

Evolution of behavior when duopolists choose prices and quantities*

Abhimanyu Khan[†] Ronald Peeters[‡]

January 19, 2012

Abstract

We study duopolistic competition in a differentiated market with firms setting prices and quantities, without explicitly imposing market clearing. Unlike the commonly adopted assumption of profit maximizing firms, we assume firm behavior to be shaped by a Darwinian dynamic: the less fitter firm imitates the fitter firm and occasionally firms may experiment with a random price and/or quantity. Our two main findings are that: (i) a market clearing outcome always belongs to the set of feasible long run outcomes, but may co-exist with non-market clearing outcomes with as well excess supply as excess demand being possible; and (ii) there exist parameter configurations for which the only feasible outcomes imply prices above monopoly level.

JEL Classification: C72, C73, D21, D43, L11, L13.

Keywords: oligopoly; bounded rationality; evolution; learning.

*We would like to thank Matthew Embrey, Jean-Jacques Herings, Rene Saran, Jörgen Weibull and the audience at HOC 2011 in Boston and SAET 2011 in Faro for useful comments and suggestions. Financial support by Meteor and the Dutch Science Foundation is gratefully acknowledged.

[†]Department of Economics, Maastricht University. E-mail: a.khan@maastrichtuniversity.nl

[‡]Department of Economics, Maastricht University. E-mail: r.peeters@maastrichtuniversity.nl

1 Introduction

The traditional literature in the field of oligopoly theory typically models firms as perfectly rational agents whose decisions are solely aimed at the maximization of profit. However, the recognition that agents might be informationally constrained and/or at most boundedly rational, has activated a growing interest in analyzing the impact that such constraints might have on decision making and market outcomes. Indeed, it is not unnatural to conceive of a situation where firms lack the relevant information for profit maximization. Hence, the tools usually used to study firm behavior lose their applicability and alternative approaches need to be adopted.

Alchian (1950) emphasizes the importance of relative performance and that relative profit might serve as the chief driver of firm behavior. If the salience of relative profit is accepted, then imitation may be a reasonable behavioral assumption. This approach agrees with the evolutionary perspective, which is built on the premise that strategies which yield a higher-than-average payoff will tend to be more attractive. Huck et al. (1999), Huck et al. (2000), Offerman et al. (2002) and Apesteguia et al. (2007) find some experimental evidence of firms imitating the fitter firm once feedback on strategies and profits is explicitly provided.

Vega-Redondo (1997) studies a situation where firms producing a homogeneous commodity compete in quantities and the market price is determined by the total quantity produced. The profits and quantities chosen by each firm are observable, and the action chosen by the most successful firm is almost always imitated. Occasionally, firms experiment and choose any quantity. As the probability of experimentation goes to zero, the Walrasian outcome prevails most of the time. Huck et al. (1999) and Apesteguia et al. (2010) provide experimental support for this finding.

Similarly, Alós-Ferrer et al. (2000) analyze, in an evolutionary setting, the long-run outcome when firms, with increasing marginal costs in a homogeneous market, set prices. Firms with the lowest price win the market. Here, analogous to Vega-Redondo (1997), prices announced by the firm and the profits realized by each are observed and the behavior of the firms is based on imitation of the successful firms. The study obtains a strict subset of the Nash equilibria of the underlying game as the long-run outcome of this dynamic process of imitation and occasional experimentation.

In this paper, we pursue a related issue. The question that we seek to address in this paper is: What is the nature of a process where firms do not choose either price or quantity but both of them simultaneously? This question is more interesting if we abandon the assumption of market clearing and then examine if it arises out of the evolutionary process. This is motivated in part by Herings (1997) and more recently by Alós-Ferrer and Kirchsteiger (2010) who show that it is possible for non-market clearing institutions to evolve.

We consider a market for a good that is differentiated horizontally. To keep things simple, we assume that there are two firms that are located at the extreme ends of a Hotelling line.¹ We impose a very basic informational setting; the only information that the firms must have pertain to is its own marginal cost and that there is a competing firm. In particular, we do not require the firms to be aware of the fact that they are in a differentiated market. Neither do they have to be cognizant of the consumers' valuation of their product or of the consumers' preferences between the two firms (i.e. the transportation cost) or the size of the market.

The firms can choose both price and quantity. Market clearing is not imposed and hence the presence of excess supply or excess demand cannot be ruled out a priori. After setting prices and quantities, the firms realize profits, all three of which are observable. The behavior that we impose on the firms is that the less successful firm immediately imitates the more successful firm in both price and quantity. This can be viewed essentially as a form of Darwinian adjustment where the more successful strategy grows in the population. In addition to this imitation dynamic, we introduce the possibility of experimentation, whereby firms choose a price or quantity that is not necessarily dictated by imitation. The possibility of such experimentation makes the system ergodic and enables us to obtain the long-run outcome. The techniques used to identify the long-run outcome are based on the concept of stochastic stability, as used by Young (1993), Kandori et al. (1993) and Ellison (2000). This process of imitation and occasional experimentation can be thought of as an application to Hamilton's theory of spite in biology (Hamilton, 1970). Via imitation, the system gravitates towards the action that yields higher relative profit, and this might occur at the expense of absolute profit.

Our results indicate that the outcome depends on the magnitude of the consumer's gross utility or valuation for the product. When the consumer valuation is lower than a certain bound, the unique stochastically stable state involves each firm behaving as a monopolist (Proposition 2). As consumer valuation increases up to another bound, we obtain a subset of prices and quantities as the long-run outcome with all prices being above the monopoly price (Proposition 3). For sufficiently high consumer valuation, any state where firms choose identical price and quantity and receive at least zero profits can be a long-run outcome (Proposition 4). In this last case, in contrast to the ones mentioned earlier, market clearing does not necessarily come about – it is possible to have long-run outcomes with persistent excess supply or excess demand.

At this juncture, we provide some preliminary intuition for the results. When consumer

¹The location of the firms is something that we take to be exogenous. Hehenkemp and Wambach (2010) show that in an evolutionary context where a good has multiple attributes, firms would choose to locate at the centre of the Hotelling line. In our paper, location is not a subject of choice but only a way to specify the demand that each firm faces. It shares no connection with the fact that rational firms may not differentiate maximally. Nonetheless, we discuss the implications of our results for homogeneous product markets in Section 5.

valuation is sufficiently low, each firm is for all practical purposes, endowed with some monopoly power. This enables monopoly pricing to be more resistant to any spiteful behavior by the competing firm. But with higher consumer valuations, this monopoly power is eroded. For an intermediate range of consumer valuation, the subset of states obtained are the ones which are most resistant to experimentation by a deviant firm. Strikingly, this goes along with prices above monopoly level. This is a reflection of the fact that the stochastically stable states are the ones most immune to any type of potential spiteful behavior by the competing firm – the high prices allow a firm to have a secure but admittedly smaller market. For sufficiently high consumer valuation, no state is more resistant to experimentation and all absorbing states are stochastically stable. Interestingly, in spite of all the emphasis on relative profit, excess supply (which represents unnecessary cost) or excess demand (which represents an opportunity to increase profit) are possible long-run features.

The paper is organized as follows. In Section 2, we present the model and the standard benchmarks. In Section 3, we present very briefly the concept of stochastic stability and related concepts and theorems we make use of later on (in particular, in the proofs). In Section 4, we present the main results of the paper. In Section 5, we discuss robustness of our results and digress on possible extensions, such as one wherein the quantity produced by the firm is not observable but the actual sales are.

2 Model, notation and standard results

We study a duopolistic market in the standard Hotelling framework, but where firms choose prices and quantities. Two firms are located at the endpoints of the unit interval $[0, 1]$ along which the consumers are uniformly distributed. We identify firms and consumers by their location on this interval – firm 0 is the firm situated at location 0 and firm 1 is the firm situated at location 1. Two firms simultaneously and independently decide on the quantity to put on the market and the price to charge for each unit.² After having observed the two prices, consumers choose from which of the two firms to acquire the product or to abstain from buying. Consumers acquire one unit of the good at most. All consumers assign a common intrinsic value of β to the good and incur a transportation cost of τ per unit of distance to the chosen firm. Hence, the utility of the consumer at location $x \in [0, 1]$ equals

$$U_x = \begin{cases} \beta - \tau \cdot x - p_0 & \text{if } x \text{ purchases from firm 0;} \\ \beta - \tau \cdot (1 - x) - p_1 & \text{if } x \text{ purchases from firm 1;} \\ 0 & \text{if } x \text{ does not consider buying any good at all.} \end{cases}$$

²In fact, the techniques used require a finite strategy space, which means the prices and quantities are chosen from a finite price and quantity grid. We continue our analysis using continuous price and quantity spaces and add clarifying footnotes regarding particular prices and quantities that should be contained in the respective grids to make the results feasible.

Consumers maximize their utility. Observe that in their decision the consumers do not take into account the quantity produced by the firms. It is possible that the demand at one firm exceeds the supply, while the other firm is left with an excess supply. We ignore the possibility of spillover demand – i.e. the possibility that a consumer who is not able to acquire the good at the preferred location to travel to the other firm (which might have excess supply). Such considerations would require additional assumptions on how goods are rationed or distributed amongst consumers by the firm that faces an excess demand. To elaborate briefly, it is possible that a consumer prefers acquiring a good at either firm to not acquiring one at all while there might be another consumer who would prefer to acquire the good only from a specific firm. The possibility of spillover demand arises only if the second consumer gets served. To analyse this, there needs to be an assumption on which consumer acquires a good in case their preferred firm faces excess demand. While it may be instructive to look at this aspect, we do not pursue it in this paper.

A firm’s sale, s_i , is given by the minimum of the supply q_i and the demand d_i that is given by the mass of consumers that decides purchasing from firm i ($i = 0, 1$). We assume an equal and constant production cost of c per unit and that the goods are perishable such that excess supply can not be utilized. As firms are aware of their production costs, we assume them not to adopt prices below unit cost. A firm’s profit is then given by:

$$\pi_i = p_i \cdot s_i - c \cdot q_i.$$

The specification of the demand of firm i depends on whether the prices give rise to all consumers deciding to purchase a good (irrespective of their success to acquire the good) or whether there is a positive mass of consumers that prefer not to purchase at given prices. In the former case the demand of the firms is determined by the location of the consumer that is indifferent between the two products: $\hat{x} = \frac{p_1 - p_0 + \tau}{2\tau}$. If $\hat{x} \geq 1$, then $d_0 = 1$ and $d_1 = 0$; if $\hat{x} \leq 0$, then $d_0 = 0$ and $d_1 = 1$; otherwise (that is, if $\hat{x} \in (0, 1)$) $d_0 = \hat{x}$ and $d_1 = 1 - \hat{x}$. In the latter case, i.e. when there is a positive mass of consumers that prefer not to purchase, the demands are given by $d_0 = \hat{x}_0$ and $d_1 = 1 - \hat{x}_1$, provided that these values are not below 0 or above 1. Here, \hat{x}_0 stands for the consumer who gets precisely zero utility if she purchases from firm 0, i.e. $\hat{x}_0 = \frac{\beta - p_0}{\tau}$; \hat{x}_1 is defined similarly, i.e. $\hat{x}_1 = 1 - \frac{\beta - p_1}{\tau}$. Note that in the former case $\hat{x}_1 \leq \hat{x} \leq \hat{x}_0$ and the firms actively compete for the indifferent/marginal consumer; in the latter case $\hat{x}_0 < \hat{x} < \hat{x}_1$ and the firms act as “local monopolists”.

Throughout the paper we assume that the consumer valuation exceeds the unit production cost: $\beta > c$. This ensures that there is scope for efficient trade. Moreover, we assume transportation costs to be of moderate size relative to the marginal cost of production: $\tau < c$. This latter assumption is purely cosmetic: it reduces the derivations, while the main insights

we derive are robust to this assumption³, and hence improves readability of the paper.

Although this paper adopts an evolutionary approach where firms are not required to have any knowledge of the demand curves and the size of the market (or, alternatively, the location of the rival firm), for the purpose of benchmarking we provide some standard solutions that assume firms to possess the relevant information to maximize profits. First, in case a firm would be a monopolist, a price of $p = \max\{\frac{\beta+c}{2}, \beta - \tau\}$ would be chosen with a quantity equal to $q = \min\{\frac{\beta-c}{2\tau}, 1\}$. Second, the Nash equilibrium prescribes the firms to produce this monopoly quantity and to charge the monopoly price in case $\beta \leq c + \tau$. When $\beta > c + \tau$, Nash equilibrium prices and quantities equal $p = c + \tau$ and $q = \frac{1}{2}$ respectively.

3 Preliminaries

In this section, we present the notion of stochastic stability and some useful theorems by Young (1993) and Ellison (2000). We restrict our attention to stochastic stability in the context of imitation and experimentation and tailor it slightly to our needs.

We focus on a two-player game, where player i chooses a strategy s_i from a finite set of strategies S_i ($i = 1, 2$). After choices have been made, the players receive their payoff that solely depends on the profile of chosen strategies $(s_1, s_2) \in S \equiv S_1 \times S_2$. In order to make imitation sensible, we restrict attention to games where the strategy space of both players is identical (i.e. $S_1 = S_2$) and the payoffs are symmetric (i.e. $\pi_1(s_1, s_2) = \pi_2(s_2, s_1)$).

3.1 Unperturbed learning dynamics and absorbing states

We consider the evolution of behavior when the less fit player (the one with the lower payoff) imitates the fitter one (the one with the larger payoff). This includes a positive probability of a player imitating the other, even if they have equal payoffs. This imitation dynamic can be captured by a Markov process (S, P) with S being the state space consisting of all strategy-profiles and P capturing the state transition dynamics. We let $P_{\omega\omega'}$ denote the probability that the system transits to the state $\omega' \in S$ when being in state $\omega \in S$.

A state $\omega \in S$ is called *absorbing* if and only if $P_{\omega\omega} = 1$ and $P_{\omega\omega'} = 0$ for all $\omega' \neq \omega$. Once a system reaches an absorbing state, it is locked in it thereafter. We denote the set of all absorbing states by $\Omega \subseteq S$. As a result of the particular imitation dynamic that we impose later on, all absorbing sets are singletons. Therefore, to our convenience, our presentation of the literature is specific to the situation at hand and ignores the possibility of non-singleton absorbing sets.⁴

³Derivations are available upon request.

⁴For a more general discourse on this, we refer to Chapter 5 of Fudenberg and Levine (1998).

The unperturbed imitation dynamic results in the system moving to an absorbing state, but it does not make a further selection amongst these states. To select from the set of all absorbing states, we perturb the imitation dynamic with the possibility of occasional experimentation. This enables us to identify the states that might be observed most of the time.

3.2 Perturbed learning dynamics and stochastic stability

Apart from the imitation dynamics, the players can experiment with any strategy. This happens with very small probability, say, ε .⁵ Let $P(\varepsilon)$ be the transition matrix associated with the probability of experimentation being equal to ε . It is possible to transit out of the absorbing states induced by imitation only with experimentation. A consequence of experimentation is that the *perturbed* process is ergodic such that it generates a unique invariant (or stationary, or equilibrium) distribution $\mu(\varepsilon)$, which is the unique solution to the system of linear equations $\mu P(\varepsilon) = \mu$. The states that are *stochastically stable* are the ones that receive a positive weight in the limiting invariant distribution, i.e. the unique invariant distribution as ε goes to zero: $\mu^* = \lim_{\varepsilon \downarrow 0} \mu(\varepsilon)$. The following result indicates that only absorbing states may be stochastically stable.

Theorem 1 (Young, 1993, Theorem 4(i)).

The set of stochastically stable states is a nonempty subset of the set of absorbing states.

The algorithm for identifying the stochastically stable states is now presented. Let the *resistance* $\rho(\omega, \omega')$ be the total number of experimentations involved in the transition from state ω to the state ω' , where $\omega, \omega' \in S$. The state space S can be represented by the vertices of a directed graph, in which $\rho(\omega, \omega')$ is the resistance of each edge. The resistance of the edge (ω, ω') is zero when the transition from state ω to state ω' is made possible entirely via the imitation dynamic while resistance of an edge emanating from an absorbing state must have a strictly positive weight. Let $r(\omega, \omega')$ be the resistance of the least resistant path among all paths between these two states.

We now consider another directed graph, where the vertices are the absorbing states $\omega \in \Omega$. A ω -tree T_ω is a directed graph with vertex set Ω such that for every $\omega' \neq \omega$ there is exactly one (cycle-free) sequence of edges (a path) from ω' to ω . Let \mathcal{T}_ω denote the set of all such ω -trees T_ω . Edges $[\omega', \omega'']$ are given a weight $r(\omega, \omega')$. The *tree-resistance* $r(T_\omega)$ of a ω -tree equals the sum of corresponding edge-resistances, i.e. $r(T_\omega) = \sum_{[\omega', \omega''] \in T_\omega} r(\omega', \omega'')$. Finally, the *stochastic potential* of $\omega \in \Omega$, denoted by $\gamma(\omega)$, is the minimal tree-resistance over all ω -trees, i.e. $\gamma(\omega) = \min_{T_\omega \in \mathcal{T}_\omega} r(T_\omega)$. A ω -tree T_ω with $r(T_\omega) = \gamma(\omega)$ is a *least resistance*

⁵In this paper, we will assume a more involved experimentation structure, as is specified in the next section. For now, let ε refer to the probability of experimentation to a random strategy.

ω -tree. The set of stochastically stable states is characterized by the set of states with minimal stochastic potential.

Theorem 2 (Young, 1993, Theorem 2).

An (absorbing) state $\omega \in \Omega$ is stochastically stable if and only if it has minimal stochastic potential, i.e. $\gamma(\omega) \leq \gamma(\omega')$ for all $\omega' \in \Omega$.

The method illustrated above requires the construction of the least resistant tree for all states so as to identify the state with the least stochastic potential. At times, it may be convenient to identify the stochastically stable state(s) by employing the concepts of radius, co-radius and modified co-radius by Ellison (2000).

The *basin of attraction* of the collection of one or more absorbing states A , is the set of initial states from which the unperturbed Markov process converges to A with probability 1, i.e. $B(A) = \{ \omega \in \Omega \mid \text{Prob}[\exists_K \forall_{k > K} \omega(k) \in A \mid \omega(0) = \omega] = 1 \}$. Then, the *radius* of A , $R(A)$, is given by the minimum number of experimentations that are necessary to leave $B(A)$ when the process starts in A .

Suppose now that the system is in a state $\omega' \notin A$. Let $z(\omega', A)$ be the minimum number of experimentations required to transit from ω' to $B(A)$. Then, the *co-radius* of the basin of attraction of A is the minimum number of experimentations needed from the state from which $B(A)$ is most difficult to access: $CR(A) = \max_{\omega' \notin A} z(\omega', A)$. Based on the concept of co-radius, the modified co-radius is defined.

Let $(\omega(1), \omega(2), \dots, \omega(K))$ be a path consisting of distinct states that starts in a state x and reaches the set A for the first time in the terminal state (A is as before, a set of absorbing states); so, $\omega(1) = x$, $\omega(k) \notin A$ for $2 \leq k \leq K - 1$ and $\omega(K) \in A$. By $S(x, A)$, we denote the sets of all such paths. Let $z(\omega(1), \omega(2), \dots, \omega(K))$ denote the total number of experimentations along the path $(\omega(1), \omega(2), \dots, \omega(K))$: $z(\omega(1), \omega(2), \dots, \omega(K)) = \sum_{k=1}^{K-1} z(\omega(k), \omega(k+1))$. Denote by A_1, A_2, \dots, A_L , the sequence of absorbing sets through which the path passes consecutively. Then, the modified cost function of the path is defined as: $z^*(\omega(1), \omega(2), \dots, \omega(K)) = z(\omega(1), \omega(2), \dots, \omega(K)) - \sum_{\ell=2}^{L-1} R(A_\ell)$. Now define $z^*(x, A) = \min_{(\omega(1), \omega(2), \dots, \omega(K)) \in S(x, A)} z^*(\omega(1), \omega(2), \dots, \omega(K))$. Finally, the *modified co-radius* of the basin of attraction of A is: $CR^*(A) = \max_{x \notin A} z^*(x, A)$.

Theorem 3 (Ellison, 2000, Theorems 1–3).

- (i) If $R(A) > CR(A)$, then the stochastically stable states are contained in A .
- (ii) If $R(A) > CR^*(A)$, then the stochastically stable states are contained in A .
- (iii) If for some absorbing state ω and some other state ω' , we have $R(\omega) > z^*(\omega', \omega)$, then $\mu^*(\omega') = 0$. If $R(\omega) = z^*(\omega', \omega)$, then $\mu^*(\omega') > 0$ implies $\mu^*(\omega) > 0$.

4 Results

Instead of assuming firms to maximize profits, we adopt the principle that firms learn by imitation and experimentation. That is, the less fit firm (i.e. the firm with the lower profit) imitates the behavior of the fitter firm (i.e. the one with the higher profit). This requires strategic decisions and the yielded profits to be perfectly observable. Once firms receive equal (and nonnegative) profits, firms can only continue learning by experimentation to alternative price–quantity pairs.

Note that for this learning dynamic, firms do not need all profit relevant information of the environment they compete in. In our setting, awareness of the size of the market, the intrinsic value consumers attach to a good and the transportation costs they are facing is not required. We even do not require the firms to be aware of their or their rival’s position on the Hotelling line. The only information we assume them to possess is their unit cost of production and the existence of a rival firm whom they can observe.

4.1 Unperturbed learning dynamics and absorbing states

We consider the following imitation dynamic. If in a particular state, a firm receives a negative profit, the rival firm is imitated in case the latter obtained a nonnegative profit. Otherwise, both firms shut down by choosing a production level of zero next period (without changing their price). If a firm receives a nonnegative profit, it imitates the rival firm in case this rival firm is strictly fitter in the current state (received a strictly higher profit); in case the rival firm is as fit, we assume the firm imitates the rival with a positive probability strictly less than one. This imitation dynamics gives rise to the following proposition.

Proposition 1. *The set of absorbing states are precisely the monomorphic states that imply non-negative profit to the firms.*

Proof. A monomorphic state (i.e. a state where both firms choose the same price and quantity) with non-negative profit is absorbing as firms receive equal profits and have the same actions – this leaves no further scope for imitation. On the other hand, if we are in a state where firms realize unequal profits, then either both of them shutdown (this occurs when both firms have negative profit) or one firm imitates the other. In either case we transit to a monomorphic state. If we are in a state where firms realize equal profit with disparate actions, the positive probability of imitation makes this state transient. ■

Without experimentation the process will converge to an absorbing state. Experimentations that give the experimenting firm the higher profit relative to its rival are needed to drive the system out of an absorbing state. The resistance of the absorbing state towards experimen-

tation provides information on the stability of an absorbing state. In the next subsection we will characterize the set of states that are most resistant towards experimentation.

4.2 Perturbed learning dynamics and stochastic stability

Before specifying the experimentation probabilities, we begin by acknowledging a conceptual issue with experimentation in multiple dimensions, which is that there is no “theory” on how experimentations happen or occur. Models that tackle this concern by endogenizing the mistake probability by assuming that agents would be less likely to experiment away from states with good realizations have been considered; eg. in Van Damme and Weibull (2002) and Blume (2003). Such an approach is unfortunately not applicable here as firms in our model do not know the demand function that they face and, hence, cannot evaluate the payoffs corresponding to other choices.

Given this background, there are two features that (we feel) might be plausible to incorporate in the specification of experimentation. Firstly, since there are two firms who make decisions independently, it might be reasonable to speculate that experimentation by one firm is more likely than experimentation by two firms. Secondly, since firms act in two dimensions, it might be hypothesized that an experimentation in one dimension is more likely than an experimentation in two dimensions. Combining these two features, we assume that the probability of experimentation by one firm in either price or quantity (say, ε) is more likely than the experimentation by a firm in price and quantity (say, ε^2). The latter in turn is more likely than two firms simultaneously experimenting in price or quantity (say, ε^3).⁶

The structure of the set of stochastically stable states is different for different values of β , the gross utility that consumers extract from consumption. We present three propositions, each of them devoted to a distinctive case with respect to the size of β .

Proposition 2. *Suppose $\beta < \tau - 2c + \sqrt{\tau^2 + 9c^2 - 2c\tau}$. The state where both firms supply the monopoly quantity $q^* = \frac{\beta - c}{2\tau}$ and set the monopoly price $p^* = \frac{\beta + c}{2}$ is the unique stochastically stable state.*⁷

Proof. Let us denote the state mentioned in the proposition by ω^* . Notice that the supposition in the proposition implies $\beta < c + \tau$, such that the quantity produced by each firm according to ω^* is less than half the market size. Our proof is constructed as follows.

⁶The results that we obtain depend partially on this assumption. This is not surprising in the light of Bergin and Lipman (1996), who show that the stochastically stable states are sensitive to the specification of how experimentations occur. Since there is no guiding principle on how experimentations occur, we adopt the system of experimentation as induced by the two features mentioned. However, we do briefly discuss in Section 5 on what would happen if we ignored the second consideration and instead said that experimentation by two firms is less likely than by a single firm (without taking into account that there are two dimensions).

⁷We need the monopoly price and quantity to be contained in the price and quantity grid respectively (recall Footnote 2).

First, we show that no experimentation in price *and* quantity can push the system outside the basin of attraction of state ω^* . We then infer that it is also not possible with *either* one price *or* one quantity experimentation. From this we conclude that the radius of ω^* is strictly larger than 2. Second, we show that the co-radius of ω^* is at most 2. Form part (i) of Theorem 3 it follows that all stochastically stable states are included in the set $\{\omega^*\}$. Next, existence leads us to conclude that ω^* is the unique stochastically stable state.

Part 1 (radius). In state ω^* the firms set price and quantity at the monopoly level and do not compete for the marginal agent. Experimentation in price by one firm leads to active competition with its rival if and only if the experimented price is less than $\frac{3}{2}\beta - \frac{1}{2}c - \tau$. Notice that any larger price leads to a decrease in the experimenting firm's profit, while it does not affect its rival's profit and will result in a return to ω^* . Suppose, therefore that the experimenting firm experiments with the price $p' < \frac{3}{2}\beta - \frac{1}{2}c - \tau$. Moreover, suppose that this firm experiments simultaneously with the optimal quantity given that price: $q' = \frac{\beta+c+2\tau-2p'}{4\tau}$. This experimentation leads to a profit of $\pi' = (p' - c)q'$ for the experimenting firm and a profit of $\pi = \frac{\beta+c}{2}(1 - q') - cq^*$ for its rival. The difference between these two profits, $\pi' - \pi$, is maximized at $p'_{max} = \frac{c+\tau}{2}$. As $\tau < c$, this price is below unit cost and it suffices to consider the case where the experimenting firm experiments with marginal cost pricing.

Suppose the experimenting firm experiments to $p' = c$ and $q' = \frac{\beta-c+2\tau}{4\tau}$. Then, $\pi' - \pi = \frac{1}{8\tau}(\beta^2 + 4c\beta - 2\tau\beta - 2c\tau - 5c^2)$, which is negative if and only if $\beta \in (\tau - 2c - \sqrt{\tau^2 - 2c\tau + 9c^2}, \tau - 2c + \sqrt{\tau^2 - 2c\tau + 9c^2})$. Since $\tau - 2c - \sqrt{\tau^2 - 2c\tau + 9c^2} < c < \beta$, we get to the condition $\beta < \tau - 2c + \sqrt{\tau^2 - 2c\tau + 9c^2}$.

We have shown that even in the best case experimentation with prices and quantities, a firm cannot obtain a situation where its profit strictly exceeds that of its rival when the condition in the proposition is met. Now, note that if it is not possible to leave ω^* with a price experimentation in combination with an optimal quantity adjustment, it is also not possible without such an adjustment. A sole quantity experimentation always decreases the profit (either because more units can be sold at a price above cost, or because more units are produced than are sold), while it doesn't affect the profit of the rival firm.

Part 2 (co-radius). From the first part of Part 1, we learn that starting from any state, if a firm experiments simultaneously to price p^* and quantity q^* , then it results in a larger profit relative to the profit of its rival. This means that, from any state, at most two experimentations are needed to get to (the basin of attraction of) ω^* . ■

Proposition 2 reveals that when the intrinsic value of the product (β) is small, the monopoly outcome is expected to prevail in the long run. In particular, the two firms are not actively competing for the same customers; the market is only partially served. Notwithstanding the market not being fully served, the stochastically stable outcome implies market clearing.

Note, however, that the restriction on the intrinsic value is that severe that the participation constraint of the consumers at the other end of the market is not satisfied for a price equal to unit cost. This hands each of the firm some market power and in this framework of imitation and experimentation, it manifests itself into monopoly prices and monopoly outputs.

Proposition 3. *Suppose $\tau - 2c + \sqrt{\tau^2 + 9c^2 - 2c\tau} \leq \beta < c + \tau$. Then, the set of stochastically stable states consists of the monomorphic states with prices $p^* \in (\bar{p}, \beta)$ and quantity $q^* = \frac{\beta - p^*}{\tau}$, where \bar{p} is the unique feasible solution to $p(\frac{c - p + \tau}{2\tau}) - c\frac{\beta - p}{\tau} = 0$.⁸*

Proof. Let Ω^* denote the set of stochastically stable states as it is claimed in the proposition. Notice that all states in Ω^* imply market clearing and in none of them firms actively compete for the marginal consumer as $\bar{p} \geq \frac{\beta + c}{2} > \beta - \frac{\tau}{2}$. Our proof is constructed as follows. First, we show that the radius of Ω^* is 2. Second, we show that the modified co-radius of Ω^* is 1. Next, it follows from part (ii) of Theorem 3 that the stochastically stable states are contained in Ω^* . Third, we show that all states in Ω^* are stochastically stable. That concludes the proof.

Part 1 (radius). We demonstrate that the radius of the set Ω^* is 2. We first show that the radius is at least 2 by showing that there exists no price experimentation or quantity experimentation that makes the system leave the basin of attraction of Ω^* .

Since supply equals demand in all states in Ω^* and the consumers' acquisition decision is independent of the quantity produced, a quantity experimentation reduces the own profit (given that the price is above marginal cost) while it leaves the rival's profit unaffected. Hence, a quantity experimentation leaves the experimenting firm worse off as compared to its rival and imitation will drive the process back to the original state.

Since firms do not compete for the marginal agent, experimentation to a larger price ($p' > p$) will not lead to a change in the rival's profit. Moreover, since the quantity is unadjusted, the cost component of the own profit does not change. The revenue component is at least as much if $p'\frac{\beta - p'}{\tau} \geq p\frac{\beta - p}{\tau}$, which is satisfied if and only if $p' \in [\beta - p, p]$; a contradiction.⁹ Experimentation to a lower price ($p' < p$) decreases the own profit, and can therefore only lead to an eruption if the rival's firm is decreased by even more. For this to be the case, the price p' must be sufficiently low to compete for the rival's consumers: $p' \leq 2\beta - p - \tau$. Such a experimentation results in a situation where the experimenting firm obtains equal or larger profit than its rival if and only if it leads to equal or larger revenue (since costs are equal): $p'\frac{\beta - p}{\tau} - \frac{p' - p + \tau}{2\tau} \geq 0$. Since the derivative of the left-hand side to p' is the constant $(2\beta - 3p)$, it suffices to check the validity of the inequality at the two extreme

⁸For this proposition, we need (i) marginal cost, (ii) the monopoly price and (iii) at least one price (say p^*) in the interval (\bar{p}, β) to exist in the price grid, while we need the monopoly quantity and the quantity demanded at p^* (i.e. $\frac{\beta - p^*}{\tau}$) to exist in the quantity grid. For the instance when there is more than one price p from (\bar{p}, β) in the price grid, we require the corresponding demand, $\frac{\beta - p}{\tau}$, to belong to the quantity grid.

⁹Note that $\beta - p < p$, since $p > \bar{p} \geq \frac{\beta + c}{2} > \frac{\beta}{2}$.

values: c and $2\beta - p - \tau$. Experimentation to $2\beta - p - \tau$ can not be disruptive for the obvious reason that no consumers are stolen away and the own profit decreases. For $p' = c$ the inequality simplifies to $p^2 - (3c + \tau)p + 2\beta c \geq 0$. The derivative of the left-hand side equals $2p - (3c + \tau)$ and is negative since $p < \beta < c + \tau < \frac{3c + \tau}{2}$.¹⁰ This means that if the inequality is not satisfied for a certain price p , then it is not satisfied for larger values of p . The price \bar{p} is precisely the price at which the inequality is binding. Hence, for prices larger than \bar{p} , an experimentation to marginal cost is not disruptive.

We can conclude that for the states in Ω^* , the radius is at least 2. All states in Ω^* are with prices above monopoly level. An experimentation by one firm to monopoly price *and* quantity gives the experimenting firm a strictly higher profit and leads to an immediate drift towards an absorbing state outside Ω^* . Thus, the radius of Ω^* is 2.

Part 2 (modified co-radius). We now show that the modified co-radius of Ω^* is 1. We do so by demonstrating that from any other state outside Ω^* , it is possible to reach any state in Ω^* with a sequence of single suitable experimentations. This implies that the modified co-radius of Ω^* is 1.

Firstly, let us suppose that the initial state is of the form $\omega^0 = (p, q, p, q)$ with $p < \beta - \frac{\tau}{2}$ such that firms compete for the marginal consumer. And, let $\omega^* = (p^*, q^*, p^*, q^*)$ be the state in Ω^* that we try to reach. If both firms earn non-negative profit state ω^0 is absorbing, otherwise firms stop producing and the process transits to state $(p, 0, p, 0)$. Irrespective of the sign of the profit, if the rival firm experiments to $q = \frac{1}{2}$, it will receive at least the same profit. This means that one experimentation is needed to get to the state $\omega^1 = (p, \frac{1}{2}, p, \frac{1}{2})$, which is an absorbing state as each firm makes non-negative profit. Next, experimentation to marginal cost implies a larger profit relative to that of the rival firm if and only if the rival firm receives a negative profit: $p \frac{c - p + \tau}{2\tau} - \frac{c}{2} \leq 0$, or simplified $(p - \tau)(p - c) \geq 0$. This inequality will always be fulfilled as $p \geq c > \tau$. Hence, one experimentation to marginal cost pricing is needed to get to the absorbing state $\omega^2 = (c, \frac{1}{2}, c, \frac{1}{2})$. Now, consider an experimentation in quantity to q^* . Note that $q^* < \frac{1}{2}$, since $p^* > \bar{p} > \beta - \frac{\tau}{2}$. Both firms earn exactly zero profit and with positive probability quantity q^* gets imitated and we are in the absorbing state $\omega^3 = (c, q^*, c, q^*)$. Finally, let one of the firms experiment with p^* . Then, this experimentation gets imitated with positive probability if and only if the resulting profit is non-negative: $p^* \frac{c - p^* + \tau}{2\tau} - c \frac{\beta - p^*}{\tau} \geq 0$. This inequality is strictly satisfied precisely for market clearing prices above \bar{p} , in particular our p^* . We have reached the desired state $\omega^* = (p^*, q^*, p^*, q^*)$. Note that all the transitions from one absorbing state to the other take place with the help of a single suitable experimentation.

Secondly, let us suppose that the initial state is of the form $\omega^0 = (p, q, p, q)$, where $p \in$

¹⁰The first inequalities follows from p being in (\bar{p}, β) , the second by assumption, and the third from the assumption that $\tau < c$.

$(\beta - \frac{\tau}{2}, \bar{p})$. In this state, the firms are not competing and they have their independent markets. If both firms are receiving negative profit, they shut down, giving rise to the absorbing state $(p, 0, p, 0)$. Irrespective of which of the two states is absorbing, let one firm experiment with the market clearing quantity $\frac{\beta-p}{\tau}$. The experimenting firm will receive the larger profit and gets immediately imitated and the process transits to the absorbing state $\omega^1 = (p, \frac{\beta-p}{\tau}, p, \frac{\beta-p}{\tau})$. Next, let one of the firms experiment with marginal cost pricing. Then, the experimenting firm receives zero profit, while its rival receives a negative profit (this is precisely what the condition $p < \bar{p}$ implies). As a result, marginal cost pricing gets imitated and we reach the state $\omega^2 = (c, \frac{\beta-p}{\tau}, c, \frac{\beta-p}{\tau})$. The sequence of single experimentations from this state to a state in Ω^* has been illustrated in the previous paragraph.

Thirdly, suppose that $\omega^0 = (p, q, p, q)$ with $p \in (\bar{p}, \beta)$ but $q \neq \frac{\beta-p}{\tau}$. Obviously an experimentation to the market clearing quantity $\frac{\beta-p}{\tau}$ is imitated and we get to a state Ω^* with a single experimentation.

Finally, from state $\omega^0 = (p, 0, p, 0)$ with $p \geq \beta$, an experimentation to marginal cost is followed with positive probability. The sequence of single experimentations from the resulting state to a state in Ω^* has been illustrated.

We conclude that it is possible to reach a state in Ω^* from any other state not belonging to Ω^* with a sequence of single experimentations. Hence, the modified co-radius of Ω^* is 1.

Part 3. We now demonstrate that all states in Ω^* are stochastically stable. Take any two states ω' and ω'' in Ω^* . From ω'' we can go to the monopoly state $(\frac{\beta+c}{2}, \frac{\beta-c}{2\tau}, \frac{\beta+c}{2}, \frac{\beta-c}{2\tau})$ with 2 experimentations (one firm experiments to monopoly price and quantity; the other follows). The monopoly state does not belong to Ω^* and it follows from Part 2 that it is possible to transit further to ω' using a series of single experimentations. Hence, $z^*(\omega'', \omega') = 2$, which is equal to $R(\omega')$, the radius of ω' (see Part 1). From Theorem 3(iii) it follows that if ω'' is stochastically stable, then so is ω' . Since, $R(\omega') = z^*(\omega'', \omega')$ for any two states ω' and ω'' in Ω^* , it follows that all states in Ω^* are stochastically stable. ■

Notice that all stochastically stable states imply market clearing and spatially segregated markets (as $\bar{p} \geq \frac{\beta+c}{2} > \beta - \frac{\tau}{2}$). The price \bar{p} in combination with quantity $\bar{q} = \frac{\beta-\bar{p}}{\tau}$ results in zero profit for a firm if the rival prices according to marginal cost. The stochastically stable states are thus precisely those states that are immune against experimentations to marginal cost pricing. The value of \bar{p} , and hence the set of stochastically stable states, depends on the consumer valuation β . As \bar{p} is increasing in β and equals to monopoly price at the lowest permissible value of β (while respecting the condition stated in the proposition), the stochastically stable states are comprised of prices higher than the monopoly price. So, where most other papers based on imitation learning lead to prices that are more competitive than the Nash prediction, we see that prices can also be higher than the Nash benchmark. Though,

in spite of the higher price, profits are lower.

The condition $\beta < c + \tau$ implies that one firm is not able to attract the full mass of consumers with marginal cost pricing. This provides both firms with total control over part of the consumers: the consumers that never consider buying the rival's product. In the monopoly state it is possible for the rival firm to experiment to a price such that consumers are poached away to such an extent that the rival ends up with a higher profit. In a stochastically stable state, firms resort to catering to the set of consumers for which it is immune against any price experimentation by the rival firm. This results in the firms producing below monopoly quantity and charging price above monopoly level. It is precisely this set of "high prices" (higher than even the monopoly price) which not only gives each firm independent markets, but under the restriction that pricing below marginal cost is not permissible, makes it impossible (evaluated in terms of relative profit) for the other firm to encroach on the other's turf. This gives rise to the emergence of niche markets that are stochastically stable.

Proposition 4. *Suppose $\beta \geq c + \tau$. Then, all absorbing states are stochastically stable.*¹¹

Proof. For the current parameter conditions every state can be disrupted by an experimentation with marginal cost. Part 2 of the proof to Proposition 3 shows this for the states where the price is less than $\beta - \frac{\tau}{2}$. For the states where the price is larger than or equal to $\beta - \frac{\tau}{2}$, an experimentation to marginal cost yields a profit of 0 for the experimenting firm and leaves the other firm with a profit of $p(\frac{c-p+\tau}{2\tau}) - c\frac{\beta-p}{\tau}$. The latter profit is decreasing in β and is therefore maximized at $\beta = c + \tau$. At $\beta = c + \tau$ this profit is negative if and only if $p < c + \tau$ or $p > 2c$. So, for prices below $c + \tau$ this profit is negative. For prices above $c + \tau$ experimentation to marginal cost works trivially as it draws away all consumers.

Next, it is possible to reach the state $(c, \frac{1}{2}, c, \frac{1}{2})$ by a quantity experimentation. So, a $(c, \frac{1}{2}, c, \frac{1}{2})$ -tree with tree-resistance $|\Omega| - 1$ is easily constructed. This is the least resistance any tree can have. Next, we show that for any absorbing state a tree with resistance $|\Omega| - 1$ can be constructed. From Theorem 2 it follows that all absorbing states are stochastically stable.

First, consider the states (c, q, c, q) , where $q < \frac{1}{2}$. All that needs to be done in this case is to reverse the (c, q, c, q) to $(c, \frac{1}{2}, c, \frac{1}{2})$ edge in the just constructed $(c, \frac{1}{2}, c, \frac{1}{2})$ -tree and deleting all edges starting from (c, q, c, q) . The modified tree has the same resistance. Thus we have a (c, q, c, q) -tree, where $q < \frac{1}{2}$, with resistance $|\Omega| - 1$. Note that, in particular, the $(c, 0, c, 0)$ -tree has resistance $|\Omega| - 1$ and is stochastically stable.

Second, consider states of the form $(p, 0, p, 0)$. We can modify the just constructed $(c, 0, c, 0)$ -tree by reversing the $(p, 0, p, 0)$ to $(c, 0, c, 0)$ edge and deleting all edges starting

¹¹The marginal cost is required to be part of the price grid.

from $(p, 0, p, 0)$. This leaves the resistance of the tree unchanged. Hence, states of that form are stochastically stable.

Third, consider states of the type (p, q, p, q) , where $p \neq c$ and q is such that both firms receive at least zero profit. To construct the (p, q, p, q) -tree we modify the corresponding $(p, 0, p, 0)$ -tree by deleting the edges starting from (p, q, p, q) and building an edge from $(p, 0, p, 0)$ to (p, q, p, q) . The constructed tree has a resistance of $|\Omega| - 1$. ■

In Figure 1 the set of stochastically stable states mentioned in the proposition is captured by the thick solid lines. In the states with prices equal to marginal cost or quantity equal to zero, the firms receive zero profit. Also the states at the curved part of the boundary imply zero profit, as this curve is composed of price–quantity pairs for which $p \cdot \frac{\beta - p}{\tau} - c \cdot q$ equals zero. In all remaining states, the firms experience a positive profit. Unlike in the previous




Figure 1: Illustration of the set of stochastically stable states.

two propositions, now there exist stochastically stable states where firms actively compete for the marginal agent. This is the case in all states with prices below $\beta - \frac{\tau}{2}$. Also, unlike in the previous propositions, markets may not clear. Only in the states along the dotted path, markets clear. In all states above this path the market operates with excess supply, whereas below this path there is excess demand.

The huge multiplicity in stochastic stable states is mainly driven by our imitation dynamic that allows an active firm (positive quantity) with zero profit to imitate an inactive firm (zero quantity). Another plausible qualification that might be reasonable to consider is the imitation of active firms only, in line with the Activity Principle of Alós-Ferrer et al. (2000). Then starting from a zero profit absorbing state with positive production, such as $(c, \frac{1}{2}, c, \frac{1}{2})$, experimentation by one firm to zero production will not induce a transition (with positive probability) to the state $(c, 0, c, 0)$ in the unperturbed (imitation) dynamic. For a fine enough grid, this would affect the result (only for Proposition 4). States with an arbitrary price and no production – i.e. $(p, 0, p, 0)$ – are now excluded from being stochastically stable. The

refinement that is obtained is, therefore, not qualitatively substantive.

An alternative assumption on the imitation dynamic that significantly reduces the multiplicity in the number of stochastically stable states is to have firms imitating the firm with larger sales in case of equal profits. This is equivalent to firms having lexicographic preferences where they primarily care about relative profit and secondary about sales. Then, $(c, \frac{1}{2}, c, \frac{1}{2})$ will be the unique stochastically stable state. In particular, all states with excess demand or supply will no longer be stochastically stable. Remarkably, the only stochastically stable state implies prices to be equal to marginal costs, in spite of the differentiated market.

5 Discussion

Proposition 4 reveals that excess supply or excess demand can be a persistent feature of a horizontally differentiated market where firm behavior is guided primarily by imitation of the better off firm. Two natural questions arise: why do firms not cut down on production when there are unsold goods? and why do firms not produce more when there is unserved demand? If we were to say that a firm adjusts accordingly on the observance of non-market clearing, then we add something apart from imitation to our framework. By consideration of such adjustments, the non-market clearing states would no longer be stochastically stable simply because they cease to be absorbing states. Then, the result that all absorbing states are stochastically stable would hold – the difference of course being that all such states are market clearing.

Next, let us comment on the fragility or robustness of the findings with respect to alternative assumptions and consider possible extensions to the model. First, consider the particular structure of imitation and experimentation that we adopted. We have assumed that the event that one firm experiments in price and quantity is infinitely less likely than one firm experimenting in price or quantity. If, instead, we would assume the likelihood of a joint price–quantity experimentation to be of the same order as a sole price or quantity experimentation, we lose Proposition 3. For values of the consumers’ gross valuation of consumption above the bound mentioned in Proposition 2, we would obtain the result in Proposition 4. To see this, one has to note the strong parallel between the resulting dynamic and a model with relative profit maximization firms. In any equilibrium with relative profit maximizing firms, firms will realize equal absolute profit, since firms can guarantee zero relative profit by choosing the same price and quantity. In case a pure strategy equilibrium exists, it will coincide with the unique stochastically stable state of Proposition 2; otherwise, no selection among absorbing states is made and we obtain the result in Proposition 4.

Second, if we would assume that firms do not imitate the rival in case of equal profits at different strategies, the non-monomorphic states of this type would become absorbing.

The set of stochastically stable states in Proposition 4 would expand accordingly, but this modification would leave Propositions 2 and 3 unaltered.

Third, our model has been built on the restriction of transportation cost being less than the marginal cost of production. A natural question is what happens with our results if transportation costs would be higher? To a large extent, the results do not change.¹² The results of Proposition 2 and 4 – that for sufficiently low values of gross valuation of consumption the monopoly outcome will prevail and for sufficiently large values there would not be a strict selection among absorbing states – would remain unaltered, though for slightly different thresholds. Our result for intermediary values of consumer valuation (Proposition 3) would slightly change. Analogously, there would be a set of stochastically stable states that implies market clearing and a segregated market. Though, now, the monopoly outcome may be an element of this set. For this to be the case the transportation cost needs to be sufficiently high.

Fourth, let us consider the implications of our assumption of differentiated products. A perfectly homogenous market is a limiting case of our model (zero transportation cost). It is easily seen that stochastic stability loses all its refinement power and all absorbing states will be stochastically stable. After all, from any state where the price is different from the marginal cost, a firm experimenting to marginal cost pricing will attract all consumers who are willing to buy at marginal cost price. This yields a profit for the experimenting firm that is at least as much as that of its rival. With positive probability the rival will imitate and, hence, it takes a single experimentation to reach a state where the price is equal to marginal cost. Next, by an argument similar to the one made in the proof of Proposition 4, it is possible to construct a path from any absorbing state to another in which the edges have a resistance of one.

Finally, we take up the issue of observability of decisions and consequences. In our model, we assumed prices, quantities and profits to be perfectly observable. It may appear more realistic to assume that, rather than the quantity produced, the quantity sold is observable. This would lead to a disparity only when the quantity produced is in excess of the sales. This in turn can have an effect only if the more successful firm produces in excess of its demand. Hence, observability of sales can only have a differential effect on Proposition 4. In Proposition 4, at a certain price p , the maximum quantity that can be produced in a stochastically stable state is $q = p \frac{\beta - p}{c\tau}$ – the quantity that yields zero profit: $p(\frac{\beta - p}{\tau}) - cq = 0$. With observability of sales, any quantity greater than $\frac{\beta - p}{\tau}$ will not be imitated. Hence, states with price above $\beta - \frac{\tau}{2}$ and quantity in the interval $(\frac{\beta - p}{\tau}, p \frac{\beta - p}{c\tau}]$ that were stochastically stable earlier, will cease to be so. In case the price is below $\beta - \frac{\tau}{2}$ and production is above $\frac{1}{2}$, both

¹²Results are available upon request.

firms observe the rival to get equal profit at lower sales and imitate to the rival's sale level with positive probability. Hence, observability of sales rules out excess supply as possible property of a stochastically stable state. The main thrust of Proposition 4 that there exist stochastically stable states that are not market clearing remains unaltered, though only with excess in demand.

References

1. Alchian A (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy* 58 (3): 211-221.
2. Alós-Ferrer C, A Ania and KR Schenk-Hoppé (2000). An evolutionary model of Bertrand oligopoly. *Games and Economic Behavior* 33 (1): 1-19.
3. Alós-Ferrer C and G Kirchsteiger (2010). General equilibrium and the emergence of (non) market clearing trading institutions. *Economic Theory* 44 (3): 339-360.
4. Apesteguia J, S Huck and J Oechssler (2007). Imitation: Theory and experimental evidence. *Journal of Economic Theory* 136 (1): 217-235.
5. Apesteguia J, S Huck, J Oechssler and S Weidenholzer (2010). Imitation and the evolution of Walrasian behavior: Theoretically fragile but behaviorally robust. *Journal of Economic Theory* 145 (5): 1603-1617.
6. Bergin J and B Lipman (1996). Evolution with state-dependent mutations. *Econometrica* 64 (4): 943-956.
7. Blume L (2003). How noise matters. *Games and Economic Behavior* 44 (2): 251-271.
8. Ellison G (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies* 67 (1): 17-45.
9. Fudenberg D and D Levine (1998). *Theory of Learning in Games*. MIT Press, Cambridge MA.
10. Hamilton W (1970). Selfish and spiteful behavior in an evolutionary model. *Nature* 228 (5277): 1218-1220.
11. Hehenkamp B and Wambach A (2010). Survival at the center – The stability of minimum differentiation. *Journal of Economic Behavior and Organization* 76 (3): 853-858.
12. Herings P (1997). Endogenously determined price rigidities. *Economic Theory* 9 (3): 471-498.

13. Huck S, HT Normann and J Oechssler (1999). Learning in Cournot oligopoly – An experiment. *Economic Journal* 109 (454): C80-C95.
14. Huck S, HT Normann and J Oechssler (2000). Does information about competitors actions increase or decrease competition in experimental oligopoly markets?. *International Journal of Industrial Organization* 18 (1): 39-57.
15. Kandori M, GJ Mailath and R Rob (1993). Learning, mutation, and long-run equilibria in games. *Econometrica* 61 (1): 29-56.
16. Offerman T, J Potters and J Sonnemans (2002). Imitation and belief learning in an oligopoly experiment. *Review of Economic Studies* 69 (4): 973-997.
17. Van Damme E and J Weibull (2002). Evolution in games with endogenous mistake probabilities. *Journal of Economic Theory* 106 (2): 296-315.
18. Vega-Redondo F (1997). The evolution of Walrasian behavior. *Econometrica* 65 (2): 275-284.
19. Young HP (1993). The evolution of conventions. *Econometrica* 61 (1): 57-84.