Econ Theory (2015) 60:203–241
DOI 10.1007/s00199-015-0885-8

RESEARCH ARTICLE

Learning and market clearing: theory and experiments

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

Received: 3 July 2013 / Accepted: 14 May 2015 / Published online: 4 June 2015

Abstract This paper investigates theoretically and experimentally whether traders learn to use market-clearing trading institutions or whether other (inefficient) market institutions can survive in the long run. Using a framework with boundedly rational traders, we find that market-clearing institutions are always stable under a general class of learning dynamics. However, we show that there exist other, non-market-clearing institutions that are also stable. Therefore, in the long run, traders may fail to coordinate exclusively on market-clearing institutions. Using a replica-economies approach, we find the results to be robust to large market size. The theoretical predictions were confirmed in a series of platform choice experiments. Traders coordinated on platforms predicted to be stable, including market-clearing as well as non-market-clearing ones, while platforms predicted to be unstable were avoided in the long run.

Keywords Market institution · Market clearing · Coordination · Learning

JEL Classification C72 · D4 · D83

✉ Carlos Alós-Ferrer
carlos.alos-ferrer@uni-koeln.de

Georg Kirchsteiger
gkirchst@ulb.ac.be

1 Department of Economics, University of Cologne, Albertus-Magnus Platz, 50923 Cologne, Germany

2 ECARES, Université Libre de Bruxelles, Avenue F. D. Roosevelt 50, CP 114, 1050 Brussels, Belgium

1 Introduction

The formation of a market requires a group of agents, some of them willing to buy and some of them willing to sell. Preferences and cost functions are sufficient to develop a theory whether market clearing is taken as granted. Actual markets, though, are not merely characterized by demand and supply. Market exchange requires an institutional framework in which action and message sets are specified and in which a process of matching and price formation can take place.

An enormous variety of market institutions can be observed in the field, even for the same good. Financial assets are traded following many different procedures (see Madhavan 1992), e.g., call markets and continuous double auction. Real estate is sold at auctions (Quan 1994), through brokers (Zumpano et al. 1996), and by means of direct negotiations. In some countries' rental markets, established but *informal* institutions bias the market in favor of the owners (e.g., group tenant visits in Switzerland). In other countries, middlemen act as platforms which compete actively for tenants. In addition, there is always an alternative "word-of-mouth" market tenants might resort to.

These details of the market institution are consequential. In addition to theoretical and empirical evidence, there is a large body of experimental evidence in this direction.[1] Trading rules affect the efficiency of the market outcome, the convergence toward equilibrium, the volatility of the prices, and the distribution of surplus over the market participants. Given that "institutions matter" and given the competition between different market institutions, we might ask which institutions are used in the long run. What are the properties of successful institutions? Do the surviving market institutions support market clearing and efficient outcomes? Are there circumstances under which inefficient trading rules can persist, or are forces and mechanisms present that drive a market toward efficient organization?

Consumer-to-consumer internet auction platforms provide a good example of competition between different market institutions. Although eBay currently has a dominating position in this market, its competitors in the past included prominent examples as Yahoo or Amazon (for an early history of competition in online auctions, see Lucking-Reiley 2000). The trading rules of these auction platforms differed, e.g., in their ending rules and in the type of the Buy-Now option sellers could use. Experimental (Ariely et al. 2005) and theoretical analysis (Reynolds and Wooders 2009) reveals that the level of realized prices as well as efficiency are influenced by these differences. A similar conclusion can be drawn for the possibility of secret reserve prices (Bajari and Hortaçsu 2003) and for buy prices (Budish and Takeyama 2001). In the context of multi-unit auctions, Ausubel and Cramton (2002) show that uniform and pay-as-bid auctions lead to different realized prices. Hence, competing institutions differ not only in their institutional setup, but the different institutional setups lead to systematic differences in the realized prices.

The survival of a specific market institution depends on whether traders employ this institution or avoid it. The decision about the use of a particular market institution gives

[1] An overview of the classical experimental evidence on the importance of market institutions is provided by Plott (1982) (see also Holt 1995).

rise to a game that combines aspects of a coordination and a minority game. On the one hand, potential buyers and sellers have to coordinate on a particular institution in order to make mutually beneficial trade possible. On the other hand, a trader is better off the fewer competing traders opt for the same trading platform. Due to the coordination aspect, such a game exhibits a multiplicity of Nash equilibria. All the traders might coordinate on an institution that does not lead to market-clearing outcomes. They might even coordinate on an institution that leads to a Pareto-inefficient outcome. Hence, we ask under what circumstances traders will indeed *learn* to coordinate on an efficient, market-clearing institution.

To provide an answer to this question, we conducted a theoretical and experimental study of a market for a homogeneous good. Potential traders have to choose simultaneously at which institution they want to trade. They choose between a market-clearing institution and other institutions that do not lead to market clearing, but realize other prices. Traders who have chosen such an institution might obtain more favorable prices but necessarily face rationing. The theoretical part of the analysis is based on a learning model where each trader has a tendency to switch from one institution to a different one next period if another institution exhibits better current-period results. Traders evaluate the results according to evaluation functions that satisfy a number of weak behavioral assumptions, compatible with standard microeconomic models but allowing also for boundedly rational behavior. The learning model is related to stochastic models of learning in games (see Fudenberg and Levine 1998, for an overview). In particular, traders are not assumed to anticipate future prices, market clearing or otherwise. They tend to switch to strategies (institutions) which are better in the current period, without anticipating the effects of their strategy change. Within this framework, the market-clearing institution is *always* stochastically stable independently of the characteristics and the number of the other available institutions. This strong prediction, however, does not imply that only market clearing will be observed in the long run. On the contrary, we find that certain non-market-clearing institutions are also stochastically stable. Hence, the theoretical analysis suggests that in the long run market-clearing institutions will be used, but in general not exclusively.

The experimental test of this result concentrates on the learning aspect of the theoretical model. More specifically, groups of 14 subjects each had to choose between two or three institutions. In the first treatment, the payoffs were designed in such a way that subjects had to choose between a market-clearing and another stochastically stable institution. In the second treatment, the choice was between the market-clearing and a non-stable institution, while in the third treatment, the choice was between all three institutions. The results show that whenever the market-clearing as well as the other stochastically stable institution were available, there was no tendency to coordinate on a single institution. Both institutions remained active in the long run, i.e., after 90 repetitions of the game. Subjects, though, learned to avoid the non-stable institution when available. We also found strong evidence that individual traders' choice behavior was in accordance with our learning model. Overall, the experimental results confirmed the theoretical predictions.

The possibility that traders might choose between different trading institutions plays a role in several existing models (e.g., Ishibuchi et al. 2002; Kugler et al. 2006; Gerber and Bettzüge 2007). Those, however, do not investigate whether traders learn to coordinate on efficient institutions guaranteeing market-clearing prices and quantities. There also exists a large experimental literature on learning in games, but to the best of our knowledge this literature does not examine the question of on which trading platforms traders coordinate.

The theoretical analysis in the paper at hand is related to our own work on competition among simultaneously available trading institutions. Alós-Ferrer et al. (2010) consider a game among two market designers confronted with boundedly rational buyers and sellers, where all sellers are endowed with a constant-return-to-scale technology. For any given characteristics of the institutions chosen by the market designers, the game played between the buyers and the sellers is a particular case of the model considered here.[2] Alós-Ferrer and Kirchsteiger (2010) consider a related model where boundedly rational traders choose among different, possibly non-market-clearing institutions within a general equilibrium framework. This approach, however, is conceptually different from the model considered here. First, since the focus of Alós-Ferrer and Kirchsteiger (2010) is on rationing, each institution is characterized directly by a parameter determining the amount of rationing. Second, traders' behavior is modeled through probabilistic behavioral rules rather than evaluation functions. Last, neither Alós-Ferrer and Kirchsteiger (2010) nor Alós-Ferrer et al. (2010) provide an experimental test of the underlying learning approach.

The paper proceeds as follows. Next, we describe the model and its basic assumptions. In Sect. 3, we describe the learning process. Sections 4 and 5 present the stability results for market-clearing and non-market-clearing institutions, respectively. Section 6 investigates the robustness of the theoretical results with respect to market size. Section 7 presents the experimental test of our model. Section 8 concludes. Proofs are relegated to "Appendix 1," and the experimental instructions are given in "Appendix 2."

2 The model

There is a homogeneous good to be traded by a finite set I of n buyers and a finite set J of m sellers. We denote the price of the good by p.

A typical buyer will be modeled through a demand function, and a typical seller through a supply function satisfying the following assumptions.

- M1 The demand function $d : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{+\infty\}$ is continuous on $[0, +\infty)$ and strictly decreasing in p in the range where $d(p) > 0$. Further, $d(0) > 0, d(p) \geq 0$ for all $p \geq 0$, and $\lim_{p\to\infty} d(p) = 0$.
- M2 The supply function $s : \mathbb{R}_+ \to \mathbb{R}_+$ is continuous and (weakly) increasing. Further, $s(0) = 0$.
- M3 There exists a price $p > 0$ with $d(p) > 0$ and $s(p) > 0$.

We allow for instance for linear demand functions of the form $d(p) = \max(a - bp, 0)$, but also for everywhere-positive functions as $d(p) = p^{-a}$, which are extended-

[2] Alós-Ferrer et al. (2010) is an example of the "asymmetric rationality" approach. Another, recent example is the location model of technology choice by Shi (2015), where rational managers set maximum capacities and mobility constraints, and boundedly rational agents select a location and a technology.

real because $d(0) = +\infty$. Notice, though, that assumption **M1** implies $d(p) < +\infty$ for all $p > 0$.

For an individual trader, the market outcome is given by the price at which he trades and by the quantity he can trade. In order to model the learning process, we describe how buyers and sellers evaluate the market outcome. Denote by q_S the quantity sold by a typical seller, and by q_B the quantity bought by a typical buyer. The evaluations of the market outcomes, $v_B(q_B, p)$ and $v_S(q_S, p)$, depend on the quantity the traders buy and sell, respectively, and on the price p at which they trade. Hence, the evaluations (payoffs) are given by functions $v_B : \mathbb{R}^2_+ \to \mathbb{R}$ and $v_S : \mathbb{R}^2_+ \to \mathbb{R}$.

The primitives in our model are the demand, supply, and the evaluation functions. We want to emphasize that this framework is more general than the usual microeconomic approach, where demand and supply are derived from maximization of the payoffs (i.e., from utility and profit maximization). We have deliberately chosen this general framework in order to allow for the possibility that demand and supply are not based on rational choices of the agents. Furthermore, in our framework the evaluation of the market outcome, which—as explained later in detail—drives the learning process, need not be identical with consumers' utility and producers' profits. In other words, we allow for more general (even boundedly rational) modes of behavior. For example, our framework allows for producers whose supply is derived from profit maximization, but who evaluate the market outcome by the revenue raised (without taking production costs into account). Such an inconsistency between the supply behavior and the learning process (which might, e.g., be due to the different divisions within a firm deciding about quantity supplied and the market chosen) can be modeled within our approach, since such a model fulfills our core assumptions (explained below). It is worth emphasizing, however, that the usual microeconomic model of utility-maximizing consumers and profit-maximizing producers is also covered by our framework, as we will show later.

Demand and supply are given meaning by the following assumptions which relate them to the evaluation of the market outcome.

A1 In the absence of rationing, a lower price is better for buyers and worse for sellers. That is, for all p, p' with $p < p'$,

$$v_B(d(p), p) > v_B(d(p'), p') \text{ whenever } d(p) > 0,$$
$$\text{and } v_S(s(p), p) < v_S(s(p'), p') \text{ whenever } s(p') > 0.$$

A2 Given the price, traders prefer not to be rationed. That is, for all $p > 0$ and all $0 < q_B < d(p)$, $0 < q_S < s(p)$,

$$v_B(d(p), p) > v_B(q_B, p) \quad \text{and} \quad v_S(s(p), p) > v_S(q_S, p).$$

A3 Given the price, traders prefer being rationed to not being able to trade. That is, for all $p > 0$ and all $0 < q_B < d(p)$, $0 < q_S < s(p)$,

$$v_B(q_B, p) > v_B(0) \quad \text{and} \quad v_S(q_S, p) > v_S(0)$$

where $v_B(0) = v_B(0, p')$ and $v_S(0) = v_S(0, p')$ for all $p' \geq 0$ are the payoffs of not being able to trade, which we explicitly assume not to depend on (hypothetical) prices.

Essentially, these assumptions are fulfilled as long as traders focus on getting more favorable prices and dislike rationing. Of course, as the next examples show, standard models fulfill **A1–A3**, but we want to emphasize that our results depend only on these minimal properties.

Note that **A3** does not state that traders prefer trading at any price over no trade. It states that, if an institution offers a price such that the trader wishes to trade according to the demand or supply function, then the evaluation of not trading is lower than the evaluation of actually trading at that price. In other words, only if the price is low (high) enough that the demand (supply) is strictly positive, buyers (sellers) prefer rationing at a strictly positive quantity over not being able to trade at all.

Example 1 (*Utility and Profit Maximization*) A first example fulfilling all assumptions above is obtained as follows. Consider identical consumers endowed with a strictly quasiconcave, continuous, and strictly monotone utility function. Fix the prices of all goods except good 1, and denote by p_{-1} the vector of (fixed) prices of goods other than 1. Assume the (reduced) demand function for good 1, $d(p_1) = x_1(p_1, p_{-1})$, to be strictly decreasing in p_1 (ruling out that it is a Giffen good). The consumers' evaluation of the market outcome is simply given by the utility derived from this outcome. That is,

$$v_B(q_B, p_1) = u(q_B, x_{-1}(W - p_1 q_B, p_{-1})),$$

where W denotes the consumer's wealth, q_B is the quantity of good 1 actually bought by a buyer at the chosen institution (that is, taking into account possible rationing), and $x_{-1}(W - p_1 q_B, p_{-1})$ is the optimal demand for goods other than 1 given the remaining wealth and the prices p_{-1}.

Sellers are identical firms that produce good 1 with a strictly convex technology without fixed costs, leading to an increasing supply function $s_1(p_1)$. The evaluation of the market outcome is given by profits, i.e.,

$$v_S(q_S, p_1) = p_1 q_S - C(q_S),$$

where C is the cost function and q_S is the quantity of good 1 actually sold by a seller at the chosen institution.

It is easy to show that valuation and demand and supply functions constructed in this way fulfill assumptions **M1–M3** and **A1–A3**.

Example 2 (*Consumer and Producer Surplus*) Another specific way to derive valuation functions for the current model is to arbitrarily specify demand and supply functions satisfying **M1–M3**, and let the evaluation of the market outcome be the corresponding consumers' and producers' surplus. It is easy to see that valuation functions constructed in this way also fulfill assumptions **A1–A3**.

As pointed out above, we view our model as reasonably general. There are, however, a number of simplifying assumptions which are made to ensure tractability and could potentially be generalized. First and foremost, we consider a buyers–sellers model, where traders' roles are predetermined and immutable. One could alternatively consider models where traders are endowed with excess demand functions, and can hence become buyers or sellers depending on the realized prices. This is the approach taken in Alós-Ferrer and Kirchsteiger (2010). Second, within each type, traders are identical, which greatly simplifies the analysis. The basic model could, however, be extended to consider heterogeneous traders (as done, e.g., in Alós-Ferrer et al. 2010, where buyers are heterogeneous). Third, buyers and sellers are endowed with continuous demand and supply functions. This allows for a straightforward conceptualization of market clearing close to standard microeconomic models, but stands in contrast with other contributions in the theoretical and experimental literature which consider unit demand and supply.

2.1 Trading institutions

The good can be traded at different market institutions. For any institution z, denote by n_z, m_z the number of buyers and sellers active at z. If either $n_z = 0$ or $m_z = 0$, then no trade takes place at z. If $n_z, m_z > 0$, let $p^*(n_z, m_z)$ be the market-clearing price at z, i.e., $p^*(n_z, m_z)$ is the solution to

$$n_z d(p) = m_z s(p). \tag{MC}$$

Under **M1**–**M3**, for every $n_z, m_z > 0$, there exists a unique $p^*(n_z, m_z)$ solving equation (MC), and it is strictly larger than zero. Note also that the equilibrium quantity is strictly positive.[3] Moreover, the market-clearing price $p^*(n_z, m_z)$ depends only on the ratio

$$r = \frac{n_z}{m_z}$$

through the implicit equation $rd(p) = s(p)$, and hence, we can write $p^* = p(r)$. It is important to note that the function $p(r)$ is strictly increasing in r [because $d(p)$ is decreasing and $s(p)$ is increasing in p].

Institutional biases come in many flavors, and it is frequently hard to formally pin down the bias. We adopt a shortcut which nevertheless allows us to tackle a wide range of examples. Because differences in the institutional setup lead to systematic differences in the realized prices, we characterize the institutions directly by the trading price generated by the institution, but abstract from the specific rules generating this price. In order to accommodate different kinds of institutions, we give here a general definition and proceed to illustrate it presenting some families of examples. Let

[3] **M1**–**M3** imply that there exists an equilibrium and that any equilibrium price is strictly larger than zero. Because of **M3** and monotonicity of supply and demand, any equilibrium quantity is strictly positive. Then, by **M1** demand at any equilibrium is strictly decreasing, implying uniqueness.

$$S(n,m) = \left\{(n_z, m_z) \in \mathbb{N}^2 \,|1 \leq n_z \leq n,\ 1 \leq m_z \leq m\right\}$$

be the set of all feasible combinations of traders and sellers which can potentially show up at the same institution.

Definition 1 An institution is characterized by a *bias function*, $\beta_z : S(n,m) \to \mathbb{R}_{++}$ which measures the ratio between the actual price realized under that market institution, p_z, and the market-clearing price. More specifically,

$$p_z(n_z, m_z) = \beta_z(n_z, m_z) p^*(n_z, m_z).$$

We say that the institution z is *market clearing* if $\beta_z(n_z, m_z) = 1$ for all $(n_z, m_z) \in S(n,m)$. We say that it is biased in favor of the sellers, or simply that it is a *seller institution*, if $\beta_z(n_z, m_z) > 1$ for all $(n_z, m_z) \in S(n,m)$. Analogously, we say that it is biased in favor of the buyers, or simply that it is a *buyer institution* , if $\beta_z(n_z, m_z) < 1$ for all $(n_z, m_z) \in S(n,m)$.

According to this definition, trade at each institution occurs at only one particular, deterministic price. One might want to give up these assumptions. It can be shown that allowing for institutions that violate this intra-institutional "law of one price" or for institutions with stochastic prices would not change our main results.[4]

We also remark that we do not assume that institutions are systematically biased in favor of the sellers or the buyers. A given institution might yield $\beta_z(n_z, m_z) < 1$ for certain pairs (n_z, m_z), and $\beta_z(n_z, m_z) > 1$ for others. Examples are given below.

If the price is not at the market-clearing level, we assume that the quantity traded is determined by the "shorter" market side and that the other market side cannot trade as much as it wishes according to its demand or supply function. This rationing is assumed to be the same for every trader of the same market side. More specifically, denote by $Q_z(n_z, m_z)$ the overall quantity traded at z. We can now distinguish between three cases:

Case 1 $\beta_z(n_z, m_z) = 1$. In this case, the market-clearing prices and quantities are realized, and no trader is rationed. The institution is market clearing. The quantities are given by $Q_z(n_z, m_z) = m_z s(p^*(n_z, m_z)) = n_z d(p^*(n_z, m_z))$; $q_B^z = d(p^*(n_z, m_z))$; $q_S^z = s(p^*(n_z, m_z))$.

Case 2 $\beta_z(n_z, m_z) < 1$. In this case, the price is below the market-clearing price, and hence the quantity is determined by supply and buyers are rationed: $Q_z(n_z, m_z) = m_z s(p_z(n_z, m_z))$; $q_S^z = s(p_z(n_z, m_z))$; $q_B^z = \frac{m_z}{n_z} s(p_z(n_z, m_z)) < d(p_z(n_z, m_z))$.

Case 3 $\beta_z(n_z, m_z) > 1$. In this case, the price is above the market-clearing price, and hence the quantity is determined by demand and sellers are rationed: $Q_z(n_z, m_z) = n_z d(p_z(n_z, m_z))$; $q_B^z = d(p_z(n_z, m_z))$; $q_S^z = \frac{n_z}{m_z} d(p_z(n_z, m_z)) < s(p_z(n_z, m_z))$.

[4] A proof of this claim is available upon request. We have also implicitly assumed that institutions are anonymous, i.e., the bias depends only on the number of sellers and buyers operating at the institution and not on their identities. Our results remain valid if this assumption is relaxed.

In summary, given an institution z characterized by a function $\beta_z(\cdot,\cdot)$, and given $r = \frac{n_z}{m_z} > 0$ and $\beta = \beta_z(n_z, m_z)$, we can compute the seller and buyer quantities as

$$q_S^z(\beta, r) = \begin{cases} s\,(\beta \cdot p(r)) & \text{if} \quad \beta \leq 1 \\ r \cdot d\,(\beta \cdot p(r)) & \text{if} \quad \beta \geq 1 \end{cases}$$

and

$$q_B^z(\beta, r) = \begin{cases} \frac{1}{r} \cdot s\,(\beta \cdot p(r)) & \text{if} \quad \beta \leq 1 \\ d\,(\beta \cdot p(r)) & \text{if} \quad \beta \geq 1 \end{cases}$$

At this point, we have to emphasize that we do not aim to analyze how a deviation from market-clearing prices comes about. Rather, we just assume that market-clearing institutions as well as institutions preventing markets from clearing are in principle feasible. And the purpose of this paper is to investigate whether a non-market-clearing institution can survive vis-a-vis a market-clearing one.

The formulation above is general enough to encompass many familiar examples.

Example 3 (*Limit price institutions*) An institution exhibits a *price cap* if there exist $p^H > 0 \in \mathbb{R}_+$ such that, for all $(n_z, m_z) \in S(n, m)$,

$$\beta_z(n_z, m_z) \leq \frac{p^H}{p^*(n_z, m_z)}.$$

Analogously, an institution exhibits a *price floor* if there exist $p^L > 0 \in \mathbb{R}_+$ such that, for all $(n_z, m_z) \in S(n, m)$,

$$\beta_z(n_z, m_z) \geq \frac{p^L}{p^*(n_z, m_z)}.$$

Price caps are often observed in housing markets, whereas price floors are prominent in labor markets—minimum wages.

Further, an institution exhibits a *fixed price* if there exist $p^F > 0$ which is simultaneously a price floor and a price cap. Such extreme public price regulation has been often observed for basic goods like food in wartime.

Other institutions do not exhibit a direct, public price regulation. Rather, market institutions like the posted offer or the posted bid institution enhance trade at prices systematically above or below the market-clearing price. The most simple type of such institutions is the following.

Example 4 (*Constant-bias institutions*) A constant-bias institution z is characterized by a bias parameter $\beta_z > 0$, i.e., $\beta_z(n_z, m_z) = \beta_z$ for all $(n_z, m_z) \in S\,(n, m)$. Thus, we can write

$$p_z(n_z, m_z, \beta_z) = \beta_z p^*(n_z, m_z).$$

Constant-bias institutions are a simple, parametric example which will actually be enough for some of our purposes.

Example 5 (*Oligopolistic institutions*) We say that a seller institution z is *oligopolistic* if $\beta_z(n_z, m_z)$ is strictly larger than one and strictly decreasing in m_z, for any given n_z. Such institutions arise, e.g., if the price is the result of a Nash equilibrium where sellers internalize buyers' demand and compete among themselves in quantities. The intuition is simply that as more and more sellers compete (larger m_z), they lose market power and the oligopolistic price approaches the competitive one (hence, the bias approaches one).

Notice that, in this formulation, sellers' market power is embodied by the institution. The market price p is higher than the market-clearing price. Still, at that market price, sellers are rationed, i.e., sell less than $s(p)$. For instance, if the market price is the Cournot–Nash one, it is only after rationing takes place that the sellers exactly supply the Cournot–Nash quantity. The institution, hence, can be seen as a coordination or commitment device.

Remark 1 Although bias functions allow to model many classes of market institutions, we explicitly exclude situations where there is a difference between the price paid by the buyer and the one received by the seller, i.e., taxation policies or transaction fees. This is, however, an important consideration. In related work (Alós-Ferrer et al. 2010), we have considered precisely this possibility in order to model markets actively designed by market designers who receive part of the market revenues.

Remark 2 In our model, agents play two different roles. On the one hand, they are boundedly rational players in a game of institution selection. On the other hand, once an institution has been chosen, agents become traders at that institution. A crucial assumption of the model (which we will make explicit below) is that institution choice is based on past, observed outcomes, and the institutions' characteristics affect institution choice through learning only. In particular, agents will not plan to, e.g., manipulate their demand and supply functions when planning an institution choice. Obviously, for any given institutional rationing scheme, traders might have an incentive to misrepresent their true demand and supply functions. This, however, is also true for any standard Arrow–Debreu framework if demand and supply functions are interpreted as consciously known trader characteristics which can be "reported" to an auctioneer (see also Alós-Ferrer and Kirchsteiger 2010, Remark 2). We abstract from this possibility here. The simplest interpretation is that our agents are boundedly rational traders who choose trading institutions on the basis of past performance, not anticipating neither future results nor manipulation possibilities.

3 The learning process

3.1 The stage game

If more than one institution is available, traders themselves can choose the institution at which they want to be active. For example, if the price for a certain good is fixed by the state, traders might choose between the official market with the fixed price, and a black market where trade is conducted at market-clearing prices. Labor might be hired at the official market where a minimum wage legislation applies, and at a black market without a price floor. Goods might be traded at a posted offer market, where the price tends to be above the market-clearing level, and at a double auction, where the market outcome tends to coincide with the competitive equilibrium.

In this section, we explicitly model the choice between trading institutions. Our aim is to be able to predict which institution(s) will be observed to be active, and whether the outstanding importance of market-clearing institutions in economics can be justified by this choice process.

A generic trader is denoted by k, while i always denotes a buyer and j always denotes a seller. There are $Z + 1$ institutions available, $z = 0, 1, \ldots, Z$. Institution 0 is a market-clearing institution ($\beta_0 = 1$). We make no assumption over the remaining others. In particular, there might be some other, competing, market-clearing institution.

We proceed now by formulating the choice process as a game. At first, all traders choose, simultaneously and independently, the institutions at which they want to trade the good.[5] Then, for each trading institution z, the number of buyers and sellers who have opted for this institution, n_z and m_z, and the bias function β_z determine—as described in Sect. 2.1—the price and the quantity exchanged at z. This in turn determines the payoffs (evaluations) of the traders having opted for z.

This choice process has some features of a coordination game. If all traders coordinate on a particular institution, every individual trader would be worse off if he deviated to another institution, since by deviating he would lose all trading partners (see **A3**). Hence, full coordination on any institution constitutes a strict Nash equilibrium and nothing guarantees coordination on the market-clearing institution. **A3** also ensures that these full-coordination equilibria are the only symmetric pure-strategy Nash equilibria where all traders of the same type choose the same institution. On top of these symmetric equilibria, asymmetric ones might also exist in general. Asymmetric Nash equilibria, however, cannot be characterized within our rather general theoretical framework, since they depend on the number of buyers and sellers, the characteristics of the feasible trading institutions, and the exact specification of the demand, supply, and evaluation functions.[6]

[5] We abstract from multi-homing considerations here.

[6] To see that asymmetric equilibria may exist in general, we can look ahead and consider the game played by the subjects in treatment 1 of our experiment. Seven buyers and seven sellers had to choose between two trading institutions, A and B. The resulting payoffs for buyers and sellers are given in Table 5 of Appendix 2, with n and m denoting the number of buyers and sellers opting for the particular institution. A direct check of those tables shows that any strategy combination with exactly three sellers and three buyers opting for A and four buyers and four sellers opting for B constitutes a pure-strategy Nash equilibrium. The same holds for two buyers and two sellers opting for A, and for a lone buyer and a lone seller opting for A. Straightforward but lengthy computations show that there exists no other asymmetric pure-strategy Nash equilibrium; in particular, there exists no pure-strategy Nash equilibrium where more than three buyers and three sellers opt for A, except for full coordination on A.

3.2 The basic learning process

We proceed now to model the learning process. First, we define the state space Ω. A state $\omega \in \Omega$ is given by

$$\omega = (\omega_B, \omega_S) \in \{0, 1, \ldots, Z\}^n \times \{0, 1, \ldots, Z\}^m$$

That is, $\omega(k) \in \{0, 1, \ldots, Z\}$ denotes the institution chosen by trader k at state ω. Given a sample path of the dynamics, the state of the process at time t is denoted by $\omega^t = \left(\omega_B^t, \omega_S^t\right) \in \Omega$.

Since interactions are anonymous and traders are symmetric, the following notation will turn out to be convenient:

$$n_z(\omega) = |\{i \in I \,|\omega(i) = z\}|$$
$$m_z(\omega) = |\{j \in J \,|\omega(j) = z\}|$$

That is, $n_z(\omega) \in \{0, 1, \ldots, n\}$ is the number of buyers and $m_z(\omega) \in \{0, 1, \ldots, m\}$ the number of sellers choosing institution z, and $n_0(\omega) + \cdots + n_Z(\omega) = n$, $m_0(\omega) + \cdots + m_Z(\omega) = m$ hold.

The learning process is based on the implicit assumption that traders understand the strategic nature of the coordination problem. Therefore, they do not regard the situation as an individual decision problem (as they would in a reinforcement learning model). Furthermore, we assume that traders only know the prices and the quantities of currently active institutions and, hence, do not have enough information to accurately predict the outcomes in all trading institutions which are in principle feasible. Thus, they lack the information necessary to compute a best reply to the current choices of all other traders.

Suppose that a trader has the possibility to revise his choice of institution (we will specify in which form revision opportunities arrive below). What can a trader do in such a situation? From his individual (myopic) standpoint, if he considers himself to be small relative to market size, the best thing he can do is to observe the outcomes (i.e., prices and quantities) of the currently active institutions and to evaluate these outcomes through his own evaluation function. That is, he will switch to that institution whose current prices and quantities he perceives as best according to his evaluation function. A trader can perceive this behavior as approximately rational, since when he chooses a new institution, the implied changes in prices and traded quantities will most of the time be small, and hence, this behavior is close to best reply. Of course, in the current (symmetric) model, this behavior could also be interpreted as imitation of successful traders of the own market type. We want to stress, though, that the described behavior does not require the observation of payoffs achieved by other traders, but merely prices and traded quantities.

Fix a state ω. Call an institution z *active* if $m_z(\omega) > 0$ and $n_z(\omega) > 0$, and inactive if $m_z(\omega) = 0$ or $n_z(\omega) = 0$. With this notation, the considerations above are captured by the following assumption.

D0 Traders who receive the opportunity to revise observe prices and traded quantities at all active institutions. Then, they choose the institution which yields the best outcome as evaluated by their own evaluation functions and go there next period (ties broken randomly).[7]

That is, provided that trader k receives revision opportunity at period t, in period $t+1$, he will choose an institution among those that in period t were yielding the highest observed payoffs for traders of his own type. Note that an agent takes his decision for period $t+1$ given the state ω^t and the associated payoffs. This decision determines the institution chosen for period $t+1$. Combining all such decisions of the individual traders determines ω^{t+1} and, hence, the basic dynamics.

3.3 Revision opportunities

When can agents revise their choices? It is common in learning models to explicitly introduce some inertia allowing for the possibility that not all agents are able to revise strategies simultaneously. Different specifications are possible. One prominent example is *independent inertia* (e.g., Samuelson 1994; Kandori and Rob 1995), where each agent has an independent, strictly positive probability of not being able to switch. A different example is *asynchronous learning* (e.g., Binmore and Samuelson 1997; Benaïm and Weibull 2003; Blume 2003), where each period one and only one agent is able to revise, all agents having strictly positive probability of receiving the revision draw. In our case, a natural variant of this dynamics would be *asynchronous learning within types*, where in every period, only one buyer *and* one seller are selected (randomly and independently) and given the opportunity to revise.

Different specifications of how revision opportunities arrive give rise to different dynamics and often affect the results (see, e.g.,Alós-Ferrer and Netzer 2010). Rather than adopting a specific formulation, we postulate a general class of dynamics encompassing the standard examples mentioned above and many others (see Alós-Ferrer 2003 and Alós-Ferrer and Netzer 2010 for a discussion).

Let $E(k, \omega)$ denote the event that agent k receives revision opportunity when the current state is ω, and let $E^*(k, \omega) \subseteq E(k, \omega)$ denote the event that agent k is the only agent of his type (i.e., the only buyer or the only seller) receiving revision opportunity in ω. With this notation, the general class of dynamics we consider is given by the following assumptions.

D1 $\Pr(E^*(k, \omega)) > 0$ for every agent k and state ω.

Notice that **D1** implies that $\Pr(E(k, \omega)) > 0$, i.e., every agent has strictly positive probability of being able to revise at any given state. Further, since we have two clearly

[7] Inactive institutions are not even observed, since no price is even posted. Hence, in the extreme case in which all institutions are inactive, traders simply stay at their respective institutions. We find this assumption plausible in this context. Alternatively, we could assume that, if there is no activity at any of the institutions, traders switch to some other institution randomly. This assumption would make states with complete inactivity easier to leave and hence "less stable." Our results would remain unchanged with this assumption, the intuition being that such states are neither stochastically stable nor crucial for transitions to stochastically stable states.

differentiated populations, we introduce a weak form of independence between the revision opportunities in those populations (it can be thought of as an anonymity requirement).

D2 For every agent k and state ω, either $\Pr\left(E^*(k,\omega) \cap E^*(k',\omega)\right) > 0$ for any agent k' of the other type, or $\Pr\left(E^*(k,\omega) \cap E(k',\omega)\right) = 0$ for any such k'.

Assumptions **D1** and **D2** are rather general. It is easy to see that they are fulfilled by the standard types of revision opportunities mentioned above. For instance, a particular example is independent inertia, where each trader receives the revision opportunity with a fixed probability $0 < \lambda < 1$, independent across traders and periods. A different example is asynchronous learning within types, where each period only one buyer and only one seller are selected to revise. Intermediate and mixed specifications are allowed.[8] The reason we explicitly choose Assumptions **D1–D2** is that, in the literature of learning in games, predictions are not always robust to minute changes in the assumptions on the dynamics. We want to make explicit that our model is not so sensitive to the details of the dynamics (see Alós-Ferrer and Netzer 2015, for a discussion).

In our context, it is plausible that traders are more likely to revise when the perceived gains from revision are higher. For instance, one might postulate that the probability of revision increases with the difference between the payoff at the institution currently chosen by the trader and the largest payoff generated at any other institution. For the case of two institutions, this would be equivalent to the *proportional imitation rule* of Schlag (1998). Such a sensitivity of revision opportunities to payoff differences is allowed by the specification above, since the revision probability $\Pr(E(k,\omega))$ is a function of the state ω.

3.4 Stochastic stability

The dynamics described till now is a Markov chain on the (finite) state space Ω, to which standard treatment applies (see, e.g., Karlin and Taylor 1975). We refer to this dynamics as the *unperturbed process*.

Given two states ω, ω', denote by $P(\omega, \omega')$ the probability of transition from ω to ω' in one period. An *absorbing set* of the unperturbed dynamics is a minimal subset of states which, once entered, is never abandoned. An *absorbing state* is an element which forms a singleton absorbing set, i.e., ω is absorbing if and only $P(\omega, \omega) = 1$.

In general, the unperturbed process presents a multiplicity of absorbing sets. In order to select among them, and following the literature, the dynamics is enriched with a perturbation in the form of mistakes or experiments as follows. With an independent probability $\varepsilon > 0$, each agent, in each period, might make a mistake ("mutate") and simply pick an institution at random,[9] independently of other considerations. This

[8] Assumption D2 explicitly precludes revision opportunity correlations, e.g., of the form "seller 13 always gets to revise whenever buyer 3 revises, but no other seller gets the chance."

[9] We mean that an institution is picked up according to a pre-specified probability distribution having full support, for instance uniformly. It is well known that the exact distribution does not affect the stochastic-stability results, as long as it has full support.

can be interpreted literally as a decision mistake or, alternatively, as an *experiment* on the side of the agent. For instance, such an experiment might correspond to an agent being replaced by a new, inexperienced one which simply builds some arbitrary theory, or to an agent discarding past information and being attracted to a new institution after observing an institutional (unmodeled) marketing campaign. Crucially, in case of experimentation also any inactive institution is chosen with a strictly positive probability.

The dynamics with mistakes (experimentation) is called a *perturbed* learning process. Since experiments make transitions between any two states possible, the perturbed process has a single absorbing set formed by the whole state space. Such processes are called *irreducible*. An irreducible process has a unique *invariant distribution* , i.e., a distribution over states $\mu \in \Delta(\Omega)$ which, if taken as initial condition, would be reproduced in probabilistic terms after updating (more precisely, $\mu \cdot P = \mu$ where P is the matrix of transition probabilities).

For a given ε, the corresponding invariant distribution is denoted by $\mu(\varepsilon)$. The *limit invariant distribution* (as the rate of experimentation tends to zero) $\mu^* = \lim_{\varepsilon\to 0} \mu(\varepsilon)$ exists and is an invariant distribution of the unperturbed process (Kandori et al. 1993; Young 1993; Ellison 2000). The limit invariant distribution singles out a stable prediction of the unperturbed dynamics in the sense that, for any $\varepsilon > 0$ small enough, the play approximates that described by μ^* in the long run. The states in the support of μ^*, i.e., $\{\omega \in \Omega \mid \mu^*(\omega) > 0\}$ are called *stochastically stable states* or long-run equilibria. The set of stochastically stable states is a union of some absorbing sets of the original, unperturbed chain ($\varepsilon = 0$).

We will rely on the characterization of the set of stochastically stable states introduced by Kandori et al. (1993) and Young (1993) and further developed by Ellison (2000). Detailed overviews can be found, e.g., in Fudenberg and Levine (1998) or Samuelson (1997).

4 Stochastic stability of market-clearing institutions

We proceed now to analyze the complete model. A first intuition for our main results is obtained when we compare the payoffs sellers and buyers receive at simultaneously active market-clearing and non-market-clearing institutions.

Lemma 1 *Assume* **A1** *and* **A2**. *Consider any distribution of traders on any number of institutions, where both a market-clearing institution* 0 *and another institution* z *are active. Let* $p_z = p_z(n_z, m_z)$. *Then, the following holds:*

For $\beta_z(n_z, m_z) \neq 1$*: If* $v_S(q_S^0, p_0) \leq v_S(q_S^z, p_z)$, *then* $v_B(q_B^0, p_0) > v_B(q_B^z, p_z)$. *Hence, if* $v_B(q_B^0, p_0) \leq v_B(q_B^z, p_z)$, *then* $v_S(q_S^0, p_0) > v_S(q_S^z, p_z)$.

For $\beta_z(n_z, m_z) = 1$*: Either* $v_S(q_S^0, p_0) \leq v_S(q_S^z, p_z)$ *and* $v_B(q_B^0, p_0) \geq v_B(q_B^z, p_z)$, *or the reverse (weak) inequalities hold.*

Lemma 1 shows that, whenever traders of a given market side obtain larger payoffs in a biased institution than their counterparts in the market-clearing one, traders of the other market side which are active in the market-clearing institution must obtain larger payoffs than those active in the biased one. This result is crucial for the analysis of the learning model. Intuitively, it points out a reason for (some) traders to move toward the market-clearing institution in the presence of another one.

We are interested in the stability of institutions. Clearly, every *monomorphic state*, where all traders coordinate in one and the same institution, constitutes an absorbing state. These are actually the only relevant absorbing states. In principle (and particularly for dynamics with asynchronous learning), there might be non-singleton absorbing sets. However, the following proposition shows that those would be made up of states where the market-clearing institution $z = 0$ is never active.

Proposition 1 *Assume* **A1**, **A2**, *and* **A3**. *Under* **D0**, **D1**, *and* **D2**,

(i) *the absorbing states of the unperturbed dynamics are the "separated states"* ω *such that there is no active institution at all, and all monomorphic states* ω_z *characterized by* $n_z(\omega_z) = n$ *and* $m_z(\omega_z) = m$, *corresponding to coordination on a particular institution;*
(ii) *no state* ω *with* $1 \leq n_0(\omega) \leq n - 1$ *and* $1 \leq m_0(\omega) \leq m - 1$ *(i.e., where the market-clearing institution is active but not all traders of any type are in it) is part of any absorbing set of the unperturbed dynamics.*

We remark, however, that this is just an intermediate result. It is a standard fact (see the general references quoted above) that only states in absorbing sets can be stochastically stable, and hence, the last proposition restricts the class of states relevant for the analysis, but not all absorbing states will be stochastically stable. Further, the key states for the analysis below are the monomorphic states.[10] Since those states correspond to full coordination on a particular market institution, we aim to identify which monomorphic states are stochastically stable.

Definition 2 We say that an institution $z \in \{0, \ldots Z\}$ is stochastically stable if the corresponding monomorphic state ω_z characterized by

$$n_z(\omega_z) = n \quad \text{and} \quad m_z(\omega_z) = m$$

is stochastically stable.

Intuitively, a stochastically stable institution is one such that, in the long run, traders frequently coordinate on it. In principle, several institutions could be stochastically stable, but if a particular institution is not, we can assert that, in the long run, this institution will be simply not be used by traders.

Our first main result establishes that market-clearing institutions are always active in the long run.

Theorem 1 *Assume* **M1**–**M3** *and* **A1**–**A3**, *and consider any dynamics satisfying* **D0**–**D2**. *Any market-clearing institution is stochastically stable.*

[10] Separated states are also absorbing because, if all institutions are inactive, no prices are observed and traders do not switch. This is inconsequential. Separated states are extremely unstable. Specifically, they are destabilized with a single mutation, in which one trader moves to an institution containing at least one trader of the other type. By Lemma 2 in Appendix 1, the outcome of the now-active institution is better for all traders than that of the inactive institution. Hence, traders at the inactive institution will switch whenever revision opportunities arise.

This result implies that, independently of which other institutions are available, coordination on the market-clearing one will always be observed at least (a non-negligible) part of the time in the long run. It is striking that this result is completely independent of what the characteristics of other institutions are. A market-clearing institution remains stochastically stable independently of how many other institutions are available and what their characteristics are, from limit pricing to oligopolistic institutions or any conceivable alternatives. Furthermore, whenever traders coordinate on a market-clearing institution, the market is efficient.[11]

Remark 3 A common criticism on the literature of learning in games is that the speed of convergence to the predicted outcomes might depend inversely (and exponentially) on population size and hence the predictions might be irrelevant for large population sizes. This criticism does not affect our results. The technical reason [see Ellison (2000) for details] is that the number of mutations involved in the stability analysis is small (two) and independent of population size. Intuitively, the transitions that destabilize non-market-clearing institutions in favor of market-clearing ones only require a few experiments, followed by high-probability revisions where traders imitate successful behavior.

5 Stable non-market-clearing institutions

In the previous section, we have shown that market-clearing institutions are always stochastically stable. However, it turns out that there exists also stochastically stable biased institutions. Strikingly, it is possible to show that even some constant-bias institutions are stochastically stable, independently of which other institutions are available.

In general, the effects of a bias on the payoffs of the traders are ambiguous. Take as an example an active institution z where prices are higher than the equilibrium price ($\beta_z(n_z, m_z) > 1$). Recall the notation $r = n_z/m_z$ for the buyers–sellers ratio at z ($0 < r < \infty$). Compared to a market-clearing institution having exactly the same r, prices as well as quantities are unfavorable for buyers, and a further increase in β_z would lead to a further decrease in buyers' payoffs. For sellers, the situation is different. For them, prices at z are more favorable than at a market-clearing institution. This comes at the price of a decrease in the quantity sellers can sell. Therefore, the impact of a further increase of β_z on sellers' payoffs is unclear.

To build an intuition, consider the standard case with demand and supply derived from utility and profit maximization. Under standard assumptions, the price set by a cartel formed by all the sellers is strictly larger than the market-clearing price. Hence, for a given number of buyers and sellers, a small increase of β above one should be beneficial for the sellers (and, of course, detrimental for the buyers).

[11] Due to the efficiency properties of the equilibrium, we view this result as "good news." In certain contexts, however, the interpretation might be different. A black labor market might be considered as a market-clearing institution which competes with regulated labor markets. Our result might thus provide an insight into the stability of moonlighting.

Similar considerations can be made for the impact of the bias on the buyers. For prices close to the equilibrium price, the positive direct effect of a price decrease on the consumers is larger than the negative effect due to the decrease in consumed quantity.

These considerations lead to Assumption **A4** below. Given a realized bias $\beta_z = \beta_z(n_z, m_z) > 0$ and given $r = \frac{n_z}{m_z} > 0$, the payoffs for buyers and sellers at institution z can be rewritten as

$$V_B(\beta_z, r) = v_B\left(q_B^z(\beta_z, r), \beta_z \cdot p(r)\right) \quad \text{and} \quad V_S(\beta_z, r) = v_S\left(q_S^z(\beta_z, r), \beta_z \cdot p(r)\right).$$

The payoffs of, say, the buyers are given by $V_B\,(\beta_z, r)$; from the buyers' point of view, though, they depend just on the actually experienced bias and buyers–sellers ratio. The following assumption spells out the effects of small deviations of the equilibrium price from the realized one for a given ratio of buyers and sellers.

A4 For any fixed ratio of buyers and sellers r with $0 < r < \infty$, there exist $\underline{\beta}(r) < 1 < \overline{\beta}(r)$ such that $V_B(\beta, r) > V_B(1, r)$ for all $\underline{\beta}(r) < \beta < 1$, and $V_S(\beta, r) > V_S(1, r)$ for all $1 < \beta < \overline{\beta}(r)$.

This condition is immediately fulfilled if the buyer's payoff $V_B(\beta, r)$ is strictly decreasing in β at $\beta = 1$,[12] and the seller's payoff $V_S(\beta, r)$ is strictly increasing in β at $\beta = 1$.

Note that the comparison of payoffs spelled out in this assumption is fundamentally different from the results of Lemma 1. There, the comparison was between payoffs yielded by two simultaneously active institutions with different traders, while in **A4**, the comparison is implicitly between payoffs yielded by two different institutions, provided that the buyers–sellers ratio is the same in both of them.

Assumption **A4**, though, is instrumental for showing that there are some non-market-clearing institutions fulfilling a property analogous to the one spelled out in Lemma 1, i.e., there is always a reason for some traders to move toward them even in the presence of a market-clearing institution.

Definition 3 Fix the number of buyers and sellers operating on the whole market. An institution $F \neq 0$ is *favored* if, given any distribution of these traders on (only) F and the market-clearing institution 0 such that both of them are active, the following holds:

If $v_S(q_S^0, p_0) \geq v_S(q_S^F, p_F)$, then $v_B(q_B^0, p_0) < v_B(q_B^F, p_F)$ (or, equivalently, if $v_B(q_B^0, p_0) \geq v_B(q_B^F, p_F)$, then $v_S(q_S^0, p_0) < v_S(q_S^F, p_F)$).

Favored institutions are those such that a statement analogous to Lemma 1 holds for them versus the market-clearing one. This is actually enough to show that favored institutions are stochastically stable.

Theorem 2 *Assume* **M1**–**M3** *and* **A1**–**A4**, *and consider any dynamics satisfying* **D0**–**D2**. *Let $z \in \{1, \ldots, Z\}$ be any favored institution. Then, independently of which other institutions are available, z is stochastically stable.*

[12] Neither $V_B(\beta, r)$ nor $V_S(\beta, r)$ are in general differentiable at $\beta = 1$, because at this point there is a transition from rationing of the demand side to rationing of the supply side. Hence, the traded quantity as a function of β has a "kink" at $\beta = 1$.

This result shows that, potentially, there might exist stochastically stable, non-market-clearing institutions. In order to actually establish their existence, it is enough to investigate under which circumstances do such favored institutions exist.

Obviously, one can just take the maximum $\underline{\beta}(r)$ and minimum $\overline{\beta}(r)$ among all (finitely many) buyers–sellers ratio which are actually possible. The intuition would be that institutions which always yield biases between those bounds should be favored. This intuition fails, though. The problem is the following. Imagine that a biased institution z with, say, constant-bias $\beta_z < 1$ and a market-clearing one are simultaneously active. In principle, since $\beta_z < 1$, prices at z are lower than at the market-clearing institution, for a given proportion of buyers and sellers. The actual proportions at z and the market-clearing institution, though, might be so different as to offset the effect of the bias. For, since the market-clearing price is an increasing function of the buyers–sellers ratio $r_z = \frac{n_z}{m_z}$, if the ratio at the market-clearing institution, r_0, is much smaller than the one at z, r_z then the price at the former, $p(r_0)$ might be so much smaller than the (theoretical) market-clearing price at z, $p(r_z)$, that the actual price there, $\beta_z \cdot p(r_z)$, might still be larger than $p(r_0)$ even though $\beta_z < 1$.

This problem, though, might be overcome by taking tighter bounds, taking full advantage of the fact that m and n are finite. Then, one obtains the following result.

Theorem 3 *Assume* **M1**–**M3***,* **A1**–**A4***, and* **D0**–**D2**. *Fix the number of buyers n and sellers m operating on the whole market. Then, there exist $\underline{\beta}^*(n, m)$ and $\overline{\beta}^*(n, m)$ with $\underline{\beta}^*(n, m) < 1 < \overline{\beta}^*(n, m)$ such that any institution F satisfying $\underline{\beta}^*(n, m) < \beta_F(n_z, m_z) < \overline{\beta}^*(n, m)$, $\beta_F(n_z, m_z) \neq 1$ for all $(n_z, m_z) \in S(n, m)$ is favored and hence stochastically stable.*

In particular, any constant-bias institution F with $\underline{\beta}^*(n, m) < \beta_F < \overline{\beta}^*(n, m)$ is stochastically stable.

This result shows that potential favored institutions do exist[13] for any n, m, and that the vicinity of the market-clearing institution consists of such favored institutions. Those non-market-clearing institutions for which $\underline{\beta}^*(n, m) < \beta_z(n_z, m_z) < \overline{\beta}^*(n, m)$ are such that they improve one market side relative to the market-clearing institution for distribution of buyers and sellers. In other words, *for any given distribution of buyers and sellers*, such a non-market-clearing institution is favored by one market side over the market-clearing one.

The last result shows that, in general, there exist non-market-clearing institutions which do not disappear in the long run. Strikingly, this includes even some very simple institutions, characterized by a constant (if small) bias.

In order to analyze the efficiency implications of coordinating on a favored non-market-clearing institution, we have to distinguish between two different notions of efficiency. On the one hand, one might ask for Pareto efficiency. Because of **A4**, for favored institutions in the vicinity of the market-clearing institution, the outcome of full coordination on at least some favored institutions is not Pareto-inferior to full coordination on the market-clearing institutions. Coordination on a favored institution

[13] That is, there are bias functions such that, if an institution is characterized precisely by that function, it will be favored. This does not mean that we assume a favored institution always to be actually available in the market.

indeed favors those for whom the price bias is advantageous, while the other market side loses as compared to coordination on a market-clearing institution.

On the other hand, one could use the aggregate (sum) of consumers' and producers' surplus as an efficiency measure. Because of the dead-weight-loss associated with rationing, full coordination on any non-market-clearing institution results in a lower efficiency than coordination on a market-clearing institution. The expected efficiency loss due to the presence of stable non-market-clearing institutions depends crucially on the proportion of time that coordination on each stable institution occurs in the long run, i.e., on the exact shape of limit invariant distribution. Theorems 2 and 3 characterize for which institutions one can expect full coordination to occur in the long run, i.e., which states are in the support of the limit invariant distribution. These results, however, do not identify the exact shape of the limit invariant distribution. In fact, the limit invariant distribution depends crucially on the properties of the demand and supply functions and also on the properties of the learning model, in particular on the specification of revision opportunities and the details of the experimentation process. Hence, the expected level of inefficiency cannot be characterized within the general framework of our model.

6 Stable non-market-clearing institutions and the market size

Theorem 2 gives us sufficient conditions for the existence of stochastically stable institutions other than the market-clearing one. By Theorem 3, favored institutions always exist for given market size, even if institutions are simply characterized by a constant-bias parameter. However, one might ask whether it is possible that only the market-clearing institution is stable if the market becomes very large. It is indeed possible to construct examples (for particular combinations of demand and supply functions) where the set of favored institutions degenerates as market size grows; however, being favored is just a sufficient condition for stochastic stability, and hence, focusing on this property would not allow us to obtain a satisfactory answer. In this section, we investigate this question by letting the size of the market grow and by analyzing stochastic stability directly.

Specifically, we adopt a "replica economy" approach as follows. We fix an economy with n buyers and m sellers and consider the K-replicated economy formed by K copies of the initial economy, i.e., with $K \cdot n$ buyers and $K \cdot m$ sellers. By Theorem 1, the market-clearing institution remains stochastically stable for all K. We aim to show that certain non-market-clearing institutions are also stochastically stable for arbitrarily large K.

We consider slightly stronger versions of our assumptions **M1**–**M3**. The following assumptions exclude that demand and supply functions have trivial parts. Note that, if, e.g., demand might be zero at a positive price, it would still be zero for any replicated economy, and hence, the sense in which the economy becomes larger would be unclear.

M1′ The demand function $d : \mathbb{R}_+ \to \mathbb{R}_+ \cup \{+\infty\}$ is continuous and strictly decreasing, with $d(p) > 0$ for all $p \geq 0$, and $\lim_{p\to\infty} d(p) = 0$.

M2′ The supply function $s : \mathbb{R}_+ \to \mathbb{R}_+$ is continuous and (weakly) increasing. Further, $s(0) = 0$ and $s(p) > 0$ for all $p > 0$.

Note that **M1′–M2′** imply **M1–M3**. The key additional implications of these assumptions are that $\lim_{r\to+\infty} p(r) = +\infty$ and $\lim_{r\to 0} p(r) = 0$.[14]

In order to study large economies, we need to specify an additional assumption on the dynamics. The reason is that assumptions **D1** and **D2** are not tailored to the case of large economies. In particular, consider a dynamics where only one agent revises every period (which is allowed by assumptions **D1** and **D2**). As K increases, the speed of learning in this dynamics effectively converges to zero. A more reasonable dynamics would be, e.g., one ensuring that at least one agent *in each replica* receives the opportunity to revise, ensuring that the speed of learning remains constant (or at least does not vanish) as K increases.[15] The following assumption fulfills this role.

D3$_K$ For every state ω, the probability that any given set of K buyers i revise (and nobody else) is strictly positive. Analogously, the probability that any given set of K sellers revise (and nobody else) is strictly positive.

The following theorem proves existence of (constant-bias) stochastically stable non-market-clearing institutions even for those cases where the set of favored institutions degenerates.

Theorem 4 *Assume* **M1′–M2′**, **A1–A4**, **D0–D2** *and* **D3**$_K$ *for the dynamics of each K-replicated economy. Suppose z with constant β_z is a favored institution for the economy with $K = 1$. The following hold.*

(i) *If $m \leq n$ (more buyers than sellers) and $\underline{\beta}(1) < \beta_z < 1$, then there exists a K^* such that z is stochastically stable for all $K \geq K^*$.*
(ii) *If $m \geq n$ (more sellers than buyers) and $1 < \beta_z < \overline{\beta}(1)$, then there exists a K^* such that z is stochastically stable for all $K \geq K^*$.*
(iii) *If $m = n$ (equal number of buyers and sellers) and $\underline{\beta}(1) < \beta_z < \overline{\beta}(1)$, then there exists a K^* such that z is stochastically stable for all $K \geq K^*$.*

Note that the bounds $\underline{\beta}(1) < 1 < \overline{\beta}(1)$ are independent of market size (recall **A4**). Further, by Theorem 3, the set of constant-bias favored institutions for $K = 1$ include a non-negligible interval around $\beta = 1$. Therefore, the set of stochastically stable institutions does not in general shrink to the market-clearing institution when the market size increases, even if the set of favored institutions degenerates. Theorem 4 shows that, under general conditions, there will be biased stochastically stable institutions even for large market size. That is, there is no “core convergence” result in this setting. Even though the set of stochastically stable institutions will always contain the market-clearing institution, other institutions will remain active as the size of the economy grows.

[14] Recall that $p(\cdot)$ is strictly increasing in r. If $\lim_{r\to\infty} p(r) \neq +\infty$, it follows that $p(r)$ is bounded above by some $L > 0$. Since $s(\cdot)$ is increasing and $d(\cdot)$ is decreasing, it follows from $rd(p(r)) = s(p(r))$ that r is bounded above by $s(L)/d(L)$, a contradiction. Analogously, if $\lim_{r\to 0} p(r) \neq 0$, we would obtain that r is bounded below by some strictly positive $s(\varepsilon)/d(\varepsilon)$, a contradiction.

[15] This problem is well known in the stochastic approximation literature. For instance, Benaïm and Weibull (2003) assume a fixed relationship between population size and the length of a time interval to ensure that the expected time between two revision opportunities of a given individual does not grow as the population size increases.

The intuition is the following. Suppose that one market side (buyers or sellers) is overrepresented in the population. Then, this market side has less market power than the other side. If, for some reason, an institution biased in the favor of this market side attracts a few sellers and buyers in similar numbers, the overrepresented side will necessarily prefer the latter institution. Once the new institution becomes active, the fact that it is favored for $K = 1$ implies that, if the appropriate proportions of agents are present in it (for instance, if the numbers of buyers and sellers are multiples of K), in practice it will behave as a favored institution in the replicated economy. This creates positive-probability paths destabilizing the market-clearing institution.

7 Experimental analysis

In this section, we test the theoretical predictions derived from our model. In particular, we investigate whether traders use stochastically stable institutions independently of whether they are market clearing or not, i.e., independently of whether they maximize the sum of the gains from trade. On the other hand, we also check whether institutions that are not stochastically stable are abandoned in the long run.

Technically, stochastic stability entails a double limit, as time goes to infinity and as the experimentation rate vanishes. None of these limits can be reproduced in reality. Hence, it becomes especially important to test whether theoretical predictions based on stochastic stability are also relevant within reasonable time horizons and in the presence of naturally noisy human decisions.

Theorem 1 implies the experimental hypothesis that market-clearing institutions should not be avoided by traders. Further, if no other stochastically stable institution is available, we obtain the immediate prediction that convergence to full coordination on the market-clearing institution should be observed. Hence, one of our experimental treatments below will be such that only one (market clearing) institution is stochastically stable, and we expect to observe convergence to full coordination.

The translation of theoretical results based on a stochastic-stability analysis to experimental predictions, however, is not always so immediate. In particular, Theorems 1 and 2 imply that in certain settings, several market institutions are stochastically stable. This, however, does *not* literally mean that one should observe coexistence. In the theoretical limit as ε becomes close to zero, coordination on each institution should be observed for a long time. After a time, though, experimentation will induce a short transitional period leading to coordination in a different institution. Hence, most of time we will observe only one active institution. In the limit as $\varepsilon \to 0$, one should observe coordination on each institution an infinitely long amount of time "before" a switch occurs. This literal prediction of stochastic stability can of course not be observed in the laboratory.

A priori, two kinds of results in the laboratory might be seen as compatible with the existence of multiple stochastically stable institutions. Within the limited time available in the laboratory, one might observe that some groups start to converge toward coordination on one institution, while other groups start converging toward another institution. Since several institutions are stochastically stable, strictly speaking this would be compatible with the theory. Of course, universal convergence toward one of the institutions would contradict the results.

The second possible result is to observe actual coexistence. If a setting as ours is simulated for very small $\varepsilon > 0$, one will indeed observe alternation, with long time intervals spent on every stochastically stable state. However, if simulations are made with small but larger values of ε, one will typically observe relatively long periods of coexistence, as transitions remain relatively slow. That is, for noisy (realistic) environments, it is reasonable to translate the theoretical multiplicity of stochastically stable institutions (which is a double-limit prediction) as a prediction of coexistence of institutions in the short to medium run. This is the kind of phenomenon that we expected to see in the laboratory whenever several institutions are stochastically stable. Beyond this specific prediction, however, we expected to observe qualitatively different results between situations with only one stochastically stable institutions and situations with several such institutions. For instance, the theoretical predictions would be refuted if we observed universal convergence toward the market-clearing institution independently of whether other stochastically stable institutions are available or not.

7.1 The experimental design

In order to test which market institutions survive in the long run, we ran experiments where buyers and sellers had to choose between three different market institutions. The focus of the experiment was the choice of the trading platform, and not the trading behavior at a given platform. Therefore, subjects did not actually conduct trading interactions on the platforms. Rather, each subject only had to choose between the feasible platforms, and his payoff was directly determined by this choice, by his type (buyer or seller), and by the number of other buyers and sellers that opted for the same market institution. That is, the subjects played a simultaneous move game where, every period, they chose one of the possible institutions. In particular, they were not pre-assigned to any institution in the first period. Such a pre-assignment would have had the advantage of allowing us to experiment with different initial conditions, but the disadvantage of possibly inducing priming or experimenter demand effects.

We restricted the experiment on a test of the platform choice part of the theory for two reasons. First, this abstraction allows us to concentrate on the institution selection process. As already discussed, the stochastic-stability results regarding institution selection are based on a double limit, which makes testing the predictions particularly important. On the other hand, any additional "behavioral noise" in the second, trading stage makes it less likely that traders will be able to coordinate on the market-clearing institution in the first, platform choice stage. In our view, the most surprising result of the theoretical analysis is the long-term survival of non-market-clearing institutions. Hence, we want a design that does not obscure the possibility of full coordination on the market-clearing institution. While the abstraction from the trading stage limits the parallelism of the experiment to real market selection, it makes it easier to observe the benchmark case of coordination on the market-clearing platform.

The buyers' demand functions and payoffs were derived from a quasilinear utility function for two goods, $u(q_0, q) = q_0 + v(q)$ with v a strictly increasing function. Specifically, for the derivation of the numerical payoffs used in the experiment, we used $v(q) = 5q - \frac{1}{2}q^2$ (for $q \in [0, 5]$). We then replaced $q_0 = w - pq$ and used $w = 1$

to obtain the valuation $v_B(q, p) = 1 - pq + 5q - \frac{1}{2}q^2$. In order to obtain payoffs in a reasonable range for the experiment, we then applied a monotonic transformation to these values.[16] The sellers' supply and payoffs were derived from the profit function $\pi(p, q) = pq - \frac{1}{8}q^2$, i.e., sellers were producers with quadratic costs.

One platform (platform A) was market clearing ($\beta = 1$), the second one (platform B) was biased with $\beta = 0.8$, and the third one was biased with $\beta = 0.4$. The resulting payoff matrices are shown in "Appendix 2." As it can be easily checked from the payoff matrices, the sum of payoffs was maximized when all traders opted for A. But whenever the distribution of traders over the platforms was such that B was active, traders of one market side were better off at B than the traders of the same market side at any of the other platforms. So B was favored and hence stochastically stable. Examination of the payoff matrices in "Appendix 2" shows that two mutations suffice for a successful transition away from platform C, while a significantly larger number of mutations is necessary in order to reach platform C from the states where full coordination in either of the other platforms obtains. Following standard arguments, this suffices to establish that platform C is not stochastically stable (and hence not favored).

We conducted three different treatments. In treatment 1 (T1), subjects had to choose between platforms A and B. In treatment 2 (T2), traders chose between A and C, and in treatment 3 (T3), they chose between all three platforms. The theoretical model predicts that in the long run, subjects will opt for both platforms in T1 and only for platform A in T2. In T3, A and B should stay active while nobody should opt for C in the long run.

Each treatment was run with six groups of seven buyers and seven sellers each. Each subject played the game for 90 times ("periods"), during which the group composition did not change. Each subject was member of only one group. In each period, subjects had to choose between the available platforms within 30 seconds. At the end of each period, traders were informed about their own payoffs as well as about the distribution of the group members over the feasible platforms. The instructions (see "Appendix 2") avoided terms like market platform, buyer/seller, etc. Instead, it used terms like decision, Type I(II), etc. Subjects had access to the payoff tables for their type (see "Appendix 2") but not to those of the opposite type. The experiments were conducted at the University of Konstanz (Germany). The subjects were undergraduates of all fields except economics and psychology. A subject's overall payoff was the sum of the payoffs earned in all the 90 periods. The exchange rate between the ECU of the payoff matrices and Euro was 0.7 Eurocent. Overall, the average subject received 11.55 Euros. A session lasted about 70 min.

7.2 Experimental results

First, we investigate which platforms are opted for in the long run. Then, we investigate the individual decision behavior and, in particular, whether the model's assumptions on the learning process are supported by the data.

[16] The transformation was $v' = 10 + 8(\arctan(1.1(v - 9.2) - \arctan(-9.02))$. Payoffs were then rounded.

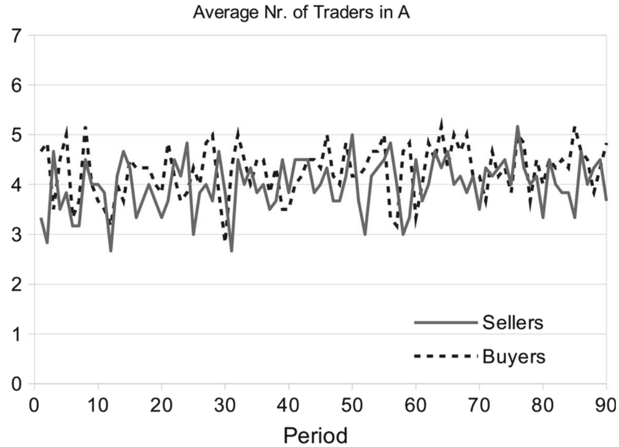

Fig. 1 Evolution of the number of traders in institution A (market clearing) in T1 (averaged across six sessions). The remaining traders are at institution B (not market clearing, but also stochastically stable)

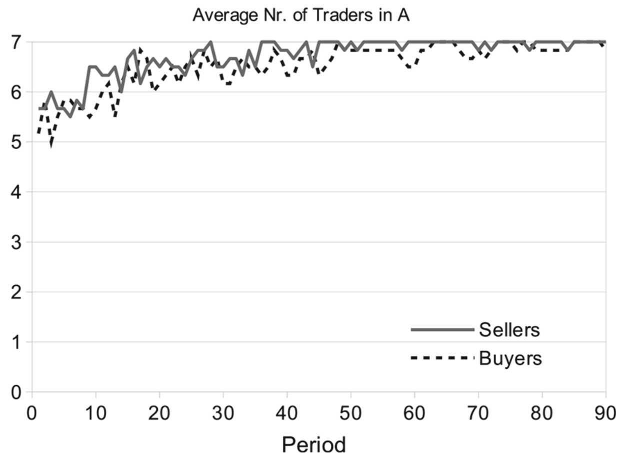

Fig. 2 Evolution of the number of traders in institution A (market clearing) in T2 (averaged across six sessions). The remaining traders are at institution C (not stochastically stable)

Figures 1 and 2 present the results of T1 and T2, respectively. The figures plot the time evolution of the number of traders in the market-clearing institution A, averaged across the six sessions of each experiment. The remaining traders are in institution B in the case of T1, and in institution C in the case of T2. The figures show a remarkable compliance with the theoretical predictions. In T1, both institutions are stochastically stable, and in the experiment, both remain active over time, with traders allocating themselves among both. In T2, only the market-clearing institution is stochastically stable, and indeed, traders quickly learn to coordinate on it and avoid the other institution.

Figure 3 presents the results of T3, where all three institutions were available. Since there is no significant difference between buyers and sellers in their choice of

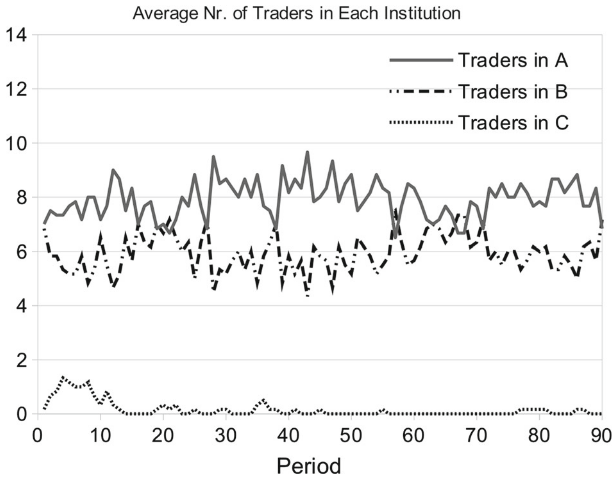

Fig. 3 Evolution of the number of traders in each institution in T3 (averaged across six sessions)

Table 1 Percentage of choices made during the last 30 periods

Group	1	2	3	4	5	6	Average all groups
T1-A	61.7	68.1	57.6	58.6	59.2	62.6	61.3
T1-B	38.3	31.9	42.4	41.4	40.8	37.4	38.7
T2-A	100.0	100.0	95.3	100.0	99.8	100.0	99.2
T2-C	0	0	4.7	0	0.2	0	0.8
T3-A	60.7	57.6	54.5	55.3	51.9	54	55.6
T3-B	38.1	42.4	45.5	44.5	48.1	46	44.1
T3-C	1.2	0	0	0.2	0	0	0.3

platform, we do not present the results for buyers and seller separately, but rather plot the average total number of traders in each of the three institutions. The results are again in agreement with the theoretical predictions. Institution C, which is not stochastically stable, is quickly abandoned in favor of the other two, stochastically stable institutions.

In each of the three treatments, at least half of the traders opt for platform A. When feasible, however (T1 and T3), platform B also remains active in the long run. But when available (T2 and T3), platform C becomes inactive during the first 15 rounds and stays empty or almost empty until the end. In summary, these observations yield:

Result 1 In the long run, traders opt for the stochastically stable platforms A and B, while platform C is avoided.

This result is not an artifact of taking the average over all groups. Rather, it can also be observed for each individual group. Table 1 presents, for all individual groups, the percentage of traders opting for the different feasible platforms during the last 30 periods. For example, in group 2 of T3, 57.6 % of the traders opted for platform A during the last 30 rounds, 42.4 % for platform B, and 0 % for platform C.

Table 2 FA denotes the benchmark payoff of full coordination on the market-clearing institution A

	Sellers	Buyers	Sum
FA	18	20	38
T1	17.01	19.62	36.63
T2	17.85	18.97	36.82
T3	16.95	19.53	36.48

In all groups, at least 50 % of the traders opted for platform A during the last 30 periods. If available, at least 30 % opted for platform B. But less than 5 % opted for platform C, and in 10 out of 12 cases, less than 1 percentage of the traders opted for C when available.

Remark 4 Full coordination on any trading institution, stochastically stable or not, is a strict Nash equilibrium of the institution choice game by virtue of Assumption **A3**. Hence, this is also true for the games tested in our experimental treatments. In particular, this implies that such profiles are trembling-hand perfect (and also proper) and, hence, we cannot obtain finer predictions from the standard game-theoretical refinement literature for our setting. One could ask whether quantal response equilibria (QRE) could be used to fit the data. However, QRE for non-negligible noise parameters (as used in the literature; see McKelvey and Palfrey 1995; Goeree et al. 2005) are completely mixed profiles, which would not explain convergence to full coordination in treatment 2. Likewise, one might ask whether observed play can be approximately described as mixed-strategy Nash equilibria. This interpretation, however, would not be compatible with the differences in results between treatments 1 and 2. A similar point concerns asymmetric (pure) equilibria. In T1, observed average play is close to a split with four agents of each type at institution A, and the rest at institution B. This profile, however, is not a Nash equilibrium, while profiles with strictly less than four traders of each type at institution A are (recall Footnote 6).

To evaluate the efficiency loss due to the existence of alternative institutions, we can take the average payoffs over all 90 periods and all sessions of the different treatments (see Table 2).

The absolute efficiency level depends crucially on the size of the market and on the number and the properties of the available institutions. Compared to full coordination on the market-clearing institution, both types of traders are made worse off by the availability of other institutions. This inefficiency is due to the fact that learning is not instantaneous, and some traders stay away from the market-clearing institution at least in the first periods. The patterns of payoffs in the different treatments are as predicted. Sellers are best off at T2, while the availability of the stable, non-market-clearing platform B with $\beta < 1$ in T1 and T3 hurts them. Buyers are better off in T1 and T3 than in T2. Taking the session averages, we find that these results are statistically significant. With respect to overall efficiency, the sum of payoffs is slightly higher in T2 than in the other treatments. This difference is not significant. It is, however, much more pronounced if one takes only the last 30 periods. In this case, the average sums of the payoffs are 36.87 (T1), 37.69 (T2), and 36.77 (T3). Taking the session averages

Table 3 Switches and imitation switches

Treatment	Switches	%	Imitation switches	%
T1	2315	31.0	1529	65.9
T2	512	6.8	267	52.1
T3	2587	34.6	1624	62.8

The percentage of switches is over the total number of decisions from period 2 onward. The percentage of imitation switches is over the total number of switches

of the last 30 periods, the sum of payoffs is significantly higher in T2 than in the other treatments.

Overall, Result 1 and the observations above provide a strong support for the main predictions of the theoretical model, namely that market-clearing as well as other stochastically stable institutions will be used in the long run, while other institutions will be avoided. To further investigate the reasons for this result, we take a closer look at the individual behavior. In particular, we investigate the behavior of traders who change the platform from one period to the next ("platform switching").

In Table 3, we provide the number of switches observed for the different treatments as absolute numbers and as percentage of the number of total decisions. We also look at the number of cases where subjects switch to an institution which gave the traders of their own type the highest possible payoff in the last period, i.e., switches consistent with our model (fourth column). Since in these cases the subjects act as if they imitate the most successful last period choice, we call them *imitation switches*. Table 3 also provides the percentage of imitation switches over all switches (fifth column).

Table 3 shows that in T1 and T3, subjects switch in about 30–35 % of all possible cases (the maximum possible number of switches per treatment is 7476, 89 periods times 14 subjects times 6 groups). In T2, however, subjects switch in only 6.8 % of all cases. In T1 and T3, 60–65 % of all switches were consistent with our learning models, while in T2, the respective percentage is 52 %.

Taking the group percentages of imitation switches, we can test the null hypothesis that these percentages are equally likely to be strictly above 50 % as weakly below 50 %, i.e., there is no tendency of imitation switching. In all six groups of T1, the percentage of imitation switching was between 62 and 70 %. In T3, these percentages were between 59 and 65 %. Therefore, a binomial test shows that for T1 and T3, the null hypothesis has to be rejected at any significance level. For T2, we find one group with 50 % imitation switching, and five groups with percentages between 51 and 55 %. In this case, the null hypothesis also has to be rejected at the 5 % level.

This provides evidence consistent with institution switching in favor of institutions with maximum observed payoffs.[17] Overall, we observe a tendency toward imitation switches in all three treatments, but this tendency is weaker in T2 than in T1 and T3. This difference is not surprising. Given the near perfect coordination on the unique

[17] That is, we obtain evidence in favor of the relevance of *observed past performance*, which is explicitly *not* compatible with purely forward-looking behavior. One natural interpretation of the behavior underlying the observed phenomena is imitation learning. However, in our context, it is also natural to postulate a form of reinforcement learning where the payoffs in other institutions are taken as observable *bygone* payoffs.

stochastically stable platform A in T2, there were only very few possibilities where an imitation switch was possible at all. This implies that a smaller percentage of switches should be imitation switches, since other reasons for a switch (e.g., experimentation) should be of comparable strength across all treatments. In summary:

Result 2 Individual traders tend to switch to a platform which in the last period gave the highest payoff to traders of their own type. This tendency is stronger in T1 and T3 than in T2.

To further investigate individual behavior, the learning model must be further specified. In particular, as mentioned in Sect. 3.3, we hypothesized that the likelihood of a revision would depend on the observed payoff differences between the own and other institutions. Hence, we test a learning model where the revision probability is strictly increasing in the difference between the highest last-period-own-type payoff and the last-period-own payoff.

Denote by Δ the difference between the highest last-period-own-type payoff and the last-period-own payoff. For the case of two platforms, s denotes a dummy which simply takes value 1 if a switch to the other platform occurs. For T3, the definition of s is more involved, because many more possibilities exist. We define s as a variable which takes the value 1 if either last period's platform did not deliver maximal payoffs and a switch to the last-period-best among the other two platforms occurred, or last period's platform did deliver maximal payoffs and a switch to some other platform occurred. This definition is the natural generalization of the dummy variable for the two-platform case. The logic is as follows. Consider first the case where $\Delta > 0$, that is, last period's platform did not deliver the highest payoffs. In the two-platform case, the decision consistent with our basic decision rule involves a switch, i.e., $s = 1$. In the three-platform case, $s = 1$ indicates again the choice consistent with the basic decision rule, which corresponds to a switch to the appropriate platform, but not to the third one. In the case $\Delta = 0$, the decision consistent with our basic decision rule in the two-platform case is to stay, i.e., $s = 0$. In the three-platform case, $s = 0$ again indicates the choice consistent with the basic decision rule. The main difference between the dummy variables in the two- and three-platform cases is that, when $\Delta > 0$, with three platforms, a value of $s = 0$ might indicate either that the agent did stay in his previous platform (which might correspond to either inertia or a mistake), or also a switch to a "third platform," which is neither his previous one nor the one which delivered highest payoffs. Switches of the last type can obviously not occur in the two-platform case. But in T3, only 85 decisions (out of 7476) were of this type.

Since each trader has to decide 89 times whether to switch or not, we have a strongly balanced panel data set. We conduct a probit regression with random effects with s as dependent variable. The most important independent variable is Δ, and to allow for nonlinearities, we include Δ^2. We also include the period, a type dummy, and dummies for the groups. Since all group dummies are insignificant except for one group in T3, they are not reported in Table 4.

The regressions deliver the following main result.

Result 3 In all three treatments, the switching probability is strictly increasing in the difference between highest last-period-own-type payoff and the last-period-own payoff.

Table 4 Regression results

	T1	T2	T3
Δ	0.0289	0.2601	0.0252
	(0.000)	(0.000)	(0.000)
Δ^2	−0.0004	−0.0087	−0.0004
	(0.006)	(0.000)	(0.000)
Period	−0.0038	−0.0251	−0.0050
	(0.000)	(0.000)	(0.000)
Type	−0.0484	−0.1412	−0.0276
	(0.674)	(0.222)	(0.735)
Constant	−0.6442	−0.9573	−0.5140
	(0.000)	(0.000)	(0.001)
No. obs.	7476	7476	7476
Log likelihood	−4239.4064	−11,151.2425	−4496.9168
Chi^2	0.0000	0.0000	0.0000

Random effects probit regressions on switches agreeing with the decision rule. Entries in brackets are p-values

As can be seen from Table 4, in all three treatments, the impact of Δ on s is positive and highly significant (p-values are shown in brackets below the corresponding coefficients). The negative coefficient of Δ^2 shows that the marginal impact of Δ is decreasing, but it remains positive for all feasible levels of Δ. That is, in accordance with our theoretical model, the likelihood of a switch to the last period's best platform is indeed increasing in the difference between the highest and the own payoff of the last period.

The period variable is significantly negative in all three treatments. This implies that the likelihood of a switch decreases over time. One could interpret this as an indication that not only the last-period experience determines the switching behavior. Rather, previous experiences also matter, and therefore in later periods, the last-period experience has a smaller impact than in earlier periods.

In general, T1 and T3 deliver quite similar results. The size of the coefficients of T2 differs substantially from those of T1 and T3. This indicates again clear but unsurprising differences between T2 and the other treatments. In T3, platform C is essentially disregarded by the traders from the very beginning, so the actual choice for the traders is between A and B, as in T1. In T2, of course, the choice is between A and C and convergence to complete coordination on A occurs, which is not observed in the other treatments.

8 Conclusions

We have presented a model where traders can choose among different trading institutions and asked whether they will learn to coordinate on an institution that guarantees market clearing. Under a general class of learning dynamics, we find that the market-clearing institution is always stochastically stable. It is, however, not necessarily the *only* one. We also find non-market-clearing institutions that are stochastically stable under general conditions, even if the market becomes large. As a result, coordination on market-clearing institutions will be often observed as the result of learning, but other institutions might also survive in the long run.

Stochastic stability, however, involves a double limit as time goes to infinity and the probability of mistakes goes to zero. Neither of these limits corresponds to a realistic situation. In order to test for the relevance of our theoretical results, we conducted a laboratory experiment on platform selection based on the structure of our model. The results are remarkably in agreement with the qualitative content of the theoretical predictions. Traders quickly learn to avoid institutions which are not stochastically stable, while all stochastically stable institutions remain active. Whenever a stochastically stable institution is confronted with another institution which fails this criterion, the outcome is a sharp selection result, as predicted by the model, even though the length of the experiment was relatively short and decision errors were not influenced in any way. Whenever two institutions are stochastically stable, the qualitative prediction is that both institutions are equally stable and no quick convergence to either one should be observed for a finite time interval and non-negligible noise level. This qualitative prediction is readily observed in the data. Hence, we can conclude that the predictions of our model are also relevant for actual decision-making. Further, the analysis of individual-level data also supports our behavioral assumptions and, in particular, the hypothesis that switching probabilities increase in the payoff difference between other institutions and the one currently used by a trader.

Acknowledgments We thank two anonymous referees, Ana B. Ania, Johannes Buckenmaier, Luis Corchón, Michihiro Kandori, Akihiko Matsui, Alex Possajennikov, and Jörgen Weibull for helpful comments and suggestions.

Appendix 1: Proofs

Proof (Proof of Lemma 1) Suppose $\beta_z(n_z, m_z) < 1$. Then, buyers, but not sellers, are rationed at institution z. We have that $q_S^0 = s(p_0)$ and $q_S^z = s(p_z)$. Suppose now that $v_S(q_S^0, p_0) \leq v_S(q_S^z, p_z)$. By **A1**, we must have that $p_0 \leq p_z$.

Then, again by **A1**, $v_B(q_B^0, p_0) \geq v_B(d(p_z), p_z)$. Since $d(p_z) > q_B^z$ (buyers are rationed), **A2** implies that $v_B(d(p_z), p_z) > v_B(q_B^z, p_z)$.

The case $\beta_z(n_z, m_z) > 1$ is analogous.

If $\beta_z(n_z, m_z) = 1$, no traders are rationed, and the analogous arguments follow with weak inequalities (**A2** does not apply since there is no rationing). □

We now prove some auxiliary lemmata. The first shows that traders in an inactive institution will always prefer any active institution.

Lemma 2 *Suppose that institution z is inactive and institution z' is active. If traded quantities are not zero at z', any trader in z or z' strictly prefers the outcome of institution z' to that of z. If traded quantities are zero at z', traders in z or z' are indifferent between the two institutions.*

Proof Since $m_z = 0$ or $n_z = 0$, traded quantities are zero at z. Since both the bias $\beta_{z'}(n_{z'}, m_{z'})$ and the market-clearing price $p^*(n_{z'}, m_{z'})$ are strictly positive, it follows that $p_{z'}(n_{z'}, m_{z'}) > 0$. If $d(p_{z'}(n_{z'}, m_{z'})) > 0$, the claim then follows from assumption **A3**. If $d\left(p_{z'}(n_{z'}, m_{z'})\right) = 0$, the claim follows from our explicit assumption as part of **A3** that evaluations do not depend on hypothetical prices. □

Lemma 3 *Assume* **A1**–**A3**. *Under* **D0**–**D2**, *given any state ω with $n_0(\omega) \geq 1$ and $m_0(\omega) \geq 1$, there exists a finite, positive probability path of the unperturbed dynamics leading from ω to the state ω_0 with $n_0(\omega_0) = n$ and $m_0(\omega_0) = m$.*

Proof Consider any institution $z \neq 0$, which is chosen by some traders in state ω. If $n_z(\omega) = 0$ or $m_z(\omega) = 0$, by Lemma 2, we can build a positive-probability path to a new state where no trader is at institution z. Hence, without loss of generality, suppose that $n_z = n_z(\omega) > 0$ and $m_z = m_z(\omega) > 0$.

If $\beta_z(n_z, m_z) \neq 1$, it follows from Lemma 1 that in state ω, at least one of the two types of traders strictly prefers the market-clearing institution. Let k be a trader of that type who is at the non-market-clearing institution z. It might happen that k prefers a third institution to the market-clearing one, but certainly will not stay in z if given revision opportunity. Further, by Assumption **D1**, there is strictly positive probability that k is the only trader of his type obtaining revision opportunity. Consider the paths where this event happens, and let k' denote a trader of the *other* type (i.e., not of the same type as k) who, in state ω, is in the same non-market-clearing institution z. Consider now the event that only k and k' get revision opportunity.

If this event has positive probability, then (if it occurs) k' may or may not change institution, but k will, switching to the market-clearing or another institution. If the probability of k and k' being the only revising traders is zero, by Assumption **D2**, no agent of the same type as k' will revise this period, and hence k will change institution but no other agent will. In any case, the process reaches a state with strictly less traders at institution z than there were in ω, but at least the same traders in the other institutions (and, in particular, the market-clearing one). If $\beta_z(n_z, m_z) = 1$, Lemma 1 yields weak preferences. The argument above applies again, because by Assumption **D0** ties are broken randomly, i.e., if a trader weakly prefers another institution to his current one, there is a maybe small but positive probability that he switches away.

Repeating this argument, we will reach a state ω' with either $n_z(\omega') = 0$ or $m_z(\omega') = 0$. From this state, all remaining traders will leave institution z as above (by Lemma 2). Hence, we reach a state where strictly less institutions are chosen than in ω.

Repeating this procedure, we will reach a state where only two institutions are chosen by traders, and one of them will necessarily be the market-clearing one. Applying again the same argument (using Lemma 1) shows that we can construct a positive-probability path to ω_0, where 0 is the only active institution. □

Proof (Proof of Proposition 1) The states given in (i) are obviously absorbing because, in the absence of experimentation, traders will never switch to unobserved institutions. To see that there are no other absorbing states, suppose that there are traders of the same type in at least two different institutions. Since necessarily one of those institutions is yielding (weakly) higher payoffs than the other, and under Assumption **D1** there is positive probability that one of the traders not in that institution is given revision opportunity, there is a positive probability transition to a different state, a contradiction. Finally, if there is only one active institution but the state is not monomorphic, the result follows from Lemma 2. Part (ii) follows immediately from Lemma 3. □

The remaining proofs rely on the characterization of the set of stochastically stable states introduced by Kandori et al. (1993) and Young (1993), and on the concept of radius and coradius developed by Ellison (2000). Given two absorbing sets A and B, let $c(A, B) > 0$ (referred to as the *transition cost* from A to B) denote the minimal number of mistakes in a positive-probability path starting in an element of A and leading to an element in B. The following lemma contains all the results on stochastic stability that we require for the analysis. Its proof is a straightforward application of Ellison (2000, Theorems 1 and 3) and is analogous to the proof of Lemma 2 in Alós-Ferrer and Kirchsteiger (2010); hence, we omit it here.

Lemma 4 *Let A be an absorbing set and define the* Radius *of A by*

$$R(A) = \min\{c(A, B)\,|\,B \text{ is an absorbing set, } B \neq A\}$$

and the Coradius *of A by*

$$CR(A) = \max\{c(B, A)\,|\,B \text{ is an absorbing set, } B \neq A\}$$

Then:

(i) *If $R(A) \geq CR(A)$, the states in A are stochastically stable.*
(ii) *If $R(A) > CR(A)$, the only stochastically stable states are those in A.*
(iii) *If the states in an absorbing set B are stochastically stable and $R(A) = c(B, A)$, the states in A are also stochastically stable.*

Proof (Proof of Theorem 1) We have to show the stochastic stability of the state ω_0. If there is any other market-clearing institution, the conclusion follows by renaming. First, notice that, by Lemma 2, no monomorphic state can be left with less than two mutations unless the traded quantity is zero. Since traded quantities at a market-clearing institution are never zero, it follows that $R(\{\omega_0\}) \geq 2$.

Consider any state in any absorbing set other than $\{\omega_0\}$. Notice that two mutations (to the market-clearing institution) suffice to reach a state ω with $n_0(\omega) \geq 1$ and $m_0(\omega) \geq 1$. By Lemma 3, there is a positive-probability path of the unperturbed dynamics (i.e., requiring no further mutations), leading to ω_0. This shows that $CR(\{\omega_0\}) = 2$ (the equality follows because two mutations are required to leave any other monomorphic state). The result follows from Lemma 4(i). □

Proof (Proof of Theorem 2) Let ω_z denote the monomorphic state corresponding to coordination on institution z. We know from Theorem 1 that ω_0 is stochastically stable. By definition of a favored institution, we see that if exactly two mutations to institution z occur at state ω_0, we reach a state where at least one type of traders strictly prefer that institution. Analogously to the proof of Lemma 3 (through repeated application of Definition 3), from this state, there exists a positive-probability path involving no further mutations which leads to state ω_z. From the proof of Theorem 1, we already know that it is possible to make the opposite transition with exactly two mutations (but no less). Thus, we obtain that $c(\{\omega_0\}, \{\omega_z\}) = 2 = R(\{\omega_z\})$, and the result follows from Lemma 4(iii). □

Proof (Proof of Theorem 3) For any given n, m, let

$$R(n,m) = \left\{ \frac{a}{b} \,|a = 1, \ldots, n, \text{ and } b = 1, \ldots, m \right\}$$

be the set of feasible buyers–sellers ratios and define

$$\underline{\beta}(n,m) = \max_{r \in R(n,m)} \underline{\beta}(r) \qquad \text{and} \qquad \overline{\beta}(n,m) = \min_{r \in R(n,m)} \overline{\beta}(r).$$

Note that for any given number of buyers and sellers, there exists only a finite number of values r can take. Hence, by **A4**, $\underline{\beta}(n,m) < 1 < \overline{\beta}(n,m)$. For any given n, m, define $T(n,m)$ as the set of all pairs (r_0, r_z) such that $r_0 = \frac{n_0}{m_0}$ and $r_z = \frac{n_z}{m_z}$ with $n_0, n_z \in \{1, \ldots, n\}$ and $m_0, m_z \in \{1, \ldots, m\}$ such that $n_0 + n_z = n$ and $m_0 + m_z = m$. In other words, $T(n,m)$ is the set of all pairs of buyer–seller ratios which are feasible when exactly two institutions are simultaneously active. Finally, define

$$\underline{\beta}^*(n,m) = \max \left\{ \underline{\beta}(n,m), \max \left\{ \frac{p(r_0)}{p(r_z)} \,|(r_0, r_z) \in T(n,m) \text{ and } r_0 < r_z \right\} \right\}$$

$$\overline{\beta}^*(n,m) = \min \left\{ \overline{\beta}(n,m), \min \left\{ \frac{p(r_0)}{p(r_z)} \,|(r_0, r_z) \in T(n,m) \text{ and } r_0 > r_z \right\} \right\}$$

Notice that $\underline{\beta}^*(n,m)$ (and analogously $\overline{\beta}^*(n,m)$) is well defined because $T(n,m)$ is finite and $r_0 < r_z$ implies $\frac{p(r_0)}{p(r_z)} < 1$. Clearly, $\underline{\beta}(n,m) \leq \underline{\beta}^*(n,m) < 1 < \overline{\beta}^*(n,m) \leq \overline{\beta}(n,m)$.

Consider an institution F such that $\underline{\beta}^*(n,m) < \beta_F(n_z, m_z) < \overline{\beta}^*(n,m)$ for all feasible n_z, m_z. We prove that F is favored.

We want to show that, whenever $v_S\left(q_S^0, p_0\right) \geq v_S\left(q_S^F, p_F\right)$, then $v_B\left(q_B^0, p_0\right) < v_B\left(q_B^F, p_F\right)$. Let $\beta_F = \beta_F(n_z, m_z)$ be the realized bias at institution F. Suppose $\beta_F < 1$. Then, buyers, but not sellers, are rationed at F. We have that $q_S^0 = s(p_0)$ and $q_S^F = s(p_F)$. Suppose now that $v_S(q_S^0, p_0) \geq v_S(q_S^F, p_F)$. By **A1**, and we must have that $p_0 \geq p_F$.

Suppose that $r_0 \geq r_F$. Then, $p_0 = p(r_0) \geq p(r_F)$ and, by **A1**, $v_B(q_B^0, p_0) = V_B(1, r_0) \leq V_B(1, r_F)$. By **A4**, $V_B(1, r_F) < V_B(\beta_F, r_F) = v_B\left(q_B^F, p_F\right)$ and the claim follows.

Suppose now that $r_0 < r_F$. Then, $p_0 = p(r_0) < p(r_F)$. If, as assumed, $v_S\left(q_S^0, p_0\right) \geq v_S\left(q_S^F, p_F\right)$, then $p_0 \geq p_F = \beta_F \cdot p(r_F)$ by **A1**. It follows that $\beta_F \leq \frac{p(r_0)}{p(r_F)}$, a contradiction with $\beta_F > \underline{\beta}^*(n,m)$. Hence, $v_S\left(q_S^0, p_0\right) < v_S\left(q_S^F, p_F\right)$.

The case $\beta_F > 1$ is analogous. □

The following Lemma is used in the proof of Theorem 4.

Lemma 5 *Assume* **A1** *and* **A4**. *If* $\underline{\beta}(1) < \beta_B < 1 < \beta_S < \overline{\beta}(1)$, *then*

(i) *if* $m \leq n$ *(more buyers than sellers), in a state where an equal, strictly positive number of sellers and buyers are at an institution* z *with constant* $\beta_z = \beta_B$ *and the remaining traders are at a market-clearing institution (and there are traders of both types in the latter), buyers strictly prefer* z*;*

(ii) *if* $m \geq n$ *(more sellers than buyers), in a state where an equal, strictly positive number of sellers and buyers are at an institution* z *with constant* $\beta_z = \beta_S$ *and the remaining traders are at a market-clearing institution (and there are traders of both types in the latter), sellers strictly prefer* z.

Proof We will show part (i). Part (ii) is analogous. Let $0 < \ell < \max(m, n)$ be the number of sellers and buyers at the alternative institution z. Since $m \leq n$, we have that $m - \ell \leq n - \ell$ and hence

$$r = \frac{n - \ell}{m - \ell} \geq 1$$

That is, there are (weakly) more buyers than sellers at the market-clearing institution. By **A4**, since $\underline{\beta}(1) < \beta_B < 1$,

$$V_B(\beta_B, 1) > V_B(1, 1)$$

and, by **A1**,

$$V_B(1, 1) = v_B \left(q_B^z(1, 1), p(1)\right) \geq v_B \left(d(p(r)), p(r)\right)$$

because $q_B^z(1, 1) = d(p(1))$ and $p(1) \leq p(r)$ since $r \geq 1$ and p is increasing in r. Hence,

$$V_B(\beta_B, 1) > v_B \left(d(p(r)), p(r)\right)$$

which proves the claim, because $v_B \left(d(p(r)), p(r)\right)$ is the buyers' payoff at the market-clearing institution, and $V_B(\beta_B, 1)$ is the payoff of the buyers at the non-market-clearing institution with $\beta_z = \beta_B$. □

Proof (Proof of Theorem 4) We will show part (i). Part (ii) is analogous, and part (iii) follows from (i) and (ii). Let ω_1 denote the monomorphic state corresponding to coordination on the buyers' institution z. By hypothesis, z is a favored institution for the economy with $K = 1$. Further, we know from Theorem 1 that ω_0 is stochastically stable. In order to show stochastic stability of ω_1 for large K, by Lemma 4(iii), it is enough to show that two mutations at ω_0 from the market-clearing institution 0 to z suffice for a transition.

Let a buyer and a seller mutate from 0 to z. Then, $r_z = 1$ and $r_0 = \frac{Kn-1}{Km-1} \geq 1$. By Lemma 5(i), the mutant buyer is strictly better off. By **D3**$_K$, let K buyers revise, including the mutant, and follow him to z (so exactly $K - 1$ buyers switch). Now, $r_z = \frac{K}{1}$ and $r_0 = \frac{(n-1)K}{Km-1}$.

Since $\lim_{r\to\infty} p(r) = +\infty$ under **M1′**–**M2′**, $r_z \to \infty$ as $K \to \infty$, and $r_0 \to \frac{n-1}{m}$ (finite), even though $\beta_z < 1$ it follows that there exists K^* such that, for all $K \geq K^*$, $p_z = \beta_z p^*(r_z) > p^*(r_0) = p_0$.

Sellers are not rationed at 0 (by definition), and they are not rationed at z either since $\beta_z < 1$. Hence, by **A1**, sellers are strictly better off at z since they face a higher price. By $\mathbf{D3}_K$, there is positive probability that K sellers, including the lone seller already in z, revise and move to z from 0, while no other trader receives revision opportunity. In the new state, we have $r_z = \frac{K}{K} = 1$ and $r_0 = \frac{(n-1)K}{(m-1)K} = \frac{(n-1)}{(m-1)}$. By Lemma 5(i), $V_B(\beta_z, r_z) > V_B(1, r_0)$ and buyers at z are strictly better off. By $\mathbf{D3}_K$, there is positive probability that K buyers from 0 revise and follow them to z.

We know that the market institution z is favored for the economy with $K = 1$. That means that one market side is better off at z for all prices resulting from population proportions *which are feasible in the economy with* $K = 1$ (recall the construction of the set $T(m, n)$ in the proof of Theorem 3). In the state we have just reached, we have $r_z = \frac{2K}{K} = 2$ and $r_0 = \frac{(n-2)K}{(m-1)K} = \frac{n-2}{m-1}$, which are feasible population proportions in the economy with $K = 1$.

We conclude that one market side is better off at z. Let K traders of the appropriate market side switch (using $\mathbf{D3}_K$). The new population distribution is always a multiple of K for each trader type and each institution; hence, we can apply the fact that z is favored in the economy with $K = 1$ again. Proceeding iteratively, eventually we reach a state where a complete market side is at institution z. By **A3**, we can complete the transition by moving groups of K traders of the other market side to z until the market-clearing institution becomes empty. □

Appendix 2: Experimental instructions and payoff matrices

The instructions and the control questionnaire below are translated from German into English as literally as possible. These instructions were distributed to Type I traders (i.e., sellers) of treatment 3 (choice between three platforms). The instructions for Type II traders (buyers) were symmetric (of course with the appropriate payoffs in the examples). The instructions for T1 and T2 were similar, with the only difference that all the references to choice C were deleted.

Each participant was provided with the appropriate payoff tables for the institutions used in the experiment. Table 5 displays the payoffs obtained by buyers and sellers at each of the three institutions used in the experiments. Within each table, each row corresponds to the number n of buyers present at the institution, each column to the number m of sellers.

Instructions: Type I

The experiment you are about to participate in is part of a research project on decision behavior. The instructions are simple, and if you read them carefully and make appropriate decisions, you can earn a considerable amount of money.

The revenues made during the experiment are counted in ECU (“experimental currency units”). After the end of the experiment all the revenues you made during the experiment will be added up and paid to you in cash. For every ECU you will receive 0.7 Eurocent.

Table 5 Experimental payoff tables

Buyers									Sellers							
$n\backslash m$	0	1	2	3	4	5	6	7	$n\backslash m$	1	2	3	4	5	6	7
Institution A: Market clearing																
									0	10	10	10	10	10	10	10
1	10	20	30	31	32	32	32	32	1	18	12	11	11	10	10	10
2	10	12	20	28	30	31	31	32	2	32	18	14	12	12	11	11
3	10	11	13	20	27	29	30	31	3	47	25	18	15	13	12	12
4	10	11	12	14	20	26	28	30	4	60	32	23	18	16	14	13
5	10	10	11	13	15	20	25	27	5	72	40	27	21	18	16	15
6	10	10	11	12	13	16	20	24	6	82	47	32	25	21	18	16
7	10	10	11	11	12	14	16	20	7	91	54	37	29	23	20	18
Institution B: Stochastically stable but not market clearing ($\beta = 0.8$)																
									0	10	10	10	10	10	10	10
1	10	23	30	31	32	32	32	32	1	15	12	11	10	10	10	10
2	10	13	23	29	30	31	31	31	2	24	15	13	12	11	11	11
3	10	11	15	23	28	29	30	31	3	34	20	15	13	12	12	11
4	10	11	13	16	23	27	29	30	4	42	24	18	15	14	13	12
5	10	11	12	14	18	23	26	28	5	50	29	21	17	15	14	13
6	10	10	11	13	15	18	23	26	6	56	34	24	20	17	15	14
7	10	10	11	12	13	16	19	23	7	62	38	27	22	19	17	15
Institution C: Not stochastically stable ($\beta = 0.4$)																
									0	10	10	10	10	10	10	10
1	10	12	14	15	16	17	17	17	1	11	10	10	10	10	10	10
2	10	11	12	13	14	15	15	16	2	14	11	11	10	10	10	10
3	10	11	12	12	13	14	14	14	3	16	12	11	11	11	10	10
4	10	11	11	12	12	13	13	14	4	18	14	12	11	11	11	11
5	10	11	11	11	12	13	13	13	5	20	15	13	12	11	11	11
6	10	10	11	11	12	12	12	13	6	22	16	14	12	12	11	11
7	10	10	11	11	11	12	12	12	7	23	17	14	13	12	12	11

n stands for the number of buyers, m for the number of sellers at the institution. The actual tables used in the experiment were less condensed. In particular, n and m were replaced with explicit descriptions of the form "Number of participants of Type I/II who choose A/B/C" in front of the rows or above the columns, as appropriate. The tables are permuted so that for every participant, the number of participants of his or her type varied across rows

Any communication between the participants is strictly forbidden.

In every round of the experiment you have to choose between three different options, A, B, or C. To do so you click on the appropriate button for decision A, B, or C. Then you confirm your decision by clicking on the OK button. In each round you have half a minute to make this decision.

In this experiment there are two types of participants, participants of Type I and participants of Type II. You are of Type I. All in all, there are 7 participants of Type

I and 7 participants of Type II within the group relevant for you. You will not be informed about the identity of the other group members, and the other group members will not be informed about your identity.

The revenues you make in a round depend on the number of other group members of each type who make the same decision as you. Assume that you have made decision A, three other Type I participants have made the same decision, and 5 Type II participants have chosen A, too. In this case 4 Type I participants and 5 Type II participants have chosen A. As you can see from the attached revenue matrix, your revenues are 21 ECU in this case.

Another example: You have chosen B, one other Type I participant has made the same decision, and 3 Type II participants have chosen B, too. In this case your revenues are 20 ECU.

After all members of your group have made a decision, you will be informed about the number of participants of each type that have chosen A, B, and C. You will also be informed about their revenues, and about the sum of revenues you have earned so far in the whole experiment.

After that, a new round will start, in which you will have to decide between A, B, and C, again. Overall there will be 90 rounds.

Control Questionnaire

1. Suppose that you have made decision B, 2 other Type I participants have made the same decision, and 2 Type II participants have also chosen B. What are your revenues?
   Suppose that 2 Type I participants and 3 Type II participants have chosen A. What are the revenues of those Type I participants who have chosen A?
   Suppose the remaining 2 Type I participants and the remaining 2 Type II participants have chosen C. What are the revenues of those Type I participants who have chosen C?
2. Suppose that you have made decision C, 3 other Type I participants have made the same decision, and 2 Type II participants have also chosen C. What are your revenues?
   Suppose that 1 Type I participant and 2 Type II participants have chosen B. What are the revenues of the Type I participant who has chosen B?
   Suppose the remaining 2 Type I participants and the remaining 3 Type II participants have chosen A. What are the revenues of those Type I participants who have chosen A?

References

Alós-Ferrer, C.: Finite population dynamics and mixed equilibria. Int. Game Theory Rev. **5**(3), 263–290 (2003)

Alós-Ferrer, C., Kirchsteiger, G.: General equilibrium and the emergence of (non) market clearing trading institutions. Econ. Theory **44**, 339–360 (2010). doi:10.1007/s00199-009-0466-9

Alós-Ferrer, C., Netzer, N.: The logit-response dynamics. Games Econ. Behav. **68**, 413–427 (2010)

Alós-Ferrer, C., Netzer, N.: Robust stochastic stability. Econ. Theory **58**(1), 31–57 (2015). doi:10.1007/s00199-014-0809-z

Alós-Ferrer, C., Kirchsteiger, G., Walzl, M.: On the evolution of market institutions: the platform design paradox. Econ. J. **120**(543), 215–243 (2010)
Ariely, D., Ockenfels, A., Roth, A.: An experimental analysis of ending rules in internet auctions. RAND J. Econ. **36**, 890–907 (2005)
Ausubel, L., Cramton, P.: Demand Reduction and Inefficiency in Multi-unit Auctions. Mimeo, New York (2002)
Bajari, P., Hortaçsu, A.: The winner's curse, reserve prices, and endogenous entry: empirical insights from eBay auctions. RAND J. Econ. **34**, 329–355 (2003)
Benaïm, M., Weibull, J.: Deterministic approximation of stochastic evolution in games. Econometrica **71**, 878–903 (2003)
Binmore, K., Samuelson, L.: Muddling through: Noisy equilibrium selection. J. Econ. Theory **74**, 235–265 (1997)
Blume, L.: How noise matters. Games Econ. Behav. **44**, 251–271 (2003)
Budish, E., Takeyama, L.: Buy prices in online auctions: irrationality on the internet. Econ. Lett. **72**, 325–333 (2001)
Ellison, G.: Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. Rev. Econ. Stud. **67**, 17–45 (2000)
Fudenberg, D., Levine, D.: The Theory of Learning in Games. MIT Press, Cambridge (1998)
Gerber, A., Bettzüge, M.O.: Evolutionary choice of markets. Econ. Theory **30**, 453–472 (2007). doi:10.1007/s00199-005-0063-5
Goeree, J.K., Holt, C.A., Palfrey, T.R.: Regular quantal response equilibrium. Exp. Econ. **8**, 347–367 (2005)
Holt, C.: Industrial Organization: A Survey of Laboratory Research. In: Kagel, J., Roth, A. (eds.) Handbook of Experimental Economics. Princeton University Press, Princeton (1995)
Ishibuchi, H., Oh, C., Nakashima, T.: Designing a Decision Making System for a Market Selection Game. Mimeo, New York (2002)
Kandori, M., Rob, R.: Evolution of equilibria in the long run: a general theory and applications. J. Econ. Theory **65**, 383–414 (1995)
Kandori, M., Mailath, G.J., Rob, R.: Learning, mutation, and long run equilibria in games. Econometrica **61**, 29–56 (1993)
Karlin, S., Taylor, H.M.: A First Course in Stochastic Processes, 2nd edn. Academic Press, San Diego (1975)
Kugler, T., Neeman, Z., Vulkan, N.: Markets versus negotiations: an experimental analysis. Games Econ. Behav. **56**, 121–134 (2006)
Lucking-Reiley, D.: Auctions on the internet: what's being auctioned, and how? J. Ind. Econ. **48**(3), 227–252 (2000)
Madhavan, A.: Trading mechanisms in securities markets. J. Finance **47**(2), 607–641 (1992)
McKelvey, R.D., Palfrey, T.R.: Quantal response equilibria for normal form games. Games Econ. Behav. **10**(1), 6–38 (1995)
Plott, C.: Industrial organization theory and experimental economics. J. Econ. Lit. **20**, 1485–1587 (1982)
Quan, D.C.: Real Estate Auctions: A Survey of Theory and Practice. J. Real Estate Finance Econ. **9**, 23–49 (1994)
Reynolds, S., Wooders, J.: Auctions with a buy price. Econ. Theory **38**, 9–39 (2009). doi:10.1007/s00199-006-0182-7
Samuelson, L.: Stochastic stability in games with alternative best replies. J. Econ. Theory **64**, 35–65 (1994)
Samuelson, L.: Evolutionary Games and Equilibrium Selection. MIT Press, Cambridge (1997)
Schlag, K.: Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. J. Econ. Theory **78**, 130–156 (1998)
Shi, F.: Long-run technology choice with endogenous local capacity. Econ. Theory **59**(2), 377–399 (2015). doi:10.1007/s00199-014-0838-7
Young, P.: The evolution of conventions. Econometrica **61**, 57–84 (1993)
Zumpano, L.V., Elder, H.W., Baryla, E.A.: Buying a house and the decision to use a real estate broker. J. Real Estate Finance Econ. **13**, 169–181 (1996)