

Learning to Trade in an Unbalanced Market

Florian Hauser and Marco LiCalzi

Abstract We study the evolution of trading strategies in double auctions as the size of the market gets larger. When the number of buyers and sellers is balanced, Fano et al. [2] show that the choice of the order-clearing rule (simultaneous or asynchronous) steers the emergence of fundamentally different strategic behavior. We extend their work to unbalanced markets, confirming their main result as well as that allocative inefficiency tends to zero. On the other hand, we discover that convergence to the competitive outcome takes place only when the market is large and that the long side of the market is more effective at improving its disadvantaged terms of trade under asynchronous order-clearing.

1 Introduction

Recently, Fano et al. [2] have studied the evolution of trading strategies for a double auction when the number of traders increases. They provide two main results. First, the competitive outcome obtains under different market architectures, provided that the size of the market is sufficiently large. Second, the choice of the order-clearing rule affects trading behavior. Under simultaneous order-clearing, marginal traders learn to act as *price takers* and make offers equal to their valuations or costs. Under asynchronous order-clearing, the intramarginal traders learn to act as *price makers* and make offers equal to the competitive price.

An important feature of their study is the assumption that buyers and sellers populate the market in equal numbers; that is, the market is *balanced*. Moreover, using a simulative approach, agents' learning is modeled by means of a genetic algorithm. Curiously enough, although it is well known that “traders on the long side of a market wind up holding the short end of the stick” [8], the agent-based literature has paid scant attention to the study of unbalanced markets; however, see Gode and Sunder [5] or Anufriev et al. [1] for the special case of one seller, one intramarginal buyer and n extramarginal buyers.

This paper explores what happens in double auctions when we remove the assumption that the market is balanced. A second methodological contribution is that we model agents' learning by means of genetic programming, extending the reach of genetic algorithms. We confirm and extend the main results in Fano et al. [2] when the market is balanced.

Florian Hauser

Dept. Banking and Finance, Universität Innsbruck, e-mail: florian.hauser@uibk.ac.at

Marco LiCalzi

Dept. Management, Università Ca' Foscari Venezia, e-mail: licalzi@unive.it

More importantly, we discover two novel effects when the market is unbalanced. Under simultaneous order-clearing, the convergence towards price-taking behavior expands from the marginal traders to the intramarginal traders on the long side, due to the competitive pressure. A second result is about the interaction of strategic behavior and order-clearing rule. When traders are forced to act as price takers, the disadvantageous terms of trade for the long side are worse under asynchronous order-clearing because the price is not uniform across all transactions. However, when we let traders optimize their strategies for the protocol in use, asynchronous order-clearing improves the average transaction price for the long side. In short, by forcing a uniform price over all transaction, simultaneous order-clearing dampens the effects due to strategic behavior.

2 The Model

There are $n = b + s$ traders, where b are buyers and s are sellers. Each trader wishes to maximize expected profits and is in the market to exchange at most one unit of a generic good per round. 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 uniform distributions on $[0, 1]$. Profits are $v - p$ for buyers and $p - c$ for sellers, where p is the price at which a transaction occurs.

We speak of intramarginal and extramarginal buyers (sellers) according to their position on the demand (supply) function with respect to the market-clearing price(s). We refine this qualitative distinction into a complete ordering, by defining 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 ($c = 1$). Hence, stronger intramarginal traders have valuations (or costs) lying farther away from the market-clearing price(s).

Agents trade using one of two distinct trading protocols, whose main difference is in the nature of the order-clearing rule. Under simultaneous order-clearing, we have a *call market* where the offers made by buyers and sellers are aggregated to form their demand and supply functions and all transactions take place at a unique price (that we choose to be the mid-value of the interval of market-clearing prices).

Under asynchronous order-clearing, we have a *continuous double auction* where agents arrive in random order and sequentially submit binding¹ 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 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.

¹ See LiCalzi and Pellizzari [6] on the role of an assumption of binding offers.

Agents from a pool of N traders interact repeatedly and anonymously. In each *round*, we randomly draw (without replacement) b buyers and s sellers from the pool of available agents and let them visit the market ($b + s = n \ll N$). Hence, the agents active in the market during a specific round change randomly over time. This increases variety and learning opportunities.

We constrain traders to make offers that are *individually rational*: no buyer can bid more than his valuation and no seller can ask less than her cost. We evolve traders' strategies using genetic programming, henceforth nicknamed GP for brevity. This optimization method is an alternative approach to the genetic algorithm used in Fano et al. [2]. The design of our GP routines follows closely the standard tree-based approach described in Koza [4]. Evolution takes place sequentially over at least 2500 optimization steps.² In each step, one trader is randomly selected from the pool and his trading function is optimized by GP. As the complexity of the optimization problem is rather low, we choose conservative parameters for the GP algorithm.

An initial population of 100 trading strategies is generated by the ramped-half-half method with a maximum number of 100 nodes. Input data for a strategy is v or c , as well as a set of constants (0.01, 0.1, 2, 3, 7). To modify the input data, GP can refer to a set of standard arithmetic functions (+, -, \times , \div , exp, ln). The initial population of strategies is improved by an evolutionary process over 50 generations. To generate a member of a new generation, GP selects successful strategies from the previous generation based on a tournament of size 4 and applies crossover and mutation with probability 0.8 and 0.2, respectively. The fitness function is the average profit of a strategy over 30 rounds. The best strategy from the last generation is selected and the trader stays with it until he is chosen for a further optimization step.

We report results over a fixed experimental design with nine cells, where buyers are on the longer side of the market. Each cell corresponds to a different traders' population, where the number of sellers is equal to $s = 1$, $s = 5$, and $s = 50$, against a number of buyers equal to $b = s$, $b = 3s$, and $b = 5s$. We usually arrange the nine cells in a 3×3 matrix, where the columns correspond to $s = 1$ (left), $s = 5$ (center) and $s = 50$ (right), and the rows to $b = s$ buyers (top), $b = 3s$ (middle), and $b = 5s$ (bottom). In each experiment, the pool of agents learns to trade using the GP as described above. At the end of the learning process, for each cell we calculate results over 1000 additional rounds of trading.

3 Convergence to the Competitive Outcome

The *competitive price* obtains under simultaneous order-clearing when all agents truthfully report their costs and their valuations. We denote this by S-TT for brevity.

² We go up to 6000 steps for large markets with many agents.

(The rest of our shorthand notation is “A” for the asynchronous order-clearing rule and “GP” for the trading strategies developed by genetic programming.)

Generally speaking, we find that the trading price rapidly stabilizes around the competitive price both for S-GP and A-GP. This convergence is faster under S-GP. To exemplify, we compare in [Figure 1](#) the average trading price (in black) and its standard deviation (in grey) for the case $s = 50$ and $b = 5s$ as a function of the optimization steps. (The statistics are computed over 1000 rounds of trading while freezing the GP process.) It is apparent that, compared to S-GP, it takes more time under A-GP before enough learning takes place to stabilize prices.




Fig. 1 Average price (black) and standard deviation (grey) under simultaneous (left) and asynchronous (right) order clearing for the case $s = 50$ and $b = 5s$.

More detailed information appears in [Table 1](#), where for each cell we report in the first line the (average) competitive equilibrium price computed using S-TT (left) and the average trading price under S-GP (center). The second line in each cell gives the average trading price under A-TT (left) and A-GP (center). To control for the accumulation of orders in the book, we also give the average closing price on the right of the second line in each cell, using data from the last transaction in each round. Averages are taken over all the transactions from 1000 rounds per at least 10 distinct simulations with different random seeds.

	$s = 1$			$s = 5$			$s = 50$		
$b = s$	0.498	0.496		0.501	0.501		0.500	0.499	
	0.495	0.494	0.494	0.503	0.500	0.499	0.500	0.500	0.500
$b = 3s$	0.601	0.698		0.690	0.723		0.746	0.748	
	0.638	0.630	0.630	0.735	0.705	0.710	0.807	0.744	0.745
$b = 5s$	0.647	0.774		0.763	0.806		0.829	0.831	
	0.716	0.703	0.703	0.818	0.789	0.794	0.882	0.829	0.830

Table 1 Each cell exhibits the average transaction price for S-TT (top left), S-GP (top center), A-TT (bottom left), A-GP (bottom center), as well as the average closing price for A-GP (bottom right).

The close alignment of the values in the cells of the first line (where $b = s$ and $s = 1, 5, 50$ as we move rightward) confirms the result in Fano et al. [2]: in a balanced market, the evolution of trading strategies stabilizes prices around the competitive price p^* , particularly when the market grows large. On the other hand, the second and third horizontal lines (with $b = 3s$ and $b = 5s$, respectively) show a mismatch between the price under the baseline case of S-TT (top-left) and the prices under S-GP or A-GP in the central column of the four cells with $s = 1$ or $s = 5$ in the bottom-left corner. This mismatch is supported for any practical level of confidence by a two-sided paired t -test. (By a *practical* level of confidence we mean a p -value lower than 10^{-5} .)

We conclude that there is no convergence of the realized prices to the competitive price p^* under price-taking behavior when the market is unbalanced ($b = 3s, 5s$) and the market size is not large ($s = 1, 5$). When the market is relatively small, the disadvantage borne by the long side turns out to lead to prices higher than p^* under either simultaneous or asynchronous order-clearing. This effect is novel with respect to what is reported by Fano et al. [2] under the assumption of balanced markets.

A careful examination of Table 1 reveals a second novel effect. Comparing the two prices in the central column of each cell for the six unbalanced markets, we find that the average transaction price under S-GP is systematically higher than the average transaction price under A-GP. For five of the six cells, this difference is statistically significant for any practical level of confidence based on a two-sided paired t -test. (The exception is the case $s = 50, b = 3s$ where we find a p -value of 0.004.) In other words, when we take into account that agents try and optimize their trading strategies with respect to the protocol in use, the long side of the market finds the asynchronous order-clearing rule more beneficial.

We come back to this issue in the next sections, but the main intuition is the following. The simultaneous order-clearing forces a uniform price for all traders on the long side, while the asynchronous rule allows different agents to trade at different prices. When both sides of the market optimize their strategies, no one on the long side can earn any advantage under a uniform price. On the other hand, the price dispersion allowed by the asynchronous rule can be put at use to mitigate the disadvantage borne by traders on the long side.

Contrast this with the case where we assume that agents stick to TT trading. Comparing values across the first column in each cell of the six unbalanced markets, we see that simultaneous order-clearing is preferred by the long side. When all agents trade truthfully, asynchronous order-clearing puts the long side at a disadvantage because it increases the competitive pressure each buyer must bear; see Gjerstad [3].

Clearly, the strategic behavior of traders interacts differently with the order-clearing rule when we move from TT to GP. Comparing S-TT and S-GP in the first row of each cell of the unbalanced markets, we see that under simultaneous order-clearing trading based on optimized strategies puts the long side to a disadvantage. On the other hand, comparing A-TT and A-GP in the second row of each cell of the unbalanced markets, we find that asynchronous order-clearing improves the terms of trade for the long side.

Finally, we note that, under both GP and TT, the difference between the average trading price under both the S or the A treatment shrinks as the size of the market grows: expanding the market size swamps the differences in prices because it increases competition among traders on the same side and thus reduces the strategic advantage that either leg of the market may hold.

4 The Evolution of Strategic Behavior

We report the outcome of our simulations about the evolution of strategic behavior in two distinct sections. The first one deals with simultaneous order-clearing. The second one discusses the asynchronous case.

4.1 Simultaneous Order Clearing

Figure 2 reports in black the average trading functions evolved by GP for the nine cells in our experimental design, as well as a sample of the actual trading functions in grey. The horizontal dashed lines denote the average theoretical price (grey), and the average price over all trades (black).

Given our assumption of individual rationality, buyers' functions lie in the lower triangle below the diagonal and sellers' functions in the upper triangle. For visual clarity and computational speed, we assume that an extramarginal agent picks truth-telling over another strategy whenever both options are optimal; e.g., since none of the individually rational bids from a buyer with valuation $v < 0.25$ leads to trading, we let him pick truth-telling. Hence, the initial (final) segment of buyers' (sellers') trading functions is actually *on* the diagonal.

A comparison of the average trading functions across the first line (where $b = s$) confirms the result in Fano et al. [2]: when the size of a balanced market grows, the marginal agents learn to make an offer equal to their valuation/cost and their trading strategies approach price-taking. Additional evidence can be gathered by the actual trading functions depicted in grey. Broadly speaking, their dispersion provides a visual representation of the level of variety in the learning of individual trading functions, and thus on the strength of the evolutionary pressure insisting on them. (Recall that by assumption the dispersion for very weak traders ends up being set to zero.)

As we move from left to right in the first row, the dispersion around the trading functions of the marginal traders shrinks to zero: there is a strong evolutionary pressure on all marginal traders to learn price-taking. The residual dispersion in learning for $s = 1$ and $s = 50$ has different explanations. When $s = 1$, we have a bilateral monopoly with one buyer and one seller where there is very little selection pressure due to lack of competition. When $s = 50$, under simultaneous order-clearing the uniform price p^* is set by the marginal traders; once these have learned to be price-




Fig. 2 Average trading functions under simultaneous order-clearing (solid black). The horizontal dashed lines denote the average competitive price (grey), and the average price over all trades (black). Columns are $s = 1$ (left), $s = 5$ (center), $s = 50$ (right); rows are $b = s$ (top), $b = 3s$ (middle), $b = 5s$ (bottom).

takers, a bid above p^* or an ask below p^* is equally profitable for the other traders. Therefore, the main evolutionary pressure experienced by the intramarginal buyers (sellers) is to make offers above (below) p^* ; accordingly, we see a lot of dispersion but the intramarginal trading functions never escape such regions. Finally, the only profitable alternative to offering p^* is for buyers (sellers) to bid (ask) slightly above (below) p^* and extract better terms of trade, although this entails a significant risk of no trade. So the actual trading functions oscillate between price-taking and an offer close to p^* .

Consider now the unbalanced markets in the second and third row. Since $b > s$, buyers represent the long side of the market. A comparison of the average trading functions shows that, when the size of the market grows, the trading strategies converge to price-taking both for the marginal sellers and for a large interval of buyers around the marginal ones: that is, price-taking behavior spreads from the marginal traders to most agents on the long side, due to the competitive pressure. Accordingly, the dispersion of intramarginal buyers' trading functions shrinks to zero (except for the strongest ones) while this is not the case for intramarginal sellers. It is apparent that "traders on the long side of a market wind up holding the short end of the stick" [8].

4.2 *Asynchronous Order Clearing*

Figure 3 reports the average trading functions evolved by GP for the nine cells in our experimental design, using the same format as in Figure 2.

A comparison across the first line (where $b = s$) confirms the result in Fano et al. [2]: when the size of a balanced market grows, the intramarginal agents learn to make a constant offer equal to the competitive price and their trading strategies approach price-making. The dispersion of trading functions makes it clear that the evolutionary pressure applies to all intramarginal traders. The same observation holds when markets are unbalanced.

Combining the results about the effects of the order-clearing rule over the evolution of strategic behavior, we can say that two forces are at work. The first one was pointed out by Fano et al. [2]: the simultaneous rule pushes marginal traders towards price-taking, while the asynchronous rule attracts intramarginal traders towards price-making. The second novel effect is that, when the market is unbalanced and order-clearing is simultaneous, the competitive pressure spreads the price-taking behavior on the long side to all the intramarginal traders.

Different reasons underline these two effects. The first one stems from the asynchronous rule leading to non-uniform prices across trades: this increases the risk of trading using price-taking, and allows many more learning opportunities to run away from it. The second effect is due to the increase in competitive pressure among traders on the long side. As discussed in Sect. 3, a side effect of the asynchronous rule in allowing better learning is to help traders on the long side of small unbalanced markets to achieve better terms of trade than under simultaneous order-clearing. A similar point is made by Pouget [7]), who compares the Walrasian auctioneer trading protocol against the call auction and finds that in the latter one agents may fail to learn the competitive equilibrium if they use reinforcement learning.




Fig. 3 Average trading functions under asynchronous order-clearing (solid black). The horizontal dashed lines denote the average competitive price (grey), and the average price over all trades (black). Columns are $s = 1$ (left), $s = 5$ (center), $s = 50$ (right); rows are $b = s$ (top), $b = 3s$ (middle), $b = 5s$ (bottom).

5 Allocative Efficiency

Our last batch of results concerns allocative efficiency, that we define as the ratio between the total surplus realized by traders and the theoretical maximum surplus. (This latter one is computed as the surplus realized under S-TT.) We confirm the result from Fano et al. [2] that the allocative inefficiency tends to zero as the size of the market grows regardless of the order-clearing rule, and we find that it holds also for unbalanced markets.

Ceteris paribus, the allocative efficiency is higher under the S treatment than under the A treatment for all cells in our experimental design. Moreover, as a function of the optimization steps, the efficiency remains virtually constant under simultaneous order-clearing while it takes several optimization steps before it levels off under the asynchronous rule. This qualitative difference conforms with our former observation that the effects of a simultaneous order-clearing rule swamp those associated with the evolution of trading strategies, while these latter ones carry much greater impact under the asynchronous rule. To exemplify, [Figure 4](#) depicts the allocative efficiency (in black) and the buyers’/sellers’ surplus ratio (in grey, to be discussed below) for the case $s = 50$ and $b = 5s$ as a function of the optimization steps. (The statistics are computed over 1000 rounds of trading after freezing the GP process.)




Fig. 4 Efficiency (black) and buyers’/sellers surplus ratio (grey) under simultaneous (left) and asynchronous (right) order clearing for the case $s = 50$ and $b = 5s$.

For a more complete description, we compare the surplus realized under S-GP and under A-GP. Using the surplus realized under S-TT as benchmark, we look at the ratios for the realized surplus in the two GP treatments over the benchmark. [Table 2](#) provides for each cell of our experimental design the average values for such two quantities, computed over at least 10 distinct simulations. We report the ratio of surplus for S-GP over S-TT on the left and for A-GP over S-TT on the right. Apart

	$s = 1$		$s = 5$		$s = 50$	
$b = s$	0.785	0.797	0.971	0.888	0.996	0.968
$b = 3s$	0.863	0.824	0.982	0.923	0.998	0.980
$b = 5s$	0.898	0.854	0.982	0.941	0.998	0.984

Table 2 Average ratios for S-GP vs. S-TT (left) and for A-GP vs. S-TT (right).

from the case where $s = b = 1$ (for which the distinction between simultaneous and asynchronous order-clearing is of little consequence), allocative efficiency is higher under S than under A and grows to 1 as the size of the market increases. This claim is supported for any practical level of confidence by a two-sided paired t -test. The

lack of balance among the two sides of the market does not seem to carry discernible effects, except for the increase in allocative efficiency due to the higher number of traders available in the unbalanced markets.

Clearly, allocative efficiency is affected both by the protocol and by the trading strategies adopted by the agents. For instance, since the ratios on the left of each cell of [Table 2](#) are lower than 1, we can conclude that the strategic behavior developed under GP reduces the efficiency that could be achieved if everyone would make truthful offers. Unfortunately, it is well known that the pursuit of individual profits may hurt the common welfare.

For the case of asynchronous order-clearing, we attempt to disentangle the effects of protocol and behavior in [Table 3](#) where we report the (average) ratio of surplus for A-TT over S-TT on the left and for A-GP over A-TT on the right. The ratio

	$s = 1$		$s = 5$		$s = 50$	
$b = s$	1.000	0.797	0.807	1.100	0.702	1.379
$b = 3s$	0.844	0.975	0.825	1.119	0.875	1.119
$b = 5s$	0.820	1.042	0.864	1.089	0.937	1.050

Table 3 Average ratios for A-TT vs. S-TT (left) and for A-GP vs. A-TT (right).

between A-TT and S-TT measures the loss in allocative efficiency when traders using truth-telling switch from simultaneous to asynchronous order-clearing. Given that this latter rule destroys uniform pricing and adds noise due to the random order of arrival, we expect to see allocative inefficiency.

The second value on the right of each cell gives the (average) ratio between the surpluses realized under A-GP and A-TT. It measures the improvement in allocative efficiency brought over by optimizing trading strategies with respect to the protocol. Since the ratio is greater than 1 in all markets with $s \neq 1$, the pursuit of individual profits is socially beneficial under asynchronous order-clearing because it reduces both missed trades and the number of transactions involving extramarginal traders. Contrasting this with the case of simultaneous order-clearing, we conclude that the social benefits of strategic behavior may depend on protocolary details.

Our last comparison concerns the ratio between the surplus realized by the long side and the surplus realized by the short side. This should be (approximately) 1 in a balanced market and strictly lower than 1 in an unbalanced market. [Table 4](#) reports the (average) ratio of buyers' to sellers' surplus under S-TT on the left, under S-GP at the center, and under A-GP on the right. As expected, the ratio is close to 1 for all balanced markets in the first row. Moreover, it is exactly 1 for S-TT with $s = 1$ because the simultaneous order-clearing rule sets the clearing price exactly midway between the best bid and the best ask, so the two agents who complete a transaction must split the surplus equally between themselves.

In all other cases, it is apparent that the buyers on the long side of the market capture a share of the allocative efficiency much lower than the sellers on the short side. Three major effects are worth noting. First, increasing the market size makes the long side worse off by increasing the competitive pressure among buy-

	$s = 1$			$s = 5$			$s = 50$		
$b = s$	1.000	1.012	1.044	1.001	0.994	0.993	1.000	1.004	1.001
$b = 3s$	1.000	0.433	0.591	0.542	0.401	0.428	0.347	0.341	0.348
$b = 5s$	1.000	0.302	0.422	0.413	0.257	0.274	0.212	0.205	0.209

Table 4 Average ratios for buyers' to sellers' surplus under S-TT (left), S-GP (center), and A-GP (right).

ers. Second, compared to the S-TT benchmark, optimizing trading strategies under simultaneous order-clearing further deteriorates the long side's position because the sellers on the short side learn to exploit their relative advantage. Finally, switching to asynchronous order-clearing under optimized strategies improves the long side's performance because the random order of arrival for traders weakens the bargaining position of sellers. These three effects closely match the observations in Sect. 3, in so far that better terms of trades for the buyers on the long side correspond to capturing a larger share of the realized surplus.

References

1. Anufriev M, Arifovic J, Ledyard J, Panchenko V (2010), Efficiency of continuous double auctions under individual evolutionary learning with full or limited information. Working Paper, CeNDEF, University of Amsterdam
2. Fano S, LiCalzi M, Pellizzari P (2011), Convergence of outcomes and evolution of strategic behavior in double auctions. *Journal of Evolutionary Economics*, forthcoming
3. Gjerstad S (2007), The competitive market paradox. *Journal of Economic Dynamics and Control* 31:1753-1780
4. Koza JR (1992), *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. The MIT Press, Cambridge, Massachusetts
5. Gode DK, Sunder S (1997), What makes markets allocationally efficient? *Quarterly Journal of Economics* 112:603-630
6. LiCalzi M, Pellizzari P (2008), Zero-intelligence trading without resampling. In: Schrederseker K, Hauser F (eds) *Complexity and Artificial Markets*, 3-14. Springer, Berlin Heidelberg New York
7. Pouget S (2007), Adaptive traders and the design of financial markets. *The Journal of Finance* 52:2835-2863
8. Taylor CR (1995), The long side of the market and the short end of the stick: Bargaining power and price formation in buyers', sellers', and balanced markets. *Quarterly Journal of Economics* 110:837-855