Experimental Economics (2022) 25:203–228
https://doi.org/10.1007/s10683-021-09710-1

ORIGINAL PAPER

Do traders learn to select efficient market institutions?

Carlos Alós-Ferrer[1] · **Johannes Buckenmaier**[1] · **Georg Kirchsteiger**[2]

Received: 11 November 2020 / Revised: 29 January 2021 / Accepted: 15 March 2021 /
Published online: 16 April 2021
© The Author(s) 2021

Abstract

When alternative market institutions are available, traders have to decide both where and how much to trade. We conducted an experiment where traders decided first whether to trade in an (efficient) double-auction institution or in a posted-offers one (favoring sellers), and second how much to trade. When sellers face decreasing returns to scale (increasing production costs), fast coordination on the double-auction occurs, with the posted-offers institution becoming inactive. In contrast, under constant returns to scale, both institutions remain active and coordination is slower. The reason is that sellers trade off higher efficiency in a market with dwindling profits for biased-up profits in a market with vanishing customers. Hence, efficiency alone might not be sufficient to guarantee coordination on a single market institution if the surplus distribution is asymmetric. Trading behavior approaches equilibrium predictions (market clearing) within each institution, but switching behavior across institutions is explained by simple rules of thumb, with buyers chasing low prices and sellers considering both prices and trader ratios.

Keywords Market selection · Market clearing · Posted offer market · Constant returns to scale · Experiment

JEL classification D4 · D83 · L1

✉ Carlos Alós-Ferrer
carlos.alos-ferrer@econ.uzh.ch

Johannes Buckenmaier
johannes.buckenmaier@econ.uzh.ch

Georg Kirchsteiger
georg.kirchsteiger@ulb.be

[1] Zurich Center for Neuroeconomics (ZNE), University of Zurich, Zurich, Switzerland

[2] ECARES, Université Libre de Bruxelles, Bruxelles, Belgium

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

1 Introduction

When investigating the functioning of markets one is puzzled by the huge variety of market institutions used to trade similar or even the same good (see e.g. Madhavan 1992; Quan 1994; Zumpano et al. 1996). Extensive research has revealed that the institutional rules governing trade have a substantial impact on prices, quantities, and the resulting efficiency of markets (see, among many others, Lucking-Reiley 2000; Reynolds and Wooders 2009; Ariely et al. 2005; Ausubel et al. 2014; Genesove and Hansen 2016). This in turn led to substantial research on the evolution of market rules (see e.g. Rud and Rabanal 2018). Starting with the seminal contribution of Hayek (1967) it has been argued that traders, if left to their own devices, choose the most efficient trading institution that the good at hand allows for.[1] This article aims to test this conjecture experimentally. Specifically, we set out to investigate whether efficient institutions are indeed used in the long run, while inefficient ones fade as they are eventually avoided by the traders.

Take a situation where different trading platforms are feasible for trading a particular good. Each trader has to choose which one he/she wants to use.[2] An individual buyer or seller will not choose a platform where he/she cannot find a trading partner, even if the institutional setup of the avoided platform would lead to very efficient trading outcomes. Hence, traders face a (partial) coordination game when choosing on which platform they want to trade.[3] After choosing the trading institution, i.e. for a given distribution of buyers and sellers over the feasible platforms, the institutional setups–together with demand and supply resulting from the platform choice–will lead to more or less efficient trading outcomes. Hence, the whole process is characterized by two steps: First, traders have to choose the trading platform (where to trade). Then, they conduct their trades with the partners available at their platform (how much and at which price to trade).

In our experiment each subject is either a buyer or a seller. In the first stage of each experimental round, each subject had to choose individually between a posted offer (PO; with sellers posting prices) institution and an open double auction (DA, Plott and Gray 1990; "open" meaning that there was no auctioneer, thus multiple prices could be realized).[4] In the second stage of each round, subjects could trade multiple units of a homogeneous good according to the rules of the chosen institution with those potential trading partners who had opted for the same institution. At the end of each round, all the participants received aggregate information about what happened at both institutions, enabling them to make an informed decision on where to trade in the next round: number of buyers and sellers, distribution of

[1] This "Hayekian" position can also be discerned in the writings of Klein and Alchian (1978) and Williamson (1975).

[2] In our notation sellers are female and buyers are male.

[3] Coordination games have been widely studied both theoretically and experimentally (see, e.g. Young 1993; Van Huyck et al. 1997). To the best of our knowledge, though, these investigations have not been connected to the functioning of markets.

[4] We speak of "market institutions" for consistency with the previous literature. Other authors have used the term "market formats" (e.g. Rud and Rabanal 2018).

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

trading prices (and their average), and number of trades per buyer and per seller at each institution. Then a new round started, where first all subjects had to choose between the DA and the PO. After that, they traded on the chosen platform anew, received information about the outcome of both institutions, and so on.

For each trade conducted, the seller received the price paid by the buyer, but in turn she had to pay production costs that were determined by the experimental design. Hence, her net earnings from a trade were the difference between the price and production costs. For a buyer, the net earnings from a trade were the difference between the resale values, which were determined by the design of the experiment, and the price paid.[5] That is, in each given round, the actual, endogenously-induced supply and demand at a given platform was determined by the production costs and resale values of the sellers and buyers who opted for that platform. To keep the demand and supply functions constant, and since our focus is on which institutions survive, the induced values were the same for all buyers and the production costs were the same for all sellers.

We chose DA and PO as institutional alternatives since it is very well documented that, when functioning in isolation, DAs are highly efficient, whereas POs tend to create inefficiencies and are biased towards sellers. Theoretically, for example, Rustichini et al. (1994) show that, in double auctions, possible inefficiencies due to traders' strategic misreporting vanish quickly as the number of traders increase. Empirically, it has been repeatedly shown that DAs induce a very quick convergence of prices and quantities to market clearing levels (see, e.g., the seminal article of Plott and Smith 1978a). As a consequence, the DA has been found to be the most efficient market institution when efficiency is measured by the sum of all gains from trade reaped by buyers and sellers (as it is usually done in the context of market experiments). In fact, DA is routinely used in market experiments as a proxy for competitive markets (e.g., Plott 2000; Crockett et al. 2011; Gillen et al. 2020). On the other hand, in the PO prices and quantities typically show a much slower convergence to the market clearing level, with prices typically converging from above (see Plott and Smith 1978a). For given demand and supply, the efficiency level is considerably higher for a DA than for a PO. If efficiency drives the selection of market institutions, we should expect that traders coordinate on the DA, and that the PO eventually falls into disuse.

The experimental results indicate that the validity of this prediction depends crucially on the properties of the distribution of production costs. If the production technology underlying the sellers' supply function exhibits decreasing returns to scale, traders indeed learn to coordinate on the efficient DA, with PO becoming mostly inactive in finitely many rounds. In contrast, if production displays constant returns to scale, resulting in a flat induced supply function, coordination on DA

[5] In order to obtain a decreasing demand function allowing for an unambiguous equilibrium prediction, the highest resale value was used for the calculation of the net earnings of the buyer's first trade, the second highest resale value for the calculation of the earnings resulting from the second trade, etc.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

happens more slowly and both institutions remain active until the last round.[6] This difference might be caused by the fact that a Pareto-dominance relation between the two institutions depends on the distribution of the production costs. For a production technology with constant returns to scale, sellers' earnings are (nearly) zero at an institution like the DA that induces prices which are at (or very close to) the market clearing level. Hence, one market side, i.e. the sellers, has a strong incentive to try to coordinate traders on an institution like the PO, where prices are above the market clearing level. To put it differently, if an institution yields outcomes which Pareto-dominate those of competing ones, it should be expected that it will clearly and quickly drive the latter out of the market. However, switching from coordination on the less efficient PO to coordination on the more efficient DA does not constitute a Pareto-improvement if sellers produce under constant returns to scale. Our results thus suggest that efficiency alone, as measured by the sum of the gains from trade, might not be enough to guarantee that a market institution drives competing ones out of the field. The reason for this is that sellers face a tradeoff between a more efficient institution which brings their profits down and an inefficient, low-volume one which however typically yields higher prices. In the absence of Pareto-dominance, the distribution of gains from trade, and not only aggregate efficiency, becomes consequential for market selection and survival.[7]

Our data allows us to look both at actual trading behavior within an institution and the previous decision of which institution to trade in. Regarding the former, we observe that aggregate behavior, as reflected by actual prices, does converge to the theoretical market-clearing benchmark both in PO and in DA (even though we use an open DA implementation without an auctioneer). Regarding the choice of market institution, we find that simple rules of thumb are sufficient to capture most behavior. Buyers seem to mostly chase after low (past) prices, switching to the institution where observed prices were better for them, along the lines of previously-postulated behavioral rules based on past performance only (e.g. Huck et al. 1999; Offerman et al. 2002; Bosch-Domènech and Vriend 2003). Sellers' behavior appears to be more complex, but is well-explained by the combination of a similar rule which points toward high prices, and a complementary rule which focuses on favorable trader ratios. In our setting, where all buyers have the same resale values and all

[6] This result is similar in flavor to those of Alós-Ferrer and Kirchsteiger (2015, 2017). In these theoretical articles we found that traders who have to learn which platform to use are more likely to coordinate on a market clearing institution if sellers have an increasing supply function. See also Alós-Ferrer et al. (2010) for implications for market design.

[7] Naturally, one could also consider a mirror-image situation where the inefficient institution favors buyers, e.g. using a posted-bid institution (PB) where buyers post bids. By symmetry we would expect symmetric results to the ones obtained for the choice between the DA and PO: If the demand is perfectly elastic and supply is increasing, the market-side "disadvantaged" by the market clearing equilibrium (in this case the buyers) would resist full coordination on the DA, and the PB would remain active. On the other hand, for decreasing demand and increasing supply we would also expect full coordination on the DA. One could expect this convergence to be even faster when the supply is perfectly elastic. Finally, if the choice is between DA and PO, but the demand is perfectly elastic and the supply is increasing, we would also expect a very quick convergence to DA since the inefficient institution does not favor the disadvantaged side in this case.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

sellers have the same cost functions, equilibrium prices are a function of trader ratios, hence the latter rule might reflect forward-looking behavior.

An extensive experimental literature, going back to Chamberlin (1948) and Smith (1962), has analyzed the empirical properties of market institutions for given demand and supply when viewed in isolation (for an overview see Plott and Smith 1978b, Part 1).[8] Surprisingly, however, very few contributions have combined the experimental investigation of actual trading behavior with the analysis of the *choice* of the trading institution for a given good. One notable exception is Campbell et al. (1991), which investigated the endogenous choice between a computerized double-auction market (with an auctioneer) and (illegal) off-floor trading in the context of stock markets, and its impact on the bid-ask spread. The latter was implemented as direct negotiations, specifically traders could submit direct offers for blocks of three units to their two neighboring traders of the other market side. Kugler et al. (2006) compare direct negotiations and centralized markets in an experimental setting where each trader can trade a single unit only. Trade can be conducted in two alternative institutions. The first captures centralized markets through a sealed-bid double auction with a single market clearing price (as in Campbell et al. 1991, computed by an auctioneer). The second captures direct negotiations through bilateral matching, where the matching was designed as to maximize trade. Inefficiency might appear due to asymmetric distribution of traders across the two institutions, but their focus was not on the comparison of institutions in terms of efficiency. Rather, they focus on heterogeneous traders with different values and find that different types of traders generally prefer different market mechanisms. Their results show an unraveling of direct negotiations, which is led by higher-value traders (buyers with high resale values or sellers with low costs), who learn faster to coordinate on the centralized market. Last, the theoretical work of Alós-Ferrer and Kirchsteiger (2015) included an experiment on platform selection but adopted a reduced-form payoff table for actual trade, that is, traders did not make actual trading decisions. Alós-Ferrer and Kirchsteiger (2015) focused on the selection of market-clearing institutions (vs. institutions with price biases) and showed that certain alternative institutions could survive in the long run when players followed simple behavioral rules of thumb.

2 The experiment

A total of $N = 480$ subjects (258 females, mean age 25.3 years, $SD = 6.9$) participated in 15 experimental sessions. Subjects were recruited from the student population of the University of Cologne (excluding students majoring in psychology) via ORSEE (Greiner 2015). All interactions took place via a custom-made market interface (see Sect. 2.2) programmed in zTree (Fischbacher 2007). There were two treatments, DRS and CRS (see Sect. 2.3), with 15 groups (240 subjects) randomly assigned to each of them and to the buyer or seller role (see Table 1).

[8] Williams et al. (2000) examined experimental double-auction markets in a setting where participants traded two different goods concurrently, with a different market for each good.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

2.1 Procedures

A session comprised 32 subjects randomly allocated into two groups of 16 subjects each. Subjects interacted only with other subjects within their own group. That is, each group constitutes an independent observation. Within each group half of the subjects were randomly assigned the role of a buyer and half were assigned the role of a seller. Roles remained fixed throughout the experiment.

Before the beginning of the experiment, subjects received written instructions describing the course of the experiment and the market selection stage (see Online Appendix). These general instructions were also read aloud by the experimenter. After that, subjects received specific written instructions according to their assigned role (buyer or seller) with a detailed description of how trade was conducted at each of the two platforms. Subjects then answered four control questions to ensure their understanding of the experimental environment.

The instructions also contained a detailed description of the payment procedure, which was carried out truthfully. At the end of the experiment a subject's earnings from each round were added up and converted to euros at a rate of €1.5 for 100 experimental currency units. In addition subjects received a show up fee of €4 leading to an average total remuneration of 23.68 EUR. Sessions lasted about 105 minutes on average.

2.2 Design and market interface

Each population of traders consisted of 8 buyers and 8 sellers. Roles were assigned randomly at the beginning of the experiment and stayed fixed throughout. All market interactions took place via a custom-made interface programmed in zTree (Fischbacher 2007).[9] There were 25 trading rounds, each consisting of two sequential stages. In the first stage (Market Selection), each trader chose individually between two market platforms, a double auction (DA) and a posted offer (PO) institution. In the second stage (Trading), subjects could trade at the selected platform according to the rules of that institution, interacting with all traders who selected the same market. That is, the number of traders within each market was determined endogenously and, hence, so was demand and supply. The number of buyers and sellers within the selected market was visible throughout the trading stage. In any given round traders had to commit to the selected market and could only trade within that market, but they could freely switch markets between rounds.

We use a multi-unit setup, specifically, each trader can trade up to 6 units of a homogeneous good. Buyers received exogenously-given resale values for each unit bought. Sellers faced exogenously-given production costs for each unit sold.[10]

[9] See the Online Appendix for screenshots.

[10] In this sense, preferences in our experiment were induced. Crockett et al. (2020) make the point that, in market experiments where traded goods involve exogenous risk, it might be important to differentiate between induced and elicited preferences. However, the trading stage of our experiment involved no exogenous risk and followed the tradition of Chamberlin (1948), Smith (1962), and many others.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Table 1 Summary of the experiment

Treatment	# Buyers	# Sellers	# Markets	# Sellers per market	# Buyers per market
DRS	120	120	15	8	8
CRS	120	120	15	8	8

Buyers and sellers were homogeneous, that is, resale values and production costs were the same for all buyers and sellers, respectively. For each trade conducted, the seller received the price at which the trade took place and had to pay the corresponding production costs, whereas buyers had to pay the price and received the corresponding resale value. Endowments and capacities were reset each period, that is, the experiment involved stationary repetition as standard in market experiments.

In each session, the trading rounds were preceded by two trial rounds intended to give subjects the opportunity to familiarize themselves with the interface and the rules of trade at both platforms. The trial periods featured no market selection stage. Instead, in the first trial round 4 buyers and 4 sellers were randomly assigned to each platform, and in the second trial round platform allocation was reversed.

Market Selection In the first stage, traders chose between a DA and a PO institution. In all but the first round, subjects received aggregate information regarding the performance of both platforms (irrespective of which platform they had selected previously) as well as a summary of their own performance in the previous round. Specifically, aggregate information was presented in a table showing the number of buyers and sellers, the average trading price, and the number of trades per buyer and seller separately for each institution. Additionally, subjects were provided with a price-quantity histogram for each market indicating the number of units traded at each price. Own performance was summarized via a second table listing the prices of all units traded by the subject and the overall profit in that trading round. On the same screen subjects then could choose one of the two platforms, neutrally labeled "Market A" and "Market B," via a button press. Subjects had 30 seconds for this choice and the remaining time was shown on screen. If a subject failed to select a platform within the time limit, he/she could not trade in the next trading round.[11] If in a given round a platform was selected only by buyers or only by sellers, it remained inactive and traders were informed of this fact on screen.

Double auction institution Figure 1 displays a graphical illustration of the interface for DA (for buyers). In DA all traders could simultaneously make offers and accept open offers by traders from the other market side. That is, there was no auctioneer or market-maker, and all offers were visible for all traders. Traders could offer any number of units $q \in \{1, \ldots, q_r\}$, with $q_r \leq 6$ being the remaining number of units that could be traded, at any price p by submitting an offer of the form (p, q), subject to the constraint of not incurring a net loss within a trading period.[12] At any

[11] This happened for 22 of the 480×25 market-selection decisions, i.e. around 0.2% of the time.

[12] The range of feasible offers that could be submitted and the feasible active offers that could be accepted by each trader were limited in such a way that net within-period losses were not possible. That is, single units could be traded at a loss as long as this was compensated by profits from other trades in

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

point in time a trader could have at most one active offer. However, offers could be withdrawn or replaced by new offers at any time. A submitted offer remained active until it was withdrawn or replaced by a new offer. Further, traders could always accept active offers by the other market side. All traders could see all open offers from both buyers and sellers (see Figure 1). Offers from the own and the other market side were always shown on the left and right parts of the screen, respectively. Each offer (p, q) was represented by q boxes showing the price p with each box representing an offer to buy/sell one unit at that price (see Figure 1). Traders could accept offers from the other market side, that is buy/sell one unit at the shown price, by clicking on the corresponding box. Traded units were indicated by crossed-out boxes and remained visible until the end of the trading round. Each subject could trade up to a maximum of 6 units and the remaining number of units that could be traded at any point was shown on screen. Additionally, buyers, respectively sellers, were shown the resale values, respectively production costs, for each of the six units as well as the price of already-traded units. Subjects could leave the trading stage by pressing the "leave market stage" button, in which case any still-active offer was withdrawn. The maximum duration of a trading round at DA was 90 seconds, with the remaining time shown on screen. Alternatively, the market platform was closed if no further trade was possible (for example because all buyers had left the trading stage or all sellers were out of stock).

Our design follows the classical Multiple Unit Double Auction (MUDA) design of Plott and Gray (1990), with the exception that in their design only the highest-price bid and the lowest-price ask are shown (replacing previous ones), while in our design subjects observe all offers. Subjects were able to see all transactions in real-time, but no graphical representation or order book were provided.

Posted offer institution In PO buyers and sellers moved sequentially in two distinct stages. In the first stage, sellers simultaneously submitted offers of the form (p, q) proposing to sell q units at price p.[13] Each seller was only able to submit a single offer (p, q) (similarly to DA, where each trader could have at most one active offer at any point in time). Sellers were shown their production costs for all six units, but received no information about the offers made by other sellers during this stage. This first stage lasted for a maximum of 20 seconds, or, alternatively until all sellers had committed to an offer. At the end of this stage the offers made by all sellers were collected and the second stage began. In that stage, buyers moved sequentially in a randomly-determined order. When a buyer's turn came, he observed all offers made in the first stage, as well as all units bought by previous buyers, and could buy units by clicking on still-available offers. As in the DA case, available offers were indicated by numbered boxes and already-traded units by crossed-out boxes (as in

Footnote 12 (continued)
the same trading period. Subjects received a detailed explanation of this procedure in the written instructions.

[13] As in the DA, feasible offers were limited to (p, q) pairs such that, if all offered units were traded, no net within-period loss for the seller would result.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

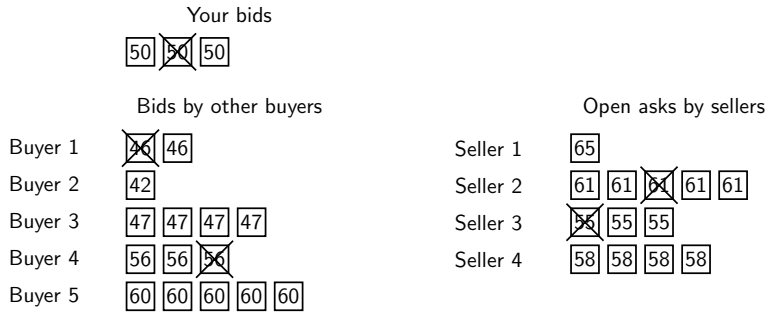


Fig. 1 Schematic representation of the screen layout for DA (buyers)

the right part of Fig. 1).[14] A buyer's turn ended after 20 seconds or once no further units were available. The trading round ended once all buyers had had their turn, or no further units were available.

2.3 Treatments

In a given round, the resale values and production costs of the traders that opted for a particular institution determined the induced supply and demand at that institution. There were two treatments, DRS and CRS, which correspond to decreasing and constant returns to scale for the supply side of the market, respectively. Specifically, the treatments differed in the resale values and production costs as given in Table 2.[15] In both treatments the demand function induced via buyers' resale values was decreasing in the price. This was accomplished by setting decreasing resale values for each unit traded by a buyer. In treatment DRS, sellers faced increasing production costs (decreasing returns to scale), hence the induced supply function was increasing in the price. In treatment CRS, sellers faced a constant production cost instead (constant returns to scale), and as a consequence the induced supply function was flat. In each round a common shock X_t was added to all these baseline resale values and production costs, so that the exact numbers varied from round to round, but induced demand and supply were unaffected.

2.4 Hypotheses

Given demand and supply, the DA typically leads to higher efficiency than the PO. Intuitively, if selection of market institutions is driven only by efficiency, over time traders should learn to coordinate on the efficient institution. However, efficiency is only concerned with overall trader surplus but not with the distribution of these gains

[14] Similarly to the DA case, buyers could only accept offers not leading to a net within-period loss.
[15] The trial periods used different resale values and production costs.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

from trade between buyers and sellers. The latter introduces a tradeoff between efficiency and profits for one market-side that could prevent coordination on the efficient institution.

In treatment DRS, the distribution of profits is "symmetric" in the sense that for prices at the market clearing level both buyers and sellers make positive profits. Consequently, in this treatment the aforementioned tradeoff is expected to be absent and convergence on the efficient institution should occur.

Hypothesis 1 In DRS, traders learn to coordinate on DA over time with PO becoming mostly inactive in the final rounds.

In contrast, treatment CRS presents exactly an "asymmetric" setting where market power is biased toward buyers in the sense that at prices at the market clearing level sellers make zero profits and buyers reap all the gains from trade. Hence, sellers have to choose between an efficient institution where their profits are (close to) zero and an inefficient institution which, due to its structure, typically leads to higher prices and seller profits. This tradeoff might hinder coordination on DA leading to slower convergence and survival of the inefficient PO institution.

Hypothesis 2 In CRS, coordination on DA is slower than in DRS with both institutions remaining active until the last round.

3 Equilibrium prices and market-clearing benchmark

In this section, we consider the theoretical benchmark where markets clear, taking as given an allocation of traders to institutions, and derive equilibrium market prices and total trader surplus, which we will use as a benchmark to measure efficiency.

Fix an institution, and consider the population of traders that has chosen to trade at that institution. This population consists of n buyers and m sellers (with $n, m \in \{0, \dots, 8\}$ in the experiment). Each trader can trade up to Q units ($Q = 6$ in the experiment) of a single homogeneous good. Denote the price of the good by $p \geq 0$. A typical buyer is characterized by a vector of resale values $(r_1, \dots, r_Q)$ with $r_1 \geq \cdots \geq r_Q > 0$ where r_k is the buyer's resale value for the kth unit. A typical seller is characterized by a vector of production costs $(c_1, \dots, c_Q)$ with $0 < c_1 \leq \cdots \leq c_Q$ where c_k is the seller's production cost for the kth unit. Assume that $r_1 \geq c_1$, that is, beneficial trade is possible for any $p \in [r_1, c_1]$. For a given price p, the resale values induce a (weakly) decreasing demand function given by

$$d(n, p) = nk \text{ where } k = \max\left\{0, \sup\left\{k \in \{1, \dots, Q\} \;\middle|\; p \leq r_k\right\}\right\}.$$

That is, demand is given by the largest number of units that the n buyers are willing to purchase (implicitly assuming, for simplicity, that traders are willing to trade when exactly indifferent). Specifically, if $r_Q < p < r_1$, k is the unique number such that $r_{k+1} < p \leq r_k$. If $p \leq r_Q$, then $k = Q$ and demand is given by nQ (the capacity constraint becomes binding). If $p > r_1$, the set over which the supremum is taken

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Table 2 Baseline resale values and production costs for treatments DRS and CRS

Unit	DRS		CRS	
	Resale Values	Costs	Resale Values	Costs
1	80	15	80	40
2	65	35	65	40
3	60	40	60	40
4	55	50	55	40
5	40	60	40	40
6	35	65	35	40

above is empty (so the supremum is formally $-\infty$), and hence $k = 0$ and demand is nonexistent. The induced demand function is illustrated in Fig. 2 for both treatments in the experiment, for the case where all 8 buyers are at the same institution.

Analogously, for a given price p, the production costs induce a (weakly) increasing supply function given by

$$s(m,p) = mk \text{ where } k = \max\left\{0, \sup\left\{k \in \{1,\dots,Q\} \;\middle|\; p \geq c_k\right\}\right\}.$$

That is, supply is given by the largest number of units that the m sellers are willing to sell. Specifically, if $c_1 < p < c_Q$, k is the unique number such that $c_k \leq p < c_{k+1}$. If $p \geq c_Q$, then $k = Q$ and supply is given by mQ (the capacity constraint becomes binding). If $p < c_1$, $k = 0$ and supply is nonexistent. The induced supply function is illustrated in Fig. 2 for both treatments in the experiment, for the case where all 8 sellers are at the same institution.

In our discrete setting, demand and supply are step functions and hence equilibrium prices might correspond to an interval instead of a unique value. Let $\underline{p}(m,n) = \sup\{p \mid d(n,p) \geq s(m,p)\}$ and $\bar{p}(m,n) = \inf\{p \mid d(n,p) \leq s(m,p)\}$. It follows that $\underline{p}(m,n) \in \{-\infty, c_1, \dots, c_Q\}$ and $\bar{p}(m,n) \in \{+\infty, r_1, \dots, r_Q\}$. Further, note that $\underline{p}(m,n) \leq \bar{p}(m,n)$, and, since $r_1 \geq c_1$, at least one of $\bar{p}(m,n)$ and $\underline{p}(m,n)$ must be finite.

The set of equilibrium prices $p^*(m,n)$ is given by

$$p^*(m,n) = \begin{cases} [\underline{p}(m,n), \bar{p}(m,n)] & \text{if } \underline{p}(m,n),\ \bar{p}(m,n) \text{ are finite} \\ \{r_1\} & \text{if } \bar{p}(m,n) = \infty \text{ (that is, } n \geq mQ) \\ \{c_1\} & \text{if } \underline{p}(m,n) = -\infty \text{ (that is, } m \geq nQ). \end{cases}$$

Figure 2 illustrates $p^*(m,n)$ for both treatments in the experiment, for the particular case $m = n = 8$. Note that $\bar{p}(m,n) = \infty$ implies that $n \geq mQ$ and demand always exceeds supply, hence r_1 is the unique equilibrium price at which trade actually occurs. Analogously, if $\underline{p}(m,n) = -\infty$, $m \geq nQ$, supply always exceeds demand, and c_1 is the unique equilibrium price at which trade occurs.

In our experiment, the resale values were given, and the production costs varied across treatments (DRS and CRS). Tables B.1 and B.2 in the Online Appendix report the equilibrium price intervals for both treatments. In the analysis below, and as a first

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

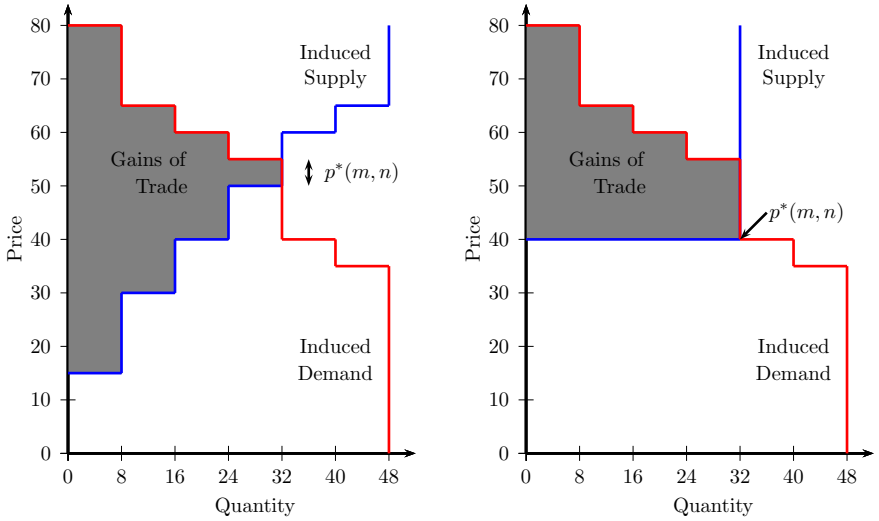


Fig. 2 Induced demand and supply functions for both treatments (DRS left, CRS right) in case of full coordination on the respective institution

measure of the efficiency of the institution, we will report the distance from actually-realized trading prices per round to the benchmark equilibrium price intervals $p^*(m,n)$.

As a second measure of efficiency, we consider the largest achievable total trader surplus for a given number of buyers and sellers. Let the mQ resale values of the m buyers be $R_1 \geq \cdots \geq R_{mQ}$ and let the nQ costs of the sellers be $C_1 \leq \cdots \leq C_{nQ}$. The *maximum gains of trade* are defined as the largest-possible total trader surplus, that is

$$S^*(m,n) = \sum_{i \in I}(R_i - C_i) \text{ where } I = \left\{ i \;\middle|\; R_i > C_i, i \leq \min\{nQ, mQ\} \right\}.$$

Figure 2 illustrates $S^*(m,n)$ for both treatments in the experiment, for the particular case $m = n = 8$. The values of $S^*(m,n)$ for DRS and CRS are given in Tables B.3 and B.4 in the Online Appendix, respectively. In the analysis below, we will compare the actually-realized total trader surplus (defined as the sum of the differences between resale value and production cost for all actually-traded units) to $S^*(m,n)$.

4 Results: Market Selection

4.1 Decreasing returns

We first consider treatment DRS, where the supply side faced decreasing returns to scale ($N = 240$ in 15 market observations). The top-left panel of Figure 3 shows the average number of buyers and sellers at DA for each period. We see a clear

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

convergence toward DA over time.[16] That is, as time goes by, more and more traders of each market side choose the DA. This is confirmed in a series of linear random effects regressions on the number of buyers/sellers at DA per period, as reported in Table 3.[17] The variable RatioDA is the (lagged) buyer-seller ratio at DA. The regressions show that buyers are attracted by a low buyer-sellers ratio, but the latter has no effect on sellers' behavior.

The bottom-left panel of Fig. 3 shows the fraction of markets where DA and PO are active over time. Across all 15 markets the DA is active in all but one period for a single market, that is, on average the DA is active 99.7% of the time. In contrast, the PO is only active 48% of the time. The average fraction of periods in which DA is active is significantly larger than the average fraction of periods in which PO is active (Wilcoxon signed-rank test, WSR: $N = 15$,[18] $z = 3.352$, $p < 0.001$). We obtain a similar result for the average number of buyers or sellers, that is, on average there are more buyers and more sellers at DA than at PO (WSR, buyers: $z = 3.352$, $p < 0.001$; sellers: $z = 3.352$, $p < 0.001$).

To confirm that traders increasingly coordinate on DA over time, we compare coordination at the beginning against coordination at the end of the experiment. To that end, we exclude the first period and split the remaining 24 periods in three parts of 8 periods each (see the Online Appendix for an alternative specification). We first consider the fraction of time DA or PO are active in part 1 vs. part 3. Because DA is essentially active in all periods, there is no difference for DA between parts 1 and 3 (WSR, $z = -1.000$, $p = 0.317$). On the other hand, PO is active more often in part 1 (64.2% of the time) than in part 3 (28.3%; WSR, $z = 3.048$, $p = 0.002$). A similar picture emerges when looking at the number of buyers and sellers at DA, which is significantly larger at the end of the experiment compared to the beginning (WSR, buyers: $z = -3.311$, $p < 0.001$; sellers: $z = -3.411$, $p < 0.001$).

4.2 Constant returns

We now consider treatment CRS ($N = 240$ in 15 different market observations). The top-right panel of Figure 3 illustrates the evolution of the number of buyers and sellers at DA over the course of the experiment in CRS. We again observe convergence toward DA, although it appears to be slower than in the case of DRS (see next subsection for a treatment comparison).[19] Table 4 displays a series of linear random effects regressions of the number of buyers/sellers at DA on period and shows that traders move toward DA over time.[20] Buyers are attracted by a low buyer-seller ratio, and the opposite is true for sellers.

[16] Figure A.1 in the Online Appendix displays the analogous graphs for each market session.
[17] Fractional logit regressions yield the same results. See Table A.1 in the Online Appendix.
[18] All WSR tests in Sects. 4 and 5 refer to paired comparisons within a treatment and hence have $N = 15$. Thus, we omit the number of observations below.
[19] Figure A.2 in the Online Appendix displays the analogous graphs for each market session.
[20] Fractional logit regressions yield the same results. See Table A.2 in the Online Appendix.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

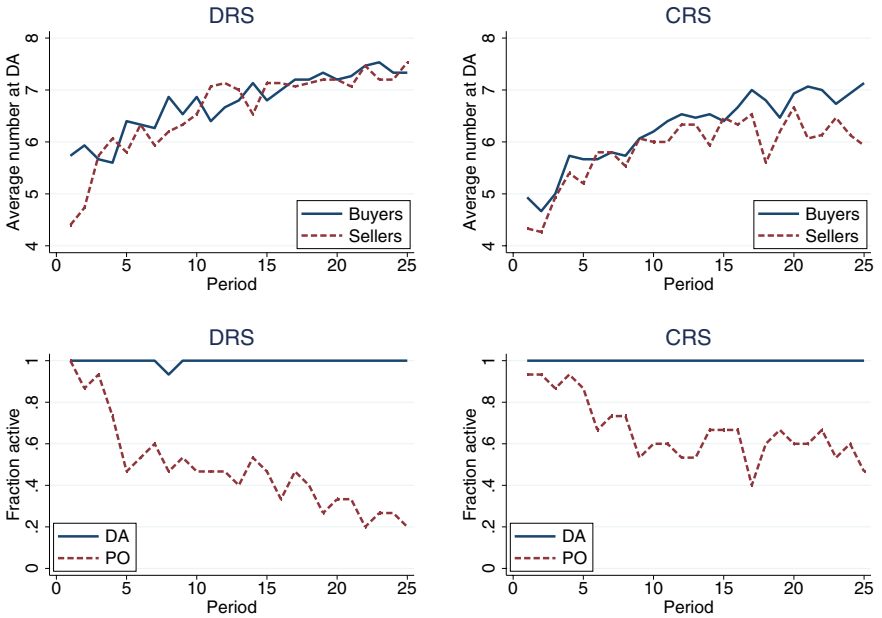


Fig. 3 Top: Average (over 15 markets) number of buyers/sellers at DA per period for DRS (left) and CRS (right). Bottom: Average (over 15 markets) fraction of markets where DA/PO is active per period for DRS (left) and CRS (right)

Table 3 Linear random effects regressions for DRS

	Number of buyers at DA		Number of sellers at DA	
	1	2	3	4
Period	0.0734***	0.0636***	0.0959***	0.0864***
	(0.0074)	(0.0074)	(0.0088)	(0.0090)
RatioDA		−0.5988***		0.1776
		(0.1308)		(0.1373)
Constant	5.8000***	6.5959***	5.3580***	5.3413***
	(0.3470)	(0.2954)	(0.3790)	(0.3387)
R^2 (overall)	0.1193	0.1652	0.1663	0.1194
Observations	375	359	375	359
Markets	15	15	15	15

Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

The bottom-right panel of Figure 3 shows the fraction of markets in CRS where DA and PO are active over time. DA is always active, whereas PO is only active 66.4% of the time, which is significantly less often than DA (WSR: $z = 3.333$, $p < 0.001$). Buyers and sellers favor DA over PO, that is, on average there are both

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Table 4 Linear random effects regressions for CRS

	Number of buyers at DA		Number of sellers at DA	
	1	2	3	4
Period	0.0882***	0.0836***	0.0622***	0.0548***
	(0.0078)	(0.0078)	(0.0101)	(0.0107)
RatioDA		−0.5331***		0.3603**
		(0.1127)		(0.1558)
Constant	5.1153***	5.7994***	5.0507***	4.7695***
	(0.3315)	(0.2999)	(0.3159)	(0.3279)
R^2 (overall)	0.1629	0.1870	0.0715	0.0495
Observations	375	360	375	360
Markets	15	15	15	15

Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

more buyers and more sellers at DA (WSR, buyers: $z = 3.352$, $p < 0.001$; sellers: $z = 3.352$, $p < 0.001$).

As in the previous treatment, we compare coordination in parts 1 (periods 2–9) and 3 (periods 18–25). We first consider the fraction of time DA or PO are active in part 1 vs part 3 of CRS. Because in CRS DA is always active in all periods, there is no difference for this platform. On the other hand, PO is active more often in part 1 (78.3% of the time) than in part 3 (59.2%; WSR, $z = 2.094$, $p = 0.036$). The analogous statement holds for the number of buyers and sellers at DA, which is significantly larger in part 1 compared to part 3 (WSR, buyers: $z = -3.401$, $p < 0.001$; sellers: $z = -2.902$, $p = 0.004$).

4.3 Treatment comparison

We now compare the results of treatment CRS with those of the DRS treatment. Examination of Figure 3 suggests that convergence toward coordination on DA is slower in the case of constant returns to scale. A treatment comparison shows that, indeed, there are on average more buyers and more sellers at DA in DRS than in CRS (Mann-Whitney-U test, MWU, buyers: $N = 30$,[21] $z = 1.660$, $p = 0.097$; sellers: $z = 2.158$, $p = 0.031$).[22]

The differences are also reflected in the comparison between the beginning and the end of the experiment. Although there is no difference in activity for DA or PO in part 1 (MWU, DA: $z = -1.000$, $p = 0.317$; PO: $z = -1.272$, $p = 0.204$), in part 3 PO is marginally more active in CRS than in DRS (MWU, $z = -1.721$, $p = 0.085$).

[21] All MWU tests in sects. 4 and 5 refer to comparisons across treatments and hence have $N = 30$. Thus, we omit the number of observations below.

[22] PO is also active more often in CRS (66.4% of the time) than in DRS (48.0%), but the difference fails to reach significance (MWU, $z = -1.519$, $p = 0.129$).

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Activity is, however, a somewhat coarse measure, hence we turn to the number of buyers and sellers at DA as a more fine-grained measure of convergence. With this measure, we find that already at the beginning of the experiment more buyers and sellers coordinate on DA in DRS compared to CRS (MWU; buyers, $z = 1.787$, $p = 0.074$; sellers, $z = 1.808$, $p = 0.071$). The difference subsists at the end of the experiment (buyers, $z = 1.660$, $p = 0.097$; sellers, $z = 2.926$, $p = 0.003$).

Table 5 displays the results of linear regressions comparing the number of buyers and sellers in the first and last parts of the experiment (rounds 2-9 and 18-25, respectively) across treatments. That is, the dummy CRS takes the value 1 for market observations in that treatment. We observe that the number of sellers at DA is significantly smaller for CRS market observations, but only in part 3.[23]

5 Results: Efficiency and the Gains from Trade

The previous section has concentrated on the market selection decision and shown that traders learn to gradually coordinate on the DA platform, although convergence is stronger under decreasing returns to scale. This section examines trading decisions within each institution, and asks whether actual trade is efficient and which market side is able to reap larger shares of the gains from trade.

5.1 Trading Efficiency and Market-Clearing

We consider two dimensions of efficiency of an institution. The first is the difference between the actual price and the equilibrium price interval $p^*(m,n)$ given that n buyers and m sellers are currently present at that institution. The second is the fraction of realized gains (trader surplus) relative to the maximal possible gains of trade $S^*(m,n)$.

Figure 4 shows the evolution of the average price at both institutions relative to the equilibrium price, for both treatments. For illustrative purposes, what is plotted is the difference $p - p^*$, where p is the realized average price and p^* is the equilibrium price, if the latter is unique. If $p^*(m,n)$ is a proper interval, we plot the difference with the closest value in that interval.

To evaluate how close realized prices were to equilibrium ones, we consider the distance, i.e. the (average across markets of the) absolute value of the average difference each period. In DRS, realized prices where closer to equilibrium in DA (average distance 2.86) than in PO (9.45; WSR, $z = -3.294$, $p < 0.001$). This is also true in CRS, although the difference is smaller. In this treatment, the average distance between realized and equilibrium prices was 7.20 for DA and 9.25 for PO (WSR, $z = -2.499$, $p = 0.013$). Comparing treatments, for DA the distance to equilibrium prices is larger in CRS than in DRS (MWU, $z = -4.500$, $p < 0.001$), whereas we find no significant difference across treatments for PO ($z = 0.477$, $p = 0.633$).

[23] Fractional logit regressions yield the same results. See Table A.3 in the Online Appendix.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Table 5 Regressions for comparison DRS vs CRS in first and last part

	Number of buyers at DA		Number of sellers at DA	
	Rounds 2-9	Rounds 18-25	Rounds 2-9	Rounds 18-25
Period	0.1536***	0.0365**	0.1849***	0.0278
	(0.0351)	(0.0178)	(0.0412)	(0.0268)
CRS	−0.6583	−0.4500	−0.5167	−1.1000**
	(0.4058)	(0.4318)	(0.4146)	(0.4613)
Constant	5.3554***	6.5484***	4.8746***	6.6528***
	(0.3550)	(0.4677)	(0.3593)	(0.6150)
R^2 (overall)	0.0864	0.0346	0.0752	0.1238
Observations	240	240	240	240
Markets	30	30	30	30

Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Figure 5 displays the evolution of the average fraction of the maximal gains of trade which were actually realized in both treatments and for both institutions. In DRS the double auction realizes on average 95.8% of the maximum possible gains of trade, whereas in PO realized gains are only 71.2%. The difference is highly significant (WSR, $z = 3.408$, $p < 0.001$). In CRS, DA realizes 98.1% of all possible gains, while in PO realized gains are only 83.5%. Again, the difference is highly significant (WSR, $z = 3.408$, $p < 0.001$). Realized gains are higher in CRS than in DRS both for DA (MWU, $z = -4.500$, $p < 0.001$) and for PO (MWU, $z = -2.344$, $p = 0.019$).

In summary, we find that DA is able to realize a larger fraction of the maximal gains of trade than PO in both treatments, coming very close to the theoretical maximum. However, the difference is larger in the DRS treatment. To see this, we compare the difference in differences and find that the relative advantage of DA over PO is larger in DRS (24.0%) compared to CRS (14.6%; MWU, $z = 1.970$, $p = 0.049$).

5.2 The Distribution of Gains from Trade

The previous subsection has examined efficiency at the level of an institution. In this subsection, we briefly compare the gains from trade separately for buyers and sellers, which illuminates the underlying dynamics. Figure 6 displays the per capita gains from actual trade over time in both treatments and both institutions, separately for buyers and for sellers (that is, the figure plots total gains of a market side divided by the number of actually-present traders from that side, averaged across active markets). Whereas in DRS the gains from trade remain of a comparable magnitude for both market sides and for each given market institution, with sellers sightly above buyers, in CRS a large gap opens in DA, with buyers' gains steadily rising while those from sellers decline (the average difference between per capita buyers' and sellers' gains is -2.0 ECUs in part 1 and

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

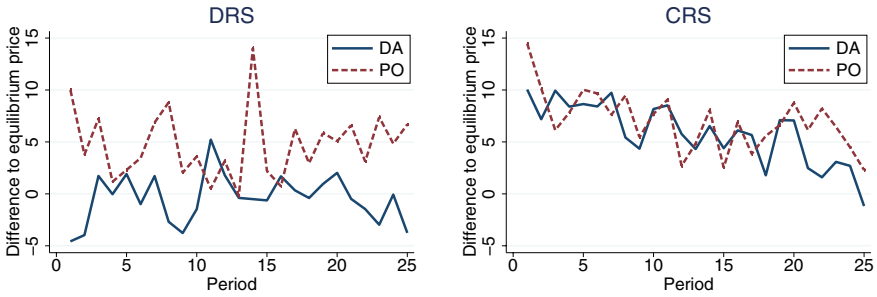


Fig. 4 Average (over 15 markets) of the differences between average realized prices and the equilibrium price ($p - p^*$, or the difference with the closest value in the interval $p^*(m, n)$) at PO vs. DA per period

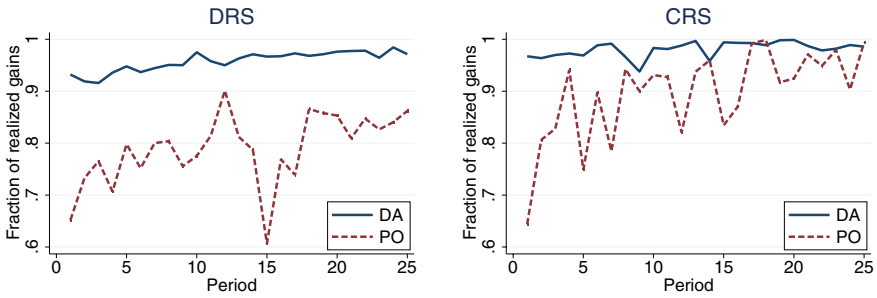


Fig. 5 Average (over 15 markets) realized gains (as fraction of maximal possible gains) at PO vs. DA per period

52.9 in part 3; WSR, $z = -3.408$, $p < 0.001$). This gap also exists but is smaller in the case of PO (the corresponding average difference in part 3 for PO is 21.6, which is smaller than the one of DA; WSR, $z = -2.840$, $p = 0.005$).

The interpretation is as follows. Over time, more traders move toward the double auction institution (Fig. 3, top-right). In this institution, provided enough sellers are present, the price should converge to market clearing, which under constant returns to scale implies zero profits for the sellers (Fig. 4, right). Convergence, however, takes time, and even more so in a setting where traders can switch away from a given institution. This means that profits remain positive for sellers even as more of them them flock to the double auction institution and the price in the latter drops. Hence, even as the price becomes relatively close to marginal cost and buyers' are able to reap larger and larger shares of the gains from trade, the institution remains comparatively attractive for sellers. Sellers are effectively facing a tradeoff between an institution (PO) that favors them but

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

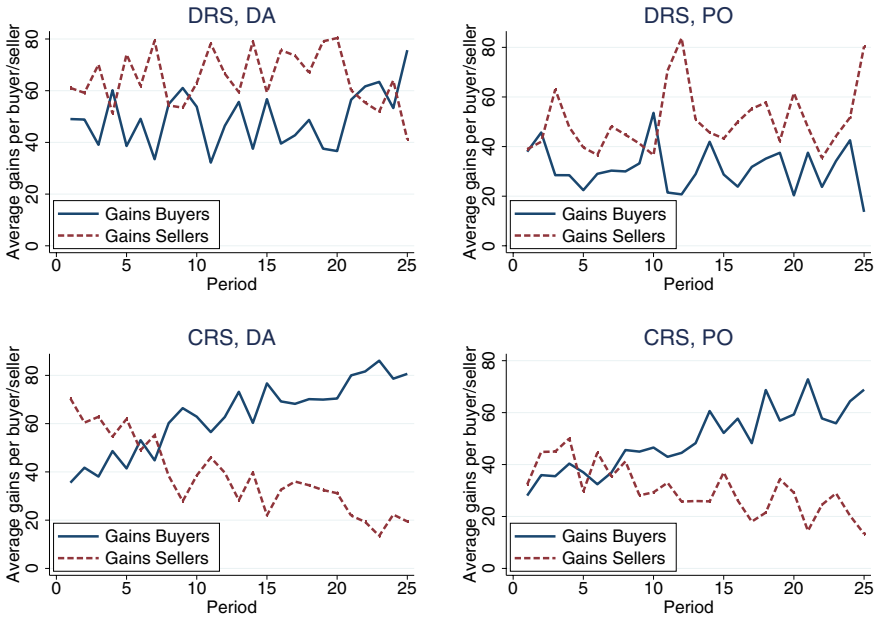


Fig. 6 Average (over active markets) gains of buyers and sellers per period

is gradually becoming empty, and a more efficient institution (DA) where their profits might eventually vanish but remain positive within a finite horizon.

6 Results: Switching Behavior

In this section we look at individual behavior regarding market selection. In particular, we look at market switching behavior, that is, the traders' decision whether to leave the market they are currently at and select the other market for the next period.[24] A *(market) switch* is a market choice where a subject selects a market for the next round that is different from the market he is currently at. A total of 19.8% of all market-choice decisions were switches, 21.1% in the case of sellers and 18.4% in the case of buyers (obviously excluding the first period).

The previous literature on experimental markets has often looked at simple behavioral rules of thumb (e.g., imitation or myopic best reply) within single markets. In particular, and following the theoretical model of Vega-Redondo (1997), a series of experimental Cournot oligopoly markets examined the role of imitation in symmetric settings where one market side is summarized by an aggregate demand

[24] The analysis of individual-level trading decisions conditional on a selected market could run into endogeneity issues. This is not a concern here, because in this section we concentrate on the decision of which market to choose, and not on subsequent trading decisions given that choice.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

function and all subjects play the role of producers. Those experiments generally found evidence in favor of imitative rules when information was centered on individual actions and profits, with possible shifts toward more forward-looking rules as myopic best reply when more information on the market structure was provided (e.g., Huck et al. 1999; Offerman et al. 2002; Apesteguía et al. 2007, 2010), although results depend on the number of players, the length of the experiment, and other factors (Bosch-Domènech and Vriend 2003; Huck et al. 2004; Friedman et al. 2015; Oechssler et al. 2016). For instance, Huck et al. (2002) suggested that their data would be consistent with a mixture of best reply and imitation. Those experiments, however, all consider production behavior within a single market institution. The exception is Alós-Ferrer and Kirchsteiger (2015), who conducted an experiment on platform selection where, however, actual trade followed a reduced-form approach (that is, traders did not make actual trading decisions). Alós-Ferrer and Kirchsteiger (2015) found evidence in favor of switching behavior reflecting which institutions had the maximum observed payoffs for the trader's market side, which can be seen as a form of imitation.

Following this literature, we hypothesized that simple rules of thumb would be able to explain switching behavior in our data.[25] We consider three behavioral rules that specify switches of a particular type depending on different (observable) market outcomes. A buyer (seller) following a *Best-Price* rule chooses the institution with lowest (highest) observed price in the current period, effectively chasing after the best observed price. To prevent trivial cases, the rule also prescribes to avoid institutions where no trade occurred, hence no price was realized. A buyer (seller) following a *Best-Average-Price* rule chooses the institution with the lowest (highest) average price in the current period, hence taking into account that a single price might not be representative of the institution (trivial cases are avoided as in Best-Price). A more forward-looking rule is to focus on the trader ratio as a predictor of future performance. In particular, in our setting, the equilibrium market price is a function of the trader ratio, and hence focusing on the latter could be seen as an attempt to predict the former, and in this sense this rule is closer to myopic best reply. A buyer following a *Best-Ratio* rule chooses the institution with lowest ratio of buyers to sellers in the current period (and follows the sellers if they are all at the same institution). Analogously, a seller following this rule chooses the institution with lowest ratio of sellers to buyers in the current period (and follows the buyers if they are all at the same institution).[26]

[25] The link to imitation is particularly interesting, because certain forms of this behavioral rule (Schlag 1998; Alós-Ferrer and Schlag 2009) give rise to aggregate-level dynamics closely related to the well-known replicator dynamics of evolutionary game theory (see Hofbauer and Sigmund 1998, Chapter 8). This link, however, has rarely been made explicit for the case of market selection. An exception is Rud and Rabanal (2018), who conduct a simulation analysis of market selection with trader heterogeneity relying on individual-level replicator dynamics to determine the time allocation across different market types.

[26] Note that the highest price and the highest average price are displayed explicitly in the market selection stage, whereas the ratio of buyers to sellers is not shown explicitly but can be computed fairly easily as both the number of buyers and sellers at each institution is shown on screen.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

The upper part of Table 6 shows the percentage of decisions that are consistent with each rule's prescriptions, in addition to the percentage of switching decisions, separately for buyers and sellers. All three rules explain a large percentage of decisions, with the simplest one, BestPrice, explaining the most both for buyers (81.0%) and for sellers (75.1%), but with the most sophisticated, BestRatio, coming close for sellers (74.3%). As decision inertia might explain a part of the non-switching decisions, the lower part of the table shows the fraction of *switches* that are consistent with following each rule, separately for trader type and treatment. BestRatio explains most switches (between 73.3% and 78.6%) for both trader types and in both treatments.

Given these results, we further analyze the predictive power of BestPrice and BestRatio in a series of random effects probit regressions (Tables 7 and 8). The dependent variable is switching behavior (a dummy variable taking the value 1 if there was a switch), but the regressors are continuous variables reflecting the price or ratio differences. The regression hence considers stochastic versions of the rules where a switch is more likely the larger the price difference (or ratio difference) is. The difference in the best price is coded as the difference between the currently chosen institution and the other institution for buyers and conversely for sellers. The difference in ratios is coded as the difference in buyer-seller ratios between the current institution and the other one for buyers, and the analogous difference in seller-buyer ratios for sellers. To avoid the natural (nonlinear) asymmetry of ratios (all ratios toward one side are condensed between 0 and 1, while ratios toward the other side are in principle unbounded), we consider the logarithm of the ratio. Last, to make the coefficients of the different rules comparable, we divide each of the two resulting variables by their respective empirical standard deviation.[27] We remark that price differences and log of ratios are undefined if either institution is empty in a given period and hence the regressions must exclude those observations. However, in those cases all rules would make the same prediction, thus they play no role to compare the rules (the regressions, however, underestimate the predictive power of the rules for that reason).

The regressions show that both BestPrice and BestRatio are predictive of switches for both trader types if considered separately. However, when taken together, in the case of buyers only BestPrice remains significant. That is, buyer behavior appears to be usefully summarized by chasing after low prices. In contrast, for sellers both rules remain significantly predictive when considered jointly. That is, sellers seem to take into account both past high prices and favorable trader ratios, suggesting a stronger focus on forward-looking evaluations. Notably, this observation holds for both treatments, DRS and CRS.[28]

[27] That is, we essentially standardize the difference variables, dividing the values by the standard deviation in the 30×24 dataset containing the price and trader-ratio differences in the experiment. We do not subtract the empirical mean because the value zero has a specific meaning (same prices or same ratios in both institutions).

[28] The Online Appendix (Tables A.4 and A.5) reports a robustness analysis controlling for the institution that a trader was at in the previous period. All previous conclusions remain unchanged, but the results also suggest that BestPrice remains predictive independently of the institution, while BestRatio is predictive only for traders who were at DA. That is, traders who just traded in a posted offers setting focus on prices but not on buyer-seller ratios.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Table 6 Consistent decisions and consistent switches

	Switches	Consistent decisions					
		BestPrice		BestAvgPrice		BestRatio	
Seller	21.1%	75.1%		69.4%		74.3%	
Buyer	18.4%	81.0%		77.7%		67.2%	
Consistent switches		BestPrice		BestAvgPrice		BestRatio	
		DRS	CRS	DRS	CRS	DRS	CRS
Seller		61.9%	62.4%	63.6%	65.7%	77.1%	78.6%
Buyer		59.0%	59.6%	62.1%	66.5%	76.9%	73.3%

Notes BestPrice, BestAvgPrice, BestRatio indicate that institution with best price, best average price (for own market side), or best buyer/seller ratio, respectively, was chosen. Consistent switches are switches in line with the respective rule

Table 7 Random effect probit regression on switches for buyers

Switch	DRS			CRS		
	1	2	3	4	5	6
DiffBestPrice	0.268^{***}		0.266^{***}	0.358^{***}		0.365^{***}
	(0.026)		(0.028)	(0.038)		(0.041)
DiffLogRatio		0.197^{***}	0.006		0.148^{***}	−0.016
		(0.041)	(0.046)		(0.035)	(0.040)
Period	-0.032^{***}	-0.035^{***}	-0.032^{***}	-0.044^{***}	-0.048^{***}	-0.043^{***}
	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.005)
Constant	-0.229^{***}	-0.396^{***}	-0.232^{**}	-0.293^{***}	-0.324^{***}	-0.288^{***}
	(0.089)	(0.095)	(0.091)	(0.080)	(0.081)	(0.081)
Observations	1350	1406	1350	1873	1929	1873

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 8 Random effect probit regression on switches for sellers

Switch	DRS			CRS		
	1	2	3	4	5	6
DiffBestPrice	0.243^{***}		0.174^{***}	0.251^{***}		0.130^{***}
	(0.035)		(0.039)	(0.033)		(0.038)
DiffLogRatio		0.278^{***}	0.190^{***}		0.284^{***}	0.231^{***}
		(0.043)	(0.048)		(0.032)	(0.038)
Period	-0.044^{***}	-0.039^{***}	-0.041^{***}	-0.026^{***}	-0.020^{***}	-0.022^{***}
	(0.007)	(0.007)	(0.007)	(0.005)	(0.005)	(0.005)
Constant	-0.375^{***}	-0.497^{***}	-0.448^{***}	-0.226^{***}	-0.353^{***}	-0.308^{***}
	(0.089)	(0.089)	(0.091)	(0.068)	(0.068)	(0.069)
Observations	1349	1405	1349	1877	1933	1877

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

The striking asymmetry in behavioral drivers between buyers and sellers is natural. In principle, in DA and under decreasing returns to scale, both trader types are treated symmetrically. However, PO is biased in favor of the sellers, and buyers are essentially passive in this latter institution. Thus, there is an institutional asymmetry which creates a tradeoff for sellers which is absent for the case of buyers. On the one hand, sellers would like to reap the high profits associated with larger prices typical of PO. On the other hand, the inefficiency of the latter institution attracts less buyers, thus leading to worse ratios for the sellers. This also holds empirically for constant returns to scale, even though in this case coordination on DA would make the sellers' profits disappear entirely. The latter observation is less surprising in view of previous sections, since the slow speed of convergence guarantees that profits in DA, although dwindling, remain positive for sellers.

It is also interesting to note that previous empirical work on experimental markets (within a fixed institution, and with respect to trading behavior) has often suggested behavioral heterogeneity, with some subjects relying more on imitative rules and others on myopic best reply (Huck et al. 1999, 2002), or behavior reflecting more or less sophisticated rules depending on available information (Offerman et al. 2002). In our markets, all traders have the same information, and they are randomly allocated to the role of buyers or sellers. Thus, the difference in behavior that we observe between buyers and sellers is causally induced by their different roles (and, specifically, the presence of an institution which treats them differently), and does not arise from cognitive or informational differences.

7 Conclusion

Most goods can be traded in different market institutions or platforms. The economics literature has intensively studied trading when the market institution is fixed, but market selection has received comparatively little attention. We study both, market selection itself, and how the presence of alternative institutions affects trade. We pit a double-auction environment, where buyers and sellers can freely trade without an auctioneer, against a posted-offers institution, where buyers are constrained to accept or reject the offers made by sellers. Although the former has been empirically shown to induce quick convergence to market clearing in isolation, the latter embodies a bias in favor of the sellers, and thus efficiency arguments suggest that traders should learn to coordinate on the double auction.

We observe coordination toward the double-auction institution, which is faster when sellers face decreasing returns to scale compared to when they are endowed with constant per-unit production costs. That is, under decreasing returns to scale, traders learn to coordinate on the efficient double auction, and the posted-offers institution becomes inactive in finitely many rounds. In contrast, under constant returns to scale, although traders tend to concentrate on the double auction, both institutions remain active until the last round. The reason is that, in the latter case, sellers try to stave off (but ultimately are not able to prevent) coordination in the double auction and the resulting convergence to market clearing prices (which would leave them with no profits).

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Within institutions, and even though traders might "jump boat" at any point, trading behavior approaches equilibrium predictions in terms of prices and gains of trade. Switching behavior across institutions, however, can be explained by simple rules of thumb. Those are different across trader types: buyers seem to simply chase after low observed prices, while sellers consider both high observed prices and trader ratios. The latter might be seen as a myopic predictor of future market prices and hence reveal a more forward-looking mindset. Because traders were randomly allocated to the roles of buyers or sellers, the differences in behavior were causally induced.

Our work empirically demonstrates that efficiency alone might not suffice for a market institution to drive competing ones out of the field. Under decreasing returns to scale, fast coordination occurs because the double auction Pareto dominates the alternative, in the sense that both trader types can reap larger gains from trade. Under constant returns to scale, gains from trade are monopolized by buyers, and sellers' earnings approach zero as prices approach the market clearing level. Thus, the distribution of gains from trade is unequal and sellers do not benefit from coordination on the efficient institution. Although, in our data, this translated into a slowdown of convergence, in general one might speculate that once profits become small enough, sellers might, in some cases and for some parameter constellations, successfully manage to turn the tide and preserve biased, inefficient institutions in the market. This might contribute to explain why multiple, possibly-inefficient institutions typically coexist for single, given goods in actual markets.

Supplementary information The online version contains supplementary material available at https://doi.org/10.1007/s10683-021-09710-1.

Funding Open Access funding provided by Universität Zürich. Financial support from the German Research Foundation (DFG) through project AL-1169/5-1 is gratefully acknowledged.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

Alós-Ferrer, C., & Kirchsteiger, G. (2015). Learning and Market Clearing: Theory and Experiments. *Economic Theory, 60*(2), 203–241.

Alós-Ferrer, C., & Kirchsteiger, G. (2017). Market Selection by Boundedly-Rational Traders under Constant Returns to Scale. *Economics Letters, 153,* 51–53.

Alós-Ferrer, C., Kirchsteiger, G., & Walzl, M. (2010). On the Evolution of Market Institutions: The Platform Design Paradox. *Economic Journal, 120*(543), 215–243.

Alós-Ferrer, C., & Schlag, K. (2009). Imitation and Learning. In P. Anand, P. Pattanaik, & C. Puppe (Eds.), *The Handbook of Rational and Social Choice*. Oxford: Oxford University Press.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Apesteguía, J., Huck, S., & Oechssler, J. (2007). Imitation-Theory and Experimental Evidence. *Journal of Economic Theory, 136,* 217–235.
Apesteguía, J., Huck, S., Oechssler, J., & Weidenholzer, S. (2010). Imitation and The Evolution of Walrasian Behavior: Theoretically Fragile But Behaviorally Robust. *Journal of Economic Theory, 145,* 1603–1617.
Ariely, D., Ockenfels, A., & Roth, A. (2005). An Experimental Analysis of Ending Rules in Internet Auctions. *RAND Journal of Economics, 36,* 890–907.
Ausubel, L. M., Cramton, P., Pycia, M., Rostek, M., & Weretka, M. (2014). Demand Reduction and Inefficiency in Multi-Unit Auctions. *Review of Economic Studies, 81*(4), 1366–1400.
Bosch-Domènech, A., & Vriend, N. J. (2003). Imitation of Successful Behaviour in Cournot Markets. *Economic Journal, 113*(487), 495–524.
Campbell, J., LaMaster, S., Smith, V. L., & Van Boening, M. (1991). Off-Floor Trading, Disintegration, and the Bid-Ask Spread in Experimental Markets. *Journal of Business, 64,* 495–522.
Chamberlin, E. H. (1948). An Experimental Imperfect Market. *Journal of Political Economy, 56*(2), 95–108.
Crockett, S., Friedman, D., & Oprea, R. (2020). Naturally Occurring Preferences and General Equilibrium: A Laboratory Study. *International Economic Review*, forthcoming.
Crockett, S., Oprea, R., & Plott, C. R. (2011). Extreme Walrasian Dynamics: The Gale Example in the Lab. *American Economic Review, 101*(7), 3196–3220.
Fischbacher, U. (2007). z-Tree: Zurich Toolbox for Ready-Made Economic Experiments. *Experimental Economics, 10*(2), 171–178.
Friedman, D., Huck, S., Oprea, R., & Weidenholzer, S. (2015). From Imitation to Collusion: Long-Run Learning in a Low-Information Environment. *Journal of Economic Theory, 155,* 185–205.
Genesove, D. & Hansen, J. (2016). The Role of Auctions and Negotiations in Housing. CEPR Discussion Paper 11392.
Gillen, B. J., Hirota, M., Hsu, M., Plott, C. R., & Rogers, B. W. (2020). Divergence and Convergence in Scarf Cycle Environments: Experiments and Predictability in the Dynamics of General Equilibrium Systems. *Economic Theory*, forthcoming.
Greiner, B. (2015). Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE. *Journal of the Economic Science Association, 1,* 114–125.
Hayek, F. (1967). *Studies in Philosophy*. London: Politics and Economics. Routledge & Paul Kegan.
Hofbauer, J., & Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge, UK: Cambridge University Press.
Huck, S., Normann, H.-T., & Oechssler, J. (1999). Learning in Cournot Oligopoly - An Experiment. *Economic Journal, 109,* C80–C95.
Huck, S., Normann, H.-T., & Oechssler, J. (2002). Stability of the Cournot Process – Experimental Evidence. *International Journal of Game Theory, 31,* 123–136.
Huck, S., Normann, H.-T., & Oechssler, J. (2004). Two Are Few and Four Are Many: Number Effects in Experimental Oligopolies. *Journal of Economic Behavior and Organization, 53*(4), 435–446.
Klein, B., C, R., & Alchian, A. (1978). Vertical Integration, Appropriable Rents, and the Competitive Contracting Process. *Journal of Law and Economics, 21*(2), 297–326.
Kugler, T., Neeman, Z., & Vulkan, N. (2006). Markets versus Negotiations: An Experimental Investigation. *Games and Economic Behavior, 56*(3), 121–134.
Lucking-Reiley, D. (2000). Auctions on the Internet: What's Being Auctioned, and How? *Journal of Industrial Economics, 48*(3), 227–252.
Madhavan, A. (1992). Trading Mechanisms in Securities Markets. *Journal of Finance, 47*(2), 607–641.
Oechssler, J., Roomets, A., & Roth, S. (2016). From Imitation to Collusion: A Replication. *Journal of the Economic Science Association, 2,* 13–21.
Offerman, T., Potters, J., & Sonnemans, J. (2002). Imitation and Belief Learning in an Oligopoly Experiment. *Review of Economic Studies, 69*(4), 973–997.
Plott, C. R. (1982). Industrial Organization Theory and Experimental Economics. *Journal of Economic Literature, 20,* 1485–1587.
Plott, C. R. (2000). Market Stability: Backward-Bending Supply in a Laboratory Experimental Market. *Economic Inquiry, 38,* 1–18.
Plott, C. R., & Gray, P. (1990). The Multiple Unit Double Auction. *Journal of Economic Behavior and Organization, 13,* 245–258.
Plott, C. R., & Smith, V. L. (1978a). An Experimental Examination of Two Exchange Institutions. *Review of Economic Studies, 45,* 133–153.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press

Plott, C. R., & Smith, V. L. (Eds.). (1978b). *Handbook of Experimental Economics Results,* (Vol. 1). Amsterdam: North Holland.
Quan, D. C. (1994). Real Estate Auctions: A Survey of Theory and Practice. *Journal of Real Estate Finance and Economics, 9,* 23–49.
Reynolds, S., & Wooders, J. (2009). Auctions with a Buy Price. *Economic Theory, 38,* 9–39.
Rud, O. A., & Rabanal, J. P. (2018). Evolution of Markets: A Simulation with Centralized, Decentralized and Posted Offer Formats. *Journal of Evolutionary Economics, 28*(3), 667–689.
Rustichini, A., Satterthwaite, M. A., & Williams, S. R. (1994). Convergence to Efficiency in a Simple Market with Incomplete Information. *Econometrica, 62*(5), 1041–1063.
Schlag, K. H. (1998). Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-Armed Bandits. *Journal of Economic Theory, 78*(1), 130–156.
Smith, V. L. (1962). An Experimental Study of Competitive Market Behavior. *Journal of Political Economy, 70*(2), 111–137.
Van Huyck, J., Battalio, R. C., & Rankin, F. (1997). On the Origin of Convention: Evidence from Coordination Games. *Economic Journal, 107*(442), 576–597.
Vega-Redondo, F. (1997). The Evolution of Walrasian Behavior. *Econometrica, 65*(2), 375–384.
Williams, A., Smith, V. L., Ledyard, J., & Gjerstad, S. (2000). Concurrent Trading in Two Experimental Markets with Demand Interdependence. *Economic Theory, 16,* 511–528.
Williamson, O. (1975). *Markets and Hierarchies: Analysis and Antitrust Implications*. New York: Free Press.
Young, H. P. (1993). The Evolution of Conventions. *Econometrica, 61,* 57–84.
Zumpano, L. V., Elder, H. W., & Baryla, E. A. (1996). Buying a House and the Decision to Use a Real Estate Broker. *Journal of Real Estate Finance and Economics, 13,* 169–181.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

https://doi.org/10.1007/s10683-021-09710-1 Published online by Cambridge University Press