

On the Evolution of Market Institutions: The Platform Design Paradox*

CARLOS ALÓS-FERRER[†] GEORG KIRCHSTEIGER[‡] MARKUS WALZL[§]

January 2006

ABSTRACT. This paper analyzes a situation where market designers create new trading platforms and traders learn to select among them. We ask whether “Walrasian” platforms, leading to market - clearing trading outcomes, will dominate the market in the long run. If several market designers are competing, we find that traders will learn to select non-market clearing platforms with prices systematically above the market-clearing level, provided at least one such platform is introduced by a market designer. This in turn leads all market designers to introduce such non-market clearing platforms. Hence platform competition induces non-competitive market outcomes.

Keywords: Market Institutions, Evolution of Trading Platforms, Learning, Asymmetric Rationality.

JEL Classification: C72, D4, D83, L1.

1. INTRODUCTION

Markets are not only characterized by demand and supply, but also by the rules that govern the trading process. The “institutional” framework determines the set of market participants, their available options, and the matching and information structure of the market. In reality we observe a huge variety of different market frameworks even for trading the very same good. Real estate, for example, is traded at auctions as well as by personal bargaining. There is also a large amount of evidence (mainly experimental) that these characteristics are crucial for the resulting trading outcome and for the realized prices (for an overview see e.g. Plott [28] or Holt [20]). While double auctions typically tend to generate market clearing prices and quantities, posted offer markets establish prices that tend to be above the market clearing level, whereas the prices on posted bid markets seem to be below the Walrasian level (see e.g. Plott and Smith[29]). As a consequence, some gains of trade are not realized on these trading platforms, and inefficiencies occur due to the design of the trading platform.

*We are grateful to seminar participants in Berlin, UL Brussels, University College London, Louvain-la-Neuve and Maastricht for helpful comments. The first author acknowledges funding by the Austrian Science Fund (FWF) under Project P18141-G09, the second author by the Banque Nationale de Belgique, and the third author by the Dutch Science Foundation (NWO).

[†]Department of Economics, University of Konstanz (Germany). Email: Carlos.Alos-Ferrer@uni-konstanz.de

[‡]ECARES, Universite Libre de Bruxelles (Belgium). Email: gkirchst@ulb.ac.be

[§]Department of Economics and METEOR, Maastricht University (The Netherlands). Email: m.walzl@algec.unimaas.nl

Since the institutional framework of a market matters, it is important to analyze how market institutions evolve and what properties the emerging trading platforms have. In particular, one might investigate the circumstances under which market institutions emerge that enforce market clearing outcomes. In order to answer this question, it is useful to distinguish between two aspects of the evolution of market institutions, namely the emergence of new institutions and the selection between existing institutions. New market institutions can either be introduced on purpose by a market designer, or be the (unintended) by-product of the actions of the traders. In what follows we will focus on market platforms introduced on purpose.¹ If a trading platform is introduced by a (profit maximizing) market designer who demands user fees, the emergence of new platforms and the selection among existing ones are closely interlinked. The market designer will try to introduce a new platform with characteristics that attract many traders. This attractiveness in turn determines the long run survival of the platform. In this paper we analyze this interplay between the creation of new and the selection among existing trading platforms, and we investigate the characteristics of the resulting platforms with respect to their ability to achieve market clearing outcomes.

A good example for what we have in mind are Business to Business (B2B) trading platforms (for an analysis of B2B e-commerce see Lucking-Reiley and Spulber [26]). The last decade has seen a proliferation of B2B platforms, and despite the burst of the internet bubble there were more than 1000 B2B marketplaces active in Europe in 2003 (see European Commission [15]). While most of the attention is devoted to e-marketplaces targeting consumers (like e-bay or Yahoo), about 95% of the e-commerce is actually B2B (see United Nations [35]). In 2004 B2B had an estimated volume of \$1 trillion (see The Economist [9]). Contrary to Business to Consumers or Consumer to Consumer platforms, large quantities of relatively standardized products are traded at B2B exchanges. On these platforms agents seem to act either as buyers or as sellers, but not as both (see European Commission [15]). A large part of B2B e-commerce is done via platforms operated by buyers or sellers themselves. In these cases the trading platforms can be viewed as a cost-reducing way of organizing the bargaining and purchasing process. The situation is different, if the trading platform is operated by a third, independent party. Such platforms are run by market designers who receive trading fees from users. The size of this part of the B2B e-commerce is large, too. Of all firms active on B2B platforms, about one third operates on such neutral platforms (see European Commission [15]).

In this paper, we develop a framework to analyze the interplay between the design of such trading platforms and their selection by traders. The platforms are created by profit maximizing and risk-neutral market designers. The designers compete with each other through platform designs. In particular, we allow that each designer demands a trading fee, but also that the platform is designed with a systematic price-bias, above or below the market clearing price. Hence, the designer decides

¹For an analysis of markets as a by-product of traders' actions, see Kirchsteiger et al.[24].

upon the share of the surplus created through trade at the platform but can also favor one type of trader with the introduction of a price bias. This model of market designs captures crucial features of trading platforms, and highlights the following trade-off for the designer. Any bias reduces the surplus generated at the platform (and thereby *ceteris paribus* the revenue for the designer), but may also make it more attractive for the favored type of trader, which in turn may enhance the platform's survival probability.

To analyze this trade-off, we model competition between two market designers and compare the results of this setting with the benchmark case of a monopolistic market designer. After the platforms have been designed, each trader decides on which platform to be active (for the monopolistic case, there is of course no real choice - traders just trade at the only existing platform). The role of the trader (buyer or seller) is exogenously given. Sellers are assumed to be firms with a constant returns to scale production technology². Buyers are characterized by their demand functions, and might be either consumers or other firms. For given platform characteristics, the selection by traders gives rise to a coordination game. If each trader opts for a particular platform and gains from trade are strictly positive, no trader has an incentive to deviate from this platform - independently of the design alternatives offered by competing platforms. However, traders might *learn* to coordinate on a *particular* platform. Following the game-theoretic learning literature (see Young [36], Kandori et al. [22], or Ellison [11]), we use a Markovian model to analyze the platform choice of the traders. We assume that the traders' behavior depends on the market outcomes generated by the different platforms and thereby on the characteristics of all feasible platforms. We are interested in the long term properties of this learning process, i.e. in its limit invariant distribution. This limit invariant distribution in turn determines the payoffs of the market designers. Hence, we establish a link between designer-revenues and the characteristics of *all* feasible platforms.

For the case of competing platforms we find that - in the long run - traders will always coordinate on a platform with prices above the market clearing level, provided that such a platform has been introduced by at least one designer. This forces both designers to introduce platforms that are not market clearing, but that have a price bias in favor of the sellers. On the other hand we find that a monopolistic designer will always introduce a market clearing platform. Therefore competition at the designers' level turns out to be detrimental for a competitive outcome at the traders' level. We regard this result as paradoxical.

The present paper is related to three strands of the literature. First, since we investigate the role of trading platforms with exogenously given buyers and sellers, our paper is to some extent related to the two-sided markets literature (see Rochet and Tirole [31] for an overview). This literature is based on the assumption of network-externalities. It analyzes the impact of these externalities and of platform competition

²In Appendix B we investigate the robustness of our results with respect to decreasing returns to scale production.

on the structure of the fees demanded by the market designers (see e.g. Armstrong [3], Belleflamme and Toulemonde [4], Caillaud and Julien [8], or Rochet and Tirole [30]). In contrast, we want to investigate whether traders learn to coordinate on market clearing trading platforms, if such platforms are feasible. Therefore we explicitly model the learning behavior of the traders, whereas the two-sided market literature assumes rational traders. Further, we ask whether platform competition induces market designers to establish platforms with characteristics that achieve market clearing outcomes. Consequently, we abstract from any network externalities that are not internalized by the price at which trade takes place. In our model trading fees demanded by the market designers are neutral insofar as the market outcome is only influenced by the total fee imposed on both market sides, but not on the distribution of the fees on the two market sides.³

Second, our paper is also related to the literature on competition between exogenously given trading institutions. Ellison and Fudenberg [13] and Ellison et al. [14] analyze under which circumstances different market institutions can coexist in equilibrium. Due to their different research questions these papers do not allow for institutions with systematic price biases. Kugler et al. [25] and Neeman and Vulkan [27] investigate the case of centralized versus decentralized trading institutions. All of these papers they rely on the assumption of rational traders, and do not allow for learning. In terms of traders' behavior, the learning model of Gerber and Bettzüge [19] is relatively close to our paper. But since they focus in the possibility of multiplicity of active trading platforms, they consider neither non-market-clearing platforms nor market designers. Closest related to the paper at hand is the work by Alós-Ferrer and Kirchsteiger [2], which also analyzes the learning behavior of traders who face the choice between different, not necessarily market clearing platforms. That paper, however, deals only with the selection among different, exogenously given institutions and does not consider competition between market designers.

In our model rational market designers are confronted with learning traders.⁴ Hence, our paper belongs to a small but growing literature that we would like to call "asymmetric rationality," where fully rational firms or otherwise sophisticated agents are confronted with a population of boundedly rational ones. The basic motivation is that consumers and small traders do not have the resources to obtain all the relevant information and fully optimize their behavior, often relying on behavioral rules of thumb instead. However, large firms, market designers, etc. can be taken as comparatively sophisticated. Gabaix and Laibson [18], Hopkins [21], and Spiegler [34] apply this approach to the analysis of industries facing boundedly rational consumers. See Ellison [12] for an overview of this literature.

The paper is organized as follows. Section 2 presents the basic model. Section 3 discusses the traders' platform choice of the traders. Section 4 analyzes the design

³Rochet and Tirole [31] define two-sided markets by the non-neutrality of the fees. In their terminology we model a one-sided market.

⁴In Appendix B we show the robustness of our results with respect to learning designers.

of the platform. Section 5 concludes. Some proofs are relegated to Appendix A. In Appendix B we analyze the robustness of our results with respect to boundedly rational designers and with respect to decreasing returns to scale production.

2. THE MODEL

2.1. Market platforms' design. A homogenous good is traded at alternative market platforms, which are set up by market designers. For simplicity, we restrict our attention to two competing market designers (referred to as competitive market design). As a benchmark, we will also analyze the case where only one market designer can set up a trading platform (referred to as monopolistic market design).

Before any trade takes place, the market designers decide upon the set of trading rules under which their respective platforms operate, and the trading fees they demand from the traders. We do not aim at a complete description of the different sets of rules the designers can introduce. Rather, we characterize them by their ability to establish market clearing. Market designers may choose to design platforms such that market clearing is guaranteed, or they may pick platforms where the price is systematically biased above or below the market clearing price. Denote by p_i^* the market clearing price if at least one seller and at least one buyer choose this platform and by $\beta_i > 0$ the bias of platform $i = 1, 2$. The actual price at which trade takes place at platform i is then given by $p_i = \beta_i p_i^*$. If platform i is not market clearing (i.e. $\beta_i \neq 1$) the quantity traded is determined by the short market side, and traders on the long market side are rationed. Sellers are rationed equally if $\beta_i > 1$. We do not specify any rationing rule for the buyers.

The common set of feasible biases is assumed to be a finite, regular grid $B = \{\beta_{\min}, \beta_{\min} + \delta, \dots, 1, \dots, \beta_{\max} - \delta, \beta_{\max}\}$, where $0 < \beta_{\min} < 1 < \beta_{\max}$ and δ is the step of the grid. We assume explicitly that $1 \in B$, and refer to the feasible platform with $\beta = 1$ as the market-clearing platform. $|B|$ denotes the number of feasible biases.

After the platforms are set up, traders will use their experience and observations to eventually learn which platform to use. Formally, we will analyze a learning process with an infinite number of trading rounds. The designers' long-run payoffs are the expected per round charges. Furthermore, we assume that the charges of designer i depend on the revenue generated by trade on i 's platform.⁵ Denote by f_i the trading fee demanded by designer i , and by ER_i the expected per round revenue generated on platform i . Then market designer i 's profits are given by $\pi_{D,i} = f_i ER_i$. The set of feasible fees is the same for both designers. Again for simplicity we assume that it is given by a finite, regular grid $F = \{f_{\min}, f_{\min} + \gamma, \dots, f_{\max} - \gamma, f_{\max}\}$, where $0 < f_{\min} < f_{\max} < 1$.⁶ $|F|$ denotes the number of feasible fees.

The trading fee can be imposed on the sellers' side, on the buyers' side, or divided between both sides. However, the market clearing price, the realized price at which

⁵Our results would not change if we assume quantity-dependent charges instead of revenue-dependent charges. The same holds for fixed participation fees.

⁶The assumption that the fees are strictly positive can be justified by (unmodelled) setup costs for the market designers.

trade is conducted, and the traded quantities depend only on the total fee, and not on the distribution of the fee over the two market sides. Buyers at platform i pay p_i for each unit, market designers receive $f_i p_i$, and sellers ultimately receive $(1 - f_i) p_i$. Hence, we do not need to specify on which market side the fee is imposed.

The characteristics of a platform i are denoted by $s_i = (\beta_i, f_i)$, and the set of feasible characteristics by $S = B \times F$.

2.2. Traders. The good is supplied by a finite set M of profit maximizing firms that use a constant returns to scale technology. We assume that there are at least two sellers at the market, i.e. $|M| > 1$. Each seller uses the same constant returns to scale technology with marginal costs of $c > 0$.⁷ When deciding about the supply, a seller takes into account the trading fee of the platform at which he is operating. Hence, sellers will supply a strictly positive but finite quantity if and only if the price net of trading fee is equal to c .

As we will see, the assumption of a constant returns to scale technology allows us to derive results for a very general class of learning models. That is, by focusing on the constant returns to scale case, we will obtain results that are robust to the details of the learning process. In the appendix we illustrate that for strictly decreasing returns to scale the results depend on the details of the learning model. In particular, the results of the constant returns to scale case can be replicated also for strictly decreasing returns to scale, but not for the whole class of learning models we analyze here.

The good is demanded by a finite set N of buyers with $|N| > 1$. Each buyer $n \in N$ is endowed with a demand function $d_n(p)$ which might be different for different buyers. All the demand functions are assumed to be strictly decreasing in p . Furthermore, $0 < d_n(p) < \infty$ for all p, n . To avoid discontinuities in the designers' profit functions we also assume that $\lim_{p \rightarrow \infty} p d_n(p) = 0$ for all $n \in N$.⁸

We call a platform *active* if both sellers and buyers are present and positive quantities are traded, and *inactive* if not. The presence of both types of traders does not ensure that the platform is active. Due to the assumption of a constant returns to scale technology the market clearing price of a platform i where both sellers and buyers are present (and, in particular, of an active platform) is given by $p_i^*(s_i) = \frac{c}{1-f_i}$. The realized price at which trade is conducted on platform i is then

$$p_i(s_i) = \beta_i \frac{c}{1-f_i}.$$

Note that if $\beta_i < 1$, the net price received by the sellers is below the marginal costs. Hence, supply is zero, and platform i is inactive.

⁷The assumption of identical sellers might seem restrictive at the first sight. Within our framework firms without access to the lowest cost technology cannot sell anything on the market, anyhow. Hence, our assumptions only rule out the case where exactly one firm has access to the production technology with the lowest costs.

⁸Our results do not depend on the assumption that the value of demand goes to zero when the price approaches infinity. However, the presentation is simplified by this assumption.

Denote by N_i the set of buyers who choose platform i , and by M_i the set of sellers who choose platform i . Then, platform i is active if and only if $|N_i| > 0$, $|M_i| > 0$, and $\beta_i \geq 1$. Let

$$D_{N_i}(p) = \sum_{n \in N_i} d_n(p)$$

denote the total demand at platform i . The quantities traded by a buyer $n \in N_i$, $q_{n,i}(N_i, M_i, s_i)$, and by a seller $m \in M_i$, $q_{m,i}(N_i, M_i, s_i)$, are given by

$$q_{n,i}(N_i, M_i, s_i) = \begin{cases} d_n(\beta_i \frac{c}{1-f_i}) & \text{if } i \text{ is active,} \\ 0 & \text{otherwise,} \end{cases}$$

$$q_{m,i}(N_i, M_i, s_i) = \begin{cases} \frac{1}{|M_i|} D_{N_i} \left(\beta_i \frac{c}{1-f_i} \right) & \text{if } i \text{ is active,} \\ 0 & \text{otherwise.} \end{cases}$$

In the single-designer case, traders cannot choose between different platforms, but have to use platform i . Hence, $N_i = N$, $M_i = M$, and the market outcome is only determined by the platform characteristics s_i .

If there is competition between market designers, trade can take place at different platforms, and the outcome depends also on the way traders learn which platform to use. This learning process is driven by the market outcomes of both platforms (see above), and by the individual evaluations of these outcomes. For the latter part note that if *buyers* trade strictly positive amounts, they are strictly better off than without trade. Hence, inactive platforms are worse for buyers than active ones. Furthermore, whenever a buyer trades a strictly positive quantity, he is not rationed at all. Hence, it is natural to assume that buyers' evaluation of active platforms is monotonically decreasing in the price. Therefore, buyers' evaluation of platform i could be represented e.g. by⁹

$$\pi_{n,i}(s_i) = \begin{cases} \frac{1}{p_i} = \frac{1-f_i}{\beta_i c} & \text{if } i \text{ is active,} \\ 0 & \text{otherwise} \end{cases}$$

Hence, $\pi_{n,i}(s_i) > 0$ for any active platform i . If both platforms are active (i.e. positive amounts are traded),

$$p(s_i) < p(s_j) \iff \pi_{n,i}(s_i) > \pi_{n,j}(s_j).$$

Clearly, platform i is active if and only if both types of traders are present and $\beta_i \geq 1$. Both platforms are active whenever $N_i, N_j \neq \emptyset$, $M_i, M_j \neq \emptyset$ and $\beta_i, \beta_j \geq 1$. In particular, if $\beta_i = \beta_j = 1$ and $f_i < f_j$, then $\pi_{n,i}(s_i) > \pi_{n,j}(s_j)$.

⁹We do not use this particular representation. If demand is derived from utility maximization, though, the realized (indirect) utility must be (a strictly monotone transformation of) this payoff function.

The sellers' evaluation of the platforms are determined by the respective profits. An inactive platform gives of course zero profits. Furthermore, whenever $\beta_i > 1$, sellers trading on platform i are on the long market side, and hence equally rationed. Hence, the sellers' evaluation of platform i is given by

$$\pi_{m,i}(N_i, M_i, s_i) = \begin{cases} \left[\frac{1}{|M_i|} D_{N_i} \left(\beta_i \frac{c}{1-f_i} \right) \right] (\beta_i - 1)c & \text{if } i \text{ is active,} \\ 0 & \text{otherwise.} \end{cases}$$

Note that as long as both platforms are active, the outcome of a non-market clearing platform is always better for the sellers than the outcome of a market clearing platform. That is, for all f_i, f_j ,

$$N_i, N_j \neq N, M_i, M_j \neq M \text{ and } \beta_i > \beta_j = 1 \Rightarrow \pi_{m,i}(N_i, M_i, s_i) > \pi_{m,j}(N_i, M_i, s_i) .$$

3. THE TRADERS' PLATFORM CHOICE

First, market designers choose the characteristics of their platforms. If there is only one market designer, traders' choice is trivial - they simply opt for the existing platform. With more than one market designer, traders have to choose between the two platforms. For any given s_i, s_j the choice of platform constitutes a coordination game. If all traders choose platform i , no trader has an incentive to deviate to the other platform j . Furthermore, if β_i and β_j are strictly larger than 1, full coordination on any platform is even a strict Nash equilibrium. Hence, nothing guarantees coordination on any particular platform, and therefore traders have to learn which platform to use.

3.1. The Learning Process. In order to model the learning process, we (implicitly) assume that a trader does not only take into account his own experiences (as he would in e.g. a reinforcement learning model). He observes the prices and the quantities of both platforms (including the observation of the inactiveness of a platform), and takes this information into account. We also assume that an individual trader does not have enough information on other traders or is not able to perform all the necessary computations in order to predict the future behavior of the other traders. Hence, an individual trader cannot accurately predict the future outcomes of the platforms. Furthermore, traders also lack the capability necessary to always compute an exact (but myopic) 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 evaluate the outcomes of both platforms, and switch to the other platform if he perceives the other platform's outcome as better. A trader can perceive this behavior as approximately rational, since when he switches, 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, this 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 any evaluation conducted by other traders, but merely the observation of prices and traded quantities of both platforms.

We proceed now to model the learning process. The state space is given by $\Omega = \{1, 2\}^{|N|} \times \{1, 2\}^{|M|}$. A typical state ω specifies which trading platform is chosen by each buyer and each seller - trader k 's platform choice is denoted by $\omega(k) \in \{1, 2\}$. The following notation will prove convenient:

$$\begin{aligned} N_i(\omega) &= \{n \in N \mid \omega(n) = i\} \\ M_i(\omega) &= \{m \in M \mid \omega(m) = i\} \end{aligned}$$

i.e. $N_i(\omega) \subseteq N$ is the set of buyers who are on platform i in state ω , and $M_i(\omega) \subseteq M$ the set of sellers who are on platform i in state ω . Of course, for $j \neq i$, $N_i(\omega) = N \setminus N_j(\omega)$ and $M_i(\omega) = M \setminus M_j(\omega)$.

The state of the process at time $t = 0, 1, 2, \dots$ is given by $\omega(t) \in \Omega$. That is, $\omega(t)(k) \in \{1, 2\}$ denotes the platform chosen by trader k at time t .

Unperturbed Learning Process. We first concentrate on the unperturbed learning process, where traders switch platform only because of learning, but not because of experimentation (experimentation is introduced in section 3.1 below). If an agent is able to revise his choice for a given period $t + 1$, he takes the new market outcomes of both platforms in period t and evaluates them. As explained above, we postulate the following learning rule:

Assumption A: A trader, who gets the opportunity to revise, observes the outcomes of both platforms in the last period. He then chooses the platform whose outcome he evaluates as best. In case of indifference, he randomizes the choice, with both platforms chosen with strictly positive probability. Choice probabilities may depend on the outcomes but not on the platforms' names.¹⁰

That is, provided that trader k receives revision opportunity at period t , he will choose the platform with a period $t - 1$ outcome that he evaluates highest. If, by chance, the outcomes of both platforms are equally evaluated, players are indifferent and will select one platform at random. For instance, in the case in which one platform is inactive and the other is active but yields exactly zero profits for the sellers, sellers randomize among them.

But when are agents allowed to 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 (or, for instance, accounting for idiosyncratic switching costs). Different specifications of how revision opportunities arrive give rise to different dynamics and often affect the results. Rather than adopting a

¹⁰In other words, we assume platform symmetry, i.e. the trader's choice does not depend on the identity of the platform but on its observable features - prices and quantities. In case of indifference, both choice probabilities have to be strictly positive but might still depend on other observable characteristics as e.g. the number of traders at each platform.

specific formulation, here we follow Alós-Ferrer and Kirchsteiger [2] and 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 $E(k, \omega)$ 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 .

Assumption B1: $\Pr(E^*(k, \omega)) > 0$ for every agent k and state ω .

Notice that assumption B1 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 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 B2: For every agent k and state ω , either $\Pr(E^*(k, \omega) \cap E^*(k', \omega)) > 0$ for any agent k' of the other type, or $\Pr(E^*(k, \omega) \cap E(k', \omega)) = 0$ for any such k' .

This assumption explicitly excludes non-anonymous situations where, say, whenever seller number 17 gets the opportunity to revise, buyers 3 and 6 also get the opportunity to revise. Assumptions B1 and B2 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. Samuelson [32] or Kandori and Rob [23]) or assumed to arrive in an asynchronous way (a case we term non-simultaneous learning; see e.g. Blume [7], Binmore and Samuelson [6] or Benaïm and Weibull [5]).¹¹ That is, our formulation covers the following standard examples.

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 state (inertia). Obviously, $\Pr(E^*(n, \omega)) = \rho(1 - \rho)^{n-1} > 0$ for any buyer n , and analogously for sellers, hence verifying B1. B2 follows from independence: $\Pr(E^*(n, \omega) \cap E^*(m, \omega)) = \Pr(E^*(n, \omega)) \cdot \Pr(E^*(m, \omega)) > 0$ for any buyer n and any seller m .

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, $\Pr(E^*(k, \omega)) = \frac{1}{|N|+|M|}$ for any trader k (verifying B1), and $\Pr(E^*(n, \omega) \cap E(m, \omega)) = 0$ for any pair of buyers and sellers (verifying B2).

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

¹¹The reason we explicitly choose Assumptions B1, B2 is that, in the literature of learning in games, many models are not robust to minute changes in the dynamic assumptions. We want to make explicit that our model is not so sensible to the details of the dynamics.

(randomly and independently) and given the opportunity to revise. Assumption B1 holds because $\Pr(E^*(n, \omega)) = \frac{1}{|N|} > 0$ for any buyer n and $\Pr(E^*(m, \omega)) = \frac{1}{|M|} > 0$ for any seller m . Assumption B2 holds by independence.

Note that the second part of Assumption B2 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 B2 would imply B1.

The specification above allows for more general learning processes than those described by independent inertia or non-simultaneous learning. Since the revision probability $\Pr(E(k, \omega))$ is a function of the state ω , it might depend e.g. on the difference between the evaluation of the outcomes of both platforms (so that unsatisfied traders are more likely to revise), or on idiosyncratic characteristics of the currently chosen platform.

Assumptions A, B1, and B2 define a stationary Markov chain on the (finite) state space Ω . Given two states $\omega, \omega' \in \Omega$, denote by $P^0(\omega, \omega')$ the probability of transition from ω to ω' in one period for the unperturbed learning process, i.e. for the process without experimentation. The transition matrix of the process is given by $P^0 = [P^0(\omega, \omega')]_{\omega, \omega' \in \Omega}$.

To fully characterize the unperturbed learning dynamics, it is useful to summarize the basic results of Markov chains. An *absorbing set*¹² 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 if $P^0(\omega, \omega) = 1$. States that are not in any absorbing set are called *transient*.

Every absorbing set of a Markov chain P 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*,¹³ 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.

¹²Also called recurrent communication class or limit set.

¹³Loosely 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.

A Markov chain is *ergodic* if it has a unique absorbing set, and *irreducible* if that absorbing set coincides with the full state space Ω . For an ergodic chain, the (unique) invariant distribution constitutes the long-run prediction, 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.

Depending on the characteristics of the platforms, the unperturbed dynamics can be ergodic, or irreducible, or might exhibit a multiplicity of absorbing sets.

Lemma 1. *Assume A, B1, and B2.*

- (a) *If $\beta_i > 1$ and $\beta_j > 1$, the only absorbing sets of the unperturbed dynamics are the two singletons made of monomorphic states ω_i^* ($i = 1, 2$) such that $N_i(\omega_i^*) = N$, $M_i(\omega_i^*) = M$.*
- (b) *If $\beta_i > 1$ and $\beta_j \leq 1$, the unperturbed dynamics is ergodic and the only absorbing set is the singleton made of the monomorphic state ω_i^* .*
- (c) *If $\beta_i \leq 1$ and $\beta_j \leq 1$, the unperturbed dynamics is irreducible, i.e. the whole state space Ω forms an absorbing set.*

Proof. See Appendix. ■

The previous Lemma shows that depending on the properties of the platforms there is more than one absorbing set. In order to select among them, we study in the next section the stability properties of the platforms with respect to experimentation.

Perturbed Learning Process. Following the literature on this type of learning models, we proceed to study stochastic stability with respect to experimentation. To do so, the dynamics is enriched with a perturbation in the form of experiments (or mistakes) as follows. With an independent, small probability $\varepsilon > 0$, each agent, in each round, might experiment (or make a mistake or “mutate”), and simply pick a platform at random,¹⁴ independently of other considerations.

The dynamics with experimentation is called *perturbed* learning process. Its transition matrix is denoted by P^ε . Since experiments make transitions between any two states possible, the perturbed process has a single absorbing set formed by the whole state space (i.e. the process is 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^0 (see e.g. Kandori et al. [22], Young [36], or Ellison [11]).

¹⁴We 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, and does not depend on ε .

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 the 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 Kandori et al. [22] and Young [36]. Detailed overviews can be found e.g. in Ellison [11], Fudenberg and Levine [16] or Samuelson [33].

Given two absorbing sets X and Y , let $c(X, Y) > 0$ (referred to as the *transition cost* from X to Y) denote the minimal number of experiments necessary for a direct transition from X to Y , i.e. a positive probability path starting in an element of X and leading to an element in Y , which does not go through any other absorbing set.

In our case, the stochastic stability analysis is simple because of Lemma 1.

Theorem 2. *Assume A, B1, and B2.*

- (a) *If $\beta_i > 1$ and $\beta_2 > 1$, only the two monomorphic states can be stochastically stable. Further, ω_i^* is stochastically stable if and only if $c(\omega_i^*, \omega_j^*) \geq c(\omega_j^*, \omega_i^*)$.*
- (b) *If $\beta_i > 1$ and $\beta_j \leq 1$, the only stochastically stable state is the monomorphic state ω_i^* .*
- (c) *If $\beta_1 \leq 1$ and $\beta_2 \leq 1$, all states in Ω are stochastically stable.*

Proof. (b) and (c) are trivial, since the set of stochastically stable states is always non-empty, only states in absorbing sets can be stochastically stable, and all states in the same absorbing set are simultaneously either stochastically stable or not. Part (a) follows from the characterization quoted above for the simple case with only two absorbing states. ■

3.2. The long run trading patterns. We now proceed to analyze the stochastic stability of platforms depending on the values of the parameters $s_i = (\beta_i, f_i)$ and $s_j = (\beta_j, f_j)$.

As a benchmark, we start with the case of identical platform design. To analyze platforms with identical characteristics ($s_i = s_j$), we observe that, for every state $\omega \in \Omega$ we can uniquely define a so-called *mirror state* $\tilde{\omega}$ by changing the platform affiliation of all traders, that is, $\tilde{\omega}$ is the only state such that $M_j(\tilde{\omega}) = M_i(\omega)$ and $N_j(\tilde{\omega}) = N_i(\omega)$. Then,

Lemma 3. *Suppose $s_i = s_j$. Then, $\mu^*(\omega) = \mu^*(\tilde{\omega}) \forall \omega \in \Omega$.*

Proof. Follows directly from $P(\omega, \omega') = P(\tilde{\omega}, \tilde{\omega}') \forall \omega, \omega' \in \Omega$ which holds for $s_i = s_j$ due to platform symmetry (recall Assumption A). ■

Theorem 2 already identifies the set of stochastically stable states whenever at least one platform i has a price bias $\beta_i \leq 1$. Hence we are left with design configurations s_i and s_j where both price biases favor sellers (i.e. $\beta_i, \beta_j > 1$). There, Theorem 2 indicates that full coordination on both platforms are the only possible stochastically stable states. To further pin down the long-run trading pattern, it proves useful to distinguish the two platforms with respect to their prices. We start with the platform which implements a (weakly) lower price.

Lemma 4. *Suppose $\beta_i, \beta_j > 1$ and $p_i = \frac{\beta_i c}{1-f_i} \leq \frac{\beta_j c}{1-f_j} = p_j$. Then, $c(\omega_i^*, \omega_j^*) \geq 2 = c(\omega_j^*, \omega_i^*)$ and ω_i^* is stochastically stable.*

Proof. If $p_i \leq p_j$, buyers (weakly) prefer platform i to platform j whenever it is active. Accordingly, ω_i^* can be reached from ω_j^* with just two mutations. To see this, suppose the system is in state ω_j^* . A buyer and a seller switching to platform i due to experimentation induce trade at platform i . B1 and B2 guarantee that there is a positive probability that first all buyers at platform j get the opportunity to revise their platform choice (and switch to i because of the lower price or because of indifference) and subsequently all sellers at platform j can revise and also switch to platform i as this is now the only platform that generates positive profits for them. Hence, $c(\omega_j^*, \omega_i^*) = 2$. Clearly $c(\omega_i^*, \omega_j^*) \geq 2$ as at least one seller and one buyer have to switch to platform j to induce trade. Then, Theorem 2(a) implies stochastic stability of ω_i^* . ■

We now turn to the stochastic stability of a platform with a strictly higher price (provided that the biases of both platforms favor sellers). While all our previous results did not depend on the modelling details such as (i) absolute population size of buyers and sellers, (ii) the relative size of these populations, (iii) the heterogeneity of buyers, (iv) the price elasticity of demand, (v) the grid size δ , and (vi) details of the learning process (e.g. adjustment speed, asymmetries between buyers and sellers), these details do matter now as the following results illustrate.

Lemma 5. *Suppose $\beta_i, \beta_j > 1$, $p_i = \frac{\beta_i c}{1-f_i} < \frac{\beta_j c}{1-f_j} = p_j$, so that ω_i^* is stochastically stable.*

- (a) *in a dynamics with independent inertia, ω_j^* is also stochastically stable (i.e. $c(\omega_i^*, \omega_j^*) = 2$) if and only if there is at least one buyer $\tilde{n} \in N$ such that*

$$d_{\tilde{n}} \left(\frac{\beta_j c}{1-f_j} \right) (\beta_j - 1) \geq \frac{1}{|M| - 1} D_{N \setminus \{\tilde{n}\}} \left(\frac{\beta_i c}{1-f_i} \right) (\beta_i - 1).$$

- (b) *in a dynamics with non-simultaneous learning, ω_j^* is also stochastically stable (i.e. $c(\omega_i^*, \omega_j^*) = 2$) if and only if there is at least one buyer $\tilde{n} \in N$ such that*

$$\frac{1}{|M| - 1} d_{\tilde{n}} \left(\frac{\beta_j c}{1-f_j} \right) (\beta_j - 1) \geq D_{N \setminus \{\tilde{n}\}} \left(\frac{\beta_i c}{1-f_i} \right) (\beta_i - 1).$$

Proof. See Appendix. ■

Note that the condition in Lemma 5(a) is violated whenever e.g. buyers are identical, $|N| \geq |M|$, and the price elasticity of demand is sufficiently high. In contrast, the condition can be satisfied for $\beta_j > \beta_i$ whenever e.g. buyers are sufficiently heterogeneous (i.e. $\exists \tilde{n} \in N$ such that $d_{\tilde{n}}(p) \gg d_n(p) \forall n \neq \tilde{n}$), or buyers are identical and $|M| \gg |N|$, or $d(p)$ is sufficiently inelastic.

The stochastic stability of ω_j^* is harder to establish if the dynamics is slow as e.g. under non-simultaneous learning. The condition in Lemma 5(b) is violated whenever e.g. buyers are identical and the price elasticity of demand is sufficiently high (in contrast to the case of independent inertia, this holds independently of the sizes of populations $|M|$ and $|N|$). The condition can be fulfilled for $\beta_j > \beta_i$ if e.g. buyers are sufficiently heterogeneous or buyers are identical and demand is sufficiently inelastic.

Remark 1. *The analysis in the previous lemmata establishes stochastic stability through transition paths involving at most two simultaneous mutations. In the terminology of Ellison [11], this implies that the estimated time of first arrival, i.e. the estimated time until a stochastically stable state is first observed, is of order ε^{-2} . Thus the speed of convergence is relatively high. The number of mutations needed for the relevant transitions does not increase with the population size. Hence, our dynamics escape the well-known critique that for large populations the long run may actually be “too long” to be relevant (see Kandori et al. [22] or Ellison [10]).*

3.3. Platform Revenues and Designers’ Profits. Till now we have analyzed the learning dynamics of the traders and the resulting long run pattern of trades. Next turn to the revenues generated by the platforms, which in turn determine the profits of the market designers.

When analyzing the market designers’ choice of the characteristics of the trading platforms we will assume that platform designers are long-lived, patient, and (relatively) rational agents when compared with individual buyers or sellers. Hence, the designers take only the long term expected revenues into account, and ignore revenues made during the adjustment process to the limit invariant distribution.¹⁵ Given the platform characteristics $s = (s_i, s_j)$, the long-run expected revenues per round $ER_i(s)$ depend on the limit invariant distribution. The profits of designer i are given by $\pi_{D,i}(s) = f_i ER_i(s)$ implying that $\pi_{D,i}(s) \geq 0$ for all s .

Consider first a platform i with $\beta_i < 1$.

Lemma 6. *Suppose $\beta_i < 1$. Then $\pi_{D,i}((\beta_i, f_i), s_j) = 0$ for all feasible f_i, s_j .*

¹⁵Otherwise, the payoffs of a market designer would not only depend on the characteristics of both platforms, but also on the initial distribution of the traders over the platforms. In the absence of a plausible theory on the initial distribution, the results would be arbitrary. Further, as pointed out in Remark 1, convergence to full coordination is fast, and hence the assumption is, to some extent, justified.

Proof. Follows immediately from the fact that trade is never possible on platforms with $\beta_i < 1$. ■

Hence, we are left with platform configurations (s_i, s_j) where both platforms have a price bias weakly bigger than one. In this case expected revenues at platform i depend not only on the design of this platform but also on the design of the other platform as the following results indicate.

Lemma 7. Consider a platform configuration $s = (s_i, s_j)$ with $s_i = (\beta_i, f_i)$, $s_j = (\beta_j, f_j)$ and prices $p_i = \beta_i \frac{c}{1-f_i}$, $p_j = \beta_j \frac{c}{1-f_j}$.

- (a) If $s_i = s_j$, then, $\pi_{D,k}(s) = \frac{1}{2} f_k p_k D_N(p_k) > 0$ for $k = 1, 2$.
- (b) If $\beta_i = \beta_j = 1$ and $f_i < f_j$, then $f_k p_k D_N(p_k) > \pi_{D,k}(s_i, s_j) > 0$ for $k = 1, 2$.
- (c) If $\beta_i > 1$ and $\beta_j \leq 1$, then $\pi_{D,i}(s_i, s_j) = f_i p_i D_N(p_i)$ and $\pi_{D,j}(s_j, s_i) = 0$.
- (d) If $\beta_i, \beta_j > 1$, $p_i \leq p_j$, and $c(\omega_i^*, \omega_j^*) > 2$, then $\pi_{D,i}(s_i, s_j) = f_i p_i D_N(p_i)$ and $\pi_{D,j}(s_i, s_j) = 0$.
- (e) If $\beta_i, \beta_j > 1$, $p_i \leq p_j$, and $c(\omega_i^*, \omega_j^*) = 2$, then $\pi_{D,k}(s_i, s_j) = \mu^*(\omega_k^*) f_k p_k D_N(p_k) > 0$ for $k = 1, 2$.

Proof. (a) follows directly from Lemma 3 as the price and the traded quantity in ω at platform i are identical to price and quantity in $\tilde{\omega}$ at platform j . To see (b), recall from Theorem 2(c) that every $\omega \in \Omega$ is stochastically stable if $\beta_i = \beta_j = 1$. Hence, there is a strictly positive probability for trade (i.e. positive revenues) at each platform. (c) follows directly from Theorem 2(b). (d) and (e) follow directly from Theorem 2(a) and Lemma 4. ■

4. THE PLATFORM DESIGN

We now compare the design choices by a monopolistic designer and by two competing designers. Observe for reference that, as shown in the last Lemma 7, the designers' expected revenues are always weakly positive.

4.1. Monopolistic Market Design. As a benchmark, we briefly consider the case where only one platform is available, with characteristics $s = (\beta, f)$. In this case traders have no choice but to use this platform. Designer's profits are given by:

$$\pi_D(s) = \begin{cases} f \frac{\beta c}{1-f} D_N\left(\frac{\beta c}{1-f}\right) & \text{if } \beta \geq 1 \\ 0 & \text{otherwise.} \end{cases}$$

What is the profit maximizing platform (β^*, f^*) ? Note first that for $\beta < 1$ the profits are zero, whereas for $\beta \geq 1$ and for $0 < f < 1$ the profits are strictly positive. Hence, $\beta^* \geq 1$. Now assume for a moment that β and f are continuous variables with

$f \in [0, 1]$ and $\beta \in (0, \infty)$. Denote $p = \frac{\beta c}{1-f}$ and recall that $\lim_{p \rightarrow \infty} d_n(p)p = 0$ for all $n \in N$. Hence, it must hold that $0 < f^* < 1$.

Differentiating the designer's profits yields (for $\beta \geq 1$)

$$\begin{aligned}\frac{\partial \pi_D}{\partial f}(\beta, f) &= pD_N(p) + f \frac{\partial p}{\partial f} [D_N(p) + pD'_N(p)] \\ \frac{\partial \pi_D}{\partial \beta}(\beta, f) &= f \frac{\partial p}{\partial \beta} [D_N(p) + pD'_N(p)]\end{aligned}$$

where $\frac{\partial p}{\partial f} = \frac{\beta c}{(1-f)^2} > 0$ and $\frac{\partial p}{\partial \beta} = \frac{c}{1-f} > 0$. Let the optimal price be $p^* = \frac{\beta^* c}{1-f^*}$. Since $0 < f^* < 1$, the first order conditions for the designer's optimum imply that $\frac{\partial \pi_D}{\partial f}(\beta^*, f^*) = 0$, thus $D_N(p^*) + pD'_N(p^*) < 0$. This implies that, $\frac{\partial \pi_D}{\partial \beta}(\beta^*, f^*) < 0$, which leads to the conclusion that the designer's profits are maximized at the corner solution $\beta^* = 1$.

Of course, in our model β and f are not continuous variables. However, if the grid of feasible fees is fine enough, the optimal fee approximates the one of the continuous case, and hence the optimal β is 1 also in the discontinuous case. Hence we conclude that a monopolistic market designer would introduce a market clearing platform, because such a platform allows him to reap the highest profits.

The intuitive reason for this result is as follows. Suppose revenues $pD_N(p)$ are maximized at price p^* . Note that this price can be attained with different (β, f) combinations and that $p^* = \frac{\beta c}{1-f}$ is increasing both in β and f . Since the monopolist designer's profits are $f p D_N(p)$, he will try to reach p^* with that (β, f) combination that has the highest fee, and hence the lowest possible $\beta \geq 1$.

4.2. Competitive Market Design. In order to reflect that platform designers are “more rational” than individual buyers and sellers, we simply consider them rational players in the normal-form game defined by these payoff functions. That is, both designers choose their platforms simultaneously and payoffs are given as above. We also allow designers to use mixed strategies, i.e. choose a probability distribution over S rather than picking up a particular characteristic for sure.

Denote by σ_i the (mixed) strategy of designer i . The expected payoff of a designer is

$$\pi_{D,i}(\sigma_i, \sigma_j) = \sum_{s_j \in S} \sum_{s_i \in S} \sigma_j(s_j) \sigma_i(s_i) f_i ER_i(s_i, s_j).$$

Further, $\pi_{D,i}(s_i, \sigma_j)$ denotes the expected payoff if designer i chooses s_i for sure and j chooses the probability distribution σ_j .

We will now analyze the Nash equilibria of this game, implicitly assuming that designers are rational in the game theoretical sense. That is, we consider asymmetric rationality where designers are comparatively more sophisticated than traders.¹⁶

¹⁶Our main results also hold if we consider boundedly rational market designers as shown in Appendix B.1.

The sets of pure strategies of designer i and j are given by $S_i = S_j = B \times F$. Since these sets are finite, a Nash equilibrium of the designers' game always exists (possibly in mixed strategies). To characterize these equilibria, we need the following Lemma.

Lemma 8. *Let (σ_i^*, σ_j^*) be a Nash equilibrium (possibly in mixed strategies). Then, for any pure strategy $s_i = (\beta_i, f_i)$ of player i such that $\sigma_i^*(s_i) > 0$, it holds that $\beta_i \geq 1$.*

Proof. Assume to the contrary that there exists a pure strategy $\bar{s}_i = (\bar{\beta}_i, \bar{f}_i)$ with $\sigma_i^*(\bar{s}_i) > 0$ and $\bar{\beta}_i < 1$. By Lemma 6 this pure strategy gives designer i a profit of zero against all strategies of j . Hence, $\pi_{D,i}(\bar{s}_i, \sigma_j^*) = 0$, and, since σ_i^* is an equilibrium strategy, $\pi_{D,i}(\sigma_i^*, \sigma_j^*) = 0$.

Suppose that, in equilibrium, j chooses only platforms with $\beta_j < 1$. That is, $\beta_j < 1$ for all $s_j = (\beta_j, f_j) \in S$ with $\sigma_j^*(s_j) > 0$. In this case, if designer i chooses with certainty a platform s'_i with $\beta'_i > 1$, Lemma 7(c) implies that $\pi_{D,i}(s'_i, \sigma_j^*) = f'_i \frac{\beta'_i c}{1-f'_i} D_N \left(\frac{\beta'_i c}{1-f'_i} \right) > 0$. Since $\pi_{D,i}(\sigma_i^*, \sigma_j^*) = 0$, this contradicts that (σ_i^*, σ_j^*) is a Nash equilibrium.

Thus, there must exist an \bar{s}_j with $\bar{\beta}_j \geq 1$ such that $\sigma_j^*(\bar{s}_j) > 0$. Then, if designer i deviates to the pure strategy $s'_i = \bar{s}_j$, $\pi_{D,i}(s'_i, \sigma_j^*) = \sigma_j^*(\bar{s}_j) f_i ER_i(s'_i, \bar{s}_j) + \sum_{s_j \in S \setminus \bar{s}_j} \sigma_j^*(s_j) \pi_{D,i}(s'_i, s_j)$. Since by Lemma 7(a) $ER_i(s'_i, \bar{s}_j) > 0$, we conclude that $\pi_{D,i}(s'_i, \sigma_j^*) > 0$, again contradicting that (σ_i^*, σ_j^*) is a Nash equilibrium. ■

Hence, we have shown that, in equilibrium, only platforms weakly biased in favor of the sellers can be chosen. We now want to show that, actually, in any equilibrium, both designers will introduce platforms that lead to prices strictly above the market clearing level - platforms that lead to market clearing prices will not be designed in equilibrium. This results holds as long as the grid of possible biases is fine enough.

Theorem 9. *Let (σ_i^*, σ_j^*) be a Nash equilibrium (possibly in mixed strategies). For any pure strategy $s_i = (\beta_i, f_i)$ of player i such that $\sigma_i^*(s_i) > 0$, it holds that $\beta_i > 1$ if δ is sufficiently small.*

Proof. By the previous lemma, only platforms with $\beta \geq 1$ will be designed in equilibrium. Assume by contradiction that there exist some pure strategies $s_i = (\beta_i, f_i)$ with $\sigma_i^*(s_i) > 0$ and $\beta_i = 1$. Denote a strategy of this type by $\bar{s}_i = (1, \bar{f}_i)$ and let $\bar{p} = \frac{c}{1-\bar{f}_i}$. Denote the carrier or support of σ_j^* by $C(\sigma_j^*) = \{s_j = (\beta_j, f_j) \in S \mid \sigma_j^*(s_j) > 0\}$.

Suppose that, for all $s_j \in C(\sigma_j^*)$ we actually had that $\beta_j > 1$. This implies by Lemma 7(c) that $\pi_{D,i}(\sigma_i^*, \sigma_j^*) = \pi_{D,i}(\bar{s}_i, \sigma_j^*) = 0$. Take any $s'_i = s'_j \in C(\sigma_j^*)$. By Lemma 7(a), and recalling that $\pi_{D,i}(s_i, s_j) \geq 0$ for all s_i, s_j , we obtain that

$$\pi_{D,i}(s'_i, \sigma_j^*) \geq \sigma_j^*(s'_j) f'_j ER_j(s'_i, s'_j) > 0.$$

Hence, player i would have an incentive to deviate from σ_i^* , a contradiction.

We conclude that there exists some $s_j \in C(\sigma_j^*)$ with $\beta_j = 1$. Let $C_1(\sigma_j^*) = \{s_j = (\beta_j, f_j) \in C(\sigma_j^*) \mid \beta_j = 1\}$. Notice that, since $\bar{s}_i = (1, \bar{f}_i)$, we have by Lemma 7(c) that $\pi_{D,i}(\bar{s}_i, s_j) = 0$ for all $s_j \in C(\sigma_j^*)$ with $\beta_j > 1$. Then, by Lemma 7(a) and (b),

$$\pi_{D,i}(\bar{s}_i, \sigma_j^*) < \sum \{ \sigma_j^*(s_j) \bar{f}_i \bar{p} D_N(\bar{p}) \mid s_j \in C_1(\sigma_j^*) \}$$

However, for any s'_i with $\beta'_i > 1$ and $f'_i = \bar{f}_i$,

$$\pi_{D,i}(s'_i, \sigma_j^*) \geq \sum \{ \sigma_j^*(s_j) \bar{f}_i \beta'_i \bar{p} D_N(\beta'_i \bar{p}) \mid s_j \in C_1(\sigma_j^*) \}$$

due to Lemma 7 (c) (the inequality follows from the fact that $\pi_{D,i}(s'_i, s_j) \geq 0$ for all s_j). This latter expression is continuous in β'_i . Thus, for β'_i approaching one from above, $\pi_{D,i}(s'_i, \sigma_j^*) > \pi_{D,i}(\bar{s}_i, \sigma_j^*) = \pi_{D,i}(\sigma_i^*, \sigma_j^*)$. Hence, if the grid is fine enough¹⁷ player i has an incentive to deviate from σ_i^* to an institution with $\beta'_i > 1$ but close enough to 1. A contradiction. ■

We have thus established the paradoxical result that competition among platform designers will induce them to select biased platforms which implement non-competitive market outcomes. In general, nothing more can be said about the specific characteristics of the Nash equilibria. A brief examination of Lemma 5 and Lemma 7(d) and (e) should convince the reader that a full characterization of the Nash equilibria will depend on the exact shape of the limit invariant distribution, and not only on its support. This distribution in turn depends on the details of the dynamics, e.g. whether learning opportunities arise simultaneously among traders or asynchronously. In contrast, the last theorem holds for any specification of the learning dynamics satisfying assumptions B1 and B2.

Still, one might suspect that competition leads to platforms close to the market clearing one, i.e. to platforms with $\beta_i = 1 + \delta$. If this suspicion would be correct, the chosen platforms would nearly resemble market clearing ones as long as the grid of feasible biases is fine enough. The next proposition, however, shows that this suspicion is in general wrong. To keep things simple consider identical buyers with a demand function $d(p)$ and denote the price elasticity of demand by $\epsilon_p(p) = -\frac{pd'(p)}{d(p)}$.

Proposition 10. *Assume independent inertia, identical buyers and $|M| = |N|$. Suppose that the grids B and F are fine enough and if ϵ_p is not much larger than one. Then,*

- i) *there exists no Nash equilibrium (σ_i^*, σ_j^*) (possibly in mixed strategies) where both designers introduce only platforms with $\beta_i = \beta_j = 1 + \delta$;*
- ii) *there exists no pure strategy Nash-equilibrium (s_i^*, s_j^*) of the design game with $\beta_i^* = \beta_j^* = 1 + \delta$.*

¹⁷The grid can be assumed to be *ex ante* fine enough by a uniform continuity argument.

Proof. See Appendix. ■

This proposition shows that equilibrium designs (beyond the features highlighted in Theorem 9) are in general rather sensitive to details of the economy and the learning process. For example it can also not be excluded that designers choose "near market clearing" platform characteristics for some specifications of the learning dynamics. In general, however, this will not prove optimal.

5. DISCUSSION

We have shown that if several trading platforms are available, traders will learn to coordinate on a platform with prices systematically above the market clearing level, if such a platform is feasible. This forces competing market designers to create such non-market clearing platforms. On the other hand a monopolistic market designer will always introduce a market clearing platform in order to maximize his profits. Hence, platform competition induces non-competitive market outcomes.

We view this result as a paradox, since it amounts to the statement that competition among market designers might induce non-competitive outcomes on the platforms. Of course, though, this result depends on several assumptions we have made.

First, we have assumed sellers to be producers endowed with a technology with constant returns to scale. Although this is a focal, economically meaningful case, it clearly simplifies the analysis and allows for a clear-cut derivation of the results. Under production technologies exhibiting decreasing returns to scale, the results are not so sharp and a characterization of the limit invariant distribution requires both a further specification of the learning behavior of the traders and a further specification of demand and supply. In Appendix B we exhibit an extended example with decreasing returns to scale where our main result still holds. It shows, however, that the optimality of a price bias is no longer independent of details like learning velocities. In particular, if one side of the market learns much faster, it might pay off to bias prices to favor this party. Nonetheless, this clearly illustrates that the scope of the paradox identified here goes beyond the constant returns to scale case.

Second, we have focused on an asymmetric rationality model where market designers are sophisticated when compared to traders. Furthermore, by focusing on long-run profits we have implicitly assumed that it is much more difficult for designers to change the properties of their platforms than for traders to switch trading platforms. In Appendix B we analyze the case of boundedly rational designers who have to learn how to design a platform through a regular (trial-and-error) design revision process. Our main result - the emergence of non-market clearing institutions - carries over to such a setting.

These robustness checks show that neither the assumption of constant returns to scale technology nor that of rational designers *drive* our results. Rather, it is indeed platform competition that leads to the emergence of non-market clearing trading platforms.

REFERENCES

- [1] Alós-Ferrer, C. (2003), Finite Population Dynamics and Mixed Equilibria, *International Game Theory Review* 5 (3), 263-290.
- [2] Alós-Ferrer, C. and G. Kirchsteiger (2003), Does Learning lead to Coordination in Market Clearing Institutions?, Working Paper 0320, Department of Economics, University of Vienna.
- [3] Armstrong, M. (2002), Competition in Two-sided Markets, mimeo.
- [4] Belleflamme, P. and E. Toulemonde (2004), Emergence and Entry of B2B Marketplaces, mimeo.
- [5] Benaïm, M. and J. Weibull (2003), Deterministic Approximation of Stochastic Evolution in Games, *Econometrica* 71, 878-903.
- [6] Binmore, K. and L. Samuelson (1997), Muddling Through: Noisy Equilibrium Selection, *Journal of Economic Theory* 74, 235-265.
- [7] Blume, L. (1995), The Statistical Mechanics of Best-Response Strategy Revision, *Games and Economic Behavior* 11, 111-145.
- [8] Caillaud, B. and B. Jullien (2002), Chicken and Egg: Competition among Intermediation Service Providers, mimeo.
- [9] The Economist May 13th (2004), E-commerce takes off.
- [10] Ellison, G. (1993), Learning, Local Interaction, and Coordination, *Econometrica* 61, 1047-1071.
- [11] Ellison, G. (2000), Basins of Attraction, Long-Run Stability, and the Speed of Step-by-Step Evolution, *Review of Economic Studies* 67, 17-45.
- [12] Ellison, G. (2005), Bounded Rationality in Industrial Organization, mimeo.
- [13] Ellison, G., and D. Fudenberg (2003), Knife-edge or Plateau: When Do Markets Models Tip?, *Quarterly Journal of Economics* 118, 1249-1278.
- [14] Ellison, G., D. Fudenberg, and M. Möbius (2004), Competing Auctions, *Journal of the European Economic Association* 2, 30-66.
- [15] European Commission, Enterprise Directorate General (2003), Report of the Expert Group on B2B Internet Trading Platforms, Final report, mimeo.
- [16] Fudenberg, D., and D. Levine (1998), *The Theory of Learning in Games*. The MIT Press, Cambridge, Massachusetts.

- [17] Freidlin, M.I. and A.D. Wentzell (1984), *Random Perturbations of Dynamical Systems*, Springer Verlag, New York.
- [18] Gabaix, X., and D. Laibson (2005), *Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*, forthcoming in *Quarterly Journal of Economics*.
- [19] Gerber, A., and M.O. Bettzüge (2002), *Evolutionary Choice of Markets*, forthcoming in *Economic Theory*.
- [20] Holt, C. (1995), *Industrial Organization*, in Kagel, J. and A. Roth (eds): *Handbook of Experimental Economics*, Princeton University Press, Princeton.
- [21] Hopkins, E. (2005), *Adaptive Learning Models of Consumer Behavior*, mimeo.
- [22] Kandori, M., G. Mailath, and R. Rob (1993), *Learning, Mutation, and Long-Run Equilibria in Games*. *Econometrica* 61, 29-56.
- [23] 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.
- [24] Kirchsteiger, G., Niederle, M., and J. Potters (2005), *Endogenizing Market Institutions: An Experimental Approach*, *European Economic Review* 49(7), 1827-1853..
- [25] Kugler, T., Neeman, Z., and N. Vulkan (2006), *Markets versus negotiations: An experimental investigation*, forthcoming in *Games and Economic Behavior*.
- [26] Lucking-Reiley, D. and Spulber, D. (2001), *Business to Business Electronic Commerce*, *The Journal of Economic Perspectives* 15, 55-68.
- [27] Neeman, Z. and N. Vulkan (2002), *Markets versus Negotiations: The Predominance of Centralized Markets*. mimeo.
- [28] Plott, C. (1982), *Industrial Organization Theory and Experimental Economics*, *Journal of Economic Literature* 20, 1485-1587.
- [29] Plott, C. and V. Smith (1978), *An Experimental Examination of Two Exchange Institutions*, *Review of Economic Studies* 45, 133-153.
- [30] Rochet, J. and J. Tirole (2003), *Platform Competition in Two-sided Markets*, *Journal of the European Economic Association* 1, 990-1029.
- [31] Rochet, J. and J. Tirole (2004), *Two-sided Markets: An Overview*, mimeo.
- [32] Samuelson, L. (1994), *Stochastic Stability in Games with Alternative Best Replies*. *Journal of Economic Theory* 64, 35-65.

- [33] Samuelson, L. (1997), *Evolutionary Games and Equilibrium Selection*. The MIT Press, Cambridge, Massachusetts.
- [34] Spiegler, R., *Competition over Agents with Boundedly Rational Expectations*, *Theoretical Economics*, forthcoming.
- [35] UNCTAD (2002), *e-Commerce and Development Report 2002*, mimeo.
- [36] Young, P. (1993), *The Evolution of Conventions*. *Econometrica* 61, 57-84.

APPENDIX

A. PROOFS

A.1. Proof of Lemma 1. We first prove the following preliminary claim. The monomorphic state ω_i^* can be reached with positive probability from any state where platform $j \neq i$ is inactive. To see this, simply notice that, by Assumptions A and B1, there is positive probability that traders will switch away from the inactive platform j until ω_i^* is reached. An immediate consequence of this claim is that, from the “cross-states” where both platforms are inactive, i.e. all buyers are at one platform and all sellers are at the other platform, both monomorphic states can be reached.

(a) Let $\beta_1 > 1$ and $\beta_2 > 1$. First note that the monomorphic states are absorbing, because at the corresponding platform both buyers and sellers make strictly positive profits and the other platform is inactive. Thus traders stay at the active one.

To complete the proof, it is enough to show that there exists a positive probability path from any state to some of the monomorphic states. By the preliminary claim, this is true if any platform is inactive. Consider, thus, an arbitrary $\omega \in \Omega$ such that $1 \leq |N_i(\omega)| \leq |N| - 1$ and $1 \leq |M_i(\omega)| \leq |M| - 1$.

Without loss of generality, suppose $p(s_i) \leq p(s_j)$. Buyers weakly prefer platform i to platform j . By B1 and B2, there is positive probability that all buyers at j receive revision opportunity in successive periods and only sellers at j receive revision opportunity. Hence, buyers will switch away from j and no new seller will switch to j . Thus, either the monomorphic state ω_i^* is eventually reached, or j becomes inactive, leading again to ω_i^* with positive probability.

(b) Let $\beta_j \leq 1$ and $\beta_i > 1$. Provided i is active, sellers strictly prefer it to platform j , regardless of prices. As above, the monomorphic state ω_i^* is absorbing. The monomorphic state ω_j^* is not. To see this, note that sellers always receive zero profits at j and will switch to i with positive probability even if profits there are zero (the latter due to randomization in case of indifference). Thus, starting at ω_j^* , with positive probability j becomes inactive such that ω_i^* is reached with positive probability.

By the preliminary claim, from any state where platform j is inactive the monomorphic state ω_i^* can be reached. If i is inactive, then ω_j^* is reached, but then the dynamics leads to ω_i^* with positive probability.

Consider an arbitrary ω such that $1 \leq |N_i(\omega)| \leq |N| - 1$ and $1 \leq |M_i(\omega)| \leq |M| - 1$. By B1 and B2, there is positive probability that all sellers at j receive revision opportunity in successive periods and only buyers at j receive revision opportunity. Hence, sellers will switch away from j and no new buyer will switch to j . Eventually, either the monomorphic state ω_i^* is reached, or j becomes inactive (which again leads to ω_i^* with positive probability).

(c) Last, suppose $\beta_i \leq 1$ and $\beta_j \leq 1$. If both $\beta_i < 1$ and $\beta_j < 1$, both platforms are always inactive. Hence every trader randomizes when given revision opportunity, thus the process is obviously irreducible.

If $\beta_j < 1$ and $\beta_i = 1$, sellers always receive zero profits and always randomize. Buyers, though, switch away from the inactive platform j whenever platform i is active and they receive revision opportunities, while they randomize whenever i is inactive. Thus, from any state, there is always positive probability to reach the monomorphic state ω_i^* . From this state, any state in

$$\Omega_i^B = \{\omega \mid N_i(\omega) = N\}$$

can be reached. Further, since sellers simply randomize, they might eventually all switch to platform j , reaching the cross state in which all sellers are at j and all buyers are at i . As long as all sellers are in j , buyers merely randomize. That is, any state in

$$\Omega_j^S = \{\omega \mid M_j(\omega) = M\}$$

can be reached. Consider an arbitrary state ω' with $|N_i(\omega')| > 0$. This state can be reached from a state $\omega \in \Omega_j^S$ with $N_i(\omega) = N_i(\omega')$ by letting sellers switch to i and giving revision opportunity only to buyers at i , who will stay there. Finally, consider a state ω' with $N_i(\omega') = \emptyset$, i.e. $N_j(\omega') = N$. This state can be reached from the monomorphic state $\omega_j^* \in \Omega_j^S$ by letting sellers randomize appropriately, because buyers who receive revision opportunity are indifferent (both platforms are inactive) and will stay at j with positive probability.

In summary, any state can be reached from ω_i^* , and the latter can be reached from any state. This completes the argument if $\beta_j < 1$ and $\beta_i = 1$.

The only remaining case is $\beta_i = \beta_j = 1$. In this case, sellers always obtain zero profits and hence randomize. Buyers in turn want to move away from inactive platforms. If both platforms are inactive, they randomize, too. However, since sellers simply randomize, it is clear that, from an arbitrary state, any platform might eventually become inactive by lack of sellers. It follows from the preliminary claim that, from any state, both monomorphic states can be reached.

From the monomorphic states, since sellers still simply randomize, all states at Ω_i^B and Ω_j^B (see above) can be reached - in particular the cross states. Then, buyers also randomize. As long as sellers stay concentrated at one platform, buyers strictly prefer that platform if it became active. That is, we can construct positive probability paths to any state in Ω_j^S and Ω_i^S (see above). Again, since sellers merely randomize,

we can construct positive probability paths to any state ω' with $0 < |N_i(\omega')| < |N|$ by letting sellers switch appropriately and giving revision opportunity only to buyers who are already at i . This completes the argument if $\beta_i = \beta_j = 1$. ■

A.2. Proof of Lemma 5. It has been shown in the proof of Lemma 4 that $c(\omega_j^*, \omega_i^*) = 2$ if $\beta_i, \beta_j > 1$ and $p_i \leq p_j$. Hence, a necessary and sufficient condition for the stochastic stability of ω_j^* is $c(\omega_i^*, \omega_j^*) = 2$.

Since $p_i < p_j$, buyers never switch to platform j as long as i is active. Hence, ω_j^* has to be reached through switching of *all* sellers to platform j and a subsequent switch of all buyers to the only remaining active platform.

In case (a), if the unperturbed dynamics is described by independent inertia, there is a positive probability that all sellers at platform j simultaneously receive the opportunity to revise. If one seller and buyer $\tilde{n} \in N$ are already present at platform j , sellers will indeed switch to j if $d_{\tilde{n}}\left(\frac{\beta_j c}{1-f_j}\right)(\beta_j - 1) \geq \frac{1}{|M|-1} D_{N \setminus \{\tilde{n}\}}\left(\frac{\beta_i c}{1-f_i}\right)(\beta_i - 1)$. Hence, this condition proves to be sufficient for the stochastic stability of ω_j^* . To see that the condition is also necessary suppose that it is violated. Then no seller will switch to j after one seller and any buyer \tilde{n} induced trade on this platform. As a consequence, more than 2 mutations are needed to reach ω_j^* .

In case (b), under non-simultaneous learning there is only one trader in any period who has the opportunity to revise its platform choice. Suppose one seller and buyer \tilde{n} switch to platform j by mutation. Suppose furthermore that in the subsequent rounds only sellers and buyers at platform i receive the opportunity to revise their decision (this happens with strictly positive probability due to B1 and B2). If $\frac{1}{|M|-1} d_{\tilde{n}}\left(\frac{\beta_j c}{1-f_j}\right)(\beta_j - 1) \geq D_{N \setminus \{\tilde{n}\}}\left(\frac{\beta_i c}{1-f_i}\right)(\beta_i - 1)$ it follows that $\frac{1}{|M_j|} d_{\tilde{n}}\left(\frac{\beta_j c}{1-f_j}\right)(\beta_j - 1) \geq \frac{1}{|M|-|M_j|} D_{N \setminus \{\tilde{n}\}}\left(\frac{\beta_i c}{1-f_i}\right)(\beta_i - 1)$ for all M_j with $1 \leq |M_j| \leq |M| - 1$. Hence, sellers prefer platform j whenever it is active and there are at least one and less than $|M|$ sellers already there. This implies that there is a positive probability path with just two mutations from ω_i^* to ω_j^* where first all sellers move to platform j and subsequently all buyers switch to j as it is the only active platform. Hence, the condition displayed in the Lemma is sufficient for the stochastic stability of ω_j^* . To see that it is also necessary, suppose it is not fulfilled. Then a seller at platform i prefers to stay there if all other sellers are at platform j together with any buyer \tilde{n} . Under non-simultaneous learning this implies that at least a third mutation is needed to reach ω_j^* , which implies that this state cannot be stochastically stable by Theorem 2(a). ■

A.3. Proof of Proposition 10. Note first that Lemma 5 (a) implies that if $\frac{\beta_i c}{1-f_i} < \frac{\beta_j c}{1-f_j}$, ω_j^* is stochastically stable iff

$$d\left(\frac{\beta_j c}{1-f_j}\right)(\beta_j - 1) \geq d\left(\frac{\beta_i c}{1-f_i}\right)(\beta_i - 1). \quad (*)$$

Proof of i) Assume to the contrary for all platform characteristics in the support of σ_i^* and σ_j^* , $\beta = 1 + \delta$. Denote by \bar{f}_i the highest fee of a platform in the support of σ_i^* , and by \bar{f}_j the highest fee of a platform in the support of σ_j^* . Without loss of generality assume that $\bar{f}_j \geq \bar{f}_i$. We can distinguish between three cases:

a) $\bar{f}_j > \bar{f}_i$: Condition (*) shows that full coordination on platform $\bar{s}_j = (1 + \delta, \bar{f}_j)$ is not stochastically stable vis a vis any platform characteristics in the support of σ_i^* . Hence, this strategy earns designer j zero profits, and since it is assumed to be in the support of j 's equilibrium strategy, j 's equilibrium profits would be zero. But j could always guarantee himself a strictly positive profit by playing the same (possibly mixed) strategy as i . Hence, case (a) is inconsistent with Nash equilibrium.

b) $\bar{f}_i = \bar{f}_j > f_{\min}$. Condition (*) shows that full coordination on platform $\bar{s}_i = (1 + \delta, \bar{f}_i)$ is not stochastically stable vis a vis any platform characteristics in the support of σ_j^* but platform $\bar{s}_j = (1 + \delta, \bar{f}_j)$. Furthermore, Lemmata 2 and 3 imply that $\mu^*(\omega_i^*) = \mu^*(\omega_j^*) = \frac{1}{2}$ if \bar{s}_i is chosen by i and \bar{s}_j is chosen by j . Therefore,

$$\pi_{D,i}(\bar{s}_i, \sigma_j^*) = \sigma_j^*(\bar{s}_j) \frac{1}{2} \bar{f}_i \frac{(1 + \delta)c}{1 - \bar{f}_i} |N| d \left(\frac{(1 + \delta)c}{1 - \bar{f}_i} \right).$$

But choosing the alternative platform design s'_i with $f'_i = \bar{f}_j - \gamma$, and $\beta'_i = 1 + \delta$ implies that $\frac{1 + \delta}{1 - f'_i} c < \frac{1 + \delta}{1 - \bar{f}_j} c$ and $d \left(\frac{1 + \delta}{1 - f'_i} c \right) \delta < d \left(\frac{1 + \delta}{1 - \bar{f}_j} c \right) \delta$. Hence, again by (*) $\mu^*(\omega_i^*) = 1$ if s'_i is chosen by i and \bar{s}_j is chosen by j , which yields

$$\pi_{D,i}(s'_i, \sigma_j^*) \geq \sigma_j^*(\bar{s}_j) (\bar{f}_i - \gamma) \frac{(1 + \delta)c}{1 - \bar{f}_i + \gamma} |N| d \left(\frac{(1 + \delta)c}{1 - \bar{f}_i + \gamma} \right).$$

If the grid of F is fine enough, i.e. if γ is small enough, this implies

$$\pi_{D,i}(s'_i, \sigma_j^*) > \pi_{D,i}(\bar{s}_i, \sigma_j^*)$$

Hence, case (b) is inconsistent with Nash equilibrium.

c) $\bar{f}_i = \bar{f}_j = f_{\min}$ - both designers choose the platform characteristics $\bar{s}_i = \bar{s}_j = (1 + \delta, f_{\min})$ for sure. Then Condition (*) guarantees the existence of a $\beta'_j > 1 + \delta$ and a $f'_j > f_{\min}$ such that platform j is stochastically stable vis a vis \bar{s}_j if the grid F is sufficiently fine, i.e. if γ is sufficiently small. Furthermore, if ϵ_p is not much larger than 1, designer j 's profits from full coordination on his platform with design s'_j , i.e. $|N| f'_j \frac{\beta'_j c}{1 - f'_j} d \left(\frac{\beta'_j c}{1 - f'_j} \right)$, is strictly larger than the respective profit from choosing \bar{s}_j , i.e. $|N| \bar{f}_j \frac{(1 + \delta)c}{1 - \bar{f}_j} d \left(\frac{(1 + \delta)c}{1 - \bar{f}_j} \right)$. It remains to show that no decrease in $\mu^*(\omega_j^*)$ overcompensates this effect. To see this suppose that $\beta'_j > \bar{\beta}_j = 1 + \delta$ in such a way

that $d(p'_j)(\beta'_j - 1) > |M||N|d(p_i)\delta$ (feasible if δ is sufficiently small). Then sellers prefer platform j with characteristics s'_j whenever it is active, while buyers prefer platform i with characteristics \bar{s}_i . For a learning dynamics with independent inertia and with $|M| = |N|$ this establishes symmetry of the transition matrix P such that $\mu^*(\omega_i^*) = \mu^*(\omega_j^*) = 1/2$. Therefore, choosing in this case $\beta'_j > 1 + \delta$ and $f'_j > \bar{f}_j$ does not reduce $\mu^*(\omega_j^*)$ while it strictly increases revenue in ω_j^* . Hence, \bar{s}_j with $\bar{\beta}_j = 1 + \delta$ and $\bar{f}_i = f_{min}$ can not be a best response to $\bar{s}_i = (1 + \delta, f_{min})$.

Proof of ii) Follows immediately from i). ■

B. ROBUSTNESS OF THE RESULTS

The results in the paper have been derived under two crucial assumptions. First, designers are assumed to be rational while traders are not (asymmetric rationality). Even though this assumption seems to be justified in a wide range of applications, one might be interested in the robustness of our results with respect to the (bounded) rationality of designers. We will discuss the case of learning designers in section B.1. Second, we assumed that sellers have a constant-returns-to-scale technology. In section B.2 we will analyze an example with decreasing returns that illustrates the robustness of our findings.

B.1. Boundedly Rational Designers. To account for learning designers, we have to extend the state space by the feasible design configurations, and we have to redefine the (unperturbed learning process). All other assumptions remain unchanged.

The *state space* is given by $\Omega = \{1, 2\}^n \times \{1, 2\}^m \times S^2$. A state $\omega \in \Omega$ denotes the location of buyers and sellers and the design of both platforms. Traders learn as specified in Assumption A. The learning process of designers is defined as follows.

Assumption C A designer who gets the opportunity to revise, observes the revenues and designs of platforms in the last period. If the platforms have different designs, he chooses the design which has led to a higher revenue (*Imitation*). In case of identical revenues designers randomize their choice, with both designs chosen with strictly positive probability. Choice probabilities may depend on the outcomes (or designs) but not on the platform's name. If designs are identical and a designer obtains zero revenues he randomizes over all possible design alternatives in the next round (*Innovation*).¹⁸

As for the opportunities to revise we invoke Assumptions B1 and B2 (on the enlarged state space and for three instead of two different types of players) as detailed in the paper.¹⁹ To illustrate the robustness of our results with respect to boundedly

¹⁸For simplicity, we assume that designers randomize over S with full support. Our results would carry over to more realistic treatments where designers use more sophisticated innovation methods.

¹⁹Note that this specification also allows for different learning speeds for traders and designers, respectively. Our model covers, for instance, the likely situation that buyers and sellers revise with a larger probability than designers.

rational designers, we prove the counterpart of Lemma 1 in the modified learning model.

Lemma 11. *Assume assumptions A, B1, B2, and C. Then the absorbing sets are all sets of the form $A_{i,\beta,f} = \{\omega \in \Omega \mid N_i(\omega) = N, M_i(\omega) = M, \beta_i = \beta, f_i = f\} \forall \beta > 1, f \in F, \text{ and } i = 1, 2.$*

Proof. We proceed in two steps. First, we show that $A_{i,\beta,f}$ is indeed absorbing for any i, f , and $\beta > 1$. Second, we show that there is always a positive probability path from any $\omega \in \Omega \setminus \cup_{i,f,\beta>1} A_{i,\beta,f}$ to an $\omega' \in \cup_{i,f,\beta>1} A_{i,\beta,f}$.

Fix any i, f and $\beta > 1$, and consider any $\omega \in A_{i,\beta,f}$. As $\beta > 1$ buyers and sellers receive strictly positive payoffs at i , and no type of trader wants to switch to the inactive platform j . According to Assumption C designer i will not change the design of his platform while designer j will innovate (randomize over S) if $s_i = s_j$ in the last period, and he will imitate (choose $s_j = s_i$) otherwise. Hence, once the system reached a state in $A_{i,\beta,f}$ buyers, sellers, and designer i will not alter their choices and designer j will alternate between $s_j = s_i$ and a random draw out of S - the system will not leave $A_{i,\beta,f}$.

Now fix a state $\omega \in \Omega \setminus \cup_{i,f,\beta>1} A_{i,\beta,f}$ and denote the designs of platform i and j in this state by $s_i = (\beta_i, f_i)$ and $s_j = (\beta_j, f_j)$, respectively. We distinguish four cases. (i) $\beta_i, \beta_j > 1$, (ii) $\beta_i > 1$ and $\beta_j \leq 1$, (iii) $\beta_i = \beta_j = 1$, and (iv) $\beta_i, \beta_j < 1$. As for (i) and (ii) it has already been shown in the proof of Lemma 1 that for fixed platform designs there is a positive probability path of the unperturbed dynamics to a monomorphic state ω_i^* with $\beta_i > 1$ from any non-monomorphic state or any monomorphic state ω_j^* with $\beta_j \leq 1$. Due to Assumptions B1 and B2 there is also positive probability that only designers with highest revenues receive revision opportunities (and do not change design) until a state ω' with $N_i(\omega') = N, N_j(\omega') = M$ is reached. Now consider (iii) $\beta_i = \beta_j = 1$. Recall from the proof of Lemma 1 that a cross-state will be reached with positive probability (for unchanged designs). In such a state designers randomize such that there is a positive probability path to a cross-state with $\beta_i > 1$, and a subsequent positive probability path to a state in $A_{i,\beta,f}$ (due to randomizing traders). Finally consider (iv) - $\beta_i, \beta_j < 1$. In such a state traders randomize and a cross-state will be reached with positive probability. In a cross state designers randomize such that a cross-state with $\beta_i > 1$ is reached. As traders continue to randomize a state in A_{i,β_i,f_i} can be reached. ■

Hence, in every absorbing set all traders are located at a platform i with $\beta_i > 1$ while designers of the other platform randomize over $S = B \times F$ (i.e. unsuccessfully innovate). As transition costs between any two absorbing sets are 1 (starting in a state in $A_{i,\beta,f}$ just one experiment by designer i is required to end up in any other $A_{i',\beta',f'}$) all absorbing states characterized in Lemma 11 are also stochastically stable. This establishes Theorem 2 for boundedly rational designers.

We conclude that boundedly rational platform designers exhibit qualitatively the same behavior as rational ones: Platform competition forces them to introduce only non-market clearing platforms.

B.2. Decreasing Returns to Scale. Consider the following example. Two identical sellers are producing with costs $c(q) = \frac{1}{2}q^2$ for q units of a commodity. For given prices (p_i) and fees (f_i) at a platform i , their profit is $\pi_{m,i}(q, p_i) = (1 - f_i)p_i q - \frac{1}{2}q^2$ and maximization leads to the supply function $s(p_i) = (1 - f_i)p_i$. Two identical buyers, - each with income of one unit - consume q units of the commodity traded at the platforms and x units of a second commodity which price is normalized to 1. The buyer's utility is given by $\pi_n(q, x) = 2\sqrt{q} + x$ and solving for the optimal consumption bundle for a given price at the respective platform p_i , yields the buyer's demand function $d(p_i) = 1/p_i^2$. Equating demand and supply gives the market clearing price at platform i in state ω $p_i^*(\omega) = r_i(\omega)^{1/3}(1 - f_i)^{-1/3}$ (with $r_i(\omega) = \frac{|N_i(\omega)|}{|M_i(\omega)|}$). Traders' and designers' profits depend on state and design and are calculated the same way as before. For our purposes it suffices to note that sellers are not rationed whenever $\beta_i \leq 1$ and their corresponding profit $\pi_{m,i}(M_i(\omega), N_i(\omega), s_i) = \frac{1}{2}(1 - f_i)^{4/3}\beta_i^2 r_i^{2/3}$ is increasing in β_i and decreasing in f_i . Analogously, sellers are rationed for $\beta_i > 1$ and profits amount to $\pi_{m,i}(M_i(\omega), N_i(\omega), s_i) = (1 - f_i)^{4/3}\frac{1}{\beta_i} r_i^{2/3}(1 - \frac{1}{2\beta_i^3})$ which is also monotonically decreasing in f_i but reaches a (global) maximum at $\beta_i = 2^{1/3}$.²⁰

Full coordination on any platform will be a singleton-absorbing set (both types of traders get strictly positive profits on any active platform). In fact, it is easy to see - as in the proof of Lemma 1 - that these states are the only absorbing sets. Moreover, it can be shown that these monomorphic states are also stochastically stable for a wide range of design configurations. Note that not only the support of the stochastically stable distribution μ^* , but also its values for the different (monomorphic) states matter for the designers' profits. Hence, this setting can not be analyzed with the tools discussed before. To obtain μ^* in this example, we have to make use of Lemma 3.1 from Freidlin and Wentzell ([17])²¹. They derive the limit invariant distribution μ^* from an analysis of graphs in the state space Ω that are induced by the stationary Markov-process P as follows. Fix an $\omega \in \Omega$. An ω -tree T is a spanning tree in Ω such that for every vertex $\omega' \neq \omega$ there exists a unique directed path from ω' to ω . Let \mathcal{T}_ω be the set of all ω -trees and define $q\omega \equiv \sum_{T \in \mathcal{T}_\omega} \prod_{(\omega', \omega'') \in T} P_{\omega', \omega''}$ (i.e. $q\omega$ is the product of all transition probabilities on a given ω -tree summed over all ω -trees). Then Lemma 3.1 of Freidlin and Wentzell [17] states that $\mu^*(\omega) = \frac{q\omega}{\sum_{\omega \in \Omega} q\omega}$. In the limit of small mutation probability ($\epsilon \rightarrow 0$) $\mu^*(\omega)$ is determined by those ω -trees that contain the smallest possible number of mutations necessary to form a spanning tree in Ω . Hence, we will restrict ourselves to those trees.

For expositional ease we further specify the learning model and substitute Assumptions B1 and B2 by

Assumption D: In every round any seller $m \in M$ is allowed to revise his location decision with probability $\rho_S \in (0, 1)$ while every buyer is allowed to revise with probability $\rho_B \in (0, 1)$.

²⁰We assume for simplicity from now on that δ is such that $2^{1/3} \in B$.

²¹For a detailed explanation see Kandori et al. [22] or Young [36].

Then a least mutation tree in $\mathcal{T}_{\omega_i^*}$ has edges as depicted in the following table.

ω'	P_{ω', ω_i^*}
(0,0)	$\epsilon^2 P_{(1,1), \omega_i^*}$
(0,1)	$\epsilon(1 - \rho_B) P_{(1,1), \omega_i^*}$
(0,2)	$\epsilon \rho_S$
(1,0)	$\epsilon(1 - \rho_S) P_{(1,1), \omega_i^*}$
(1,1)	$P_{(1,1), \omega_i^*}$
(1,2)	ρ_S
(2,0)	$\epsilon \rho_B$
(2,1)	ρ_B

The respective table for ω_j^* can be derived by a permutation of indices B and S . Hence, $\frac{\mu^*(\omega_j^*)}{\mu^*(\omega_i^*)} = \rho_B^2 \rho_S^2 (1 - \rho_B)(1 - \rho_S) \left(P_{(1,1), \omega_j^*} \right)^4 / \rho_B^2 \rho_S^2 (1 - \rho_B)(1 - \rho_S) \left(P_{(1,1), \omega_i^*} \right)^4$. Note that this expression is continuous in ρ_S and ρ_B . This leads to the following useful result (as sellers and buyers are identical profits only depend on the *number* of sellers and buyers at a platform).

Lemma 12. *Suppose $\pi_{m,i}(1, 1, s_i) > \pi_{m,j}(1, 1, s_j)$. Then for every $\epsilon > 0$ there is a $\bar{\rho}_S < 1$ such that $\mu^*(\omega_j^*) < \epsilon$ for all $\rho_S > \bar{\rho}_S$.*

Proof. First recall that the monomorphic states are the only absorbing sets. Moreover, it is easy to see that ω_i^* is the only stochastically stable state if and only if $\pi_{m,i}(1, 1, s_i) > \pi_{m,j}(1, 1, s_j)$ and $\pi_{n,i}(1, 1, s_i) > \pi_{n,j}(1, 1, s_j)$ (as it then needs more than 2 mistakes to get from ω_i^* to ω_j^*). If $\pi_{m,i}(1, 1, s_i) > \pi_{m,j}(1, 1, s_j)$ and $\pi_{n,i}(1, 1, s_i) = \pi_{n,j}(1, 1, s_j)$, $P_{(1,1), \omega_i^*} = \rho_S(1 - \rho_B)\rho_B + \frac{1}{2}\rho_B(1 - \rho_S)\rho_S + \frac{1}{2}\rho_B\rho_S$ while $P_{(1,1), \omega_j^*} = \frac{1}{2}\rho_B(1 - \rho_S)\rho_S$. Hence, $P_{(1,1), \omega_j^*}/P_{(1,1), \omega_i^*}$ approaches zero if $\rho_S \rightarrow 1$ such that $\lim_{\rho_S \rightarrow 1} \mu^*(\omega_j^*) = 0$.

If $\pi_{m,i}(1, 1, s_i) > \pi_{m,j}(1, 1, s_j)$ and $\pi_{n,i}(1, 1, s_i) < \pi_{n,j}(1, 1, s_j)$, $P_{(1,1), \omega_i^*} = \rho_S(1 - \rho_B)\rho_B$ while $P_{(1,1), \omega_j^*} = \rho_B(1 - \rho_S)\rho_S$. Hence, $P_{(1,1), \omega_j^*}/P_{(1,1), \omega_i^*}$ again approaches zero if $\rho_S \rightarrow 1$ such that $\lim_{\rho_S \rightarrow 1} \mu^*(\omega_j^*) = 0$. ■

Intuitively, if sellers learn much faster than buyers, only the platform that offers higher revenues to sellers will survive with a positive probability if both platforms are active. This induces the following strict Nash-Equilibrium.

Proposition 13. *There exists a $\bar{\rho}_S < 1$ such that for all $\rho_S > \bar{\rho}_S$ it holds: The combination of platforms (s_i^*, s_j^*) with $f_i^* = f_j^* = f_{\min}$ and $\beta_i^* = \beta_j^* = 2^{1/3}$ is a strict Nash equilibrium.*

Proof. As seller's profits decrease in f_i and reach their global maximum in $\beta = 2^{1/3}$ it is clear that $\pi_{m,i}(1, 1, s_i) < \pi_{m,i}(1, 1, s_i^*)$ for any $s_i \neq s_i^*$. But then Lemma 12 indicates that for any $\epsilon > 0$ there is a $\bar{\rho}_S$ such that $\mu^*(\omega_i^*, s_i', s_j^*) < \epsilon$ for all $\rho_S > \bar{\rho}_S$. Hence, if ϵ is chosen small enough, it holds for all $\rho_S > \bar{\rho}_S$ that

$$\pi_{D,i}(s'_i, s_j^*) < \pi_{D,i}(s_i^*, s_j^*). \quad (1)$$

In this case, (s_i^*, s_j^*) is a strict Nash equilibrium. ■

In general, if ρ_S is sufficiently large, it can be shown that $\beta_i \leq 1$ is not chosen by any designer in any pure strategy equilibrium. Moreover, in any mixed strategy equilibrium (σ_1^*, σ_2^*) there is at least one designer i where $s_i \in C(\sigma_i^*)$ implies that $\beta_i > 1$.²²

The previous proposition shows that also for decreasing returns to scale competition between market designers might lead to the design of non-market clearing institutions. This result, however, is not independent of the details of the learning model and the demand and supply conditions.

²²A proof of these claims and more detailed exposition of the material discussed in the Appendix is available on request.