under simultaneous order-clearing, as n increases, only marginal traders learn to be *price takers* and make offers equal to their valuations/costs. Under asynchronous order-clearing, as n increases, all intramarginal traders learn to be *price makers* and make offers equal to the competitive equilibrium price. The nature of the order-clearing rule affects in a fundamental way what kind of strategic behavior we should expect to emerge.

Keywords: trading protocols, asymptotic equivalence, learning, genetic algorithms.

JEL Classification Numbers: D44, D82; C63, C72

Correspondence to:

Marco LiCalzi	Department of Applied Mathematics
	University of Venice
	Dorsoduro 3825/e
	30123 Venice, Italy
Phone:	[+39] 041-234-6925
Fax:	[+39] 041-522-1756
E-mail:	licalzi@unive.it

* We thank the audience of the workshop on "Evolution and market behavior in economics and finance" held at the Scuola Superiore Sant'Anna for their comments. The second and third author acknowledge financial support from MIUR under grants 2007EENEAX and 2007TKLTSR.

1 Introduction

The double auction institution encompasses different families of trading protocols that gather buyers and sellers in a single exchange market; see Friedman (1993). The two prominent instances are the call market and the continuous double auction, that differ in whether orders are cleared simultaneously or asynchronously. Both kinds of protocols have been widely studied with a three-pronged approach based on analytical derivations, laboratory or field experiments, and computer simulations.

While the k-double auction has emerged as the standard model for the call market, there is no dominant paradigm for the continuous double auction. The sheer number of possible variants impairs the emergence of a single representative format. There is a vast literature scattered over different formulations, using a devilish variety of assumptions on the richness of the strategy space. This may depend on information as diverse as the past history of orders or transactions, the current status of the book, the timing at which an offer is made, and another myriad of protocolary details such as the option to cancel and resubmit orders or the obligation to submit price-improving offers (a.k.a. as "NYSE spread-improvement rule"). In order to make progress, it is necessary to give up on generality.

We study a simple model for the double auction where the strategies of the traders depend only on their private types. This simplification allows to provide a unified model for the k-double auction and the continuous double auction, where the only difference is in the order-clearing rule. For the special case where the market has only one buyer and one seller, we prove that the equilibrium strategies coincide but the equilibrium outcome is different. This provides a sharp illustration of the direct effects of the order-clearing rule on the performance of a trading protocol.

Our main objective is the study of which trading strategies emerge as strategically plausible when the number of traders increases. This plausibility encompasses two requirements: the profile of strategies must be (close to) an equilibrium and it must be the outcome of an evolutionary process (e.g., learning) that justifies its prominence. (In other words, we do not assume that "all equilibria are created equal".) Our approach is to let agents use genetic algorithms to maximize individual profits and coevolve a profile of (possibly randomized) trading strategies. This evolutionary approach circumvents many of the computational difficulties that affects the search for equilibrium strategies in current models of continuous double auctions.

In the long run, as evolution takes place, we show that agents learn to (approximately) play equilibrium strategies. Our first result is consistent with the asymptotic approach to the "equivalence principle" (see Aumann, 1987), according to which increasing the number of agents diminishes their strategic influence and lead the system towards the competitive outcome. Regardless of the order-clearing rule, as the market grows in size, allocative inefficiency tends to zero and performance converges to the competitive outcome. For the k-double auction, this result is magisterially exemplified in Rustichini et al. (1994) that proves the stronger statement that traders' equilibrium offers asymptotically converges to truth-telling: they are willing to accept any price below (above) their valuations (costs) so they act as perfectly competitive price-takers.

Our main results concern the evolution of strategic behavior when the market grows in size. In a nutshell, we show that the asymptotic emergence of the competitive outcome does not imply that all traders should learn to be price takers. Moreover, the order-clearing rule affects in a fundamental way what kind of strategic behavior we should expect to emerge. Under simultaneous order-clearing, as n increases, only marginal traders learn to be *price takers* and make offers equal to their valuations/costs. This implies that the asymptotically unique equilibrium strategies derived in Rustichini et al. (1994) are unlikely to be learnable. Under asynchronous order-clearing, as n increases, all intramarginal traders learn to be *price makers* and make offers equal to the competitive equilibrium price.

There is a growing literature on the application of evolutionary processes to the study of market protocols and trading strategies. MacKie-Mason and Wellman (2006) provides a survey centered around the general field of computational market design, where both protocols and strategies may be allowed to change. This paper exploits genetic algorithms to evolve and compare strategic behavior across two specific formats for the double auction. To the best of our knowledge, Dawid (1999) is the first explicit application of genetic algorithms for the derivation of strategic behavior in a double auction market. Phelps et al. (2006) provides an evolutionary comparison between three families of strategies: it concludes that, as the market grows in size, truth-telling becomes increasingly likely to emerge in a call market while it tends to disappear in a continuous double auction. Anufriev et al. (2010) applies an algorithm known as individual evolutionary learning to study the effects of public disclosure of information on the evolution of strategic behavior in a continuous double auction.

Our paper is organized as follows. Section 2 presents a unified model for the k-double auction and the continuous double auction, considered as games with incomplete information where traders' strategies depend only on their private types but different order-clearing rules apply. Section 3 describes in detail the setup for our simulations. Section 4 collects our results on the asymptotic emergence of outcome and strategic behavior for the k-double auction, when order-clearing is simultaneous. Section 5 reviews and contrasts the analogous results for the continuous double auction, when order-clearing is asynchronous. Section 6 provides an additional comparative analysis between the strategies evolved by our genetic algorithm for the continuous double auction against the current benchmark in the literature, as provided by Zhan and Friedman (2007). An appendix collects spurious material.

2 A unified model

There are many variants of the double auction; see Friedman (1993). We focus on whether orders are cleared simultaneously or asynchronously. The first case characterizes the class of trading protocols known as *call markets* or batch auctions. The second case characterizes the family of the *continuous double auctions*. We study a model that encompasses both clearing rules and make it simple to elucidate their effects on the asymptotic convergence of outcomes and the emergence of strategic behavior.

The presentation is organized as follows. The environment described in Section 2.1 collects the general characteristics of the economy, including agents' preferences and endowments. Section 2.2 summarizes the well-known model of a *k-double auction* for the call market and reviews its main properties. For our purposes, its most important feature is that it is assumes a simple strategy space for each trader. Section 2.3 describes our own model for a continuous double auction. We impose assumptions that make the strategy space of each trader as simple as in the k-double auction. Analogous simplifications underlie other models of the continuous double auction, ranging from the nonstrategic (e.g., Gode and Sunder, 1993) to the strategic ones (e.g., Zhan and Friedman, 2007). Section A in the appendix discusses a special case known as the *bilateral trading model* with one buyer and one seller, which provides a common foundation for the k-double auction and our version of the continuous double auction.

2.1 The environment

There is an equal number n of buyers and sellers. All traders wish to maximize expected profits. Each of them is in the market to exchange at most one unit of a generic good per day. Each buyer i has a private valuation v_i and each seller j has a private cost c_j. Valuations and costs are drawn from two (stochastically independent, as well as atomless and absolutely continuous) distributions F and G over the same support, which we normalize to $[0, 1]$ without loss of generality. As a special case, it is customary to assume that F and G are uniform distributions on $[0, 1]$.

When all traders are price takers, it is customary to define intramarginal and extra-marginal buyers (sellers) depending on their position on the demand (supply) function with respect to the market-clearing price(s). We follow tradition but refine this qualitative distinction into a complete ordering. Define the *strength* of a buyer with valuation v as the distance from the valuation of the weakest buyer ($v = 0$) and the strength of a seller with cost c as the distance from the valuation of the weakest seller ($c = 1$). Stronger traders have valuations (or costs) lying farther away from the market-clearing price(s).

2.2 Simultaneous order clearing

The standard model for a call market where orders are cleared simultaneously is the k-double auction; see Satterthwaite and Williams (1993). Buyers and sellers are required to submit price offers simultaneously. Each buyer declares the maximum *bid price* at which he is willing to buy and each seller issues the minimum *ask price* at which she is willing to sell. Traders decide strategically their price offers to maximize their expected payoffs. The strategy of a buyer i is a *bidding function* $\beta_i : [0, 1] \to \mathbb{R}^+$ that defines his bid $b_i = \beta_i(v_i)$ as a function of his valuation v_i. Similarly, the strategy of a seller j is an *asking function* $\alpha_j : [0, 1] \to \mathbb{R}^+$ that defines his ask $a_j = \alpha_j(c_j)$ as a function of her cost c_j.

Buyers' and sellers' offers are aggregated to form the demand and supply functions. Their intersection defines an interval $[p_1, p_2]$ of market-clearing prices. The k-double auction selects as trading price the value $p^* = (1 - k)p_1 + kp_2$, where each choice of k in $[0, 1]$ defines a different mechanism. Trade occurs among buyers who bid no less than p^* and sellers who ask no more than p^*. (Some rationing may take place at the margin, but the exact details are not relevant.) When a transaction takes place at price p between a buyer with valuation v and a seller with cost c, the payoffs are $v - p$ and $p - c$ respectively. If he does not enter into a transaction, the payoff for the trader is zero.

As discussed in Leininger et al. (1989), there exist infinitely many Bayes-Nash equilibria for the k-double auction, but two important families have been singled out. The first class collects the equilibria where the strategies satisfy a pair of differential equations; the second class is formed by the equilibria where the strategies are step functions. The literature usually restricts attention to the subset of symmetric[1] equilibria in the first class, where all buyers use the same (differentiable) bidding function β and all sellers use the same (differentiable) asking function α. This simplifies the specification of a profile of equilibrium strategies for n buyers and n sellers to a single pair (β, α).

We say that the symmetric strategy profile (β, α) is *nontrivial* if: 1) traders never play (weakly) dominated strategies; that is, $\beta(v) \leq v$ and $\alpha(c) \geq c$; 2) the set of buyers bidding $b > 0$ and the set of seller asking $a < 1$ have strictly positive probability. The first requirement upholds *individual rationality*: a buyer never bids above his value and a seller never asks below her cost. The second requirement rules out "no-trade" equilibria. The general features of a nontrivial symmetric and differentiable equilibrium of the k-double auction are stated in a well-known result from Rustichini et al. (1994).

Theorem 1 *For any nontrivial symmetric and differentiable equilibrium* (β, α)*:*

> *1) there exist values* $v^* < 1$ *and* $c^* > 0$ *such that a buyer with valuation* v *trades with positive probability if and only if* $v > v^*$ *and a seller with cost* c *trades with positive probability if and only if* $c < c^*$*;*
>
> *2)* β *and* α *are increasing over* $(v^*, 1]$ *and* $[0, c^*)$*, respectively;*
>
> *3)* $\lim_{v \downarrow v^*} \beta(v) = v^* = \lim_{c \downarrow 0} \alpha(c)$ *and* $\lim_{v \uparrow 1} \beta(v) = c^* = \lim_{c \uparrow 0} \alpha(c)$*.*

The intervals $(v^*, 1]$ and $[0, c^*]$ define the domain of serious offers where the bidding and asking functions are uniquely defined. We call *serious buyers (sellers)* those traders who are supposed to make serious bids (asks) in equilibrium.

For a given value n, in any symmetric and differential equilibrium serious buyers shade their valuations and bid $\beta(v) < v$; similarly, serious sellers markup their costs and ask $\alpha(c) > c$. This misrepresentation marks a departure from naive price-taking behavior and occurs as a result of the strategic interaction between all traders. Intuitively, when n gets large, one expect the scope for strategic misrepresentation to shrink so that the equilibrium price should tend towards the competitive value. This is formally shown in Rustichini et al. (1994) who prove that, as $n \uparrow \infty$, in any nontrivial symmetric and differentiable equilibrium the amount of strategic misrepresentation $|v - \beta(v)| + |\alpha(c) - c|$ drops to zero as $O(1/n)$ and the allocative inefficiency disappears at a rate $O(1/n^2)$. Satterthwaite and Williams (2002) additionally proves that trade in the k-double auction is worst-case asymptotically optimal within a class of plausible mechanisms.

2.3 Asynchronous order clearing

Roughly speaking, a continuous double auction works as follows. Agents sequentially submit offers on the selling and buying books. Orders are immediately executed at the outstanding price if they are marketable; otherwise, they are recorded on the books with the usual price-time priority and remain valid unless a cancellation occurs. When a transaction takes place

[1] We use a stronger notion of *role-symmetry* later in the paper.

between two traders, their orders are removed from the books and they leave the market. Hence, orders are cleared asynchronously in separate trades, usually at different prices.

The main complication associated with moving from the simultaneous clearing of the call market to the asynchronous clearing of the continuous double auction is that the latter trading protocol requires agents to make and revise choices sequentially. As usual, turning a game with simultaneous actions into one with sequential moves greatly expands the strategy space. In general, even in the simplified environment of Section 2.1, the complexity of a trader's strategy in a continuous double auction can be daunting, since he can make (or withdraw) price offers depending on the past history as well as deciding the timing of his actions. The model studied in this paper makes three appropriate assumptions and reduces the complexity of the strategy space.

First, we assume the environment from Section 2.1. This forces each trader to make only unit orders and mutes any issue about the quantity of good to be demanded or supplied. Second, we assume that the order of arrival of traders is randomly drawn according to a uniform distribution over all possible queues and that each trader gets only one chance to act. This eliminates issues of timing or order cancellation. Third, we interpret the bid b_i of a buyer i as a limit order. (Analogous assumption holds for the seller.) When his bid b_i reaches the market, the protocol compares it against the outstanding ask a: if $a \leq b_i$, then his limit order is marketable and trade occurs at price $p = a$. Otherwise, $a > b_i$; then, a limit order of b_i is stored on the buying book and awaits for a matching offer. Thus, the strategy of a buyer is a *bidding function* $\beta_i : [0, 1] \to \mathbb{R}^+$ that yields a "limit bid" $b_i = \beta_i(v_i)$. Analogously, the strategy of a seller j is an *asking function* $\alpha_j : [0, 1] \to \mathbb{R}^+$ that defines her "limit ask" $a_j = \alpha_j(c_j)$.

This setup defines a game with incomplete information, where the strategy space of the players is the same as in the k-double auction. Clearly, in general the two games underlying the trading protocols with simultaneous or asynchronous order clearing are not payoff equivalent; f.i., in the k-double auction all trade takes place at the same price while in the continuous double auction each trade carries its own price. However, surprisingly enough, Appendix A shows that they are strategically equivalent (i.e., equilibrium strategies coincide) in the special case of the bilateral trading model when $n = 1$. This provides some justification for our use of a unified model to single out the effects of the order-clearing rule.

3 Setup

This section illustrates the setup for our simulations. When obvious, we describe features only on the buyers' side because sellers are modeled symmetrically. For computational purposes, we discretize both the sets of types and strategies. Given an integer m, let $\delta = (1/m)$ be the tick size used to define an equispaced grid of buyers' valuations (and sellers' costs) over the interval $[0, 1]$. We assume that buyers' valuations are drawn from a uniform distribution on the support $\mathcal{V} = \{\delta, 2\delta, \ldots, (m-1)\delta\}$. The set of buyers' feasible offers $O = \{\delta, 2\delta, \ldots, (m-1)\delta\}$ is taken to coincide with $\mathcal{V}$. (This choice is for simplicity: we found no relevant difference in results by testing for an offers' grid finer than the valuations' grid.)

There is a pool of $2N$ potential traders. Half of them are potential (male) buyers and half are potential (female) sellers. Each of them has private information about his valuation or her cost: in game-theoretic parlance, each trader in the pool is an agent who knows his *type*. Given that there are $(m-1)$ possible valuations (costs), the pool of N potential buyers (sellers) naturally divides into $(m-1)$ *groups* of agents of the same type. In accordance with the assumption that types are uniformly distributed, each of these groups has the same cardinality r; thus, $N = r(m-1)$. For reasons that will be clear momentarily, we assume that r is a multiple of 4.

A (pure) *strategy* for a buyer with valuation v in $\mathcal{V}$ is a discrete bidding function $\beta : \mathcal{V} \to O$. Alternatively, one can think of a pure strategy as the string of (pure) *actions* taken by each possible type of a buyer. In general, we wish to allow buyers to play randomized strategies. Denote by $\Delta(O)$ the set of probability distributions over O. A *randomized strategy* $\mathbf{q} : \mathcal{V} \to \Delta(O)$ is a function that associates with each valuation v in $\mathcal{V}$ a probability distribution $\mathbf{q}(v, \cdot)$ over the bids b in O such that $\mathbf{q}(v, b) \geq 0$ and $\sum_{b \in O} \mathbf{q}(v, b) = 1$. Assuming individual rationality, we also impose $\mathbf{q}(v, b) = 0$ for all $b > v$. Hence, a strategy $\mathbf{q}$ is defined as a discrete probability distribution over the individually rational offers. Sellers' strategies are similarly defined. When it is his turn to make an offer, a buyer with valuation v uses his strategy $\mathbf{q}$ to issue a bid according to the probability distribution $\mathbf{q}(v, \cdot)$. The same applies on the sellers' side.

Agents interact repeatedly and anonymously. A *trading day* is made of several rounds. In each *round*, we randomly draw (without replacement) n buyers and n sellers from the pool of available agents and let them visit the market ($n \ll N$). Conventionally, we set the number of rounds played in a day so that on average each of the $2N$ potential traders has one chance to trade: hence, a trading day consists of about N/n rounds. (We give us a bit of slack and round down this value.)

We evolve traders' strategies using a genetic algorithm, henceforth nicknamed GA for brevity. This is a well-known and robust optimization method that we use to model how agents learn what strategies they should play. Learning is driven by the average profit of a trading strategy. Profits are $v - p$ for buyers and $p - c$ for sellers, where p is the price at which the transaction occurs. We measure the performance of an agent (a.k.a. fitness of a genotype, in GA parlance) by the average of his trading profits over the past τ days. For simplicity, we refer to a consecutive sequence of τ trading days as the *evaluation window*. Periodically (i.e., after every τ days), we update agents' strategies using a standard GA machinery based on selection, crossover and mutation. Our model is often called "4-2" in the literature, because genotypes are grouped in sets of 4 individuals and the worst couple is replaced by crossed over and (possibly) mutated copies of the best genotypes. See Ashlock (2006) for details.

More precisely, recall that the pool of N potential buyers is formed by $(m-1)$ groups of r agents with the same type. Within the group corresponding to type v, each agent holds his own bidding strategy. (Once again, analogous assumptions hold for sellers.) We apply the genetic algorithm separately for each v to mimic learning at the *interim* stage, when an agent knows his valuation v and is interested only in choosing an action appropriate for v. Thus, for each buyer i with type v, we update his randomized strategy $\mathbf{q}_i(v, \cdot)$ as follows.

We compute the fitness index as the average profit over the evaluation window for each buyer. We randomly partition the set of r buyers with valuation v into $(r/4)$ foursome groups and update each group as follows. (This is the only place where we use the assumption that

4 divides r exactly.) Let $f_1 \geq f_2 \geq f_3 \geq f_4$ be the fitness indices for the four agents in a group. (Ties are broken randomly.) The two individuals with the largest fitnesses produce by crossover two "siblings" that replace the other two individuals with the (lower) fitnesses f_3 and f_4. Crossover is applied using the standard one-point operator: given two randomized actions $\mathbf{q}(v, b)$ and $\mathbf{q}'(v, b)$, we pick a random position w and generate the siblings' actions

$$\mathbf{q}''(v, b) = \left[q(\delta), \ldots, q(w\delta), q'((w+1)\delta), \ldots, q'((m-1)\delta)\right]$$

and

$$\mathbf{q}'''(v, b) = \left[q'(\delta), \ldots, q'(w\delta), q((w+1)\delta), \ldots, q((m-1)\delta)\right].$$

Mutation is applied to the siblings' actions in two ways. First, inspired by Lettau (1997), with probability linearly decreasing in time (as measured by the cumulative number of rounds) we substitute one of the siblings' (randomized) actions with a pure action selected with uniform probability among all bids that are individually rational with respect to v. Second, again with probability linearly decreasing in time, we apply a zero-mean additive shock to one component of the other sibling's action and accordingly renormalize his randomized strategy.

This process is repeated for each v across the corresponding groups. Hence, the actions that make up buyers' strategies are separately (and simultaneously) evolved. This matches the intuition that learning takes place during the *interim* stage. The evolutionary success of an individual is based on his capability to gain higher profits with respect to other traders with the same valuation (or cost), while his overall gains depend on the actions that people are similarly evolving for different valuations and costs. Agents strive for higher profits *within* their peers but trade *across* the whole population.

The parameters used for all the simulations presented in this paper are shown in Table 1. When more than one possible value is given, boldface is used to denote the baseline: the case with $n = 10$ is used as benchmark. A full run of the genetic algorithm involves 5000 trading

Parameter	Description	Value
n	Number of active buyers (and sellers)	$\{1, \mathbf{10}, 100\}$
δ	Tick size ($= 1/m$) for grid of types	1/20
τ	Length of the evaluation window	100
r	Size of a same-type group of traders	16

Table 1: Description and value of the parameters used for the simulations.

days; given $\tau = 100$, each agent is given 50 opportunities to revise his strategy. (Convergence usually takes place within 40 revisions.) Based on n, a trading day lasts $300/n$ rounds. All the results reported below are based on 20 runs where each trader is initially endowed with a pure strategy chosen with equal probability from the set of his individually rational offers. (We have extensively tested different initializations but there are no detectable differences in our results.)

4 Simultaneous order clearing: Results

This section presents the outcome of our simulations for the case of simultaneous order-clearing. We keep this section brief, because we find it more instructive to provide more details and expand the analysis of the continuous double auction.[2] Therefore, we restrict attention only on our main focus: the evolution of strategic behavior when the market grows in size. Additional statistics and figures, such as the dynamics of transaction prices or traders' profit as learning progresses, have been curtailed for brevity.

Recall Theorem 1 in Section 2.2: strategic misrepresentation should shrink to zero as the number n of active traders on each side of the market in the market increases. Hence, as n grows, all traders should learn to be price takers. In a nutshell, we show that this is not warranted. Even a search procedure as powerful as our genetic algorithm may fail to learn "price-taking" behavior.

For each possible valuation v, there are $r = 16$ buyers. (As usual, analogous statements hold for sellers.) Each of these $i = 1, 2, \ldots, r$ buyers of type v is associated with a (possibly randomized) strategy $\mathbf{q}_i(v, \cdot)$, so our genetic algorithm maintains, compares and updates $r = 16$ different strategies for a type v. At the end of a simulation, the algorithm evolves strategies that tend to have very similar fitness indices but need not be identical. In other words, buyers of the same type v are not constrained to learn the same strategy as far as they find a way to attain similar profits.

The richness of strategic behavior that the genetic algorithm may generate can be visually appreciated using the following representation. For each v in the grid, we randomly pick one buyer i of type v and generate a bid $b_i(v)$ using his strategy $\mathbf{q}_i(v, \cdot)$. Then we join these pairs of points $(v, b_i(v))$ and obtain a bidding function that shows a realization of bids for each of the types of a buyer. This "sampled" bidding function is noisy and should not be expected to exhibit special mathematical properties, except for those implied by individual rationality: the graph of bidding (asking) functions is always below (above) true valuations (costs), represented by the bisector. Moving from left to right, Figure 1 displays a bundle of three different "sampled" bidding and asking functions for $n = 1, 10, 100$. The competitive price $p^* = 0.5$ is depicted as a dashed horizontal line. Price-taking behavior is equivalent to stating an offer equal to the true valuation (cost) so it is represented by the bisector.

Looking at the picture, a few regularities emerge very clearly. Bidding and asking function tend to be increasing in types, as it should be expected. Moreover, as we move from $n = 1$ to $n = 10$, there is a clear shift towards truth-telling behavior. Increasing the number of traders make the environment more competitive and hence introduces positive incentives for all types to learn and bid more aggressively. The more people around, the riskier it is to lie so that all traders reduce the amount of their strategic misrepresentation. This aligns perfectly with the result in Rustichini et al. (1994). But, when notching up from $n = 10$ to $n = 100$, this effect seems to evaporate. What is going on?

In order to answer, we remove the noise and look at the two "smoothed" trading functions that for each v (or c) plot the average over 20 runs of the expected bid (ask) for each group of

[2] With obvious modifications, Results 1–3 listed in Section 5 for the continuous double auction hold also for the call market.

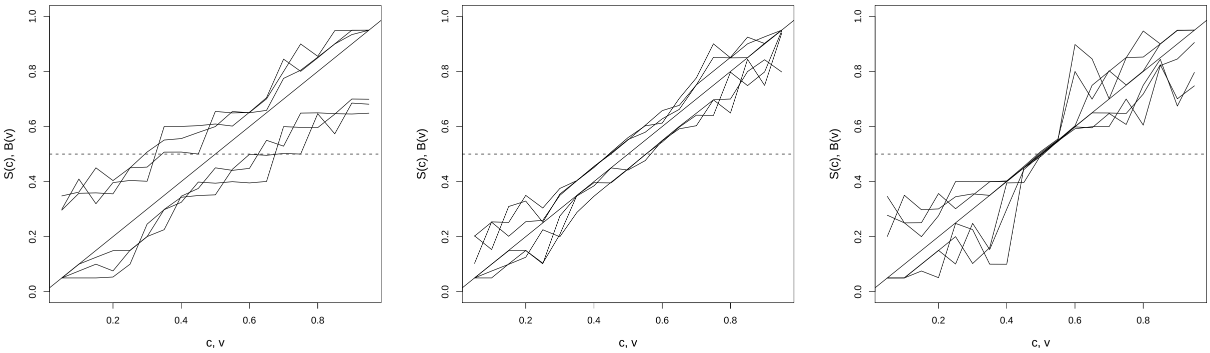

Figure 1: Bundles of realized trading functions for $n = 1, 10, 100$ (from left to right).

type v (c). Moving from left to right, these are displayed in Figure 2 for $n = 1, 10, 100$. When we go from $n = 1$ to $n = 10$, the intuition above is fully confirmed. Moving from $n = 10$ to $n = 100$, we see that it holds only for the marginal traders with valuations or costs around $p^* = 1/2$. A moment of thought explains the puzzle. Under simultaneous-order clearing, all traders trade at a single price which (as n increases) is increasingly likely to be set by the marginal traders. Bids and asks from extramarginal traders do not matter because they are extremely unlikely to trade anyway; hence, there is no sufficient push for them to learn anything. Similarly, bids and asks from deeply intramarginal traders do not matter because trade is going to occur at price p^* for any offer sufficiently away from p^*; again, there is no sufficient drift for learning to be price-takers. The market protocol is so robust that mere price-taking behavior from the marginal traders is sufficient to yield the competitive price p^*. All the same, such robustness makes it highly unlikely that non-marginal traders may learn the (unique) symmetric equilibrium.

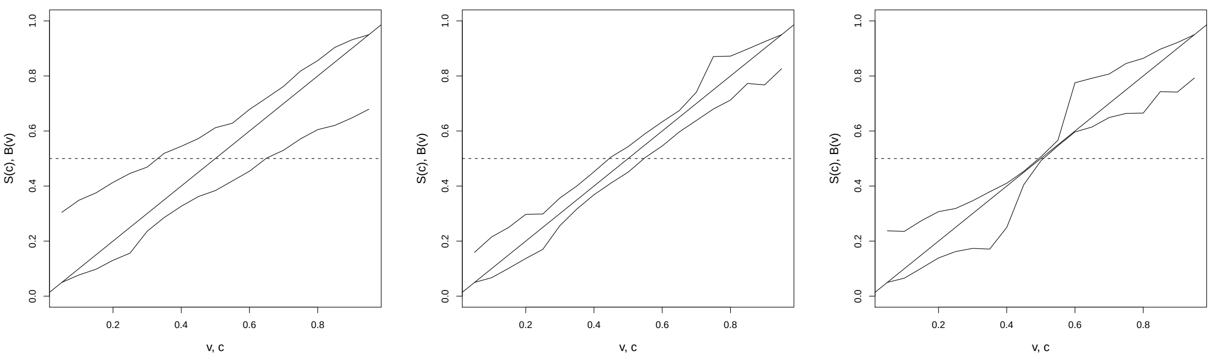

Figure 2: Average trading functions for $n = 1, 10, 100$ (from left to right).

4.1 Quality control for the GA strategies

We speak of the evolved trading strategies as an "equilibrium" profile that is learned by agents. However, it is clear by its nature that the genetic algorithm produces simply an approximation and thus we need to provide some form of quality control. A natural measure for the goodness of computational approximations to an equilibrium profile is suggested by the notion of an ε-equilibrium.

For simplicity, we recall its definition with reference to a game with complete information. Using customary notation, assume that $G = (n, S, u)$ is a game with n players, $S = S_1 \times \ldots \times S_n$ is the space of strategy profiles and u is the vector of n payoff functions for each player. As usual, assume that players maximize expected payoffs. Given $\varepsilon > 0$, we say that a (possibly randomized) strategy profile σ is an *ε-equilibrium* if, for all players i and for all strategies s_i in S_i,

$$u_i(\sigma_i, \sigma_{-i}) \geq u_i(s_i, \sigma_{-i}) - \varepsilon \tag{1}$$

Clearly, σ is a Nash equilibrium if and only if (1) holds for $\varepsilon = 0$. Therefore, if we denote by $\varepsilon^*(\sigma) \geq 0$ the least ϵ that satisfies (1), we obtain a direct estimate of the "distance" that separates σ from being an equilibrium. Roughly speaking, $\varepsilon^*(\sigma)$ is a worst-case measure for the temptation of at least one agent to break away from the strategy profile σ. The lower $\varepsilon^*(\sigma)$, the lower the push towards exploring alternative strategies. The extension of this notion to a game with incomplete information requires to check the analog of (1) for all types t_i of each player i: $u_i(\sigma_i(t_i), \sigma_{-i}; t_i) \geq u_i(s_i(t_i), \sigma_{-i}; t_i) - \varepsilon$.

The notion of ε-equilibrium applies to a given strategy profile σ. An evolutionary process like our genetic algorithm, however, is unlikely to lead always to the same strategy profile. Each run of GA, in fact, ultimately evolves strategy profiles that are similar but not necessarily identical. Hence, we need to find a measure of quality control $\varepsilon^*(\text{GA})$ taking into account that each run of GA may end up recommending slightly different strategy profiles. To this purpose, we extend the worst-case logic underlying the notion of ε-equilibrium as follows. Recall that we execute a batch of 20 runs of GA. At the end of each run $k = 1, 2, \ldots, 20$, GA evolves a strategy profile σ^k. Assuming that all other agents play their part of σ^k, we check for all types t_i of each player i the difference between his average payoff from playing what GA recommends (that is, $\sigma^k(t_i)$) and an arbitrary pure strategy:

$$\sum_{k=1}^{20} \frac{u_i(\sigma_i^k(t_i), \sigma_{-i}^k; t_i)}{20} \geq \sum_{k=1}^{20} \frac{u_i(s_i(t_i), \sigma_{-i}^k; t_i)}{20} - \varepsilon \tag{2}$$

We define $\varepsilon^*(\text{GA}) \geq 0$ as the least ϵ that satisfies (2). It provides a direct estimate of the "distance" that separates playing according to GA from being an optimal choice (against any of the pure strategies). This value is used to apply a t-test for the null hypothesis that $\varepsilon^* = 0$.

Table 2 collects and display the data used for our quality control. Each row in the top panel reports the number n of traders on each side of the market, the value of $\varepsilon^*(\text{GA})$, the type of the agent for whom the incentive to deviate is greatest, the offer that constitutes his optimal deviation and the p-value for the null hypothesis. For instance, the first row in the panel for GA reports the information that a seller with cost $c = 0.30$ can make an additional

	n	ε^*	type	offer	p-value
GA	1	0.0033	$c = 0.30$	0.45	0.25
	10	0.0060	$c = 0.25$	0.30	0.0197
	100	0.0014	$c = 0.20$	0.25	0.13
TT	1	0.0858	$v = 0.95$	0.70	10^{-12}
	10	0.0480	$v = 0.65$	0.60	0.0438
	100	0.0009	$c = 0.10$	0.40	0.27

Table 2: Quality control for the trading strategies evolved by GA in a call market.

average profit of 0.0033 by offering an ask price of 0.45 instead of complying with GA's recommendation. This additional gain is not sufficiently high to reject the null hypothesis for any reasonable level of confidence. It is apparent that $\varepsilon^*(\text{GA})$ is sufficiently close to zero for $n = 1$ and $n = 100$. For $n = 10$, $\varepsilon^*(\text{GA})$ is still negligible although statistically significant only at a confidence level above 2%. The left panel of Figure 3 provides a detailed graphical representation for the values of ε for the sellers' costs. (Due to the symmetry of our model, the graph for the buyers' valuations is analogous.) It is worth noting that the highest values of ε cluster around intramarginal traders of intermediate strength.

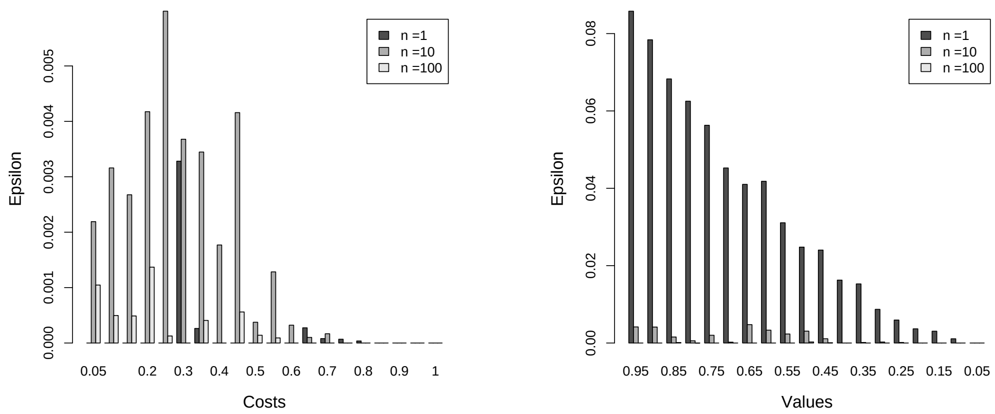

Figure 3: From left to right: values of ε^* for sellers (GA) and buyers (TT).

For comparison, the bottom panel of Table 2 provides the same information under the assumption that the strategy profiles are generated by truth-telling (TT). Consistent with Rustichini et al. (1994), the p-values for $\varepsilon^*(\text{TT})$ are sharply increasing in n: the hypothesis $\varepsilon^*(\text{TT}) = 0$ for $n = 1$ is rejected for any reasonable level of confidence; on the other hand, this hypothesis goes unscathed for $n = 100$. A simpler way to summarize the evidence is to focus on economic relevance and look at the order magnitude of ε^*: for any n under GA and for $n = 100$ under TT, this is measured in thousandths; on the other hand, for $n = 1, 10$ under TT, it is measured in hundredths. The strength of the average temptation to deviate comes out sizeably different. The right panel of Figure 3 provides details for the values of

ε for the buyers' valuations; note the different scale for the y-axis. (Sellers' costs lead to a similar graph.) For $n = 1$, there is an approximately monotonic relationship between the value of ε and the strength of a trader.

5 Asynchronous order clearing: Results

This section presents the outcome of our simulations for the case of asynchronous order-clearing. We first give evidence that our GA converges to a steady state that we dub the "equilibrium" outcome; see Dawid (1999) for a similar approach in an environment simpler than ours. Next, we describe and evaluate the trading strategies evolved by our simulations. Section 6 elaborates on how these strategies compare with another class of equilibria provided in the literature.

5.1 Transaction prices and profits

We begin with a quick look at the effects of evolving trading strategies on some fundamental parameters of the market. Our first item is the time series of transaction prices. The left panel in Figure 4 shows a typical sample of realized prices for the baseline sampled at days $t = 100, 500, 1000, 1500, 2000$ and 5000. Note that the width of the vertical bands is proportional

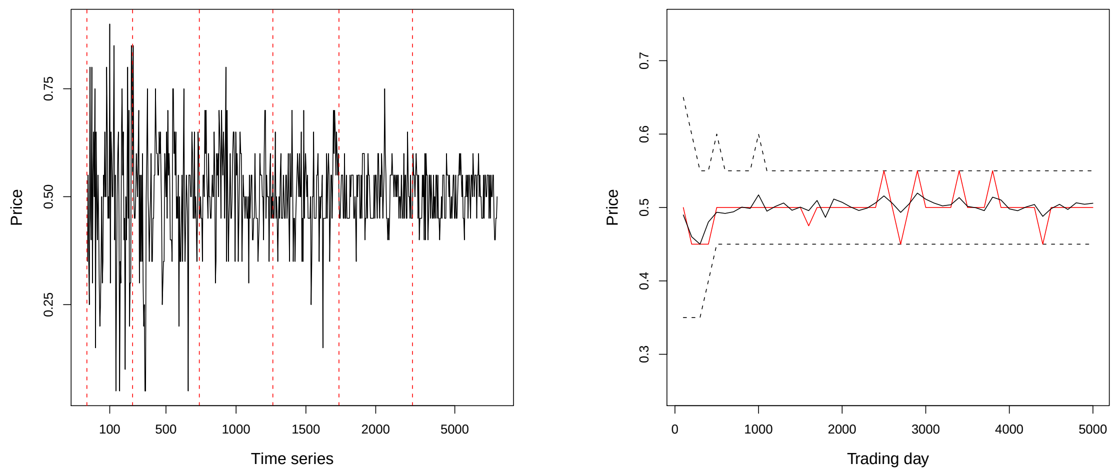

Figure 4: On the left, realized transaction prices at fixed times. On the right, average (black), median (red) and interquartile range (dashed) of the transaction prices for each period.

to the number of transactions. Therefore, as learning progresses, trading prices become less volatile and volume increases. Moreover, the transaction prices approach the competitive price p^* in the long run. The right panel exhibits the average (in black) and the median daily price (in red), together with the first and third quartile of the distribution of the intraday transaction prices.

Result 1 *In a continuous double auction under uniform priors,*[3] *the evolution of trading strategies stabilizes prices around the competitive price* p^*.

Furthermore, the amount of dispersion of prices around p^* is a decreasing function of the number n of traders in the market as shown in Table 3. This is aligned with well-known

n	μ	σ
1	0.506	0.124
10	0.497	0.053
100	0.499	0.014

Table 3: Average price μ and standard deviation σ of the intraday price for different markets.

general convergence results for trading protocols based on the alternative assumption of simultaneous order clearing; see Rustichini et al. (1994) for the k-double auction and Mendelson (1985) for Walrasian trading. On the other hand, it is understood that, *ceteris paribus*, there is higher variability in the price of the continuous double auction as a consequence of the assumptions of asynchronous order clearing and random arrival of traders.

Result 2 *If* $n_1 > n_2$, *the distribution of the transaction price* $P(n_2)$ *is more diffuse than* $P(n_1)$.

Similarly reassuring results hold for the aggregate profits. The left panel in Figure 5 shows the average daily gain per trader for a typical simulation run. Traders, who initially enter the marketplace making individually rational but otherwise random offers, coevolutively learn to extract much higher profits from trade. In turn, as shown in the right panel, this leads to an increase in the volume of transactions effected within a trading day.

We conclude that evolution leads to trading strategies that are jointly fit to maximize overall gains from trade. However, by construction, the GA does not attempt to maximize these latter ones. Each agent strives to learn and improve his private gains within the group of traders of his own type while the environment is coevolutively changing. Hence, the globally successful extraction of the trading surplus is a byproduct of the joint and unrelated maximizations of individual profits.

Figure 6 separately exhibits how the average gain per type evolves over time for buyers and sellers. The left panel shows the average profit of buyers with valuations in the set $\{0.95, 0.9, 0.8, 0.7, 0.6, 0.5, 0.45\}$ arranged from the strongest ($v = 0.95$) to the weakest ($v = 0.45$) as we move downwards. Analogously, the right panel depicts the average profit of sellers with costs in $\{0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 0.55\}$ arranged from the strongest ($c = 0.05$) to the weakest ($c = 0.55$). Profits are role-symmetric (up to the inevitable noise). Gains are increasing in valuations and decreasing in costs: stronger traders realize higher gains from trade, while marginal traders (with valuations and costs close to p^*) reap minute profits. Indeed, zooming in on the bottom lines of Figure 6 shows that gains shrink over time for the marginal types. This is due to the gradual improvement of the strategies used by the strong intramarginal traders that makes them much less susceptible to being exploited by weaker marginal agents.

[3] For brevity, we leave this qualification implicit in the statement of similar following results.

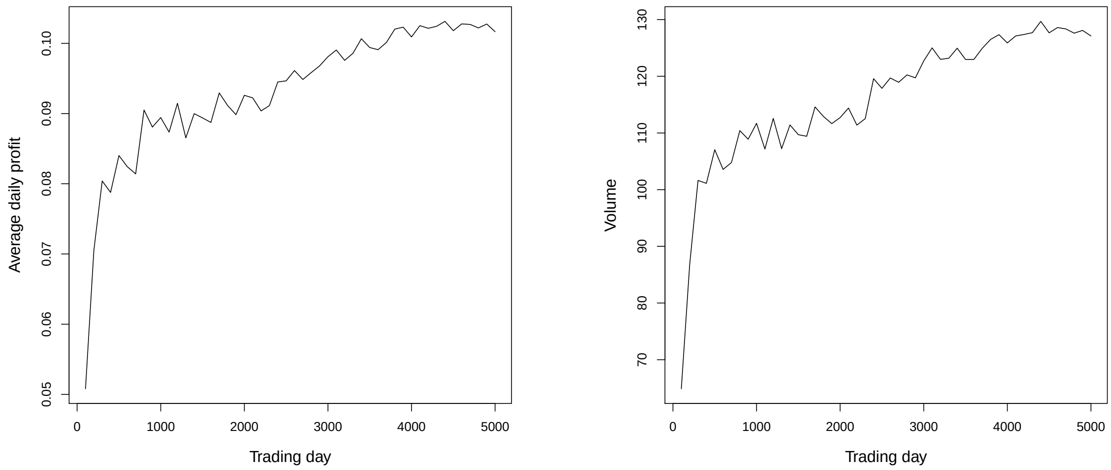

Figure 5: Traders' average profit (left) and transaction volume (right) for each period.

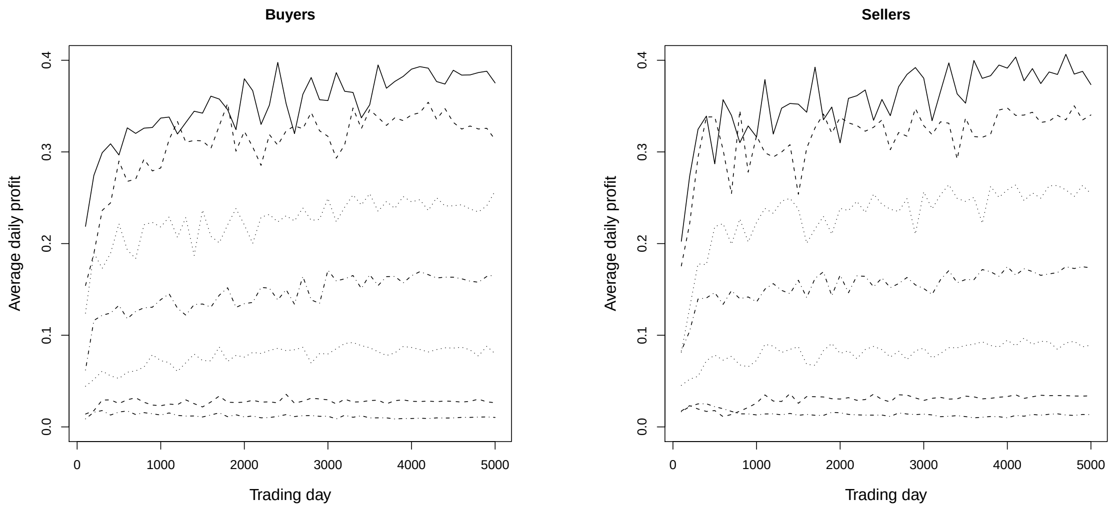

Figure 6: Average daily gains for buyers (left) and sellers (right). Groups of traders are arranged from top to bottom according to their strength.

Result 3 *The evolution of trading strategies improves aggregate profits and is more beneficial for stronger traders.*

5.2 Trading strategies

The genetic algorithm mimics an attempt to learn equilibrium strategies. Proceeding as in Section 4, we present our results. Moving from left to right, Figure 8 displays a bundle of three "sampled" bidding and asking functions for $n = 1, 10, 100$. As before, the competitive price $p^* = 0.5$ is represented by a dashed line. There is a sharp dependence of the shape of

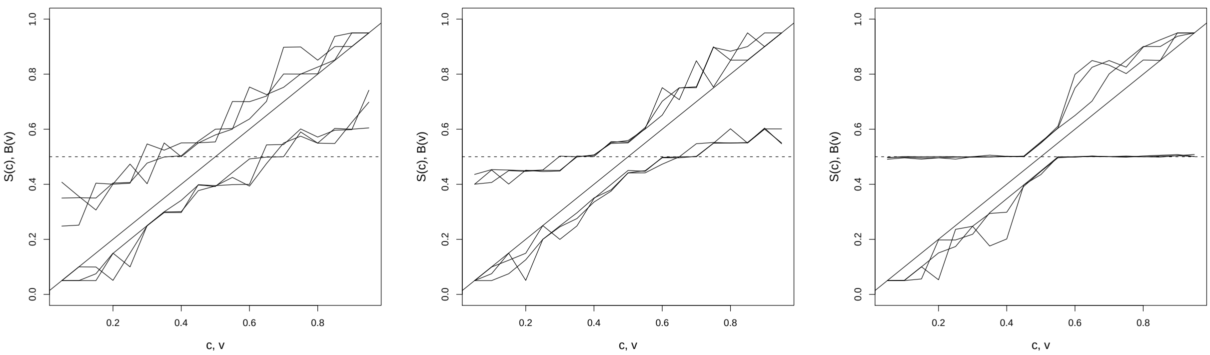

Figure 7: Bundles of trading functions for $n = 1, 10, 100$ (from left to right).

the trading strategies on the number n of agents. The pattern is even more transparent if we remove noise and look at the "average" trading functions, reported in Figure 8. As the

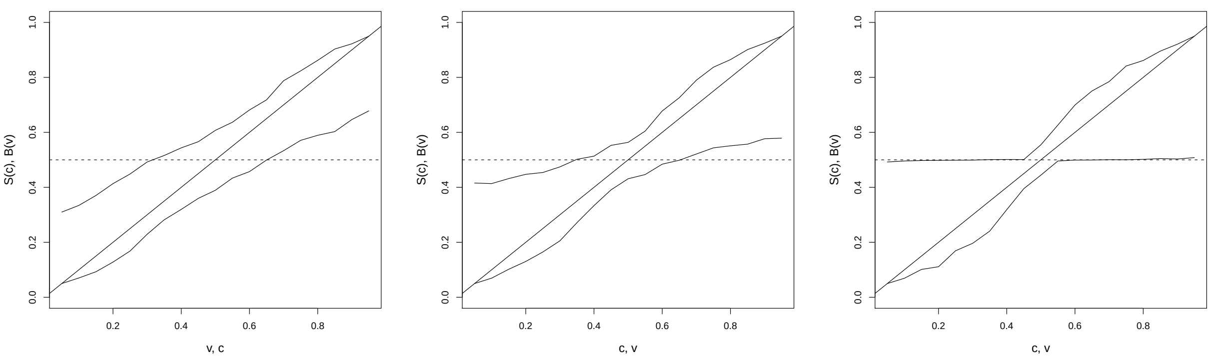

Figure 8: Average trading functions for $n = 1, 10, 100$ (from left to right).

market gets thicker, all intramarginal traders become less aggressive until (as clearly visible for $n = 100$) they almost always offer exactly the competitive price p^*. On the other hand, the variability of the bidding and asking functions for extramarginal traders is quite large and increasing in n: bids (asks) for $v < 0.5$ ($c > 0.5$) in Figure 7 fluctuate a lot. The simple explanation is that extramarginal agents trade very rarely and, consequently, get almost no chances to learn: the GA has no push to learn because it is pointlessly trying to optimize a zero-profit constant function. The contrast between the rightmost panels from Figure 2 and Figure 8 is particularly striking: under simultaneous order clearing, only marginal traders get very strong incentives to learn, while under asynchronous order clearing learning takes place for *all* intramarginal traders. Moreover, these latter ones learn to be *price-makers* and offer a price equal to the competitive price p^*.

Indeed, a closer look at the trading functions confirms that all intramarginal agents learn to play a pure strategy. Up to the inevitable noise inherent to the GA, each of them ends up making an offer equal to p^*. Figure 9 displays a typical set of genotypes (probability vectors) for the whole population when $n = 100$. The box on the left panel naturally divides into two triangles. The bottom triangle pertains to buyers; the top triangle to sellers. We describe only the buyers' side; the analog holds for the sellers'. Probabilities are color-coded with (bright) yellow and (dark) red meaning 1 and 0, respectively. The colors along the vertical segment between the x-axis and the bisector represent the randomized strategy of a buyer of type v over his bids. For instance, consider the vertical segment at $v = 0.8$: the only bid with non-null probability (yellow) is $p^* = 0.5$, while there is zero mass (red) on all the other choices in O.

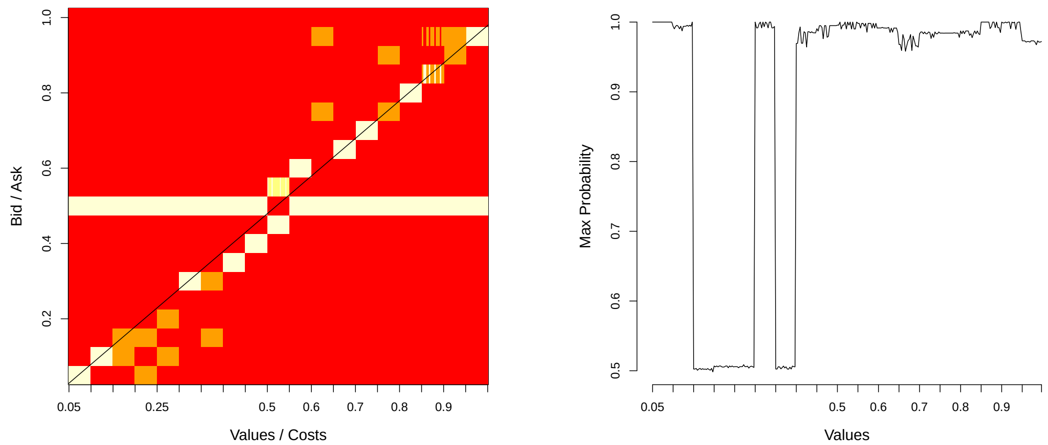

Figure 9: Left panel: randomized strategies as a color-coded map with hues representing different values of the probabilities. Right panel: plot of the modal value of the probability distribution for each type of buyer when $n = 100$.

The diffusion of the red color demonstrates that positive probability is attached to a small handful of bids and asks. For $n = 100$, the right panel of Figure 9 plots $\max_{b \in O} q_i(v, b)$ for each buyer i by adjoining values around v. Whenever this probability is close to 1, the corresponding agent is using a pure strategy. (It is understood that negligible noise is an intrinsic feature of GA.) All intramarginal traders offer the same price p^*. Extramarginal traders may end up with a nontrivial randomized strategy, but the previous discussion has explained that this is ultimately irrelevant.

Result 4 *As n increases, the trading strategies of the intramarginal agents move towards price-making: they learn to make a constant offer equal to the competitive price.*

This is our most important result. *Ceteribus paribus*, the details of the order clearing rule for the market protocols determine what kind of trading strategies we expect agents to learn. Regardless of whether order clearing is simultaneous or asynchronous, as the market grows in size, there is convergence to the competitive outcome where (almost) all trades take place at price p^*. However, the trading strategies used by agents evolve in very different directions. Under simultaneous order clearing, marginal traders make truthful offers which set up the price for everybody else and this deprives stronger intramarginal traders of the incentives to learn and be price takers. Under asynchronous order clearing, all intramarginal traders end up acting as price-makers.

6 Comparative analysis

This last section compares the equilibrium strategies evolved by the genetic algorithm against the current benchmark in the literature, as provided by Zhan and Friedman (2007) or ZF for short. They study the continuous double auction in an environment like ours, under the same assumptions of Section 2.3. However, they restrict the choice of the bidding and asking functions to subsets of functions that are easily interpreted as "markdowns" on v and "markups" on c. (For simplicity, we usually speak only of markups.) They look at three possibilities, but the leading example is the family of *standard markups* defined by

$$\beta(v) = v(1 - m_d) \qquad \text{and} \qquad a(c) = c(1 + m_u), \tag{3}$$

where $m_d, m_u \geq 0$ are the markdown and markup coefficients.

The major difference between GA and ZF is apparent. ZF search for equilibria in pure strategies within a (restricted) parametric class and do this by running an exhaustive search; GA search for equilibria in randomized strategies within a (general) nonparametric class by a non-exhaustive evolutionary approach. A second subtler but important difference is that ZF consider *ex ante* equilibria, where the optimality of a strategy is evaluated assuming that the trader does not know his own type and hence must look at the expected payoff over all his private types. Instead, following standard game-theoretic practice, we derive *interim* equilibria where each agent evaluates payoffs using knowledge of his own private type (but not others'). In other words, ZF constrains all types of a player to adopt the same markup; we allow each type of a player to pick his own markup.

Using computer simulations, ZF searches for *ex ante* equilibria[4] within the class of standard markup strategies and finds unique pairs of equilibrium coefficients: $m_d = m_u = 0.3$ for $n = 10$, and $m_d = 0.4, m_u = 0.3$ for $n = 100$. (We have independently confirmed this result.) ZF does not study the case $n = 1$; using their methodology, we obtain $m_d = m_u = 0.3$ as

[4] They look at two cases: cartels and single players. We consider only this latter, as it is more appropriate.

unique equilibrium values. (More precisely, we find that this constitutes an ε-equilibrium with $\varepsilon = 0.00002$.)

ZF considers the performance of these markup equilibrium strategies with respect to allocative inefficiency and traders' surpluses. (The *allocative inefficiency* is defined as the complement to 1 of the ratio between the realized surplus and the maximum attainable surplus.) Table 4 compares the allocative inefficiency incurred by GA and ZF in a continuous double auction. It is apparent that allocative inefficiency declines as the market grows in

n	1	10	100
GA	0.196	0.084	0.048
ZF	0.202	0.115	0.052

Table 4: Allocative inefficiencies for the continuous double auction.

size, regardless of whether trading strategies are obtained by GA or ZF.

Result 5 *Allocative inefficiency is decreasing in market size. A plausible conjecture is that it vanishes as* $n \uparrow \infty$.

A quick look confirms that the equilibrium strategies in ZF yield an allocative inefficiency slightly superior than those evolved by GA, but are still within the same order of magnitude. This may suggest that ZF is in some respect a comparable approach. We are going to argue that this impression is deceptive, because it fails to properly account for the strategic issues that underlie traders' behavior. Equilibrium strategies *per se* do not attempt to maximize overall gains from trade. Allocative efficiency is only a byproduct of the individual effort to maximize private gains. Hence, an equilibrium should be judged by its strategic plausibility rather than by its collateral social effects.

Our claim is that ZF has two serious shortcomings from a strategic point of view. The first one is that it imposes an unjustified asymmetry between buyers and sellers. Figure 10 displays the average trading gains per type using GA or ZF strategies, for $n = 1, 10, 100$. The first panel plots average gains for buyers' types under GA. (We omit the panel for sellers because it is virtually identical.) The second and third panel show the average gains

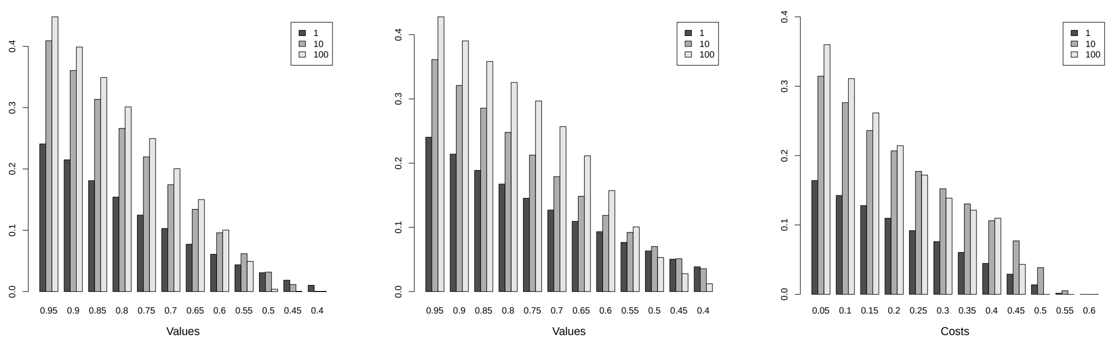

Figure 10: From left to right: average gains for buyers (GA), buyers (ZF) and sellers (ZF).

using ZF's equilibrium strategies for buyers' and sellers' types, respectively. To help a direct comparison, the scale on the y-axis is the same for the three panels and traders' types on the x-axis are always arranged by decreasing strength. As expected, stronger traders reap higher gains. However, it is apparent that ZF's trading strategies let buyers extract more surplus than sellers of equal strength, in spite of the perfect balance in the assumptions over the trading environment.

This asymmetry in outcomes is due to a flaw in the definition of the standard markup strategies. Define the strength of a bid b as its distance from the lowest bid ($b = 0$); similarly, let the strength of an ask a be its distance from the highest ask ($a = 1$). A profile of trading strategies is *role-symmetric* when the strengths of the bid and the ask issued by traders of equal strength x are the same; that is, when $\beta(x) = 1 - \alpha(1 - x)$. Cervone et al. (2009) shows that ZF's markup strategies are not role-symmetric and hence lead to equilibrium outcomes that favor buyers over sellers. To overcome this limitation, it introduces a role-symmetric formulation called *convex markup* and illustrates its advantages by means of a comparison over different market protocols.

The second shortcoming in ZF is perhaps more substantial. Similarly to Table 2 for the call market, Table 5 displays values for $\varepsilon^*(\text{GA})$ and for $\varepsilon^*(\text{ZF})$, when $n = 1, 10, 100$. For

	n	ε^*	type	offer	p-value
GA	1	0.0034	$v = 0.90$	0.65	0.62
	10	0.0054	$c = 0.15$	0.45	0.0655
	100	0.0062	$c = 0.05$	0.50	0.0131
ZF	1	0.0481	$c = 0.05$	0.20	10^{-11}
	10	0.0776	$c = 0.05$	0.40	10^{-16}
	100	0.0771	$c = 0.05$	0.50	10^{-15}

Table 5: Quality control for the trading strategies evolved by GA in a CDA.

trading strategies evolved by GA, these values are much closer to zero than for ZF. The strength of the average temptation to deviate is sizeably different. At a confidence level of 5%, for $n = 1$ and $n = 10$ we cannot reject the null hypothesis $\varepsilon^*(\text{GA}) = 0$. The same holds for $n = 100$, but at a confidence level of 1%. On the contrary, for $n = 1, 10, 100$, the null hypothesis $\varepsilon^*(\text{ZF}) = 0$ is rejected for any practical level of confidence.

Furthermore, the trader who has the highest incentive to deviate turns out to be always the strongest seller (who has cost $c = 0.05$). Besides the lack of role-symmetry hurting sellers, there is a second effect at play.[5] ZF's trading strategies systematically expect stronger sellers to ask too little. The equilibrium markup for sellers under ZF is $m_u = 0.3$ for $n = 1, 10, 100$. Hence, a seller with cost $c = 0.05$ should ask (on average) $a(0.05) = 0.05 \cdot 1.3 = 0.065$. Instead, as shown in the fourth column, the optimal ask is remarkably higher.

Moving from left to right, Figure 11 shows the values of ε for the sellers' under GA (the graph for buyers is analogous and hence omitted), as well as for buyers and sellers under ZF. After noting that the scales on the y-axes are different, we see again that there is an order

[5] In the language of mechanism design, ZF's trading strategies from (3) are not incentive-compatible for the sellers. The convex markup rule in Cervone et al. (2009) is incentive-compatible.

of magnitude of difference between the ε^* values for GA and ZF.

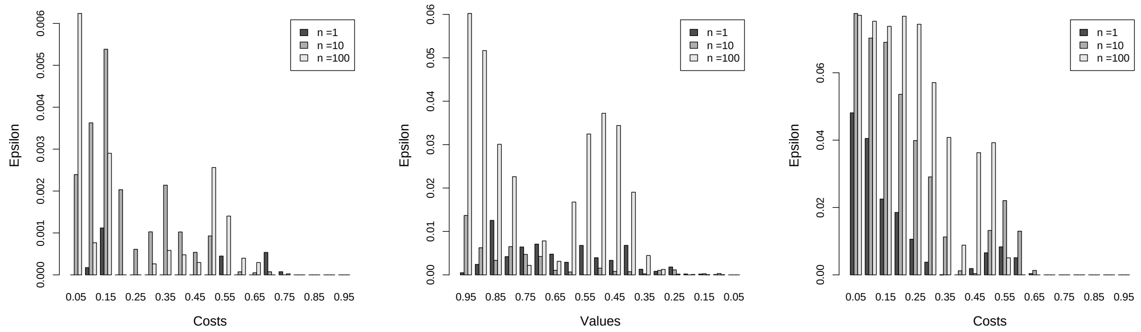

Figure 11: From left to right: values of ε^* for sellers (GA), buyers (ZF) and sellers (ZF).

6.1 Evolutionary stability

This last subsection rounds up our comparison between GA and ZF using the notion of evolutionarily stable strategies (ESS) introduced in Maynard Smith and Price (1973). Roughly speaking, a profile of trading functions is an ESS if, once it is adopted, it is not susceptible to invasion by a new strategy. This notion formalizes the "robustness" of a strategy profile as the ability to prevent the spread of competing alternatives under evolutionary pressure. We claim that ZF does not pass the test of evolutionary stability against GA.

We consider a population of agents using ZF's trading strategies and inject a small fraction of traders playing the strategies evolved by GA. We measure the average profits made by agents using the two competing strategies. For $n = 1, 10, 100$, Figure 12 depicts the level curves for the joint distribution of average daily gains obtained by traders in the invading GA population versus those collected by the ZF agents. (The fraction of invading GA traders is set equal to 1/16 and we collect realized profits over 1000 trading days.)

As n increases, it is apparent that the advantage enjoyed by GA traders gets stronger. Hence, the prospect of invadability for ZF increases when the market grows in size. The intuitive explanation is clear. The ZF equilibrium trading strategies are linear functions. As shown in Figure 7, for $n = 1$ the GA strategies can be decently approximated by a linear function. Therefore, the scope for differentiation between ZF and GA is limited. On the other hand, as n increases, the GA strategies morph towards a flat price-making offer (for the intramarginal traders) that is sharply different from ZF's prescription.

The evolutionary pressure can be quantified by looking at the (average) incremental profit per trading day reaped by the invading GA agents pitted against ZF traders. For $n = 100$, we find a value of 0.010 as sum of an average incremental profit of -0.005 for buyers and 0.025 for sellers. A "naive" statistician using a t-test for the overall value would require only about 30 trading days to reject the null hypothesis that the GA strategy is less profitable than ZF at a level of confidence of 0.1%. (Rejection would be even faster by testing only for the sellers' side.) Similarly, for $n = 10$, the evolutionary advantage to GA is 0.009 (buyers:

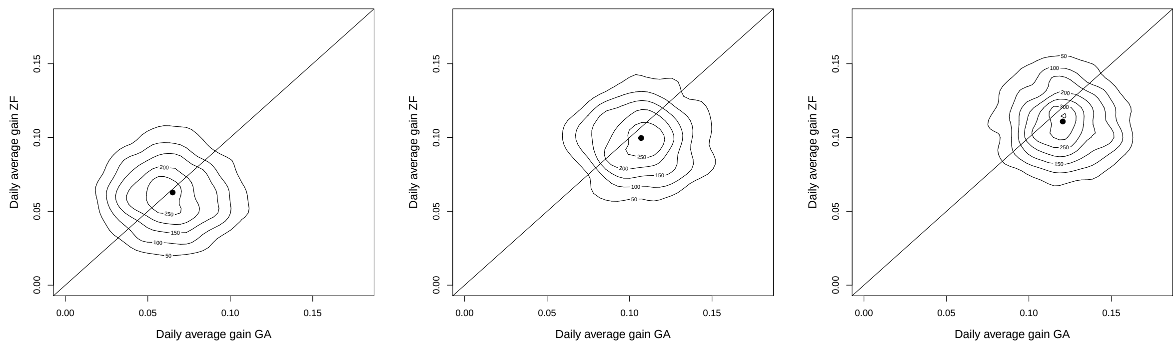

Figure 12: Joint distribution of the average daily gains for GA agents (x-axis) invading a ZF population (y-axis) for $n = 1, 10, 100$ (from left to right). The mean of the distribution is plotted as a thick point, along with the line of equal average profits.

-0.002; sellers: 0.020); a statistically significant rejection at the same level of confidence would occur after about 550 days. For $n = 1$, the average incremental profit for GA is 0.002 (buyers: -0.003; sellers: 0.007); more than 1000 trading days would be necessary for an analogous statistical rejection.

Additional intuition is gained by considering Figure 13 that for $n = 100$ separately shows the average incremental profits of buyers (on the left) as a function of their valuations and of sellers (on the right) as a function of their costs. (As before, the fraction of invading GA

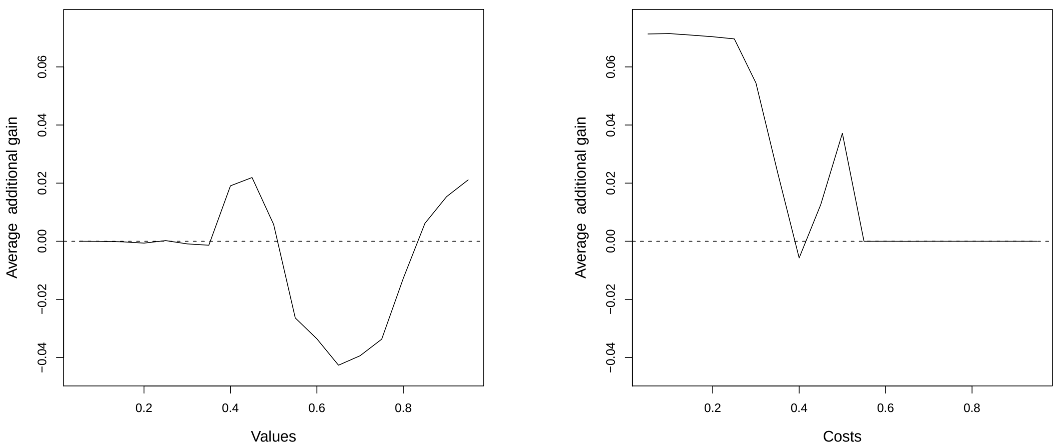

Figure 13: Average incremental profits for GA invading ZF: buyers (left) and sellers (right).

traders is 1/16 and we compare profits realized over 1000 trading days.) Marginal and very strong buyers perform slightly higher under GA, while intramarginal buyers of intermediate strength are much better off sticking with ZF. On the other hand, almost all intramarginal sellers find GA more profitable than ZF and thus the stronger evolutionary pressure to break free from ZF and shift towards GA lies on the sellers' side.

A The bilateral trading model

The bilateral trading model (with simultaneous order clearing) was introduced in Chatterjee and Samuelson (1983) and studied in Myerson and Satterthwaite (1983), spawning a long and still flourishing literature. It describes a situation where one buyer and one seller are engaged in the trade of a single object. The environment is the same as in Section 2.1, but there are only one buyer and one seller so $n = 1$. Viewed as a game with incomplete information, it is equivalent to the k-double auction described in Section 2.2 for $k = 1/2$ and $n = 1$.

An equilibrium profile (β, α) of bidding and asking functions for the bilateral trading model requires that a buyer with valuation v offers a bid $b = \beta(v)$ that solves

$$\max_b \int_0^1 \left[v - \frac{\alpha(c) + b}{2} \right] \mathbf{1}\{b \geq \alpha(c)\} \, \mathrm{d}G(c) \tag{4}$$

and a seller with cost c submits an ask $a = \alpha(c)$ that solves

$$\max_a \int_0^1 \left[\frac{a + \beta(v)}{2} - c \right] \mathbf{1}\{a \leq \beta(v)\} \, \mathrm{d}F(v), \tag{5}$$

where $\mathbf{1}\{\cdot\}$ is the indicator function.

The formal description of the bilateral trading model under asynchronous order clearing requires only the following changes. There are two equally likely queues: the buyer arrives first, or the seller arrives first. If a buyer with valuation v arrives first, he finds no outstanding ask and hence records his bid $b = \beta(v)$ on the buying book and, if trade occurs, the price is b. Similarly, if a seller with cost c arrives first, she writes her ask $a = \alpha(c)$ on the selling book and, if trade occurs, the price is a. Therefore, if trade occurs, it takes place at price b with probability $1/2$ and at price a with probability $1/2$. Recall that, under simultaneous order clearing, the transaction price is $p = (a + b)/2$. Therefore, roughly speaking, the expected value of the trading price is the same under simultaneous or asynchronous order clearing but the latter one adds some variability around it. (Formally speaking, the distribution of the trading price is a mean-preserving spread.)

Given that traders are risk neutral, this immediately translates in the strategic equivalence of the equilibria for the two models. We prove this claim by showing that the expected payoffs under any strategy profile (β, α) are the same for the two models. In the bilateral trading model under asynchronous order clearing, a buyer with valuation v who offers a bid $b = \beta(v)$ obtains a payoff

$$\max_b \frac{1}{2} \int_0^1 \big[v - \alpha(c) \big] \mathbf{1}\{b \geq \alpha(c)\} \, \mathrm{d}G(c) + \frac{1}{2} \int_0^1 \big[v - b \big] \mathbf{1}\{b \geq \alpha(c)\} \, \mathrm{d}G(c). \tag{6}$$

In the bilateral trading model under simultaneous order clearing, the payoff to a buyer with valuation v who offers a bid $b = \beta(v)$ is given in Equation (4). Clearly, (4) and (6) are identical. A similar argument applies for the seller. Therefore, the set of equilibria under arbitrary priors for the two models is the same. More generally, a similar argument applies to show the strategic equivalence of the equilibria for a k-double auction and a continuous double auction where the buyer arrives before the seller with probability k. Note that, while strategic equivalence holds, equilibrium payoffs coincide only in expectation.

References

[1] M. Anufriev, J. Arifovic and V. Panchenko (2010), "Efficiency of continuous double auctions under individual evolutionary learning with full or limited information", CeNDEF Working paper 10-01, University of Amsterdam.

[2] D. Ashlock (2006), *Evolutionary Computation for Modeling and Optimization*, Springer.

[3] R. Aumann, "Game theory", in: J. Eatwell, M. Milgate, and P. Newman (eds.), *The New Palgrave: A Dictionary of Economics*, volume 2, Macmillan, 1987, 460–482.

[4] R. Cervone, S. Galavotti, and M. LiCalzi, "Symmetric equilibria in double auctions with markdown buyers and markup sellers", in: C. Hernández, M. Posada and A. López-Paredes (eds.), *Artificial Economics*, Springer, 2009, 81–92.

[5] K. Chatterjee and W. Samuelson (1983), "Bargaining under incomplete information", *Operations Research* **31**, 835–851.

[6] H. Dawid (1999), "On the convergence of genetic learning in a double auction market", *Journal of Economic Dynamics and Control* **23**, 1545–1567.

[7] D. Friedman (1993), "The double auction institution: A survey", in: D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories and Evidence*, Cambridge (MA): Perseus Publishing, 3–25.

[8] D.K. Gode and S. Sunder (1993), "Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality", *Journal of Political Economy* **101**, 119–137.

[9] W. Leininger, P.B. Linhart, and R. Radner (1989), "Equilibria of the sealed-bid mechanism for bargaining with incomplete information", *Journal of Economic Theory* **48**, 63-106.

[10] M. Lettau (1997), "Explaining the facts with adaptive agents: The case of mutual fund flows", *Journal of Economic Dynamics and Control* **21**, 1117–1147.

[11] J.K. MacKie and M.P. Wellman (2006), "Automated markets and trading agents", in: L. Tesfatsion and K.L. Judd (eds.), *Handbook of Computational Economics*, vol. 2, Elsevier, 1381–1431.

[12] J. Maynard Smith and G.R. Price (1973), "The logic of animal conflict", *Nature* **246**, 15–18.

[13] H. Mendelson (1985), "Random competitive exchange: Price distributions and gains from trade", *Journal of Economic Theory* **37**, 254-280.

[14] R.B. Myerson and M.A. Satterthwaite (1983), "Efficient mechanisms for bilateral trading", *Journal of Economic Theory* **29**, 265–281.

[15] S. Phelps, S. Parsons, and P. McBurney (2006), "An evolutionary game-theoretic comparison of two double-auction market designs", in: P. Faratin and J.A. Rodriguez-Aguílar (eds.), *Agent-Mediated Electronic Commerce VI*, Springer, 101–114.

[16] A. Rustichini, M.A. Satterthwaite, and S.R. Williams (1994), "Convergence to efficiency in a simple market with incomplete information", *Econometrica* **62**, 1041-1063.

[17] M.A. Satterthwaite and S.R. Williams (1993), "The bayesian theory of the k-double auction", in: D. Friedman and J. Rust (eds.), *The Double Auction Market: Institutions, Theories and Evidence*, Cambridge (MA): Perseus Publishing, 99–123.

[18] M.A. Satterthwaite and S.R. Williams (2002), "The optimality of a simple market mechanism", *Econometrica* **70**, 1841-1863.

[19] W. Zhan and D. Friedman (2007), "Markups in double auction markets", *Journal of Economic Dynamics and Control* **31**, 2984–3005.