

Does Learning Lead to Coordination on Market Clearing Institutions?

Carlos Alós-Ferrer* and Georg Kirchsteiger†

September 2003

Abstract

This paper analyzes the question whether traders learn to coordinate on a trading institution that guarantees market clearing, or whether other market institutions can survive in the long run. While we find that the market clearing institution is indeed always stable under a general class of learning dynamics, it turns out that also other, non-market clearing institutions are stable. Hence, in the long run traders may fail to coordinate exclusively on market clearing institutions.

JEL-Codes: D4, D83

Keywords: Market Institutions, Coordination, Learning

1 Introduction

The formation of a market requires a group of people some of whom want to buy and some of whom want to sell. For the Walrasian theory of market clearing this is all that is needed. As soon as preferences and cost functions are specified, equilibrium analysis can be applied. Actual markets, though, are not merely characterized by demand and supply. Market exchange does not take place in an institutional void. It requires a 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. For example, the bazaar differs from financial markets in terms of the matching and bargaining process and the dissemination of information. The typical (Dutch) rules at flower auctions are different from the (English) rules at art auctions. Trading institutions not only vary across different goods, different market institutions exist even for the same good. Real estate is sold

*Department of Economics, University of Vienna (Austria). Email: Carlos.Alos-Ferrer@Univie.ac.at

†Department of Economics, University of Maastricht (The Netherlands). Email: G.Kirchsteiger@algec.unimaas.nl

both at auctions and by means of direct negotiations, and both call markets and continuous double auction markets are used to exchange financial assets.

There is some theoretical and empirical and a lot of experimental evidence indicating that the details of the market institution are consequential.¹ Trading rules affect the efficiency of the market outcome, the convergence towards equilibrium, the volatility of the prices, and the distribution of surplus over the market participants. Given that “institutions matter”, important questions are: Do the existing market institutions support market clearing equilibrium outcomes? Can we be confident that actual markets will be characterized by efficient institutions? Are there circumstances under which inefficient trading rules arise and persist, or are forces and mechanisms present that drive a market towards efficient organization?

In order to answer these questions, it is useful to distinguish between two different approaches. On the one hand, one can ask how new market institutions come into existence, and what the properties of these new institutions are.² On the other hand, we might ask which institutions survive in the long run if several institutions compete. In the paper at hand we provide an answer to the second question. Whether a specific market institution survives the competition with other existing market institutions depends on whether traders employ this institution or whether they trade under another set of rules. Hence, the (evolutionary) success of a market institution is a function of the number of traders it can attract.

The decision about the use of a particular market institution gives rise to a coordination game - potential buyers and sellers have to coordinate on a particular institution in order to make mutual beneficial trade possible. Of course, such coordination games typically exhibit a multiplicity of Nash-equilibria. There exists the possibility that in a Nash-equilibrium traders coordinate on an institution that does not lead to market clearing outcomes and that does not maximize the gains of trade. They might even coordinate on an institution that leads to a Pareto-inefficient outcome. It has been claimed that traders will eventually learn to overcome this problem and co-

¹Empirical and theoretical evidence is mostly found with respect to the impact of the design of auction rules. An empirical example of the importance of auction design in the context of internet auctions is provided by Ockenfels and Roth [20], who analyze empirically the impact of the auction design on internet auctions. A theoretical example in the context of multi-unit auctions is provided by Ausubel and Cramton [2]. For an overview of the theoretical literature of the impact of auction design see Klemperer [16]. An overview of the experimental evidence of the importance of market institutions is provided by Holt [10] and by the classical article of Plott [21].

²There exists some theoretical literature along these lines, where market institutions are modelled as networks (see Bala and Goyal [3], Jackson [12], and Kranton and Minehart [18]). Another example of this approach is Kirchsteiger, Niederle, and Potter [17], where the driving forces behind the development of new market institutions and the efficiency properties of these institutions are experimentally investigated.

ordinate on an efficient institution that guarantees market clearing outcomes (see e.g. Hayek [9]). To our knowledge this claim has never been put to a thorough investigation.³ Hence, our paper asks under what circumstances traders will indeed learn to coordinate on an efficient, market-clearing institution.

To provide an answer to this question, we use a partial equilibrium framework where we concentrate on the market institutions for trading an homogenous good. Potential traders have to choose simultaneously at which institution they want trade. We specifically postulate one market clearing institution and an arbitrary number of other, feasible institutions that do not lead to the market clearing, but realize a price below or above the market clearing level. Consequently, buyers or sellers who have chosen such an institution face rationing.

In order to model the learning process we use the stochastic stability techniques brought into the economics literature by Kandori, Mailath, and Rob [13] and Young [24] to analyze coordination games with players learning on which equilibrium to coordinate. Specifically, we postulate that traders (buyers and sellers) will decide to switch from one institution to a different one next period if they observe that the current-period results (prices and traded quantities as resulting from rationing) are better for them. Traders evaluate these 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. In particular, traders are not assumed to anticipate future prices, market-clearing or otherwise. The learning model can be interpreted as a version of Kandori, Mailath, and Rob's [13] because traders tend to switch to strategies (institutions) which are better in the current period, without anticipating the effects of their strategy change.⁴

We find two types of results. First, the market clearing institution is always stochastically stable under a general class of learning models. Second, also other, non-market clearing institutions are stochastically stable in

³There exist several papers which are similar in spirit to our paper insofar as also in these models traders can choose between different trading institutions (see e.g. Ishibuchi, Oh, and Nakashima [11], or Neeman and Vulkan [19]). Those papers, however, do not investigate whether traders learn to coordinate on efficient institutions guaranteeing market clearing prices and quantities.

⁴A conceptually related model is analyzed by Gerber and Bettzüge [8]. They postulate a finite population of traders who might choose among two identical, market-clearing asset markets. Traders have idiosyncratic preferences for the markets themselves and also perceive markets to be more attractive when the number of traders in them increases (size effect). When trading on a market, traders maximize a mean-variance utility function. Learning is modelled through Young's [24] adaptive play, i.e. traders best-reply to a sample of past play. In particular, traders are able to compute future market-clearing prices (also anticipating the effects of their own change in strategy) but are otherwise myopic. Gerber and Bettzüge find that, for a large number of traders, the only stochastically stable state is one where both markets are active and traders split equally among them.

general.

The paper proceed as follows. Next we describe the model and its basic assumptions. We proceed to present several examples and preliminary results concerning static comparisons of different institutions. In Section 3 we describe the learning process and in Section 4 we present the main results. Section 5 concludes. In the Appendix we provide a generalization of all our results to the case of institutions with stochastic prices.

2 The model

There is a homogeneous good to be traded by a set $I = \{1, \dots, n\}$ of buyers and a set $J = \{1, \dots, m\}$ of sellers. We denote the price of the good by p . The demand of a typical buyer $i \in I$ is given by a function $d(p)$ such that $d(0) > 0$, $d(p) \geq 0$ for all $p \geq 0$, and $\lim_{p \rightarrow \infty} d(p) = 0$. The function d is assumed to be continuous and strictly decreasing in p in the range where $d(p) > 0$. The supply of a typical seller $j \in J$ is given by a function $s(p)$ with $s(0) = 0$. We assume s to be continuous and strictly increasing in $p \geq 0$.

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 evaluation 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 $v_B : \mathfrak{R}_+^2 \rightarrow \mathfrak{R}$ and $v_S : \mathfrak{R}_+^2 \rightarrow \mathfrak{R}$.

The primitives in our model are the demand, supply and the payoff (evaluation) functions. We want to emphasize that this framework is much 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 more 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.

About the relation between demand, supply and the evaluation of the market outcome we make the following assumptions:

Assumption A1: In the absence of rationing, a lower price is better for buyers and worse for sellers. I.e.,

$$v_B(d(p), p) > v_B(d(p'), p') \text{ and } v_S(s(p), p) < v_S(s(p'), p') \quad (\text{A1})$$

for all p, p' with $p < p'$ and $d(p) > 0$.

Assumption A2: Given the price, traders prefer not to be rationed. I.e.,

$$v_B(d(p), p) > v_B(q_B, p) \text{ and } v_S(s(p), p) > v_S(q_S, p) \quad (\text{A2})$$

for all p and all $q_B < d(p), q_S < s(p)$

Assumption A3: Given the price, traders prefer being rationed to not being able to trade. I.e.,

$$v_B(q_B, p) > v_B(0) \text{ and } v_S(q_S, p) > v_S(0) \quad (\text{A3})$$

for all $p > 0$ and all $0 < q_B < d(p), 0 < q_S < s(p)$

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.

In the next subsection we show that these assumptions are not only very plausible at first sight, but that they are also fulfilled in standard models. In a first approach, demand and supply are derived from standard consumers' and producers' problems and the evaluation functions are obtained from the corresponding utility and profit functions. In a second approach, demand and supply are arbitrarily specified and then payoffs are taken to be consumers' and producers' surplus. Assumptions A1-A3, though, can also be fulfilled if payoffs, demand, and supply are not derived from standard problems. Take as an example a producer whose supply is derived from profit maximization, but who evaluates 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 modelled by our approach, since despite the inconsistency such a model fulfills A1-A3.

2.1 Standard Frameworks

2.1.1 Utility and Profit Maximization

Here we elaborate on the standard case of utility maximizing consumers and profit maximizing firms, because it provides insights useful for the further analysis.

Take identical consumers, each endowed with the same fixed income. The preferences of a consumer are represented by a strictly quasiconcave, continuous and strictly monotone utility function $u(x)$, where x denotes the consumption bundle. Furthermore, we assume that none of the goods is a Giffen-good. Utility maximization and price-taking behavior give us the demand functions. We are interested in the market for good 1. If we fix the prices of all goods but 1, we obtain the (reduced) demand function for good 1, $d(p)$, where p denotes the price of good 1. Consumers' evaluation

of the market outcome coincides of course with the utility derived from this outcome.

In our framework we allow for rationing, i.e. a consumer (buyer) may receive less (but not more) than the quantity specified by the demand function. In this case we can think that the consumer maximizes utility taking the rationing as an additional constraint. By the uniqueness of the solution of the utility maximization problem without rationing (which is due to strict quasiconcavity of the utility function), it is clear that for the problem with rationing the rationing constraint is binding, and the solution of this problem leads to an outcome inferior to the non-rationing case. Hence, rationing lowers the consumer's utility (and hence the valuation of the market outcome) as required by A2. Furthermore, again due to strict quasiconcavity the problem with rationing has also a unique solution. Hence rationing at a strictly positive level leads to a better outcome than no-trade, as required by A3. Finally, standard revealed-preference arguments show that if the demand for good 1 is strictly positive, any increase in p decreases utility, and hence A1 is fulfilled.

Identical firms produce good 1 with a strictly convex technology. Cost minimization allows us to derive - for given input prices - a strictly convex cost function. Maximizing profits and price taking behavior lead to the supply function $s(p)$. We assume the absence of fixed costs and that marginal costs are not bounded away from zero. This implies that $s(p)$ is strictly increasing in p as required.⁵ Firms's evaluation of the outcome coincides with the profits.

Of course, if due to rationing firms sell less than the quantity specified by $s(p)$, their profits are lower than in the non-rationing case. This confirms A2. Furthermore, rationing at a strictly positive level is better for the firms than no trade at all, as demanded by A3. Finally, in absence of rationing any increase in the output price leads to an increase in the profits of the producer for any given level of output, and the optimal adjustment of output according to the supply function can only lead to a further increase in the profits. This consideration confirms A1.

2.1.2 Consumer's and Producer's Surplus

Another specific way to derive valuation functions for the current model would be to arbitrarily specify demand and supply functions and let the evaluation of the market outcome be the corresponding consumers' and producers' surplus (denoted by u and π). While such an approach is of course

⁵This feature of the supply function guarantees that equilibrium prices and quantities are strictly positive as long as at least one buyer and one seller coordinate on an institution. Of course, there are weaker conditions also guaranteeing such an interior equilibrium outcome. Working with these weaker conditions, however, would have complicated the presentation without providing new insights.

not exempt of criticism, it is a particularly well-known case which also fulfills our assumptions.

For simplicity, let $p^R > 0$ be such that $d(p^R) = 0$ (the reservation price). Taking possible rationing (i.e. $q_B \leq d(p)$) into account, the buyer surplus is given by

$$u(q_B, p) = \int_{p_B^Q}^{p^R} d(t)dt + [p_B^Q - p] q_B$$

where $p_B^Q = d^{-1}(q_B)$. Analogously, seller's surplus with possible rationing ($q_S \leq s(p)$) is

$$\pi(q_S, p) = \int_0^{p_S^Q} d(t)dt + [p - p_S^Q] q_S$$

where $p_S^Q = s^{-1}(q_S)$. In this case, Assumption A1 holds trivially. If $p < p' \leq p^R$, then d is strictly positive in $]0, p^R[$,

$$u(d(p), p) = \int_p^{p^R} d(t)dt > \int_p^{p^R} d(t)dt = u(d(p'), p')$$

and analogously for the sellers' surplus. Assumption A2 is also easy to check. If $q_B < d(p)$, then changing the integrating variable

$$u(q_B, p) = \int_p^{p^R} d(t)dt - \int_{q_B}^{d(p)} d^{-1}(x)dx < \int_p^{p^R} d(t)dt = u(d(p), p)$$

where the strict inequality holds because $q_B < d(p)$ and d^{-1} is a strictly positive function. Further, if $0 < q_B < d(p)$, then from the same expression, and since $p_B^Q > p$, $u(q_B, p) > 0$ and $u(0, p) = 0$ for all p , verifying Assumption A3.

2.2 Trading Institutions

The good can be traded at different market institutions. For any institution z , denote by n_z, m_z the number of traders active at z . 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 our assumptions, 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. 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

$$r d(p) = s(p)$$

which yields

$$p = p(r).$$

The function $p(r)$ is strictly increasing in r (because $d(p)$ is decreasing and $s(p)$ is increasing in p).

An institution z is characterized by a bias parameter $\beta_z > 0$ which measures the difference between the actual price realized under that market institution, p_z , and the market clearing price. More specifically:

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

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 wants 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 $Q_z(n_z, m_z, \beta_z)$ the overall quantity traded at z . We can now distinguish between three cases:

Case 1: $\beta_z = 1$. In this case the market clearing prices and quantities are realized, and nobody is rationed. The institution is market clearing. The quantities are given by $Q_z(n_z, m_z, 1) = m_z \cdot s(p^*(n_z, m_z)) = n_z \cdot d(p^*(n_z, m_z))$; $q_B^z = d(p^*(n_z, m_z))$; $q_S^z = s(p^*(n_z, m_z))$

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

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

In summary, given an institution characterized by $\beta_z > 0$, and given $r = \frac{n_z}{m_z} > 0$, we can define the seller and buyer quantities as

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

and

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

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.

Remark 1 Notice that, if $m_z = 0$ or $n_z = 0$, there is no trade at that institution. We say then that the institution is inactive. Traders in an inactive institution will always prefer any active institution by Assumption A3, because in such an institution either demand (or supply) is satisfied or traders are rationed, i.e. traders buy or sell strictly less than prescribed by demand or supply.

Given an institution characterized by $\beta_z > 0$, and given $r = \frac{n_z}{m_z} > 0$, the payoffs for buyers and sellers at institution z are given by

$$V_B(\beta_z, r) = v_B(q_B^z(\beta_z, r), \beta_z \cdot p(r))$$

and

$$V_S(\beta_z, r) = v_S(q_S^z(\beta_z, r), \beta_z \cdot p(r)).$$

In general the effects of a bias on the payoffs of the traders are ambiguous. Take as an example an active institution z ($0 < r < \infty$) where prices are higher than the equilibrium price ($\beta_z > 1$). Compared to the market clearing institution with 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, however, the situation is different. For them, prices at z are more favorable than at a market clearing institution. However, this comes at the price of a decrease in the quantity sellers can sell. Therefore the impact of further increase of β_z on sellers' payoffs is unclear. To get an intuition we look again at the standard case with demand and supply derived from utility and profit maximization. Under the usual assumptions about demand and production cost structure, the monopoly price (and hence the optimal price for a cartel formed out of all sellers) is above the equilibrium price, and at the equilibrium price the profits of the monopolist are increasing in the price. Taking into account that all sellers are equal, this translates into the statement that for β_z close to one the direct, profit-increasing impact of higher prices is larger than the indirect, profit-decreasing impact of higher prices via smaller quantities. On the other hand, if the price is above the monopoly level, the negative indirect effect of a price increase dominates the positive direct effect - a further increase in the price lowers the profits because of the implied decrease in the quantity. Similar considerations can be made for the optimal price of a buyers cartel (i.e. for the monopsony price). 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. For prices below the monopsony level, however, the reverse is true: The negative indirect effect dominates. These considerations lead to:

Assumption A4: For any fixed ratio of buyers and sellers r with $0 < r < \infty$, there exist $\underline{\beta}(r) < 1 < \bar{\beta}(r)$ such that, for all $\underline{\beta}(r) < \beta < 1$,

$$V_B(\beta, r) > V_B(1, r)$$

and for all $1 < \beta < \bar{\beta}(r)$,

$$V_S(\beta, r) > V_S(1, r).$$

This condition is immediately fulfilled if the buyer's payoff $V_B(\beta, r)$ is strictly decreasing in β at $\beta = 1$,⁶ and the seller's payoff $V_S(\beta, r)$ is strictly increasing in β at $\beta = 1$.

In what follows we will have to look at the maximum bias such that one type of traders prefers the non-market clearing institution over the market clearing one *for all possible proportions of traders in both institutions*. Hence, we use the following

Definition 1 *For any given m, n , define*

$$\begin{aligned} R(m, n) &= \left\{ \frac{a}{b} \mid a = 1, \dots, n, \text{ and } b = 1, \dots, m \right\} \\ \underline{\beta}(m, n) &= \max_{r \in R(m, n)} \underline{\beta}(r) \\ \bar{\beta}(m, n) &= \min_{r \in R(m, n)} \bar{\beta}(r) \end{aligned}$$

Note that for any given number of buyers and sellers, there exists only a finite number of values r can take. Hence, by A4 $\underline{\beta}(m, n)$ and $\bar{\beta}(m, n)$ exist with $\underline{\beta}(m, n) < 1 < \bar{\beta}(m, n)$.

It will prove convenient to look at those non-market clearing institutions, which improve one market side relative to the market clearing institution for any fixed $r \in R(m, n)$. In other words *for any given ratio of buyers and sellers* such a non-market clearing institution is favored by one market side over the market clearing one. Formally, we say that an institution F is *favored* when $\underline{\beta}(m, n) < \beta_F < \bar{\beta}(m, n)$, $\beta_F \neq 1$.

By A4, potential F -institutions do exist⁷ for any m, n (and the vicinity of the market clearing institution consists of such favored institutions). To exemplify the underlying intuition, think of an institution where the price is above the market clearing level. As long as the bias is low enough so that the resulting price is below the monopoly (cartel) price, sellers profits are larger if for any given r the biased institution is used rather than the market clearing institution.

⁶Note that neither $V_B(\beta, r)$ nor $V_S(\beta, r)$ can be differentiable in general 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 $\beta = 1$.

⁷That is, there exist values of the bias parameter such that, if an institution is characterized precisely by that bias, it will be favored. This does not mean that we assume a favored institution always to be actually available in the market.

2.3 Examples

In this subsection, we present two specific examples derived from the standard formulations discussed above, and use them to exemplify Assumption A4. We will return later to these examples to illustrate our results.

Example 1 *Cobb-Douglas utility function.*

Consider an economy with two goods. We analyze trading institutions for good 1. Firms produce good 1 with the quadratic cost function $c(q_S) = \frac{1}{2}q_S^2$, where q_S denotes the quantity of good 1 produced. Denoting the price for good 1 by p , the profit function is given by $\pi(q_S, p) = pq_S - \frac{1}{2}q_S^2$. Maximizing this profit function w.r.t. q_S yields the linear supply function $s(p) = p$. Firms' valuation of the outcome coincides with the profits.

Buyers are endowed with a Cobb-Douglas utility function $u(q_B, x) = q_B x$, where q_B is the quantity of good 1 consumed and x is the consumption of the other good. Let the price for good 2 be fixed at 1, and let the income be 2 units. Then the standard consumers' problem is given by:

$$\begin{aligned} \max_{q_B, x} \quad & q_B \cdot x \\ \text{s.t.} \quad & pq_B + x = 2 \end{aligned}$$

whose solution yields the demand for good 1:

$$d(p) = \frac{1}{p}$$

Consumers' valuation of an outcome coincides of course with the utility derived from this outcome.

Suppose there are $m_z > 0$ sellers and $n_z > 0$ buyers active at institution z , where the good 1 is traded. The equilibrium price p^* is given by $p^* = p(r) = \sqrt{r}$, with $r = \frac{n_z}{m_z}$.

If the institution is market clearing, the quantity of good 1 sold by an individual seller is denoted by q_S^0 . The quantity of good 1 consumed by an individual buyer is denoted q_B^0 , and the quantity of good 2 denoted x^0 . These quantities are given by $q_S^0 = \sqrt{r}$, $q_B^0 = \sqrt{\frac{1}{r}}$, and $x^0 = 1$. Hence for the market clearing institution the payoffs are:

$$V_B(1, r) = \sqrt{\frac{1}{r}} \quad \text{and} \quad V_S(1, r) = \frac{1}{2}r$$

What happens if the institution is non-market clearing? The price of the non-market clearing institution, p_N is given by $p_N = \beta\sqrt{r}$. Which market side is rationed depends on whether β is larger or smaller than one.

We first look at the case with $\beta > 1$. Sellers are rationed in this case, and their quantities are given by:⁸

$$q_S^N(\beta, r) = rd(p_N) = \frac{r}{p_N} = \frac{\sqrt{r}}{\beta}$$

Hence, profits are given by

$$V_S(\beta, r) = q_S^N(\beta, r) p_N - \frac{1}{2} (q_S^N(\beta, r))^2 = r - \frac{1}{2} \frac{r}{\beta^2}$$

If the profits in the non-market clearing institution are larger, the following must hold:

$$1 < \frac{V_S(\beta, r)}{V_S(1, r)} = \frac{r - \frac{1}{2} \frac{r}{\beta^2}}{\frac{1}{2} r} = 2 - \frac{1}{\beta^2}$$

Since $\beta > 1$, this condition is fulfilled independently of r for all $\beta > 1$. Hence, $\bar{\beta}(r) = \infty$ for all r . Therefore, A4 is fulfilled, and $\bar{\beta}(m, n) = \infty$. This implies that any institution with $\beta > 1$ is a favored institution.

If $\beta < 1$, buyers are rationed. Since the rationing constraint is binding when utility is maximized, the quantities consumed and the utility are given by:

$$q_B^N(\beta, r) = \frac{ms(p_N)}{n} = \frac{p_N}{r} = \beta \frac{1}{\sqrt{r}}$$

$$x^N = 2 - pq_B^N(\beta, r) = 2 - \beta \sqrt{r} \beta \frac{1}{\sqrt{r}} = 2 - \beta^2$$

$$V_B(\beta, r) = \frac{1}{\sqrt{r}} \beta (2 - \beta^2)$$

If the profits in the non-market clearing institution are larger, the following must hold:

$$1 < \frac{V_B(\beta, r)}{V_B(1, r)} = \frac{\beta \sqrt{r} (2 - \beta^2)}{\sqrt{r}} = \beta (2 - \beta^2)$$

Whenever $\beta > 0.62$, this condition is fulfilled. It is again fulfilled independently of r . Hence, we can take $\bar{\beta}(r) = 0.62$ for all r . Therefore, A4 is again fulfilled, and $\bar{\beta}(m, n) = 0.62$.

This implies that any institution with $\beta > 0.62$ is a favored institution, independently of the number of sellers and buyers in the market.

Example 2 *Linear demand.*

⁸Note that in case of rationing all feasible quantities of trade are produced at marginal costs below the price sellers get in a market where they are rationed. Hence, it cannot be profit maximizing to sell less than the maximum quantity possible under rationing.

As in the previous example, sellers are producers with cost function $c(q_S) = \frac{1}{2}q_S^2$. The payoff function is given by the profits $\pi(q_S, p) = pq_S - \frac{1}{2}q_S^2$ and the supply function $s(p) = p$ is derived from profit maximization. Buyers, though, are directly endowed with a linear demand function

$$d(p) = \begin{cases} 1 - p & \text{if } p \leq 1 \\ 0 & \text{if } p > 1 \end{cases}$$

and their payoff is defined as the corresponding consumer's surplus. The market clearing price as a function of the buyers-sellers ratio $r = \frac{n_B}{m_S}$ is

$$p(r) = \frac{r}{1+r}.$$

In a non-market clearing institution with bias parameter $\beta > 1$, sellers are rationed and the quantity per seller becomes

$$q_S(\beta, r) = r \cdot d(\beta p(r)) = \begin{cases} r \left(1 - \beta \frac{r}{1+r}\right) & \text{if } \beta \leq 1 + \frac{1}{r} \\ 0 & \text{if } \beta > 1 + \frac{1}{r} \end{cases}.$$

Hence the sellers payoffs are (for $\beta > 1$)

$$V_S(\beta, r) = \pi(q_S(\beta, r), \beta p(r)) = \begin{cases} \frac{1}{2}r^2 \left(1 - \beta \frac{r}{1+r}\right) \left[\beta \left(\frac{2+r}{1+r}\right) - 1\right] & \text{if } \beta \leq 1 + \frac{1}{r} \\ 0 & \text{if } \beta > 1 + \frac{1}{r} \end{cases}$$

which in particular implies $\bar{\beta}(r) \leq 1 + \frac{1}{r}$. Since $1 + \frac{1}{r}$ converges to 1 as $r \rightarrow \infty$, it is clear that $\bar{\beta}(r) \rightarrow 1$. However, from the inequality $V_S(\beta, r) > V_S(1, r)$ for all $1 < \beta < \bar{\beta}(r)$ we can explicitly compute

$$\bar{\beta}(r) = 1 + \frac{2}{r^2 + 2r} > 1$$

and hence we see that Assumption A4 is satisfied for $\beta > 1$, although the bound $\bar{\beta}(r)$ becomes tighter as the number of buyers increases relative to the number of sellers. We will make use of this fact later.

Consider now a non-market clearing institution biased in favor of the buyers, i.e. $\beta < 1$. Buyers' payoff is computed as the consumer surplus

$$u(q_B, p) = \int_{1-q_B}^1 (1-p') dp' + [(1-q_B) - p] q_B = (1-p)q_B - \frac{1}{2}(q_B)^2$$

and buyers are rationed and receive

$$q_B(\beta, r) = \frac{1}{r} \cdot s(\beta p(r)) = \beta \frac{1}{1+r}$$

which yields (for $\beta < 1$)

$$V_B(\beta, r) = u(q_B(\beta, r), \beta p(r)) = (1-\beta)\beta \frac{1}{1+r} + \frac{1}{2} \left(\beta \frac{1}{1+r}\right)^2.$$

From the inequality $V_B(\beta, r) > V_B(1, r)$ for all $\underline{\beta}(r) < \beta < 1$ we compute

$$\underline{\beta}(r) = \frac{1}{1 + 2r} < 1$$

which shows that Assumption A4 is also fulfilled for $\beta < 1$. Notice, though, that $\underline{\beta}(r) \rightarrow 1$ as $r \rightarrow 0$, i.e. the bound becomes tighter as the number of sellers increases relative to the number of buyers.

In this example we see that, although for any given m and n there exist values of β that would give rise to favored institutions (recall that $R(m, n)$ is always finite), for any particular institution there exists a number of buyers and sellers such that the ratio becomes large enough (for $\beta > 1$) or small enough (for $\beta < 1$) such that the institution is no longer a favored one. Note that, if the number of traders (both sellers and buyers) in the market increases, the range of feasible values of the buyer-sellers ratios in $R(m, n)$ increases and hence both $\underline{\beta}(m, n)$ and $\overline{\beta}(m, n)$ converge to one. Hence, the set of favored institutions shrinks around the market clearing one.

For the example with Cobb-Douglas utility functions, the set of F -institutions is independent of the number of traders. The second example, however, shows that in general the set of F -institutions cannot be defined independently of the number of buyers and sellers operating on the market as a whole. To understand the intuition behind this, think of any example where the individual demand becomes zero above a certain price p_{\max} . Fix the number of sellers, and let the number of buyers increase. The larger the number of buyers, the flatter the market demand curve becomes, and as the number of buyers goes to infinity, the market demand curve becomes perfectly elastic at p_{\max} . Therefore, the monopoly (cartel) price as well as the market clearing price will converge to p_{\max} as n goes to infinity. Hence, while A4 implies that $\overline{\beta}(m, n)$ is always strictly larger than one, it would make no sense to assume that $\overline{\beta}(m, n)$ is bounded away from one as n approaches infinity.

2.4 Comparison of Institutions

In the following we will compare three different types of institutions. First, we will look at the market clearing institution (denoted by 0) with $\beta_0 = 1$. The second type of institutions are non-market clearing ones (denoted by N) which are characterized by $\beta_N \neq 1$. Finally, we will look at favored institutions as defined above, which of course form a subset of the set of non-market clearing institutions.

In this subsection, we compare the payoffs sellers and buyers get at different types of simultaneously active institutions. Note that this comparison, which will be determinant for the learning process, is fundamentally different than that spelled out e.g. in assumption A4, where the comparison is among

payoffs yielded by two different institutions, provided the buyers-sellers ratio was the same in both of them.

Comparing the market clearing with a general non-market clearing institution, we obtain the following general result:

Lemma 1 *Assume A1 and A2. Consider any distribution of traders on any number of institutions, where both the market clearing institution 0 and a given non-market clearing institution N (i.e. $\beta_N \neq 1$) are active. Then the following holds:*

If $v_S(q_S^0, p_0) \leq v_S(q_S^N, p_N)$, then $v_B(q_B^0, p_0) > v_B(q_B^N, p_N)$. Hence, if $v_B(q_B^0, p_0) \leq v_B(q_B^N, p_N)$, then $v_S(q_S^0, p_0) > v_S(q_S^N, p_N)$.

Proof. Suppose $\beta_N < 1$. Then buyers, but not sellers, are rationed at institution N . We have that $q_S^0 = s(p_0)$ and $q_S^N = s(p_N)$. Suppose now that $v_S(q_S^0, p_0) \leq v_S(q_S^N, p_N)$. By A1, we must have that $p_0 \leq p_N$.

Then, again by A1, $v_B(q_B^0, p_0) \geq v_B(d(p_N), p_N)$, and, by A2, $v_B(d(p_N), p_N) > v_B(q_B^N, p_N)$, because $d(p_N) > q_B^N$ (buyers are rationed).

The case $\beta_N > 1$ is analogous. ■

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 be obtaining larger payoffs than those active in the biased one. The next result shows that the analogous statement holds for favored institutions versus the market clearing one.

Lemma 2 *Assume A1 and A4. Fix the number of buyers and seller operating on the whole market. Consider any distribution of these traders on any number of institutions, where both the market clearing institution 0 and a given favored institution F (i.e. $\underline{\beta}(m, n) < \beta_F < \overline{\beta}(m, n)$) are active. Then the following holds:*

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

Proof. Suppose $\beta_F < 1$. Then buyers, but not sellers, are rationed at institution 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, we must have that $p_0 \geq p_F$.

Then, again by A1, $v_B(q_B^0, p_0) \leq v_B(d(p_F), p_F)$. Note that $v_B(d(p_F), p_F) = V_B(1, r_F)$, and $v_B(q_B^F, p_F) = V_B(\beta_F, r_F)$. Hence, by A4 and since $\beta_F > \underline{\beta}(m, n)$, $v_B(d(p_F), p_F) < v_B(q_B^F, p_F)$, and therefore $v_B(q_B^F, p_F) > v_B(q_B^0, p_0)$.

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

3 The Learning Process

If more than one institution is available, traders themselves can choose the institution at which they want to be active. In this section, we explic-

itly model this decision process. 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 institutions available, $z = 0, 1, \dots, Z$. Institution 0 is a market clearing institution ($\beta_0 = 1$), the remaining others are non-market clearing ($\beta_z \neq 1$).

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. 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 parameter β_z determine - as described in section 2.2 - the price and the quantity exchanged at z . This in turn determines the payoffs of the traders having opted for z .

It is easy to see that this choice of the institution constitutes 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, nothing guarantees coordination on the market clearing institution, and even coordination on a Pareto-inferior institution is a strict Nash-equilibrium.

As already explained, we want to analyze under which circumstances traders learn to coordinate on a market clearing institution. Hence, we model the choice of the trading institution as a learning process.

We implicitly assume 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.

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 payoff 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.

We proceed now by modelling the learning process. First we define the state space. For any point in time t , the state of the process is given by

$$\varpi(t) = (\varpi_B(t), \varpi_S(t)) \in \{0, 1, \dots, Z\}^n \times \{0, 1, \dots, Z\}^m$$

i.e. $\varpi(t)(k) \in \{0, 1, \dots, Z\}$ denotes the institution chosen by trader k at time t .

Since interactions are anonymous and traders are symmetric, it is convenient to use reduced-form state spaces. In this case, a typical state would be:

$$\omega = [(n_0, m_0), \dots, (n_Z, m_Z)]$$

where $n_z \in \{0, 1, \dots, n\}$ is the number of buyers and $m_z \in \{0, 1, \dots, m\}$ the number of sellers choosing institution z , and $n_0 + \dots + n_Z = n$, $m_0 + \dots + m_Z = m$. Let Ω denote the set of all such states.

There are three elements in a learning model which require careful consideration: how do agents revise their choices, when are they able to do such a revision, and whether mistakes are possible. We discuss now the first two elements, and postpone the discussion of the third element to section 3.2.

If an agent is able to revise his choice for a given period $t + 1$, he takes his decision given the state $\omega(t)$ and the associated payoffs. This decision determines next the institution chosen for period $t + 1$. As explained above, we postulate that traders who get the opportunity to revise observe prices and traded quantities at all institutions. Then they choose the institution which yields the best outcome as evaluated by their own payoff functions, and go there next period (ties broken randomly). 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.

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 of how revision opportunities arrive give rise to different dynamics and often affect the results. Rather than adopting a specific formulation, here, we postulate a general class of dynamics encompassing the standard examples (and many others), which are then reviewed below. See Alós-Ferrer [1] for a discussion.

Let $E(k, t)$ denote the event that agent k receives revision opportunity in period t , and let $E^*(k, t) \subseteq E(k, t)$ 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 period t .

Assumption D1: $\Pr(E^*(k, t)) > 0$ for every agent k and period t .

Notice that assumption D1 implies that $\Pr(E(k, t)) > 0$, i.e. every agent has strictly positive probability of being able to revise at any given period. Further, since we have two clearly differentiated populations, we introduce a weak form of independence between the revision opportunities in those populations (it can actually be thought of as an anonymity requirement).

Assumption D2: For every agent k and period t , either $\Pr(E^*(k, t) \cap E^*(k', t)) > 0$ for any agent k' of the other type, or $\Pr(E^*(k, t) \cap E(k', t)) = 0$ for any such k .

Assumptions D1 and D2 are rather general. They are fulfilled by the standard models considered in the literature of learning in games. In these models, revision opportunities are either modelled through independent probabilities (a case we call independent inertia; see e.g. [22] or [14]) or assumed to arrive in an asynchronous way (a case we term non-simultaneous learning; see e.g. [6] or [4]).

Independent Inertia. There is an exogenous, independent (across traders and periods) probability $0 < 1 - \rho < 1$ such that the agent does not get revision opportunity in a given period (inertia). Obviously, $\Pr(E^*(i, t)) = \rho(1 - \rho)^{n-1} > 0$ for any buyer i , and analogously for sellers, hence verifying D1. D2 follows from independence: $\Pr(E^*(i, t) \cap E^*(j, t)) = \Pr(E^*(i, t)) \cdot \Pr(E^*(j, t)) > 0$ for any buyer i and any seller j .

Note that this dynamics satisfies even stronger properties. For instance,

$$P((\cap_{i \in I} E(i, t)) \setminus (\cup_{j \in J} E(j, t))) > 0$$

i.e. there is positive probability that, in any single period, all buyers revise but no seller revises (and vice versa). We will return to this property later on.

Non-simultaneous Learning. Each period, only one agent (i.e. either a buyer or a seller) is (randomly) selected and allowed to revise his strategy. Hence, $\Pr(E^*(k, t)) = \frac{1}{n+m}$ for any trader k (verifying D1), and $\Pr(E^*(i, t) \cap E(j, t)) = 0$ for any buyer i and any seller j , and vice versa (verifying D2).

Non-simultaneous Learning within Types. In our case, it is natural to conceive a dynamics where in every period, only one buyer *and* one seller are selected (randomly and independently) and given the opportunity to revise. Assumption D1 holds because $\Pr(E^*(i, t)) = \frac{1}{n} > 0$ for any buyer i and $\Pr(E^*(j, t)) = \frac{1}{m} > 0$ for any seller j . Assumption D2 holds by independence.

Note that the second part of Assumption D2 is the one that specifically allows for dynamics where only one agent at all is allowed to revise each period (non-simultaneous learning). If this part were dropped (which still allows for non-simultaneous learning within types and independent inertia), the modified Assumption D2 would imply D1.

3.1 Unperturbed Learning Process

We refer to the dynamics as described till now as the *unperturbed* learning process, for reasons that will become apparent later. We first look at this process in order to identify properties of the model that will greatly simplify the analysis in the sequel.

The unperturbed learning process is a Markov chain on the (finite) state space Ω , to which standard treatment applies (see e.g. [15].) In order to describe the results of this model, it is useful to summarize the basic results of such dynamics for reference.

Given two states ω, ω' , denote by $P(\omega, \omega')$ the probability of transition from ω to ω' in one period. The transition matrix of the process is given by $P = [P(\omega, \omega')]_{\omega, \omega' \in \Omega}$. An *absorbing set*⁹ 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 $P(\omega, \omega) = 1$. States that are not in any absorbing set are called *transient*.

Every absorbing set of a Markov chain induces an *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, $\mu \cdot P = \mu$). The invariant distribution induced by an absorbing set $A \subseteq \Omega$ has support A . The set of all possible invariant distributions of the process is the convex hull of the invariant distributions associated to the absorbing sets. By the Ergodic Theorem, the invariant distribution associated to a given absorbing set describes the time-average behavior of the system once (and if) it gets into that class. That is, $\mu(\omega)$ is the limit of the average time that the system spends in state ω , along any sample path that eventually gets into the corresponding absorbing set. If, additionally, the absorbing set is *aperiodic*,¹⁰ then the associated invariant distribution describes also the long-run probabilities of the states in the class, $\lim_{T \rightarrow \infty} q \cdot P^T = \mu$ for all probability distributions q whose support is contained in the absorbing set. This result is referred to as the Fundamental Theorem of Markov Chains.

A Markov chain is *ergodic* if it has a unique absorbing set. The (unique) invariant distribution constitutes the long-run prediction for such a process, since it represents the limit behavior of the process independently of initial conditions. If the process is not ergodic, then several invariant distributions exist, describing the long-run behavior along different sample paths, i.e. the prediction depends on the initial conditions.

In our case the unperturbed dynamics is not ergodic, presenting a multiplicity of absorbing sets. Clearly, every “monomorphic” state where all

⁹Also called recurrent communication class or limit set.

¹⁰Loosely speaking, an absorbing set is aperiodic if it contains no deterministic non-trivial cycles. A sufficient condition for aperiodicity is that for some state ω in the set, $P(\omega, \omega) > 0$. Note also that any absorbing state is aperiodic.

traders coordinate in one and the same institution constitutes an absorbing state. In order to show that these are the only relevant absorbing states, we need the following preliminary result. Given $\omega, \omega' \in \Omega$, $\omega \neq \omega'$, a positive probability path from ω to ω' is a finite sequence of states $\omega_0 = \omega, \omega_1, \dots, \omega_{n-1}, \omega_n = \omega'$ such that $P(\omega_{k-1}, \omega_k) > 0$ for all $k = 1, \dots, n$.

Lemma 3 *Assume A1, A2, and A3. Under D1 and D2, given any state $\omega = [(n_0, m_0), \dots, (n_Z, m_Z)]$ with $n_0 \geq 1$ and $m_0 \geq 1$, there exists a positive probability path of the unperturbed dynamics leading from ω to the state $\omega_0 = [(n, m), (0, 0), \dots, (0, 0)]$.*

Proof. Consider any non-market clearing institution, $z \neq 0$, which is chosen by some traders in state ω . If $n_z = 0$ or $m_z = 0$, by Remark 1 we can build a positive probability path to a new state where no trader is at institution z . Hence, without loss of generality, suppose $n_z > 0$ and $m_z > 0$.

By Lemma 1, it follows that in state ω at least one of the two types of traders strictly prefers the market clearing institution in state ω . 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). Repeating this argument, we will reach a state with either $n_z = 0$ or $m_z = 0$. From this state, all remaining traders will leave institution z as above (by Remark 1). 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. ■

Obviously, if there exists a positive probability path from ω to ω' and ω' is an absorbing state, then ω is necessarily transient. Using this fact, we establish the following

Lemma 4 Assume A1, A2, and A3. Under D1 and D2,

- (i) the absorbing states of the unperturbed dynamics are the cross-states $\omega = [(n_0, m_0), \dots, (n_Z, m_Z)]$ such that $n_z = n$ and $m_{z'} = m$ for some $z, z' \in \{0, \dots, Z\}$. In particular, this includes all monomorphic states $\omega_z = [(0, 0), \dots, (n_z = n, m_z = m), \dots, (0, 0)]$ corresponding to coordination on a particular institution;
- (ii) no state $\omega = [(n_0, m_0), \dots, (n_Z, m_Z)]$ with $1 \leq n_0 \leq n - 1$ and $1 \leq m_0 \leq m - 1$ (i.e., where the market clearing institution is active) is part of any absorbing set of the unperturbed dynamics

Proof. Cross states as defined 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 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. Part (ii) follows immediately from Lemma 3. ■

In principle (and particularly for slow dynamics as e.g. those with non-simultaneous learning), there might be non-singleton absorbing sets. By (ii) in the last Lemma, though, those would be made up of states where the market clearing institution is never active.

3.2 Perturbed Learning Process

In order to select among the absorbing states of the unperturbed learning model, and following the literature, we proceed to study stochastic stability. 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 (experiment or mutate), and simply pick an institution at random,¹¹ independently of other considerations.

The dynamics with mistakes (experimentation) is called *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*. Hence, the perturbed process is ergodic. The corresponding (unique) invariant distribution is denoted $\mu(\varepsilon)$. The *limit invariant distribution* (as the rate of experimentation tends to zero) $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu(\varepsilon)$ exists and is an invariant distribution of the unperturbed process P (see e.g. [13], [24], or [5]).

¹¹We mean that an institution is picked up according to a pre-specified probability distribution having full support, for instance uniformly. The exact distribution does not affect the results, as long as it has full support.

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. Clearly, the set of stochastically stable states is a union of some absorbing sets of the original, unperturbed chain ($\varepsilon = 0$).

In the sequel, whenever we say absorbing sets or states, we refer to the unperturbed dynamics. Since the perturbed dynamics is irreducible, no confusion should arise.

We will rely on the characterization of the set of stochastically stable states developed by [13] and [24] and further developed by [5]. Detailed overviews can be found e.g. in [7] or [23]. In particular, 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 will require for the analysis. Its proof is a straightforward application of [5, Theorems 1 and 3] (see the Appendix).

Lemma 5 *Let A be an absorbing set and define the Radius of A by*

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

and the Coradius of A by

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

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.*

Intuitively, the radius of an absorbing set (e.g. the monomorphic state corresponding to coordination on one particular institution) is a measure of how hard it is to destabilize it, while the coradius is a measure of how difficult it might be to reach it from any different state. Part (i) says that, whenever the radius exceeds the coradius, the absorbing set (say, an institution) is easier to reach than to destabilize, and hence it is stochastically stable. If the inequality is strict, part (ii) further establishes that the institution it is necessarily the only stochastically stable one. Finally, part (iii) says that, if an institution is stable but there is another one which can be reached from the former just as easily as it is left, then the later must also be stochastically stable.

4 Stochastically Stable Institutions

We proceed now to analyze the complete model. First, notice that it is straightforward to show that the non-monomorphic cross-states (in which all institutions are inactive, i.e. no trader is actually trading) are rather unstable. Specifically, they are destabilized with a single mutation, in which one trader joins the traders of the other type. To see this, recall simply that trade is preferred to no trade (recall Remark 1), and hence 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.

Since monomorphic states correspond to full coordination on a particular market institution, we aim to identify which of those states are stochastically stable.

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

$$\omega_z = [(0, 0), \dots, (n_z = n, m_z = m), \dots, (0, 0)]$$

is stochastically stable.

Intuitively, a stochastically stable institution is one such that, in the long run, traders frequently coordinate in 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.

Theorem 1 *Under A1, A2, A3, D1, and D2, the market clearing institution is always stochastically stable.*

Proof. We have to show the stochastic stability of the state $\omega_0 = [(n, m), (0, 0), \dots, (0, 0)]$. First, notice that, by Remark 1, no monomorphic state can be left with less than two mutations. In particular, $R(\{\omega_0\}) \geq 2$.

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 $\omega = [(n_0, m_0), \dots, (n_Z, m_Z)]$ with $n_0 \geq 1$ and $m_0 \geq 1$. 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 5(i). ■

This result shows 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. The immediate question is whether other institutions can also be observed. As a first insight, we offer a condition that guarantees a negative answer.

Definition 3 *The market clearing institution is locally robust if at all states of the form $[(n-1, m-1), (0, 0), \dots, (1, 1), \dots, (0, 0)]$, both types of traders strictly prefer the market clearing institution.*

Proposition 1 *Under A1, A2, A3, D1, and D2, if the market clearing institution is locally robust, the state ω_0 is the only stochastically stable state.*

Proof. If the market clearing institution is locally robust, it is not possible to destabilize ω_0 with two mutations, and hence $R(\{\omega_0\}) > 2$, while $CR(\{\omega_0\}) = 2$ as in the proof of Theorem 1. The result follows from Lemma 5(ii). ■

Local robustness, though, might be too much to ask for. The following result shows that any favored institution will necessarily be stochastically stable. As a corollary, in the presence of such an institution local robustness must fail.

Theorem 2 *Assume A1, A2, A3, and A4, and consider any dynamics satisfying D1 and D2. Let $z \in \{1, \dots, Z\}$ be any favored institution. Then z is stochastically stable.*

Proof. 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 Lemma 2, 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 Lemma 2), 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 5(iii). ■

Theorem 2 gives us sufficient conditions for the existence of stochastically stable institutions other than the market clearing one. Recall from the example with Cobb-Douglas utility functions that the set of favored institutions can be very large and independent of the number of traders in the market, and hence the same holds for the set of stochastically stable institutions. This depends, however, crucially on the particular demand, supply and evaluation functions. In the example with linear demand and consumers surplus as evaluation, the set of favored institutions shrinks as the number of traders increases. Hence, in this example Theorem 2 delivers little information when the number of traders becomes large. The following results are significant also for those cases where Theorem 2 has no bite as the number of traders increases.

The following lemma shows that, independently of the actual size of the market, it is always possible to find a non-market clearing institution

(which needs not be favored) such that the market clearing one fails local robustness in the presence of the former. Therefore, the set of stochastically stable institutions may not shrink to the market clearing institution when the market size increases, even if the set of favored institutions degenerates with increasing market size.

Lemma 6 *Assume A1 and A4. Let $0 < \underline{\beta}(1) < \beta_B < 1 < \beta_S < \overline{\beta}(1)$. Then,*

- (i) *if $m \leq n$, in a state where all traders except for one seller and one buyer are at a market clearing institution and the two remaining traders are at an institution with $\beta_z = \beta_B$, buyers strictly prefer the latter;*
- (ii) *if $m \geq n$, in a state where all traders except for one seller and one buyer are at a market clearing institution and the two remaining traders are at an institution with $\beta_z = \beta_S$, sellers strictly prefer the latter.*

Proof. (i) Since $m \leq n$, we have that $m - 1 \leq n - 1$ and hence

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

i.e. 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

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

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

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

which proves the claim, because $v_B(d(p(r)), p(r))$ is the buyers' payoff at the market clearing institution, and $V_B(\beta_B, 1)$ is the payoff of the (only) buyer at the non-market clearing institution with $\beta_z = \beta_B$ (where there is just one buyer and one seller).

(ii) Since $m \geq n$, we have that $m - 1 \geq n - 1$ and hence $1 < \beta_S < \overline{\beta}(1)$

$$r = \frac{n-1}{m-1} \leq 1$$

i.e. there are (weakly) less buyers than sellers at the market clearing institution. By A4, since $1 < \beta_S < \overline{\beta}(1)$,

$$V_S(\beta_S, 1) > V_S(1, 1)$$

and, by A1

$$V_S(1, 1) = v_S(q_S^z(1, 1), p(1)) \geq v_S(s(p(r)), p(r))$$

because $q_S^z(1, 1) = s(p(1))$ and $p(1) \geq p(r)$ since $r \leq 1$ and p is decreasing in r . Hence,

$$V_S(\beta_S, 1) > v_S(s(p(r)), p(r))$$

which proves the claim, because $v_S(s(p(r)), p(r))$ is the sellers' payoff at the market clearing institution, and $V_S(\beta_S, 1)$ is the payoff of the (only) buyer at the non-market clearing institution with $\beta_z = \beta_B$ (where there is just one buyer and one seller). ■

This lemma has a simple interpretation. Suppose 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 an equal number of sellers and buyers, the referred side will necessarily prefer the latter institution.

Note that, since $\underline{\beta}(1) \leq \underline{\beta}(m, n)$ and $\overline{\beta}(1) \geq \overline{\beta}(m, n)$, the set of institutions fulfilling the hypothesis of the last Lemma is in general significantly larger than the set of favored institutions. In particular, the former set remains unchanged as the number of traders in the market increases, while the later can shrink. For instance, in the example with linear demand, the set of favored institutions becomes degenerate as the market size increases, while simple calculation shows that $\underline{\beta}(1) = \frac{1}{3}$ and $\overline{\beta}(1) = \frac{5}{3}$. Hence, the last lemma and the next theorem apply for all institutions with $\frac{1}{3} < \beta < \frac{5}{3}$ independently of market size.

The following theorem proves existence of stochastically stable non-market clearing institutions even for those cases where the set of favored institutions degenerates. However, we can establish this result only for dynamics with independent inertia.

Theorem 3 *Assume A1, A2, A3, and A4. Suppose we have a dynamics with independent inertia. Let $0 < \underline{\beta}(1) < \beta_B < 1 < \beta_S < \overline{\beta}(1)$ and consider three or more institutions, one of them with bias parameters 1 (the market clearing institution), β_B (the buyers' institution) and β_S (the sellers' institution). Then,*

(i) *if $m \leq n$, both the market clearing and the buyers' institution are stochastically stable.*

(ii) *if $m \geq n$, both the market clearing and the sellers' institution are stochastically stable.*

Proof. We will show part (i). Part (ii) is analogous. Let ω_1 denote the monomorphic state corresponding to coordination on the buyers' institution. We know from Theorem 1 ω_0 is stochastically stable. By part (i) of

the previous Lemma, we see that if exactly two mutations to the buyers' institution at state ω_0 occur, we will reach a state where buyers strictly prefer that institution. With independent inertia, there is positive probability that all buyers revise (hence switching to the buyer institution) and at least the mutant seller does not. Regardless of what other sellers do, next period the buyers' institution will be the only active one and all remaining traders will eventually switch to it. Hence, we have constructed a positive probability path from ω_0 to ω_1 involving only two mutations. 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 $c(\{\omega_0\}, \{\omega_1\}) = 2 = R(\{\omega_1\})$, and the result follows from Lemma 5(iii). ■

Theorem 3 presents a trade-off with respect to the result shown in Theorem 2. On the one hand, the former applies to a wider range of institutions than the latter. On the other hand, though, the price to pay is that the result is only established for a subset of the considered dynamics, namely those with independent inertia. Careful inspection of the proof of Theorem 3 shows that it is in general not possible to make a similar statement for more general dynamics. The intuition is that the hypothesis of the Theorem guarantees that as long as two traders (of different types) switch to the considered non-market clearing institution, one market side will benefit from the switch. If the dynamics is quick enough (as in the independent inertia case), all traders of that market side will switch to the non-market clearing institution and the other side will have to follow suit as the market clearing one becomes inactive. If the dynamics is slow, however (as e.g. a dynamics with non-simultaneous learning), the hypothesis of Theorem 3 is not sufficient to derive the result.

5 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. We also identify a sufficient condition for this institution to be the only stochastically stable one. However, this condition is rather strong. Hence, we also find non-market clearing institutions that are stochastically stable under rather general conditions. These conditions are of two types. The first one (favored institutions) depends on actual market size. If the market becomes very large, for certain examples the set of institutions fulfilling this condition may become degenerate. The second one is independent of market size, but guarantees stochastic stability of non-market clearing institutions only for a restricted class of learning dynamics.

In general, we conclude that coordination on market clearing institutions will be often observed as the result of learning, but other institutions might also survive.

References

- [1] Alós-Ferrer, C. (2003), Finite Population Dynamics and Mixed Equilibria, *International Game Theory Review* 5 (3), forthcoming.
- [2] Ausubel, L. and P. Cramton (2002), Demand Reduction and Inefficiency in Multi-Unit Auctions, mimeo.
- [3] Bala, V. and S. Goyal (2000), A Non-Cooperative Model of Network-Formation, *Econometrica* 68, 1181-1229.
- [4] Benaïm, M. and J. Weibull (2003), Deterministic Approximation of Stochastic Evolution in Games, *Econometrica* 71, 878-903.
- [5] Ellison, G. (2000), Basins of Attraction, Long Run Stochastic Stability, and the Speed of Step-by-Step Evolution, *Review of Economic Studies* 67, 17-45.
- [6] Friedman, J. W., and C. Mezzetti (2001), Learning in Games by Random Sampling, *Journal of Economic Theory* 98, 55-84.
- [7] Fudenberg, D., and D. Levine (1998), *The Theory of Learning in Games*. The MIT Press, Cambridge, Massachusetts.
- [8] Gerber, A., and M.O. Bettzüge (2002), Evolutionary Choice of Markets, mimeo.
- [9] Hayek, F. (1967), *Studies in Philosophy, Politics and Economics*. Routledge & Paul Kegan, London.
- [10] Holt, C. (1995), Industrial Organization: A Survey of Laboratory Research. in: Kagel, J. and A. Roth: *Handbook of Experimental Economics*. Princeton University Press, Princeton.
- [11] Ishibuchi, H., C. Oh, and T. Nakashima (2002), Designing Decision Making Systems for a Market Selection Game. mimeo
- [12] Jackson, M. (2002), The Evolution of Social and Economic Networks. *Journal of Economic Theory* 106, 265-295.
- [13] Kandori, M., G. Mailath, and R. Rob (1993), Learning, Mutation, and Long-Run Equilibria in Games. *Econometrica* 61, 29-56.

- [14] Kandori, M., and R. Rob (1995), Evolution of Equilibria in the Long Run: a general theory and applications. *Journal of Economic Theory* 65, 383-414.
- [15] Karlin, S., and H. M. Taylor (1975), *A First Course in Stochastic Processes*, 2nd edition. Academic Press, San Diego.
- [16] Klemperer, P. (1999), *Auction Theory: A Guide to the Literature*. *Journal of Economic Surveys* 13, 227-286.
- [17] Kirchsteiger, G., M. Niederle, and J. Potter (2002), Public versus Private Exchanges, mimeo.
- [18] Kranton, R. and D. Minehart (2001), A Theory of Buyer-Seller Networks. *American Economic Review* 91, 485-508.
- [19] Neeman, Z. and N. Vulkan (2002), Markets versus Negotiations: The Predominance of Centralized Markets. mimeo.
- [20] Ockenfels, A., and A. Roth (2002), Last-Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon on the Internet. *American Economic Review* 92, 1093-1103.
- [21] Plott, C. (1982), *Industrial Organization Theory and Experimental Economics*. *Journal of Economic Literature*, 20, 1485-1587.
- [22] Samuelson, L. (1994), Stochastic Stability in Games with Alternative Best Replies, *Journal of Economic Theory* 64, 35-65.
- [23] Samuelson, L. (1997), *Evolutionary Games and Equilibrium Selection*. The MIT Press, Cambridge, Massachusetts.
- [24] Young, P. (1993), The Evolution of Conventions. *Econometrica* 61, 57-84.

Appendix A: Proof of Lemma 5

Lemma 5 *Let A be an absorbing set and define the Radius of A by*

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

and the Coradius of A by

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

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. Part (ii) follows immediately from [5, Theorem 1]. Part (i) follows from part (iii) taking B to be any absorbing set containing stochastically stable states (which always exist). Part (iii) follows from [5, Theorem 1], but the result there is stronger and the relationship might not be apparent at first glance. That Theorem uses the concept of modified costs $c^*(B, A)$, which are smaller than the transition costs ([5, p. 28-9]). However, if $R(A) = c(B, A) \geq c^*(B, A)$ and states in B are stochastically stable, it cannot happen that $R(A) > c^*(B, A)$ because in such case [5, Theorem 1] also implies that no state in B can be stochastically stable. Hence, $R(A) = c(B, A)$ implies $R(A) = c^*(B, A)$, which is the hypothesis in [5, Theorem 1]. ■

Appendix B: Stochastic Institutions

Experimental evidence reveals that at least for some institutions the price at which trade takes place has a large variance. To capture this feature, the model can be generalized as follows.

For an institution z the price bias is a random variable. Its value is drawn from an arbitrary distribution over the strictly positive real numbers, with expected value $\tilde{\beta}_z$ and variance v_z . Denote a realization of this random variable by β_z . The actual price at which trade takes place at z is given by:

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

Hence, $\tilde{\beta}_z$ specifies the extent to which the price at z differs on average from the market clearing price, and v_z measures the volatility of the bias. Note that $\tilde{\beta}_z = 1, v_z = 0$ imposes market clearing.

Given a realization β_z , if the price is not at the equilibrium level, rationing takes place as in the deterministic case above.

We now proceed to sketch the analysis of this more general model, pointing only at the differences with the deterministic case. First, consider a market clearing institution 0 with $\tilde{\beta}_0 = 1, v_0 = 0$, and a *stochastically non-market clearing* institution N with either $\tilde{\beta}_N \neq 1$, or $v_N > 0$, or both. We obtain the following generalization of Lemma 1.

Lemma 7 *Under Assumptions A1, A2,*

- (i) *For all realizations of β_N it holds that: $v_B(q_B^0, p_0) \leq v_B(q_B^N, p_N) \implies v_S(q_S^0, p_0) \geq v_S(q_S^N, p_N)$.*
- (ii) *The probability of a realization of β_N such that $v_B(q_B^0, p_0) \leq v_B(q_B^N, p_N)$ does not imply $v_S(q_S^0, p_0) > v_S(q_S^N, p_N)$, is zero.*

Proof. In order to prove the lemma, we have to distinguish between three possible realizations of β_N . If either $\beta_N < 1$ or $\beta_N > 1$, the proof is identical with that of Lemma 1. If $\beta_N = 1$, neither sellers nor buyers are rationed at N . Hence, $q_B^N = d(\beta_N p^*(m_N, n_N))$. By A1, $v_B(q_B^0, p_0) \leq v_B(q_B^N, p_N)$ implies that $p_N \leq p_0$. Again by A1, this leads to $v_S(q_B^0, p_0) \geq v_S(q_S^N, p_N)$.

In all three cases, $v_B(q_B^0, p_0) \leq v_B(q_B^N, p_N) \implies v_S(q_S^0, p_0) \geq v_S(q_S^N, p_N)$ as demanded by part (i). Part (ii) is implied by the fact that $\beta_N = 1$ happens with probability zero. ■

Next we define the stochastic analog of favored institutions. We say that an institution F is *stochastically favored* if it is stochastically non-market clearing and the support of its bias parameter β_F is contained in $]\underline{\beta}(m, n), \overline{\beta}(m, n)[$. Then, we obtain a generalization of Lemma 2.

Lemma 8 *Assume A1 and A4. Fix the number of buyers and seller operating on the whole market. Consider any distribution of these traders on any number of institutions, where both the market clearing institution 0 and at a given stochastically favored institution F are active. Then the following holds:*

- (i) *If $v_B(q_B^0, p_0) \geq v_B(q_B^F, p_F)$, then $v_S(q_S^0, p_0) < v_S(q_S^F, p_F)$.*
- (ii) *The probability of a realization of β_N such that $v_B(q_B^0, p_0) \geq v_B(q_B^F, p_F)$ does not imply $v_S(q_S^0, p_0) < v_S(q_S^F, p_F)$, is zero.*

Proof. Again for realizations $\beta_F < 1$ or $\beta_F > 1$ the proof is identical with that of Lemma 2. If $\beta_F = 1$, neither sellers nor buyers are rationed at F . Hence, $q_B^F = d(\beta_F p^*(m_F, n_F))$. By A1, $v_B(q_B^0, p_0) \geq v_B(q_B^F, p_F)$ implies that $p_F \geq p_0$. Again by A1, this leads to $v_S(q_B^0, p_0) \leq v_S(q_S^F, p_F)$. This shows part (i). Part (ii) is again implied by the fact that $\beta_F = 1$ happens with probability zero. ■

The analogues of Lemmata 3 and 4 follow directly from Lemma 7. Then we obtain, with the same proof of Theorem 1:

Theorem 4 *Under A1, A2, A3, D1, and D2, the market clearing institution ($\tilde{\beta}_0 = 1, v_0 = 0$) is always stochastically stable.*

We say that the market clearing institution is *stochastically locally robust* if at all states of the form $[(n-1, m-1), (0, 0), \dots, (1, 1), \dots, (0, 0)]$, both types of traders strictly prefer the market clearing institution for all possible realizations of the bias parameter of the other active institution. Of course, this is an even stronger condition than local robustness in the deterministic case. The proof of Proposition 1 then yields

Proposition 2 *Under A1, A2, A3, D1, and D2, if the market clearing institution is stochastically locally robust, the monomorphic state ω_0 is the only stochastically stable state.*

Lemma 8, however, allows us to establish the analog of Theorem 2.

Theorem 5 *Assume A1, A2, A3, and A4, and consider any dynamics satisfying D1 and D2. Let $z \in \{1, \dots, Z\}$ be any stochastically favored institution. Then z is stochastically stable.*

Of course, in examples as the one with linear demand, as the number of traders increases the support of the bias parameter of stochastically favored institutions must become narrower, degenerating in the limit to $\beta = 1$. Again, though, we can obtain the analog of Lemma 6, leading to a generalization of Theorem 3:

Theorem 6 *Assume A1, A2, A3, and A4. Suppose we have a dynamics with independent inertia. Let B and S be two institutions such that the support of their bias parameters β_B, β_S are contained in $] \underline{\beta}(1), 1[$ and $] 1, \overline{\beta}(1)[$, respectively. Then,*

- (i) *if $m \leq n$, both the market clearing and B are stochastically stable.*
- (ii) *if $m \geq n$, both the market clearing and S are stochastically stable.*

In the deterministic case, all favored institutions fulfill the hypothesis of Theorem 3. In the stochastic case, though, the analogous statement is not true. The distribution of the bias parameter of a stochastically favored institution may assign positive probability both to a set of bias parameters larger than 1 and to another set of bias parameters smaller than 1, so that its bias works in an ambiguous way, sometimes favoring sellers and sometimes buyers. The institutions in the statement of the last theorem, though, must always be biased in favor of the same side of the market. As in the deterministic case, however, the set of such institutions remains unchanged as the market size increases.