

ON THE EVOLUTION OF MARKET INSTITUTIONS: THE PLATFORM DESIGN PARADOX*

Carlos Alós-Ferrer, Georg Kirchsteiger and Markus Walzl

We study competition among market designers who create new trading platforms, when boundedly rational 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 compete, we find that traders 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 market designers to introduce non-market clearing platforms. Hence platform competition induces non-competitive market outcomes.

Markets are not only characterised 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 that these characteristics are crucial for the resulting trading outcome and for the realised prices. Since the impact of the trading rules on market outcomes is difficult to investigate with time-series of real-life market data, for a more detailed discussion see Friedman (1993), empirical evidence mainly relies on laboratory experiments, for an overview of the evidence see, e.g., Plott (1982) and Holt (1995); in the context of financial markets see also Friedman (1993). 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 (Plott and Smith, 1978). As a consequence, some gains of trade are not realised on these trading platforms and inefficiencies occur due to the design of the trading platform. In a similar way, Dutch or first-price auctions are notorious for inducing overbidding and creating inefficient allocations compared to second-price formats (Kagel, 1995). In a field study, Roth and Ockenfels (2002) show that fixed ending-rules (‘hard-close’) in online auctions lead to late bidding (‘sniping’); see also Ockenfels and Roth (2006). A laboratory experiment by Ariely *et al.* (2005) confirmed this finding, and also showed that fixed ending-rules lead to lower revenues for the seller (and less efficient allocations) than automatic extensions of the auction (‘soft ending’). All these studies suggest that socially desirable features of market outcomes such as unbiased (market-clearing) prices and efficient allocations are rather sensitive to details of the respective market institution. Moreover, there seems to be a trade-off between efficiency and a price bias in favour of one of the market sides.

* We are grateful to Dan Friedman, Fei Shi, seminar participants in Berlin, UL Brussels, University College London, Louvain-la-Neuve and Maastricht, and an anonymous referee 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).

A remarkable example for the coexistence of a variety of trading institutions is provided by Business to Business (B2B) trading platforms, for a comprehensive analysis see, e.g., Lucking-Reiley and Spulber (2001). Recent decades have 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 (European Commission, 2003). While most of the public and scientific attention is devoted to e-marketplaces targeting consumers (like e-bay or Yahoo), about 95% of the e-commerce is actually B2B (UNCTAD, 2002). In 2004 B2B had an estimated volume of \$1 trillion (*The Economist*, 2004). In contrast with Business to Consumers or Consumer to Consumer platforms, large quantities of relatively standardised products are traded at B2B exchanges. Agents seem to act either as buyers or as sellers on these platforms but not as both (European Commission, 2003). B2B e-commerce is organised in three different ways. The predominant modus in the early days of e-commerce were platforms opened by buyers or sellers (or respective umbrella organisations). An example is MetalSite, a platform organised by steel producers that suspended operations in 2001. Currently, B2B e-commerce is typically organised either as e-procurement¹ (where sellers use standardised software and exchange opportunities offered by platforms such as Ariba or CommerceOne to design and allocate procurement contracts) or via institutions operated by a third party (this holds, e.g., for CheMatch – a trade platform for chemical products – or for a large part of the product portfolio offered at EnronOnline – a multi-commodity exchange run by Enron until 2002). Of all firms active on B2B platforms, about one third operates on such platforms run by third parties (European Commission, 2003). Both e-procurement software and market designs of third parties show a variety of institutional arrangements. EnronOnline, for instance, was organised as a posted offer market while competing platforms such as AltraEnergy (or on the chemical sector CheMatch) are exchange platforms that work like double auctions. The software solutions offered by Ariba and CommerceOne include various institutional arrangements such as Dutch auctions or proxy-bidding (with hard and soft ending rules).² Platform designs seem to exist with different propensities to generate market clearing outcomes in B2B e-commerce (see the experimental literature cited above).

Given the variety of different market institutions, and the variety of their efficiency properties, one wonders which type(s) of trading institutions will be observed in the long run. In particular, our article investigates whether institutions promoting efficient, market clearing outcomes will dominate less efficient trading platforms in the long run. To answer this question, we are led to investigate the evolution of market institutions. It is useful to distinguish between two aspects of this evolution, namely the selection between existing institutions by the traders and the emergence of new 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 focus on

¹ For a recent discussion of the adoption of e-procurement in B2B and an overview of market designs see Davila *et al.* (2003).

² Interestingly, CommerceOne applied for Chapter 11 bankruptcy in 2004 and was bought out later. In general, entry and exit are still a frequent phenomenon in the market for B2B platforms, suggesting that this industry is still at a relatively early stage of development. However, business analysts identify a development from the creation of new platforms (the common business model in the 1990s) to buy-outs, which indicates some degree of maturity (Keys, 2002).

market platforms introduced on purpose.³ If a trading platform is introduced by a market designer who demands user fees, the design of a new platform by the designer and the selection among existing ones by the traders 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 article we analyse 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.

Trading platforms are created by profit-maximising, risk-neutral market designers. The designers compete with each other through platform designs. Each designer chooses a trading fee that he demands from the traders for the use of his platform. To capture the trade-off between efficiency and a price bias for one market side that has been observed in laboratory studies (see above), we allow designers to choose platform designs with systematic price biases, above or below the market clearing price. Hence, through the trading fee, each designer decides upon his share of the surplus created through trade at his platform. But he can also favour one type of trader with the introduction of a price bias. 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 favoured type of trader, which in turn may enhance the platform's survival probability.

To analyse 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, traders decide which platform they want to be active on (for the monopolistic case, there is of course no real choice – traders just trade on the only existing platform). The role of a trader (buyer or seller) is exogenously given. Sellers are assumed to be firms with a constant returns to scale production technology.⁴ Buyers are characterised 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 no trader has an incentive to deviate from this platform – independently of the design alternatives offered by the competing platform. If traders were fully rational, we would have established a standard two stage game (Stage 1: Market Design; Stage 2: Traders' platform choice) with (network) externalities in the Stage 2-subgame. Such a game typically exhibits a multiplicity of (subgame perfect) equilibria. As in the battle of sexes, coordination of all traders on each platform is clearly an equilibrium of the 2nd stage, next to a mixed strategy equilibrium where traders are indifferent between platforms. To select among these equilibria, we drop the assumption of fully rational traders and instead assume that traders are boundedly rational but may *learn* to coordinate on a particular platform. Following the game-theoretic learning literature (Young, 1993; Kandori *et al.*, 1993; Ellison, 2000), we use a Markovian model to analyse the platform choice of the traders. We assume that the traders' behaviour depends on the market outcomes generated by the different platforms and thereby on the characteristics of all feasible platforms. We are interested

³ For an analysis of markets as a by-product of traders' actions, see Kirchsteiger *et al.* (2005).

⁴ In Appendix B we investigate the robustness of our results with respect to decreasing returns to scale.

in the long-term properties of this learning process, i.e. in its (limit) invariant distribution. This 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 designers to introduce platforms that are not market clearing but that have a price bias in favour 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 article is related to three strands of the literature. First, since we investigate the role of trading platforms with exogenously given buyers and sellers, our article is to some extent related to the two-sided markets literature; see Rochet and Tirole (2006) for an overview. This literature is based on the assumption of network externalities. It analyses the impact of these externalities and of platform competition on the structure of the fees demanded by the market designers (Armstrong, 2006; Belleflamme and Toulemonde, 2004; Caillaud and Julien, 2003; Rochet and Tirole, 2003). 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 behaviour 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 internalised 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.⁵

Second, our article is also related to the literature on competition between exogenously given trading institutions. Ellison and Fudenberg (2003) and Ellison *et al.* (2004) analyse the circumstances under which 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.* (2006) investigate the case of centralised versus decentralised trading institutions. All of these papers rely on the assumption of rational traders and do not allow for learning. In terms of traders' behaviour, the learning model of Gerber and Bettzüge (2007) is relatively close to our article. But since they focus on the possibility of multiplicity of active trading platforms, they consider neither non-market-clearing platforms nor market designers. The paper most closely related to the one at hand is that of Alós-Ferrer and Kirchsteiger (2008), which also analyses the learning behaviour 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.

⁵ Rochet and Tirole (2006) define two-sided markets by the non-neutrality of the fees. In their terminology we model a one-sided market.

In our model, rational market designers are confronted with boundedly rational, learning traders.⁶ Hence, our article belongs to a small but growing literature that we 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 optimise their behaviour, often relying on behavioural rules of thumb instead. However, large firms, market designers, etc. can be taken as comparatively sophisticated. Schlag (2004), Gabaix and Laibson (2006), Hopkins (2007) and Spiegler (2006) apply this approach to the analysis of industries facing boundedly rational consumers. See Ellison (2006) for an overview of this literature.

The article is organised as follows. Section 1 presents the basic model. Section 2 discusses the traders’ platform choice of the traders. Section 3 analyses the design of the platform. Section 4 concludes. All proofs are in Appendix A. In Appendix B we analyse the robustness of our results with respect to boundedly rational designers and with respect to decreasing returns to scale in production.

1. The Model

We study the trade of a homogenous good at different market platforms, which are set up by profit-maximising market designers. For simplicity, we restrict our attention to two competing market designers (referred to as competitive market design). As a benchmark, we also analyse the case where only one market designer can set up a trading platform (referred to as monopolistic market design). In this Section, we introduce the trading rules that are at the market designer’s disposal and analyse trade and profits for a given choice of trading rules by the designers and a given platform choice by buyers and sellers.

1.1. Market Platforms’ Design

Before trade takes place, 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 characterise 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 the actual price 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

⁶ In Appendix B we show the robustness of our results with respect to learning designers.

the grid. To understand why institutions with different price biases are feasible for the designers, recall the experimental and empirical results mentioned in the introduction. In our framework posted offer markets or first price auctions are characterised by $\beta > 1$, posted bid markets or proxy-auctions with 'hard-close' by a $\beta < 1$, while double auctions can be represented by $\beta = 1$. We refer to the platform with $\beta = 1$ as the market-clearing platform and we assume that such a platform is always feasible. $|B|$ denotes the number of feasible biases.

After the platforms are set up, traders will use their observations and experience to learn eventually which platform to use. Formally, we analyse 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 are a fixed share of the revenue generated by trade on i 's platform.⁷ 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. 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 $0 < f_{\min} < f_{\max} < 1$.⁸ $|F|$ 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 realised price at which 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 $S = B \times F$.

1.2. Traders

The good is supplied by a finite set M of at least two profit-maximising firms (sellers) that use the same constant returns to scale technology with marginal costs of $c > 0$.⁹ When deciding upon his supply, a seller takes into account the trading fee of the platform at which he operates. Hence, sellers supply a strictly positive but finite quantity if and only if the price net of trading fee is equal to c .

As shown later, 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 this case, we obtain results that are robust to the details of the learning process. In Appendix B.2. 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

⁷ Our results would not change if we assume quantity-dependent charges instead of revenue-dependent charges.

⁸ The assumption that the fees are strictly positive can be justified by (unmodelled) setup costs for the market designers.

⁹ The assumption of identical sellers might seem restrictive at the first sight. Within our framework, firms without access to the lowest cost technology would face zero market demand. Hence, our assumptions only rule out the case where exactly one firm has access to the lowest-cost technology.

returns to scale case can be replicated also for strictly decreasing returns to scale, but not for the whole class of learning models we analyse 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 $n \in N$.¹⁰

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. To see this, note that 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 is given by $p_i^*(s_i) = c/(1 - f_i)$. The realised price at which trade is conducted on platform i is then

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

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 despite both types of traders being present on platform i .

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

$$D_{N_i}(p) = \sum_{n \in N_i} d_n(p) \tag{2}$$

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

$$q_{n,i}(N_i, M_i, s_i) = \begin{cases} d_n\left(\beta_i \frac{c}{1 - f_i}\right) & \text{if } i \text{ is active,} \\ 0 & \text{otherwise,} \end{cases} \tag{3}$$

$$q_{m,i}(N_i, M_i, s_i) = \begin{cases} \frac{1}{|M_i|} D_{N_i}\left(\beta_i \frac{c}{1 - f_i}\right) & \text{if } i \text{ is active,} \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

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,

¹⁰ Our 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.

whenever a buyer trades a strictly positive quantity, he is not rationed at all. It is thus 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. by¹¹

$$\pi_{n,i}(s_i) = \begin{cases} \frac{1}{p_i} = \frac{1-f_i}{\beta_i c} & \text{if } i \text{ is active,} \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

If both platforms i and j are active (i.e. positive amounts are traded),

$$p(s_i) < p(s_j) \iff \pi_{n,i}(s_i) > \pi_{n,j}(s_j). \quad (6)$$

This implies in particular that if $\beta_i = \beta_j = 1$ and $f_i < f_j$, then $\pi_{n,i}(s_i) > \pi_{n,j}(s_j)$.

The sellers' evaluations 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 equally rationed. Hence, the sellers' evaluation of platform i is given by

$$\pi_{m,i}(N_i, M_i, s_i) = \begin{cases} \left[\frac{1}{|M_i|} D_{N_i} \left(\beta_i \frac{c}{1-f_i} \right) \right] (\beta_i - 1)c & \text{if } i \text{ is active,} \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Note that for $\beta_i > 1$ the sellers' profits are strictly positive provided that platform i is active. On the other hand, for $\beta_j = 1$, $\pi_{m,j}(N_j, M_j, s_j) = 0$ irrespective of whether platform j is active or not. So as long as there is an active non-market clearing platform, its outcome is always strictly better for the sellers than the outcome of a market clearing platform. That is, for all f_i, f_j

$$|N_i| > 0, |M_i| > 0 \text{ and } \beta_i > \beta_j = 1 \Rightarrow \pi_{m,i}(N_i, M_i, s_i) > \pi_{m,j}(N_j, M_j, s_j). \quad (8)$$

2. The Traders' Platform Choice

In our model, market designers first choose their platforms' characteristics and then buyers and sellers decide which platform to join. If there is only one market designer, traders' choices are 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. In this Section, we introduce the learning process and analyse long-run trading patterns and platform revenues for a given configuration of designs.

¹¹ We do not use this particular payoff function. If demand is derived from utility maximisation, though, the realised (indirect) utility must be a strictly monotone transformation thereof.

2.1. *The Learning Process*

We consider a social learning process defined by

- (i) the information available to each trader,
- (ii) the way traders revise their platform choices whenever they have the opportunity,
- (iii) the opportunities to revise a platform choice, and
- (iv) the way traders make mistakes when choosing a platform.

The information available to a trader is not only his own experience (as it would be in a reinforcement learning model, for instance). Rather, each trader observes the prices and the quantities of *both* platforms (including the observation of the inactiveness of a platform). 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 behaviour of the other traders. Hence, individual traders cannot accurately predict the future outcomes of the platforms. Furthermore, they 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 behaviour 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 behaviour is close to the 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 behaviour described does not require the observation of any evaluation conducted by other traders but merely the observation of prices and traded quantities in both platforms.

With this learning rule the switching decision of each trader depends on the trading outcomes of both platforms in the last trading round. These trading outcomes depend on the distribution of the traders over both platforms. Hence, the distribution of traders over platforms depends on the last period's distribution of traders over platforms. A state ω specifies which trading platform is chosen by each buyer and each seller. The state space is given by $\Omega = \{1,2\}^{|N|} \times \{1,2\}^{|M|}$, and trader k 's platform choice in state $\omega \in \Omega$ is denoted by $\omega(k) \in \{1,2\}$. The following notation will prove convenient:

$$N_i(\omega) = \{n \in N \mid \omega(n) = i\} \tag{9}$$

$$M_i(\omega) = \{m \in M \mid \omega(m) = i\} \tag{10}$$

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 . By definition, all those traders who are not on platform i have to be on the other platform j .

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 .

2.1.1. *Unperturbed learning process*

We first concentrate on the unperturbed learning process, where traders switch platforms only because of learning but not because of experimentation (experimentation is introduced in the next subsection). 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 chooses the platform whose outcome he evaluates as best. In case of indifference, he stays with his old platform.

Whenever trader k receives a revision opportunity at period t , he will choose the platform with the period $t - 1$ outcome that he evaluates highest. If, by chance, the outcomes of both platforms are equally evaluated, the trader sticks to his former platform choice. For instance, in the case in which one platform is inactive and the other is active but yields zero profits for the sellers, sellers do not switch. This assumption could be justified by small but positive switching costs.¹²

But when are agents allowed to revise their choices? It is common in learning models to introduce some inertia explicitly 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 specific formulation, here we follow Alós-Ferrer and Kirchsteiger (2008) and postulate a general class of dynamics encompassing the standard examples (and many others).¹³ This general dynamics is defined by the following assumptions.

ASSUMPTION B1. For every agent k and state ω there is strictly positive probability that agent k is the only trader of his own market side who is able to revise his platform choice.

Notice that B1 implies that every agent has a strictly positive probability of being able to revise at any given state. It also allows the revision probability to depend on the state ω and on the identity of the trader, k .

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 considered as an anonymity requirement).

ASSUMPTION B2. For every agent k and state ω , if k is the only one of his own market side who is allowed to revise, either no trader of the other market side is able to revise his platform choice, or there is a strictly positive probability for each trader of the other market side to be allowed to revise.

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

¹² See Oechssler (1997) for a discussion. In Alós-Ferrer *et al.* (2006), we investigate a different tie-breaking rule. Indifferent traders randomise their platform choice, i.e., every platform is chosen with a strictly positive but not necessarily identical probability. All our results (in particular Theorem 10) are robust towards such a modification of the tie-breaking rule.

¹³ See Alós-Ferrer (2003) for a discussion. Learning processes fulfilling B1 correspond to 'regular' learning processes in Alós-Ferrer and Netzer (2007).

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 (1994); Kandori and Rob (1995)) or assumed to arrive in an asynchronous way (a case we term asynchronous learning; see Blume (1995), Binmore and Samuelson (1997) and Benaïm and Weibull (2003)).¹⁴ That is, our formulation covers the following standard examples.

Independent Inertia. For each agent k and each state ω there is an exogenous, equal, independent and strictly positive probability $\rho < 1$ that agent k does not get a revision opportunity.

Asynchronous Learning. Each period, only one agent (i.e. either a buyer or a seller) is (randomly) selected and allowed to revise his strategy.

Asynchronous Learning within Types. In our case, it is natural to conceive of a dynamic where, in every period, only one buyer *and* one seller are selected (randomly and independently) and given the opportunity to revise.

Obviously, B1 and B2 are fulfilled by all these types of learning. The specification above allows for more general learning processes than those described by independent inertia or asynchronous learning. Since the revision probability is allowed to be a function of the state ω , it might depend, for example, 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. The transition matrix is given by $P^0 = [P^0(\omega, \omega')]_{\omega, \omega' \in \Omega}$. An *absorbing set* of the unperturbed dynamics is a minimal subset of states which, once entered, is never abandoned. An *absorbing state* is an element which forms a singleton absorbing set, i.e. $P^0(\omega, \omega) = 1$.¹⁵

As a first step in the analysis of long-run trading patterns, we determine the absorbing sets of the unperturbed learning dynamics. Depending on the design of the two platforms, there exist multiple such absorbing sets. The reason is that no trader ever switches to a platform which does not have an agent of each market side and/or has a bias below 1 and is therefore inactive. Moreover, indifferent traders do not switch. In particular, sellers never switch to a market-clearing platform as it does not offer a positive profit for them. These considerations lead to the following results.

LEMMA 1. *Assume A, B1, and B2. Let $i \neq j$, $\Omega_i^B = \{\omega | N_i(\omega) = N\}$ and $\Omega_j^B = \{\omega | N_j(\omega) = N\}$. All absorbing sets of the unperturbed dynamics are singletons. Depending on platforms' characteristics, the absorbing states are as follows.*

¹⁴ The 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 sensitive to the details of the dynamics.

¹⁵ Our analysis of Markov chains as defined by the learning dynamics uses the methods and concepts introduced in Kandori *et al.* (1993) and Young (1993). Detailed overviews can be found e.g. in Ellison (2000), Fudenberg and Levine (1998) or Samuelson (1997).

- (a) If $\beta_i > 1$ and $\beta_j > 1$, the monomorphic states ω_k^* ($k = i, j$) such that $N_k(\omega_k^*) = N, M_k(\omega_k^*) = M$ and every state in $\Omega_0 = \{\omega \mid \pi_{m,i} = \pi_{m,j}, \pi_{n,i} = \pi_{n,j}\}$.¹⁶
- (b) If $\beta_i > 1$ and $\beta_j = 1$, the monomorphic state ω_i^* , the cross-state ω_i^0 with $N_i(\omega_i^0) = N, M_i(\omega_i^0) = \emptyset$, and every state in Ω_j^B (which includes ω_j^* and ω_i^0).
- (c) If $\beta_i = \beta_j = 1$, the elements in Ω_i^B and Ω_j^B , plus, if and only if $p(s_i) = p(s_j)$, every state with two active platforms.
- (d) If $\beta_i > 1$ and $\beta_j < 1$, the monomorphic state ω_i^* , and all states in which platform i is inactive (for $\beta_j < 1$, platform j is always inactive).
- (e) If $\beta_i = 1$ and $\beta_j < 1$, the elements of Ω_i^B and all states in which platform i is inactive.
- (f) If $\beta_i < 1$ and $\beta_j < 1$, all states $\omega \in \Omega$.

2.1.2. Perturbed learning process

In order to select among the multiple absorbing states, we now turn to the analysis of the stability properties of the platforms with respect to experimentation. The dynamics are enriched with a perturbation in the form of experiments (or mistakes) in the following way. 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,¹⁷ 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) and there is a unique probability distribution over states $\mu_\varepsilon \in \Delta(\Omega)$ which, if taken as initial condition, would be reproduced in probabilistic terms after updating (more precisely, $\mu_\varepsilon P^\varepsilon = \mu_\varepsilon$). This μ_ε is called the *invariant distribution* of P^ε . For the perturbed dynamics P^ε the *limit invariant distribution* $\mu^* = \lim_{\varepsilon \rightarrow 0} \mu_\varepsilon$ exists and is an invariant distribution of the unperturbed process P^0 (see e.g. Kandori *et al.* 1993; Young 1993; Ellison 2000). It 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. Thereby $\mu^*(\omega)$ is the probability that (for small ε) the process will be in state ω in the long-run. The states in the support of μ^* , i.e. $\{\omega \in \Omega \mid \mu^*(\omega) > 0\}$ are called *stochastically stable states* or long-run equilibria. The set of stochastically stable states is the union of some absorbing sets of the original, unperturbed chain ($\varepsilon = 0$).

We call a platform *active in the long-run* if there is a positive probability for trade at this platform in the long-run, i.e., if there is a stochastically stable state with platform i being active.¹⁸

¹⁶ Ω_0 is non-empty as it always contains cross-states ω_i^0 ($i = 1, 2$) (see (b)).

¹⁷ We mean that an institution is picked up according to a pre-specified probability distribution having full support. The exact distribution does not affect the results, as long as it has full support and does not depend on ε .

¹⁸ In the following, whenever we say absorbing sets or states, we refer to the unperturbed dynamics. Since the perturbed dynamics is irreducible, no confusion should arise.

THEOREM 2. Assume A, B1, and B2.

- (a) Suppose $\beta_i < 1$. Then, platform i is not active in the long run.
- (b) Suppose $\beta_i = \beta_j = 1$. Then, platforms i and j are active in the long run.
- (c) Suppose $\beta_i > 1$ and $\beta_j \leq 1$. Then, platform i is active and j is inactive in the long run.

The intuition for this theorem is straightforward. Since there is no trade on a platform with $\beta_i < 1$, it will never be active. Furthermore, on a market clearing platform the sellers' profits are always zero. Hence, sellers do not care at which platform they are if they have to choose between two market clearing platforms. Consequently, both platforms are active in the long run if both are market clearing. Finally, a seller is never worse off at platform i with $\beta_i > 1$ than at platform j with $\beta_j \leq 1$, even if there are no buyers at i . A buyer, on the other hand, is worse off at j than at i when he finds no seller at j . So sellers have an unambiguous tendency to learn to use i , whereas buyers do not always have a tendency towards j . As a consequence, all traders will coordinate on the non-market clearing platform i in the long run. On the level of the platform design, it is thus easy to compete with a market clearing platform by introducing a platform design with a positive price bias.

2.2. The Long-run Trading Patterns

We now proceed to analyse long-run trading patterns (i.e. the stochastic stability of platforms) for a given design configuration $s_i = (\beta_i, f_i)$ and $s_j = (\beta_j, f_j)$.

As a benchmark, we start with the case of identical platform design. To analyse 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, the distribution of traders over the platforms is symmetric in the long run, i.e., $\mu^*(\omega) = \mu^*(\tilde{\omega}) \forall \omega \in \Omega$.

Theorem 2 already identifies the set of long-run active platforms (i.e. stochastically stable states with active platforms) 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 favour sellers (i.e. $\beta_i, \beta_j > 1$). There, Lemma 1(a) implies that full coordination on each platform and states with indifference of both buyers and sellers are the only candidates for stochastically stable states. To pin down stochastic stability, it proves useful to distinguish the two platforms with respect to their prices.

LEMMA 4. Suppose $\beta_i, \beta_j > 1$ and $p_i = \beta_i c / (1 - f_i) < \beta_j c / (1 - f_j) = p_j$. Then,

- (a) only monomorphic states can be stochastically stable.
- (b) ω_i^* is stochastically stable.

According to Lemma 4, the platform with trade at a lower price is always stochastically stable as it is preferred by buyers as long as it is active. The only other candidate for stochastic stability is coordination on the high price institution. 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 $\beta_i \beta_j > 1$, $p_i = \beta_i c / (1 - f_i) < \beta_j c / (1 - f_j) = p_j$, so that ω_i^* is stochastically stable.

- (a) In a dynamics with independent inertia, ω_j^* is also stochastically stable if and only if there is at least one buyer $\tilde{n} \in N$ such that

$$d_{\tilde{n}} \left(\frac{\beta_j c}{1 - f_j} \right) (\beta_j - 1) > \frac{1}{|M| - 1} D_{N \setminus \{\tilde{n}\}} \left(\frac{\beta_i c}{1 - f_i} \right) (\beta_i - 1). \quad (11)$$

- (b) In a dynamics with asynchronous learning, ω_j^* is also stochastically stable if and only if there is at least one buyer $\tilde{n} \in N$ such that

$$\frac{1}{|M| - 1} d_{\tilde{n}} \left(\frac{\beta_j c}{1 - f_j} \right) (\beta_j - 1) > D_{N \setminus \{\tilde{n}\}} \left(\frac{\beta_i c}{1 - f_i} \right) (\beta_i - 1). \quad (12)$$

The condition in part (a) is violated whenever buyers are identical, $|N| \geq |M|$, and the price elasticity of demand is sufficiently high. It can be satisfied for $\beta_j > \beta_i$ whenever buyers are sufficiently heterogeneous (i.e. $\tilde{n} \in N$ 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 asynchronous learning. The condition in part (b) is violated whenever 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 buyers are sufficiently heterogeneous or buyers are identical and demand is sufficiently inelastic.

REMARK 1. The proofs of the previous lemmata (see Appendix A) rely on transition paths involving at most two simultaneous mutations. Thus the speed of convergence is relatively high. Since the number of required mutations does not increase with population size, our dynamics escapes the well-known critique that for large populations the long run may actually be 'too long' to be relevant (Kandori *et al.* 1993; Ellison 1993). We find this point important, because real-world designers are not infinitely long-lived, and some market institutions (e.g. online platforms) appear to go out of business relatively quickly if lacking customers. Since the speed of convergence of the trader-learning process is quick, we think that our results remain relevant.

2.3. Platform Revenues and Designers' Profits

Till now we have analysed the learning dynamics of the traders and the resulting long run pattern of trades. Next we turn to the revenues generated by the platforms, which in turn determine the profits of the market designers.

When analysing 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 consider a platform profitable if it is active in the long run and they ignore revenues made during the adjustment process to the limit invariant distribution.¹⁹ 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 .*

Unsurprisingly, a platform with $\beta_i < 1$ does not generate any profit for the designer as it is always inactive. Hence, we are left with platform configurations (s_i, s_j) where both platforms have a price bias weakly larger than one. In this case expected revenues at platform i depend not only on the design of platform i 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 $p_i = \beta_i [c / (1 - f_i)]$, $p_j = \beta_j [c / (1 - f_j)]$.*

- (a) *If $s_i = s_j$ and $\beta_i \beta_j \geq 1$, 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 $\pi_{D,j}(s_j, s_i) = 0$.*

The first part of this Lemma shows that for identical platforms the designers' long-run profits are identical and strictly positive. This follows from the symmetry of the limit invariant distribution (Lemma 3). With two market clearing institutions, none of the platforms can reap all long-term revenues even if the fees differ. Finally and most importantly, when a non-market clearing and a market-clearing platform compete, the designer of the former makes strictly positive profits, whereas the profits of a designer of a market-clearing institution are zero, because all traders coordinate on the non-market clearing platform in the long run (see Theorem 2(c)).

3. The Platform Design

We now compare the design choices by a monopolistic designer and by two competing designers.

3.1. 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 and designer's profits are given by:

¹⁹ Otherwise, market designers' payoffs would depend 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.

$$\pi_D(s) = \begin{cases} f \frac{\beta c}{1-f} D_N \left(\frac{\beta c}{1-f} \right) & \text{if } \beta \geq 1 \\ 0 & \text{otherwise.} \end{cases} \quad (13)$$

It follows that a monopolistic designer introduces a market-clearing platform – as long as the grid of feasible fees is fine enough.

PROPOSITION 8. *Suppose there is only one platform available. Then, the designer chooses a market-clearing platform, i.e. $\beta^* = 1$, if γ is sufficiently small.*

The intuitive reason for this result is as follows. Suppose revenues $pD_N(p)$ are maximised at price p^* . Note that this price can be attained with different (β, f) combinations and that $p^* = \beta c / (1 - f)$ 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 $\beta \geq 1$.

3.2. 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 their payoff functions.²⁰ That is, both designers choose their platforms simultaneously and payoffs are given as in Section 2.3. The sets of pure strategies of designer i and j are given by $S_i = S_j = B \times F$. 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 i is

$$\pi_{D,i}(\sigma_i, \sigma_j) = \sum_{s_j \in S} \sum_{s_i \in S} \sigma_j(s_j) \sigma_i(s_i) f_i E R_i(s_i, s_j). \quad (14)$$

Since the sets of pure strategies are finite, a Nash equilibrium of the designers' game always exists (possibly in mixed strategies). To characterise these equilibria, we need the following Lemma.

LEMMA 9. *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 $\beta_i \geq 1$.*

Lemma 9 is an immediate consequence of the fact that there is no trade at a platform with $\beta < 1$. Therefore, only platforms weakly biased in favour of the sellers will be chosen in equilibrium. We now 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 result holds as long as the grid of feasible biases is fine enough.

THEOREM 10. *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.*

²⁰ As shown in Appendix B.1, our main results do not change if we consider boundedly rational market designers who learn the same way as traders do.

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. As we have seen, competition between a market clearing and a non-market clearing platform leads to full coordination of the traders on the latter platform (see Theorem 2(c)). Hence, designers do not introduce a market clearing platform when facing platform competition.

In general, nothing more can be said about the specific characteristics of the Nash equilibria. A brief examination of Lemma 5 should convince the reader that a full characterisation 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 hypothesis 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 hypothesis is in general false. For simplicity, consider identical buyers with a demand function $d(p)$ and denote the price elasticity of demand by $\varepsilon_p(p) = -pd'(p)/d(p)$.

PROPOSITION 11. *Assume independent inertia, identical buyers and $|M| = |N|$. If δ and γ are sufficiently small and ε_p is not much larger than one, then there exists no Nash equilibrium (σ_i^*, σ_j^*) (neither in pure nor in mixed strategies) where both designers introduce only platforms with $\beta_i = \beta_j = 1 + \delta$.*

Beyond the features highlighted in Theorem 10, equilibrium designs are rather sensitive to details of the economy and the learning process. It cannot be excluded that in equilibrium designers choose ‘near market-clearing’ platform characteristics for some specifications of the learning dynamics and/or of the demand functions. But in general the equilibrium choices are not ‘near market-clearing’ platforms.

4. 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 maximise his profits. Hence, we derive the paradoxical result that platform competition induces non-competitive market outcomes. This result could also explain why so many B2B platforms exhibit institutional designs that are notorious for biased (non-market-clearing) prices (e.g. posted offer markets, proxy auctions with ‘hard-close’ or Dutch auctions).

The results of our article depend of course on our key assumptions. First, we have focused on boundedly rational traders who choose platforms myopically. For fully rational designers and traders, our set-up would correspond to a two stage game

where, in stage 1, designers (simultaneously) choose platform designs, and in stage 2, traders coordinate on platforms. The second stage thereby resembles a coordination problem or a game with network externalities. Naturally, this structure induces a multiplicity of equilibria – in particular, there is a subgame perfect equilibrium where all traders coordinate on a market-clearing platform (with monopolistic trading fees). Our analysis of the coordination problem in stage 2 using a learning dynamics can be interpreted as an equilibrium selection device, which rejects the above-described equilibrium and selects only configurations with non-market-clearing platform designs.

Second, we assumed asymmetric rationality in the sense that designers are more sophisticated than traders. Furthermore, by focusing on long-run profits we have implicitly assumed that it is much harder for designers to change the properties of their platforms than for traders to switch trading platforms, or that platform providers are much more patient than traders. Cases such as the downfall of Enron Online or the bankruptcy of CommerceOne illustrate, however, that this assumption may not be fully justified. Sometimes platform providers indeed suffer rather quickly from a lack of traders and are removed from the market at short notice. But our modelling framework can cope with cases where platform designers and traders revise their decisions with equal speed. In Appendix B we investigate the case of boundedly rational designers who have to learn how to design a platform through a regular (trial-and-error) design revision process (i.e., designers are as myopic or impatient as traders). Our main results carry over to such a setting. Hence, while asymmetric rationality is a crucial ingredient of our model, it is not the driving force behind the emergence of non-market clearing institutions. Furthermore, our results also hold when the traders' learning process shows a relatively high speed of convergence. In real life this would imply that a market designer can enter and remain in the market with a superior design, because traders can coordinate on the respective platform at short notice. In this sense our model allows for the 'free entry' of superior platform designs and still non-market clearing institutions emerge in equilibrium.

Third, 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 less strong and a characterisation of the limit invariant distribution requires both a further specification of the learning behaviour of the traders and a further specification of demand and supply. In Appendix B we provide 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. Nonetheless, this clearly illustrates that the scope of the paradox identified here goes beyond the constant returns to scale case.

These robustness checks show that neither the assumption of constant returns to scale 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.

Appendix A. Proofs

The proofs rely on the well-known characterisation of the invariant distribution of a Markov Chain developed by Freidlin and Wentzell (1984) and its implications for stochastic stability as discussed by Kandori *et al.* (1993) and Young (1993), which we briefly review here. Lemma 3.1 from Freidlin and Wentzell (1984) states the following. Fix an $\omega \in \Omega$. An ω -tree T is a tree in Ω with root ω , i.e. a graph on Ω such that for every state $\omega' \neq \omega$ 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_ω is the product of all transition probabilities on a given ω -tree summed over all ω -trees). Then, the invariant distribution for fixed ε is given by $\mu_\varepsilon(\omega) = q_\omega / \sum_{\omega \in \Omega} q_\omega$. Relying on this result, Kandori *et al.* (1993) or Young (1993) observe that in the limit as $\varepsilon \rightarrow 0$, $\mu^*(\omega)$ is determined by those ω -trees that imply the smallest possible number of mutations necessary to form a tree in Ω . Given two states ω and ω' , let $c(\omega, \omega')$ (the *transition cost* from ω to ω') denote the minimal number of experiments necessary for a transition (or link) from ω to ω' along a positive probability path starting in ω and leading to ω' . The cost of an ω -tree is the sum of all costs along links in it. Let $\gamma(\omega)$ (the *stochastic potential* of ω) be the minimal cost of an ω -tree. A state ω is stochastically stable if and only if its stochastic potential is minimal, i.e. $\gamma(\omega) \leq \gamma(\omega')$ for all $\omega' \in \Omega$. Let A be the set of absorbing sets. If $X \in A$, all states in X have the same stochastic potential, denoted $\gamma(X)$.

Proof of Lemma 1

(a) Let $\beta_1 > 1$ and $\beta_2 > 1$. 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. Moreover, elements of Ω_0 are absorbing because traders do not switch if the respective profits are identical on both platforms. It is now enough to show that there exists a positive probability path from any other to a monomorphic state or a state in Ω_0 .

Consider $\omega \notin \Omega_0 \cup \{\omega_i^*\} \cup \{\omega_j^*\}$. At least one platform has to be active. If only platform i is active, the monomorphic state ω_i^* is reached with positive probability because buyers and sellers receive positive profits at platform i and zero profits at platform j . Hence, suppose that both platforms are active. Suppose that $p(s_i) \neq p(s_j)$ and (without loss of generality) that $p(s_i) < p(s_j)$. Then, buyers strictly 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 platform j receive revision opportunity. Then, 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. Now suppose that $p(s_i) = p(s_j)$. Buyers are indifferent and will never switch. As $\omega \notin \Omega_0$, sellers prefer one platform and there is a positive probability path to a state with an inactive platform or a state in Ω_0 .

(b) Let $\beta_j = 1$ and $\beta_i > 1$. Monomorphic states are absorbing, because at the corresponding platform both buyers and sellers make weakly positive profits and the other platform is inactive. Thus traders do not switch. Cross states are absorbing because traders do not switch if profits are identical on both platforms. Finally, sellers receive zero profits at platform j . Hence, they are indifferent between an inactive platform i and platform j , so they never switch. Buyers strictly prefer an active platform j to an inactive platform due to strictly positive profits. Therefore, all states in $\Omega_j^B = \{\omega | N_j(\omega) = N\}$ are absorbing.

It remains to show that there exists a positive probability path from any $\omega \notin \Omega_j^B \cup \{\omega_i^0\}$ to a monomorphic state or a state in Ω_j^B . In such a state ω , at least one platform has to be active. If only platform i is active, buyers and sellers strictly prefer i to j and ω_i^* is reached with positive probability. If only platform j is active, buyers strictly prefer j to i while sellers do not switch at all. Hence, a state in Ω_j^B is reached with positive probability. If both platforms are active, sellers strictly prefer i and, by B1 and B2, there is positive probability that all sellers at j but only buyers at

i receive revision opportunities in successive periods. Hence, sellers will switch away from j and no new buyer will switch to j . Thus, either the monomorphic state ω_i^* is eventually reached, or j becomes inactive.

(c) Let $\beta_i = \beta_j = 1$. Then, sellers never switch. Suppose first that $p(s_i) \neq p(s_j)$. Buyers strictly prefer an active to an inactive platform and are indifferent between two inactive platforms. Hence, states in $\Omega_i^B = \{\omega | N_i(\omega) = N\}$ and $\Omega_j^B = \{\omega | N_j(\omega) = N\}$ are absorbing. It is enough to show that there exists a positive probability path from any state outside $\Omega_j^B \cup \Omega_i^B$ to a state inside this joint set. Consider $\omega \notin \Omega_j^B \cup \Omega_i^B$. In ω , at least one platform has to be active. If only one is active, buyers strictly prefer this platform to the other. Hence, a state in Ω_i^B or Ω_j^B (depending on which was the active platform) is reached with positive probability. If both platforms are active, suppose without loss of generality that $p(s_i) < p(s_j)$. Then, buyers strictly prefer platform i and there is a positive probability path to an element in Ω_i^B .

Consider now the case $p(s_i) = p(s_j)$. States in Ω_i^B and Ω_j^B are absorbing, because buyers strictly prefer an active to an inactive platform. Also, every state with two active platforms is absorbing, because buyers are indifferent between active platforms. Consider a state ω with exactly one active platform (i , say) which is not in Ω_i^B . Then, buyers strictly prefer i to j and there is a positive probability path to a state in Ω_i^B .

(d) Let $\beta_i > 1$ and $\beta_j < 1$. Then, platform j is always inactive. If platform i is active, buyers and sellers strictly prefer platform i to platform j . Hence, ω_i^* is absorbing. If platform i is inactive, buyers and sellers do not switch at all. Hence, every state with an inactive platform i is absorbing.

To complete the proof consider a non-monomorphic state ω with an active platform i . Then, buyers and seller strictly prefer platform i to platform j and there is a positive probability path to ω_i^* .

(e) Let $\beta_i = 1$ and $\beta_j < 1$. Then, sellers never switch and platform j is always inactive. If platform i is active, buyers strictly prefer i to j . Hence, every state in Ω_i^B is absorbing. If platform i is inactive, buyers and sellers do not switch at all. Hence, every state with an inactive platform i is absorbing. Consider a state $\omega \notin \Omega_i^B$ with an active platform i . Then, buyers strictly prefer platform i to platform j and there is a positive probability path to a state in Ω_i^B .

(f) If $\beta_i, \beta_j < 1$, neither buyers nor sellers switch, hence every state is absorbing.

Proof of Theorem 2

(a) For $\beta_i < 1$, no trade occurs at platform i independent of the number of buyers or sellers at platform i .

(b) Let $X \in A$, then $\gamma(X) \geq |A| - 1$ as at least one mistake is needed for a transition between any two absorbing sets.

A transition between any two absorbing states ω, ω' with $\|M_i(\omega) - M_i(\omega')\| + \|N_i(\omega) - N_i(\omega')\| = 1$ is possible with one mistake. Also a transition from a cross-state ω_i^0 to ω_j^* is possible with one mistake: Consider ω_i^0 and a buyer who (by mistake) switches to platform j . As platform i is inactive and buyers receive strictly positive profits at platform j , by Assumptions B1 and B2 there is a positive probability path of the unperturbed dynamics to ω_j^* . Hence $\gamma(X) = |A| - 1$ for all $X \in A$ and every state in A is stochastically stable.

(c) Let $\beta_j = 1$. The absorbing states are ω_i^*, ω_i^0 and the states in Ω_j^B . A transition from ω_i^* to an element in $\Omega_j^B \cup \{\omega_i^0\}$ needs at least two mistakes because buyers and sellers receive strictly positive profits at i . Hence, $\gamma(\omega) > |A| - 1$ for $\omega \in \Omega_j^B \cup \{\omega_i^0\}$. A transition between any two states $\omega, \omega' \in \Omega_j^B$ with $\|M_j(\omega) - M_j(\omega')\| = 1$ is possible with one mistake. A transition from the cross-state ω_j^0 to ω_i^* is possible with one mistake: Consider ω_j^0 and a buyer who (by mistake) switches to platform i . As platform j is inactive and buyers receive strictly positive profits at platform i , by Assumptions B1 and B2 there is a positive probability path of the unperturbed dynamics to ω_i^* . In the same way, one can also construct a positive probability path (with one mistake by a seller) from ω_i^0 to ω_i^* . Hence, $\gamma(\omega_i^*) = |A| - 1$ and ω_i^* is the only stochastically stable state.

Let $\beta_j < 1$. Then, A consists of ω_i^* and the states without active platforms. A transition from ω_i^* to a state with two inactive platforms needs at least two mistakes because buyers and sellers receive strictly positive profits at i . Hence, $\gamma(\omega) > |A| - 1$ for $\omega \in A \setminus \{\omega_i^*\}$. A transition between any two states $\omega, \omega' \in A \setminus \{\omega_i^*\}$ with $\|M_i(\omega) - M_i(\omega')\| + \|N_i(\omega) - N_i(\omega')\| = 1$ is possible with one mistake. A transition from the cross-state ω_j^0 to ω_i^* is possible with one mistake as well, exactly as in the case $\beta_j = 1$. Hence, $\gamma(\omega_i^*) = |A| - 1$ and ω_i^* is the only stochastically stable state.

Proof of Lemma 3

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).

Proof of Lemma 4

If $p_i < p_j$, buyers strictly prefer platform i to platform j whenever it is active. Lemma 1(a) implies that the states $\omega_i^*, \omega_j^*, \omega_j^0$ and ω_j^1 form the only absorbing sets. A single experiment suffices for a transition from a cross state to a monomorphic state: consider without loss of generality ω_i^0 and suppose a buyer switches (by mistake). Then, platform j is active and platform i is not. Buyers and sellers strictly prefer platform j and there is a positive probability path to ω_j^* . In contrast, at least two experiments are necessary for a positive probability path from a monomorphic state to another monomorphic state or a cross-state: consider without loss of generality ω_i^* . As long as only one trader switches (by mistake), platform j remains inactive and platform i is strictly preferred by buyers and sellers. As a consequence, $\gamma(\omega_i^0) \geq 5, \gamma(\omega_j^0) \geq 5, \gamma(\omega_i^*) \geq 4$, and $\gamma(\omega_j^*) \geq 4$. There is a positive probability path with two experiments from ω_j^* to ω_i^* : Consider ω_j^* and suppose that a buyer and a seller switch (by mistake) to platform i . Then, platform i is active and buyers strictly prefer platform i to platform j . By Assumptions B1 and B2, there is a positive probability that only buyers and sellers at j receive revision opportunity. But then, ω_i^* is reached with positive probability. Hence, $\gamma(\omega_i^*) < 4 < 5 \leq \gamma(\omega_j^0)$ and $\gamma(\omega_i^*) < 4 < 5 \leq \gamma(\omega_j^1)$ which shows Part (a). Part (b) follows from $\gamma(\omega_j^*) \geq 4$.

Proof of Lemma 5

By the proof of Lemma 4, $\gamma(\omega_i^*) = 4$ if $\beta_i, \beta_j > 1$ and $p_i < p_j$. By Lemma 1(a), only the monomorphic states can be stochastically stable. Hence, a necessary and sufficient condition for the stochastic stability of ω_j^* is $\gamma(\omega_j^*) = 4$.

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), under independent inertia there is 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 switch to j if

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

Hence, this condition is sufficient for the stochastic stability of ω_j^* . To see that it 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 experiments are needed to reach ω_j^* .

Consider part (b) (asynchronous learning). Suppose one seller and buyer \tilde{n} switch to platform j by mutation. By B1 and B2, with positive probability in the subsequent rounds only sellers and buyers at platform i receive the opportunity to revise. If

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

it follows that

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

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. Thus, 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 asynchronous 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).

Proof of Lemma 6

Follows from the fact that trade is never possible on platforms with $\beta_i < 1$ (see Theorem 2(a)).

Proof of Lemma 7

(a) follows from Lemma 3 as the price and the traded quantity in ω at platform i are identical to those in $\tilde{\omega}$ at platform j .

(b) follows from Theorem 2(b). (c) follows from Theorem 2(c).

Proof of Proposition 8

For $\beta < 1$ the profits for a monopolistic designer 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 $\beta \in [1,\infty)$. Denote $p = \beta c / (1-f)$ 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 $\beta \geq 1$)

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

where

$$\frac{\partial p}{\partial f} = \frac{\beta c}{f(1-f)^2} > 0 \text{ and } \frac{\partial p}{\partial \beta} = \frac{c}{1-f} > 0.$$

Let the optimal price be $p^* = \beta^* c / (1-f^*)$. 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 $D_N(p^*) + p D'_N(p^*) < 0$. This implies that

$$\frac{\partial \pi_D}{\partial \beta}(\beta^*, f^*) < 0,$$

hence the designer's profits are maximised at the corner solution $\beta^* = 1$.

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.

Proof of Lemma 9

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 $\sigma_j^*(s_j) > 0$. If designer i chooses with certainty a platform s'_i with $\beta'_i > 1$, Lemma 7(c) implies that $\pi_{D,i}(s'_i, \sigma_j^*) = f'_i [(\beta'_i c)/(1 - f'_i)] D_N [(\beta'_i c)/(1 - f'_i)] > 0$. 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 $s'_i = \bar{s}_j$, $\pi_{D,i}(s'_i, \sigma_j^*) = \sigma_j^*(\bar{s}_j) f_i ER_i(s'_i, \bar{s}_j) + \sum_{s_j \in S \setminus \bar{s}_j} \sigma_j^*(\bar{s}_j) \pi_{D,i}(s'_i, s_j)$. 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.

Proof of Theorem 10

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 $\bar{p} = c/(1 - \bar{f}_i)$. Denote the carrier or support of σ_j^* by $C(\sigma_j^*) = \{s_j = (\beta_j, f_j) \in S \mid \sigma_j^*(s_j) > 0\}$.

Let $\pi_{D,i}(s_i, \sigma_j)$ denote the expected payoff if designer i chooses s_i for sure and j chooses the probability distribution σ_j . 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, 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 enough²¹ player i has an incentive to deviate from σ_i^* to an institution with $\beta'_i > 1$ but close enough to 1, a contradiction.

Proof of Proposition 11

Lemma 5 (a) implies that if $\beta_i c / (1 - f_i) < \beta_j c / (1 - f_j)$, ω_i^* is stochastically stable iff

$$d\left(\frac{\beta_j c}{1 - f_j}\right) (\beta_j - 1) > d\left(\frac{\beta_i c}{1 - f_i}\right) (\beta_i - 1). \tag{A.1}$$

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

²¹ The grid can be assumed to be *ex ante* fine enough by a uniform continuity argument.

(i) $\bar{f}_j > \bar{f}_i$: Condition (A.1) shows that full coordination on platform $\bar{s}_j = (1 + \delta, \bar{f}_j)$ is not stochastically stable *vis-à-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 \hat{f} 's equilibrium strategy, \hat{f} '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 (i) is inconsistent with Nash equilibrium.

(ii) $\bar{f}_i = \bar{f}_j > f_{\min}$. Condition (A.1) shows that full coordination on platform $\bar{s}_j = (1 + \delta, \bar{f}_j)$ is not stochastically stable *vis-à-vis* any platform characteristics in the support of σ_j^* but platform $\bar{s}_j = (1 + \delta, \bar{f}_j)$. Furthermore, Lemma 3 implies 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'_j with $f'_j = \bar{f}_j - \gamma$, and $\beta'_j = 1 + \delta$ implies that

$$\frac{1 + \delta}{1 - f'_j} c < \frac{1 + \delta}{1 - \bar{f}_j} c \text{ and } d \left(\frac{1 + \delta}{1 - \bar{f}_j} c \right) \delta < d \left(\frac{1 + \delta}{1 - f'_j} c \right) \delta.$$

Hence, again by (A.1) $\mu^*(\omega_i^*) = 1$ if s'_j is chosen by i and \bar{s}_j is chosen by j , which yields

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

If the grid of F is fine enough, i.e. if γ is small enough, this implies $\pi_{D,i}(s'_j, \sigma_j^*) > \pi_{D,i}(\bar{s}_i, \sigma_j^*)$, a contradiction with Nash equilibrium.

(iii) $\bar{f}_i = \bar{f}_j = f_{\min}$ - both designers choose $\bar{s}_i = \bar{s}_j = (1 + \delta, f_{\min})$ for sure. Then Condition (A.1) guarantees the existence of a $\beta'_j > 1 + \delta$ and a $f'_j > f_{\min}$ such that platform j is stochastically stable *vis-à-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 ,

$$\text{i.e. } |N| f'_j \frac{\beta'_j c}{1 - f'_j} d \left(\frac{\beta'_j c}{1 - f'_j} \right),$$

is strictly larger than the respective profit from choosing \bar{s}_j ,

$$\text{i.e. } |N| \bar{f}_j \frac{(1 + \delta)c}{1 - \bar{f}_j} d \left(\frac{(1 + \delta)c}{1 - \bar{f}_j} \right).$$

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 . With independent inertia and $|M| = |N|$, this establishes symmetry of the transition matrix P and hence $\mu^*(\omega_i^*) = \mu^*(\omega_j^*) = 1/2$. 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 $\bar{s}_i = (1 + \delta, f_{\min})$.

Appendix B. Robustness of the Results

We have derived our results 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 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 analyse 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). 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 according to 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 revenues differ, he chooses the design which led to a higher revenue (imitation). If both platforms generate zero profits, he randomises with positive probability in the next round over all possible design alternatives (innovation).²² In case of identical positive profits at both platforms designers stick to their former choice (inertia).*

We further assume B1 and B2 (on the enlarged state space and for three instead of two different types of players).²³ The perturbations are defined as in Section 2.1. The perturbed process is again irreducible. We now prove the counterpart of Theorem 2 in the modified learning model.

Proposition 12. Assume A, B1, B2, and C. Then, $\beta_i > 1$ for $i = 1,2$ in every stochastically stable state.

Proof. Throughout this proof, we adopt the convention that for a given platform i the other platform is denoted by $-i$.

First, observe that the monomorphic state ω_i^* with $\beta_i \geq 1$ (and $\beta_{-i} = \beta_i$) is an absorbing state of the unperturbed process. Second, each cross-state with designers randomising over all designs forms a (non-singleton) absorbing set.

A platform i with a positive number of buyers and sellers cannot have a price bias $\beta_i < 1$ in any absorbing set. To see this, consider a state ω with $N_i(\omega) \neq \phi$ and $M_i(\omega) \neq \phi$ and $\beta_i < 1$.

Case 1: Suppose $\beta_{-i} \geq 1$ and platform $-i$ is active. Then, there is a positive probability that i imitates $-i$ (inducing $\beta_i \geq 1$) without any migration of buyers and sellers. Both platforms yield positive profits and have the same design. For this case Assumption C implies that designs can only change if a cross state is reached. But a cross state (with randomising designers) forms an absorbing set. Hence, the unperturbed process never reaches a state with $\beta_i < 1$ and a positive number of buyers and sellers at i again.

Case 2: Suppose $\beta_{-i} \geq 1$ and platform $-i$ is inactive or $\beta_{-i} < 1$. Then, both platforms generate profits of zero and designers randomise. With positive probability, platform i will have $\beta_i > 1$ and generate positive profits for designers, buyers and sellers while $\beta_{-i} < 1$. Then, ω_i^* with $\beta_i > 1$ is reached with positive probability.

Moreover, a state where both platforms are active and one has bias $\beta_i > 1$ and the other bias $\beta_{-i} = 1$ cannot be part of an absorbing set. To see this, observe that from such a state there is always a positive-probability path to ω_i^* because sellers strictly prefer platform i and there is a positive probability that they are the only ones with an opportunity to revise for sufficiently many periods.

²² For simplicity, we assume that designers randomise over S with full support.

²³ This specification also allows for different learning speeds for traders and designers, allowing for example, for the likely situation where buyers and sellers revise with a larger probability than designers.

Hence, absorbing sets are of three kinds. First, states where all active platforms have price biases $\beta > 1$. Denote the set of such absorbing sets by A_0 and the set of all other absorbing sets by A_1 . Second, states where either two active platforms have price bias $\beta = 1$, or the only active platform has price bias $\beta = 1$. Last, *cross states* where no platform has a positive number of buyers and sellers.

To prove the Proposition, we compare the stochastic potential of absorbing sets in A_0 and A_1 . First, observe that it takes at least two mistakes to leave an absorbing set in A_0 . If only one platform i is active, the corresponding absorbing state will be ω_i^* . Because both types of traders receive positive profits, mistakes by both types of traders are needed to induce another active platform. It takes also more than one mistake to reach a cross state. If both platforms are active, designers generate positive profits and designs will only change to $\beta \leq 1$ if a cross state is reached (inertia and imitation do not lead to platforms with $\beta \leq 1$; and designers innovate only in cross-states). However, from a state with two active platforms, it takes more than one mistake by traders to reach a cross-state.

If there is only one active platform i , with bias $\beta_i = 1$, then platform $-i$ also has $\beta_{-i} = 1$ due to imitation. At i all buyers and at least one seller are assembled. The different absorbing states differ only with respect to the number of sellers and can be connected to a (restricted) tree with one mistake per absorbing set (see the proof of Theorem 2). This tree can be connected to the cross states with one mistake. If two active platforms have price bias $\beta = 1$, prices have to be the same at both platforms and sellers do not make any profits on both platforms. The respective absorbing sets can be connected with one mistake (i.e., a switching seller) per absorbing set and the resulting tree can again be connected to the cross states with one mistake. Hence, there is a tree with root in some cross state connecting all absorbing sets in A_1 with one mistake per absorbing set. But a cross state can be left with one mistake towards a monomorphic state in A_0 (due to innovation, there is a positive probability that there is a platform i with $\beta_i > 1$ and, with one mistake, this platform becomes active and is strictly preferred by all traders). Denote by $\tilde{\omega} \in A_0$ the state which minimises the number of mistakes needed to form a tree restricted to A_0 . For a tree in Ω , the minimal cost of a $\tilde{\omega}$ -tree is $|A_1| + x_{\tilde{\omega}}$ where $x_{\tilde{\omega}}$ is the minimal number of mistakes needed to connect all states in A_0 to the respective $\tilde{\omega}$ -subtree. Now consider a state $\hat{\omega} \in A_1$. To construct a minimal-cost $\hat{\omega}$ -tree, we take the $\tilde{\omega}$ -tree, delete the outgoing link of $\hat{\omega}$ (which has cost one), and add a (least-resistance) link from $\tilde{\omega}$ to a state in A_1 (of cost larger than one). Hence, $\gamma(\hat{\omega}) > \gamma(\tilde{\omega})$.

In every stochastically stable state traders are therefore located at a platform i with $\beta_i > 1$. We conclude that boundedly rational platform designers exhibit qualitatively the same behaviour as rational ones: Platform competition forces them to introduce non-market clearing platforms only.

B.2. Decreasing Returns to Scale

Consider the following example. Two identical sellers produce with costs given by $c(q) = (1/2)q^2$. For given prices (p_i) and fees (f_i) at a platform i , their profit is $\pi_{m,i}(q, p_i) = (1 - f_i)p_i q - (1/2)q^2$ and maximisation 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 normalised to 1. The buyers' utility is given by $\pi_n(q, x) = 2\sqrt{q} + x$ and utility maximisation for a given price p_i at the respective platform 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 $r_i(\omega) = |N_i(\omega)|/|M_i(\omega)|$). 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

$\pi_{m,i}(M_i(\omega), N_i(\omega), s_i) = (1 - f_i)^{4/3}(1/\beta_i)r_i^{2/3}[1 - (1/2\beta_i^3)]$ which is also monotonically decreasing in f_i but reaches a (global) maximum at $\beta_i = 2^{1/3}$.²⁴

For expositional ease we further specify the learning model and assume independent inertia within types. That is, 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 $\rho_B \in]0,1[$.

Full coordination on any platform is an absorbing state (both types of traders get strictly positive profits on any active platform). Analogously to Lemma 1, it is easy to see that the monomorphic states and the states in $\Omega_0 = \{\omega \mid \pi_{m,i} = \pi_{m,j}, \pi_{m,i} = \pi_{m,j}\}$ form the only absorbing sets. Moreover, only the monomorphic states can be stochastically stable. The designers' profits, though, depend not only on the support of the limit invariant distribution μ^* but also on its weights for the different (monomorphic) states. Hence, analysing this setting requires a direct application of Lemma 3.1 in Freidlin and Wentzell (1984).

To economise on notation, identify a state with a pair (k, l) where k is the number of sellers and l the number of buyers present at platform i . Let $P_i^* = P((1, 1), \omega_i^*)$ and $P_j^* = P((1, 1), \omega_j^*)$. It is easy to see that a minimal-cost tree in $\mathcal{T}_{\omega_i^*}$ is as follows. State $(1, 1)$ is connected to ω_i^* with transition probability P_i^* . States $(0, 0)$, $(0, 1)$, and $(1, 0)$ are connected to $(1, 1)$ with transition probabilities $P((0, 0), (1, 1)) = \varepsilon^2$, $P((0, 1), (1, 1)) = \varepsilon(1 - \rho_B)^2$, and $P((1, 0), (1, 1)) = \varepsilon(1 - \rho_S)^2$. All other states are directly connected to ω_i^* with transition probabilities $P((0, 2), \omega_i^*) = \varepsilon\rho_S$, $P((1, 2), \omega_i^*) = \rho_S$, $P((2, 0), \omega_i^*) = \varepsilon\rho_B$, and $P((2, 1), \omega_i^*) = \rho_B$. Hence, the product of transition probabilities in this tree is $\rho_B^2\rho_S^2(1 - \rho_B)^2(1 - \rho_S)^2P_i^*$. The probabilities for ω_j^* can be derived by a permutation of indices B and S , thus the corresponding product is $\rho_B^2\rho_S^2(1 - \rho_B)^2(1 - \rho_S)^2P_j^*$. In the limit, $\mu^*(\omega_j^*)/\mu^*(\omega_i^*)$ is determined by the quotient of these two products, which simplifies to P_i^*/P_j^* . 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 13. *Suppose $\pi_{m,i}(1, 1, s_i) > \pi_{m,j}(1, 1, s_j)$. Then for every $\kappa > 0$ there is a $\bar{\rho}_S < 1$ such that $\mu^*(\omega_j^*) < \kappa$ for all $\rho_S > \bar{\rho}_S$.*

Proof. 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 $\pi_{n,i}(1, 1, s_i) = \pi_{n,j}(1, 1, s_j)$, $P_i^* = \rho_S(1 - \rho_B)\rho_B + \frac{1}{2}\rho_B(1 - \rho_S)\rho_S + \frac{1}{2}\rho_B\rho_S$ and $P_j^* = \frac{1}{2}\rho_B(1 - \rho_S)\rho_S$. If $\pi_{n,i}(1, 1, s_i) < \pi_{n,j}(1, 1, s_j)$, $P_i^* = \rho_S(1 - \rho_B)\rho_B$ and $P_j^* = \rho_B(1 - \rho_S)\rho_S$. Hence, in both cases P_i^*/P_j^* approaches zero if $\rho_S \rightarrow 1$ and $\lim_{\rho_S \rightarrow 1} \mu^*(\omega_j^*) = 0$. 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 14. *There exists a $\bar{\rho}_S < 1$ such that, for all $\rho_S > \bar{\rho}_S$, the platform profile (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 sellers' profits decrease in f_i and reach their global maximum at $\beta = 2^{1/3}$ it is clear that $\pi_{m,i}(1, 1, s_i) < \pi_{m,i}(1, 1, s_i^*)$ for any $s_i \neq s_i^*$. Suppose both s_i^* and some institution $s_i \neq s_i^*$ are available, and let ω be the state where all traders are at s_i . By Lemma 13, for any $\kappa > 0$ there exists $\bar{\rho}_S$ such that $\mu^*(\omega) < \kappa$ for all $\rho_S > \bar{\rho}_S$. Hence, for κ small enough (and since there are finitely many strategies), we obtain $\pi_{D,i}(s_i, s_j^*) < \pi_{D,i}(s_i^*, s_j^*)$ for all $s_i \neq s_i^*$; thus (s_i^*, s_j^*) is a strict Nash equilibrium.

²⁴ We assume for simplicity from now on that δ is such that $2^{1/3} \in B$.

It can be shown that, for ρ_S large enough, $\beta_i \leq 1$ is not chosen by any designer in any pure strategy equilibrium. Moreover, in any mixed strategy equilibrium (σ_1^*, σ_2^*) there is a least one designer i where $s_i \in C(\sigma_i^*)$ implies that $\beta_i > 1$.²⁵ Note also that the condition on ρ_S is a sufficient but not a necessary one. For a smaller ρ_S we cannot characterise the limit invariant distribution and hence do not know the equilibrium behaviour of the designers. The result might hold even for $\rho_S < \bar{\rho}_S$, depending on the details of the demand and supply conditions.

These results illustrate that the ‘Platform Design Paradox’, i.e. the fact that competition between market designers might lead to the design of non-market clearing institutions, also appears in the case of decreasing returns to scale.

University of Konstanz
ECARES, Universite Libre de Bruxelles
Bamberg University

Submitted: 1 September 2006

Accepted: 2 January 2009

References

- Alós-Ferrer, C. (2003). ‘Finite population dynamics and mixed equilibria’, *International Game Theory Review*, vol. 5 (3), pp. 263–90.
- Alós-Ferrer, C. and Kirchsteiger, G. (2008). ‘Learning and market clearing’, mimeo, available at <http://www.ecares.org/kirchsteiger.html>.
- Alós-Ferrer, C., Kirchsteiger, G. and Walzl, M. (2006). ‘On the evolution of market institutions: the platform design paradox’, CEPR Discussion Paper No. 5538.
- Alós-Ferrer, C. and Netzer, N. (2007). ‘The logit-response dynamics’, TWI Research Paper Series No 28.
- Ariely, D., Ockenfels, A. and Roth, A. (2005). ‘An experimental analysis of ending rules in internet auctions’, *RAND Journal of Economics*, vol. 36, pp. 790–809.
- Armstrong, M. (2006). ‘Competition in two-sided markets’, *Rand Journal of Economics*, vol. 37, pp. 668–91
- Belleflamme, P. and Toulemonde, E. (2004). ‘Emergence and entry of B2B marketplaces’, CORE Discussion Paper No. 78.
- Benaïm, M. Weibull, J. (2003). ‘Deterministic approximation of stochastic evolution in games’, *Econometrica*, vol. 71, pp. 878–903.
- Binmore, K. and Samuelson, L. (1997). ‘Muddling through: noisy equilibrium selection’, *Journal of Economic Theory*, vol. 74, pp. 235–65.
- Blume, L. (1995). ‘The statistical mechanics of best-response strategy revision’, *Games and Economic Behavior*, vol. 11, pp. 111–45.
- Caillaud, B. and Julien, B. (2003). ‘Chicken and egg: competition among intermediation service providers’, *RAND Journal of Economics*, vol. 34, pp. 521–52.
- Davila, A., Gupta, M. and Palmer, R. (2003). ‘Moving procurement systems to the internet: the adoption and use of e-procurement technology models’, *European Management Journal*, vol. 21 (1), pp. 11–23.
- The Economist* (2004). ‘E-commerce takes off’, May 13th.
- Ellison, G. (1993). ‘Learning, local interaction, and coordination’, *Econometrica*, vol. 61, pp. 1047–71.
- Ellison, G. (2000). ‘Basins of attraction, long-run stability, and the speed of step-by-step evolution’, *Review of Economic Studies*, vol. 67, pp. 17–45.
- Ellison, G. (2006). ‘Bounded rationality in industrial organization’, in (R. Blundell, W.K. Newy and T. Persson, eds.), *Advances in Economics and Econometrics: Theory and Applications*, pp. 142–80, Ninth World Congress, Cambridge: Cambridge University Press.
- Ellison, G. and Fudenberg, D. (2003). ‘Knife-edge or plateau: when do markets models tip?’, *Quarterly Journal of Economics*, vol. 118, pp. 1249–78.
- Ellison, G., Fudenberg, D. and Möbius, M. (2004). ‘Competing auctions’, *Journal of the European Economic Association*, vol. 2, pp. 30–66.
- European Commission, Enterprise Directorate General (2003). ‘Report of the Expert Group on B2B Internet Trading Platforms’, final report
- Freidlin, M.I. and Wentzell, A.D. (1984). *Random Perturbations of Dynamical Systems*, New York: Springer Verlag.

²⁵ A proof of these claims and more detailed exposition of the material discussed in Appendix B is available on request.

- Friedman, D. (1993). 'How trading institutions affect financial market performance: some laboratory evidence', *Economic Inquiry*, vol. 31, pp. 410–35.
- Fudenberg, D. and Levine, D. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Gabaix, X. and Laibson, D. (2006). 'Shrouded attributes, consumer myopia, and information suppression in competitive markets', *Quarterly Journal of Economics*, vol. 121 (2), pp. 505–40.
- Gerber, A. and Betzùge, M.O. (2007). 'Evolutionary choice of markets', *Economic Theory*, vol. 30, pp. 453–72.
- Holt, C. (1995). 'Industrial organization', in (J. Kagel and A. Roth, eds.), *Handbook of Experimental Economics*, pp. 349–443, Princeton: Princeton University Press.
- Hopkins, E. (2007). 'Adaptive learning models of consumer behavior', *Journal of Economic Behavior and Organization*, vol. 64, pp. 348–68.
- Kagel, J. (1995). 'Auctions: a survey of experimental research', in (J. Kagel and A. Roth, eds.), *Handbook of Experimental Economics*, pp. 501–86, Princeton, NJ: Princeton University Press, Princeton.
- Kandori, M., Mailath, G. and Rob, R. (1993). 'Learning, mutation, and long-run equilibria in games', *Econometrica*, vol. 61, pp. 29–56.
- Kandori, M. and Rob, R. (1995). 'Evolution of equilibria in the long run: a general theory and applications', *Journal of Economic Theory*, vol. 65, pp. 383–414.
- Keys, S. (2002). 'Online trading platforms: to build or to buy?', mimeo, available where?.
- Kirchsteiger, G., Niederle, M. and Potters, J. (2005). 'Endogenizing market institutions: an experimental approach', *European Economic Review*, vol. 49(7), pp. 1827–53.
- Kugler, T., Neeman, Z. and Vulkan, N. (2006). 'Markets versus negotiations: an experimental investigation', *Games and Economic Behavior*, vol. 56, pp. 121–34.
- Lucking-Reiley, D. and Spulber, D. (2001). 'Business to business electronic commerce', *Journal of Economic Perspectives*, vol. 15, pp. 55–68.
- Ockenfels, A. and Roth, A. (2006). 'Late and multiple bidding in second-price internet auctions: theory and evidence concerning different rules for ending an auction', *Games and Economic Behavior*, vol. 55, pp. 297–320.
- Oechssler, J. (1997). 'An evolutionary interpretation of mixed-strategy equilibria', *Games and Economic Behavior*, vol. 21, pp. 203–37.
- Plott, C. (1982). 'Industrial organization: theory and experimental economics', *Journal of Economic Literature*, vol. 20, pp. 1485–587.
- Plott, C. and Smith, V. (1978). 'An experimental examination of two exchange institutions', *Review of Economic Studies*, vol. 45, pp. 133–53.
- Rochet, J. and Tirole, J. (2003). 'Platform competition in two-sided markets', *Journal of the European Economic Association*, vol. 1, pp. 990–1029.
- Rochet, J. and Tirole, J. (2006). 'Two-sided markets: an overview', *RAND Journal of Economics*, vol. 35(3), pp. 645–67.
- Roth, A. and Ockenfels, A. (2002). 'Last-minute bidding and the rules for ending second-price auctions: evidence from eBay and Amazon on the internet', *American Economic Review*, vol. 92(4), pp. 1093–103.
- Samuelson, L. (1994). 'Stochastic stability in games with alternative best replies', *Journal of Economic Theory*, vol. 64, pp. 35–65.
- Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*, Cambridge, MA: MIT Press.
- Schlag, K. (2004). 'Competing for boundedly rational consumers', mimeo, Universitat Pompeu Fabra.
- Spiegler, R. (2006). 'Competition over agents with boundedly rational expectations', *Theoretical Economics*, vol. 1, pp. 207–31.
- UNCTAD (2002). *E-Commerce and Development Report 2002*, Geneva: United Nations.
- Young, P. (1993). 'The evolution of conventions', *Econometrica*, vol. 61, pp. 57–84.