Discussion Paper No. 53

Imitators and Optimizers in Cournot Oligopoly

Burkhard C. Schipper*

March 2005

*Burkhard C. Schipper, Department of Economics, University of Bonn. schipper@uni-bonn.de

Sonderforschungsbereich/Transregio 15 · www.gesy.uni-mannheim.de

Universität Mannheim · Freie Universität Berlin · Humboldt-Universität zu Berlin · Ludwig-Maximilians-Universität München
Rheinische Friedrich-Wilhelms-Universität Bonn · Zentrum für Europäische Wirtschaftsforschung Mannheim

Speaker: Prof. Konrad Stahl, Ph.D. · Department of Economics · University of Mannheim · D-68131 Mannheim,
Phone: +49(0621)1812786 · Fax: +49(0621)1812785

Imitators and Optimizers in Cournot Oligopoly[*]

Burkhard C. Schipper[†]
Department of Economics
University of Bonn

this version: March 2005

Abstract

We analyze a symmetric n-firm Cournot oligopoly with a heterogeneous population of optimizers and imitators. Imitators mimic the output decision of the most successful firms of the previous round a là Vega-Redondo (1997). Optimizers play a myopic best response to the opponents' previous output. Firms are allowed to make mistakes and deviate from the decision rules with a small probability. Applying stochastic stability analysis, we find that the long run distribution converges to a recurrent set of states in which imitators are better off than are optimizers. This finding appears to be robust even when optimizers are more sophisticated. It suggests that imitators drive optimizers out of the market contradicting a fundamental conjecture by Friedman (1953).

JEL-Classifications: C72, D21, D43, L13.

Keywords: profit maximization hypothesis, bounded rationality, learning, Stackelberg.

"Men nearly always follow the tracks made by others and proceed in their affairs by imitation, even though they cannot entirely keep to the tracks of others or emulate the prowess of their models. So a prudent man should always follow in the footsteps of great men and imitate those who have been outstanding." Niccolò Machiavelli

1 Introduction

One of the most fundamental assumptions in economics is that firms maximize absolute profits. However, already Alchian (1950) suggested that firms may maximize relative profits in the long run rather than absolute profits. In contrast, Friedman (1953) argued that evolutionary selection forces favor absolute profit maximization. In particular, he postulated that, although firms may not know their profit functions, we can assume that they behave as if they maximize profits because otherwise they would be driven out of the market by firms that do behave as if they maximize profits. Koopmans (1957), p. 140, remarked that if selection does lead to profit maximization then such an evolutionary process should be part of economic modeling. Taking Koopmans' suggestion into consideration, this paper describes an attempt to prove Friedman's conjecture. This attempt failed. That is, in the model presented here it turns out that Friedman's conjecture is false.

The present paper was partly inspired by Vega-Redondo (1997).[1] He shows that, in a quantity setting symmetric n-firm Cournot oligopoly with imitators, the long run outcome converges to the competitive output if small mistakes are allowed. Imitators mimic the output of the most successful firms in the previous round. His result is in sharp contrast to optimizers, whose outputs are known to converge under certain conditions in the Cournot tatonnement to the Cournot Nash equilibrium. It seems natural to wonder what happens if imitators and optimizers are mixed together in a heterogeneous population. According

---

[1]See also related work by Schaffer (1989), Rhode and Stegeman (2001), and Alós-Ferrer, Ania, and Vega-Redondo (1999).

to Friedman, we should find that optimizers are better off than are imitators, and that consequently optimizers drive out imitators in any payoff monotone selection dynamics. However, we find that imitators are strictly better off than are optimizers, which is at first glance a rather surprising result given that imitators are less sophisticated than optimizers. In a sense, this result is reminiscent of Stackelberg behavior. That's why we name the support of the long run distribution the set of Pseudo-Stackelberg states. First, imitators and optimizers play roles analogous to those of the "independent" and the "dependent" firms respectively in von Stackelberg's (1934) work.[2] Optimizers are "dependent" since by definition they play a best reply. Imitators are "independent" because they do not perceive any influence on the price but take it as given. Note however, that they do not conform exactly to the Stackelberg conjecture. Second, analogous to the profits of von Stackelberg's independent and dependent firms, every imitator is better off than every optimizer. Finally, our analysis retains the important aspect of von Stackelberg's idea: the modeling of asymmetries and behavioral heterogeneity of firms.

Imitators and optimizers differ with respect to the knowledge required to take their decisions. Whereas for imitators it is sufficient to know the previous period's outputs of every firm and their associated profits, optimizers need to know the total output of their opponents as well as their own profit function, which involves knowing inverse demand and costs, in order to calculate the myopic best response. Imitation is often associated with boundedly rational behavior but note that imitation of successful behavior can be also viewed as a rational rule of thumb (Vega-Redondo, 1997) when firms and decision makers have difficulties in perceiving their profit functions. They can easily judge their performance relative to other firms in the industry. This might be also one reason why a part of executives' remuneration-packages is often based on the firm's stock outperforming

---

[2]It is interesting to note that von Stackelberg himself never used the word "leader" in his book but spoke of the "independent" and the "dependent" firm. Today's familiar sequential representation of the Stackelberg game is not due to von Stackelberg. The idea of a game with a first mover advantage was introduced first without reference to Stackelberg (1934) as the "majorant game" by von Neumann and Morgenstern (1944), pp. 100. I thank Prof. Selten for pointing me to the "majorant game".

the market index or similar means of relative comparison.

In the proofs of our results, we rely on two main concepts, quasi-submodularity of payoff functions and stochastic stability analysis. Quasi-submodularity (see Topkis, 1998, pp. 43) is closely related to strategic substitutes (see Bulow, Geanakoplos, and Klemperer, 1985) and the dual single-crossing property (see Milgrom and Shannon, 1994). The intuition for quasi-submodularity in our context is that if a firm prefers a larger quantity to a lower quantity for a given total market quantity, then it prefers also the larger quantity to the lower quantity for a lower total market quantity. The Cournot oligopoly satisfies this property by definition (see Lemma 1). A similar version of this property is used in modern oligopoly theory (see Vives 2000, Amir, 1996, Amir and Lambson, 2000, etc.). Vega-Redondo's (1997) result can be generalized to a class of quasi-submodular games (see Schipper, 2003).

Following Kandori, Rob, and Mailath (1993) and Young (1993), the dynamic analysis in this paper uses the concept of stochastic stability developed by Freidlin and Wentzel (1984) (see also Ellison, 2000 and others). The general idea is that mutations select among absorbing sets of the decision process such that only the most robust absorbing sets remain in the support of the limiting invariant distribution. There are several alternative interpretations of the noise in our context. First, firms are assumed to innovate with a small probability in a sense of experimenting with various output levels. Second, firms are assumed to be boundedly rational such that there is always a small positive probability of making mistakes in output decisions. Finally, every period, a small fraction of the firms is replaced by newcomers who choose their output from tabula rasa. Any of those interpretations adds some realistic feature to the model. Instead making use of the graph theoretic arguments developed by Freidlin and Wentzel (1984) as well as Kandori, Rob, and Mailath (1993) and Young (1993), we employ a simpler necessary condition for stochastic stability introduced by Nöldeke and Samuelson (1993, 1997) and Samuelson (1994). They show that a necessary condition for a state to be contained in the support of the unique invariant limiting distribution is that this state is contained in the minimal set of absorbing sets that is robust to a single mutation. Such a set is

called a recurrent set. In our main result we show that the symmetric Cournot Nash equilibrium, the only absorbing state in which optimizers are as well off as imitators, is not the unique stochastically stable state. Moreover, we also show in an example that there are assumptions on the parameters of the game such that the entire set of Pseudo-Stackelberg states is the unique recurrent set. In any case, the support of the unique limiting invariant distribution implies that imitators are strictly better off than are optimizers.

Apart from a pure theoretical interest, the analysis presented here is of practical relevance since imitation, in the form of "benchmarking" and "best practices", is widely used in today's management. Given that such imitative behavior exists among other decision rules in today's business practice, it is only natural for theorists to investigate imitation as well as the heterogeneity of decision rules.

Conlisk (1980) also analyzes a dynamic model with imitators and optimizers. However, he takes the cost of optimizing into account, and this cost is a key for obtaining his results. Our result appears to be stronger since in our work imitators are better off than are optimizers even without any optimizers' cost of sophistication. Conlisk's (1980) result has a similar flavor to Stahl (1993), who concludes using a different approach that dumb players may never die out and smart players with maintenance costs may vanish. Using a different approach, Banerjee and Weibull (1995) study optimizers and players that are programmed to actions in evolutionary symmetric 2-player games. They show that long run resting states hold a positive share of programmed players. There has been extensive research on imitation in game theory. For instance, Schlag (1998) analyzes various imitation rules in multi-armed bandit problems and shows that a certain type of imitation rule is optimal. Gale and Rosenthal (1999) study imitators and experimenters where former mimic to a certain extent the population average. Roughly they find that the population converges to the Nash equilibrium in various games with a unique equilibrium, but note that their imitators differ from ours. Kaarbøe and Tieman (1999) study imitators and myopic optimizers in strict supermodular games and find among others that the set of absorbing sets corresponds to the set of Nash equilibria. This is

in contrast with the strict submodular game studied in our paper, for which there are also other absorbing states than the Nash equilibrium. Research on Friedman's profit maximization hypothesis has been done for example by Blume and Easley (2002) and Sandroni (2000), who find support for it in a general equilibrium context. Dutta and Radner (1999) show in a model with entrepreneurs and capital markets that other behaviors than profit maximization may survive. The present paper is also related to the literature on interdependent preferences. In particular, Koçkesen, Ok, and Sethi (2002) found that players who also care about relative payoffs may have a strategic advantage in a class of symmetric games including the Cournot game. Note that imitators do care about relative payoffs since their decision rule involves a comparison of profits among firms.

The paper is organized as follows: Section 2 introduces the model and the decision rules. It is followed in section 3 by an informal discussion of candidates for solutions. Section 4 presents the results, which are subsequently discussed in the concluding section 5. All proofs are contained in the appendix. The required mathematical tools are introduced along the way.

2 Basic Model and Decision Rules

This section outlines the basic model in the spirit of Cournot (1838), pp. 79. Consider a finite number of firms $N = \{1, 2, ..., n\}$ and a market for a homogeneous good. Inverse demand is given by a function $p : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$. For every total output quantity $Q \in \mathbb{R}_+$ this function specifies the market clearing price $p(Q)$. By the assumption of symmetry, every firm $i \in N$ faces the same demand and possesses the same production technology. Hence the cost functions $c : \mathbb{R}_+ \longrightarrow \mathbb{R}_+$ are identical. For each firm it is a function of the quantity q_i it produces. Let the total output over all firms be $Q := \sum_{i \in N} q_i$. For later analysis, it will be convenient to write profits as function from the individual quantity

and the total quantity,

$$\pi_i(q_i, Q) := q_i p(Q) - c(q_i), \forall i \in N. \tag{1}$$

We restrict our analysis to a symmetric oligopoly since imitation is more reasonable if firms face similar conditions of production.

For technical reasons we assume that firms choose output from a common finite grid $\Gamma = \{0, \delta, 2\delta, ..., \nu\delta\}$, where both $\delta > 0$ and $\nu \in \mathbb{N}$ are arbitrary. This turns the strategic situation into a game with a finite action space and allows us to focus on finite Markov chains later in the dynamic analysis.

In the proofs of our results, the following observation will be crucial. This observation does not require any additional assumptions (other than Assumption 1) but is a property of the Cournot oligopoly. Therefore we introduce it here instead later in the text.

Definition 1 (Submodularity) *π_i is submodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$ if for all $q_i'' > q_i'$, $Q'' > Q'$*

$$\pi_i(q_i'', Q') - \pi_i(q_i', Q') \geq \pi_i(q_i'', Q'') - \pi_i(q_i', Q''). \tag{2}$$

It is strictly submodular if Inequality (2) holds strictly.

Assumption 1 (Strictly Decreasing Demand) *For all $Q, Q' \in \{0, \delta, 2\delta, ..., n\nu\delta\}$, if $Q' > Q$ then $p(Q') < p(Q)$.*

Lemma 1 *By Assumption 1, π_i is strictly submodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$.*

If Assumption 1 is modified such that p is weakly decreasing then π_i is submodular but not strictly submodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$.

Remark 1 (Quasi-Submodularity) *Submodularity implies that π_i is quasi-submodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$ (but not vice versa), i.e., for all $q_i'' > q_i'$, $Q'' > Q'$*

$$\pi_i(q_i'', Q'') \geq (>)\pi_i(q_i', Q'') \implies \pi_i(q_i'', Q') \geq (>)\pi_i(q_i', Q'), \tag{3}$$

$$\pi_i(q_i', Q') \geq (>)\pi_i(q_i'', Q') \implies \pi_i(q_i', Q'') \geq (>)\pi_i(q_i'', Q''). \tag{4}$$

The observation that the payoff function is quasi-submodular in the individual quantity and the total output is used later in the proofs repeatedly. Note that this property follows directly by the structure of the Cournot game. No additional assumptions on the game have to be imposed.

The dynamics of the system is assumed to proceed in discrete time, indexed by $t = 0, 1, 2, ...$ At each t the state of the system is identified by the current output schedule

$$\omega(t) = (q_1(t), q_2(t), ..., q_n(t)).$$

Thus, the state space of the system is identical to Γ^n. Associated with any such state $\omega(t) \in \Gamma^n$ is the induced profit profile $\pi(t) = (\pi_1(t), \pi_2(t), ..., \pi_n(t))$ at t, defined as follows:

$$\pi_i(t) := q_i(t)p(Q(t)) - c(q_i(t)), \forall i \in N. \tag{5}$$

Assumption 2 (Inertia) *At every time $t = 1, 2, ...$, each firm $i \in N$ has regardless of history an i.i.d. probability $\rho \in (0, 1)$ of being able to revise her former output $q_i(t-1)$.*

Note that since $0 < \rho < 1$ the process has inertia. That is, not every period all firms adjust output. The idea is that it is too costly to always adjust output. Moreover, it will become clear later on that with this assumption we rule out cycles of the best response dynamics.

The finite population of firms N is partitioned into two subpopulations of imitators and optimizers respectively. Let I be the subset of N that contains all imitators. The fraction of imitators in the population is denoted by $\theta = \frac{\sharp I}{\sharp N}$. The firms in the two subpopulations are characterized by different decision rules. The idea of a decision rule is appropriately summarized by Nelson and Winter (1982, p. 165) who write that *"...at any time, firms in an industry can be viewed as operating with a set of techniques and decision rules (routines), keyed to conditions external to the firm ... and to various internal state conditions..."* Conventional economics focuses mainly on profit maximization. However,

"benchmarking", "best practices", and other imitation rules can be found in today's management practice.

Definition 2 (Imitator) *An imitator $i \in I$ chooses with full support from the set*

$$D_I(t-1) := \{q \in \Gamma : \exists j \in N \ s.t. \ q = q_j(t-1) \ and \ \forall k \in N, \pi_j(t-1) \geq \pi_k(t-1)\}. \quad (6)$$

Definition 3 (Optimizer) *An optimizer $i \in N \backslash I$ chooses from the set*

$$D_O(t-1) := \{q \in \Gamma : q \in b(q_{-i}(t-1))\}, \quad (7)$$

with $q_{-i} := \sum_{j \in N \backslash \{i\}} q_j$ and $b : \{0, \delta, 2\delta, ..., (n-1)\nu\delta\} \longrightarrow \Gamma$ is defined to be firm's best reply correspondence

$$b(q_{-i}) := \{q_i' \in \Gamma : q_i'p(q_{-i} + q_i') - c(q_i') \geq q_i p(q_{-i} + q_i) - c(q_i), \forall q_i \in \Gamma\}. \quad (8)$$

It is assumed that initially in $t = 0$ every firm starts with an arbitrary output within the admissible domain Γ.

The imitation rule is explained as follows: Every period there exists a firm j that had the highest profit in the previous period. An imitator imitates the previous period's quantity of firm j. It is the same imitation rule as used by Vega-Redondo (1997). Definition 3 means that an optimizer sets an output level that is a best reply to the opponents' total output in the previous period. In the last section we discuss how our results generalize to more sophisticated optimizers.

The process induced by the decision rules is a n-vector discrete time finite Markov chain with stationary transition probabilities. Finiteness is provided by the finite state space Γ^n. It is a vector process since each ω is a vector in Γ^n. Due to the myopic decision rules, the process has the Markov property, namely $prob\{\omega(t+1)|\omega(t), \omega(t-1), ..., \omega(t-k)\} = prob\{\omega(t+1)|\omega(t)\}$. That is, $\omega(t)$ contains all the information needed to determine transition probabilities. Since the decision rules themselves do not change over time, the process has stationary transition probabilities $prob\{\omega'(t+1)|\omega(t)\} = prob\{\omega'(t+k+1)|\omega(t+k)\}$, $k = 0, 1, ...$

The Markov operator is defined in the standard way as the $\sharp \Gamma^n \times \sharp \Gamma^n$-transition probability matrix $P = (p_{\omega\omega'})_{\omega,\omega' \in \Gamma^n}$ with $p_{\omega\omega'} = prob\{\omega'|\omega\}$, $p_{\omega\omega'} \geq 0$, $\omega, \omega' \in \Gamma^n$ and $\sum_{\omega' \in \Gamma^n} p_{\omega\omega'} = 1$, for all $\omega \in \Gamma^n$. That is, the element $p_{\omega\omega'}$ in the transition probability matrix P is the conditional probability that the state is in ω' at $t+1$ given that it is in ω at t. According to this definition of a Markov transition matrix, probability distributions over states are represented by row vectors.

We conclude the model of the decision processes with following assumption:

Assumption 3 (Noise) *At every output revision opportunity t, each firm follows her decision rule with probability $(1-\varepsilon)$, $\varepsilon \in (0, a]$, a being small, and with probability ε she randomizes with full support Γ.*

As a matter of convention, we call a firm mutating at t if it randomizes with full support at t. The noise has a convenient technical property. Let $P(\varepsilon)$ be the Markov chain P perturbed with the level of noise ε. Then by Assumption 3, $P(\varepsilon)$ is regularly perturbed (Young, 1993, p. 70), i.e., it is an ergodic and irreducible Markov chain on Γ^n. This implies that there exists a unique invariant distribution $\varphi(\varepsilon)$ on Γ^n (for standard results on Markov processes see for example Masaaki, 1997). To put it more intuitively, the noise makes any state accessible from any other state in finite time. This is sufficient for the existence of the unique invariant distribution.

The following analysis focuses on the unique limiting invariant distribution φ^* of P defined by $\varphi(\varepsilon)P(\varepsilon) = \varphi(\varepsilon)$, $\varphi^* := \lim_{\varepsilon \to 0} \varphi(\varepsilon)$ and $\varphi^* P = \varphi^*$. In particular, the focus is on how to characterize this probability vector since it provides a description of the long run output behavior of the market when the noise goes to zero. For that reason we will refer to it also as the long run distribution. It determines the average proportion of time spent in each state of the state space in the long run, or expressed differently, the relative frequency of a state's appearance as the time goes to infinity (see Fudenberg and Levine, 1998, or Samuelson, 1997, for an introduction and discussion of this method).

3 Candidates for Solutions

In this section we informally discuss candidates for solutions. By standard results (e.g. see Samuelson, 1997, Proposition 7.4) we know that the support of the long run distribution can only contain states that are elements of absorbing sets of the unperturbed process. Therefore we consider first the case of no noise, $\varepsilon = 0$, and define an absorbing set $A \subseteq \Gamma^n$ in the standard way by

(i) $\forall \omega \in A$, $\forall \omega' \notin A$, $p_{\omega \omega'} = 0$, and

(ii) $\forall \omega, \omega' \in A$, $\exists m \in \mathbb{N}$, m finite, s.t. $p_{\omega \omega'}^{(m)} > 0$, $p_{\omega \omega'}^{(m)}$ being the m-step transition probability from ω to ω'.

Vega-Redondo (1997) showed that a homogeneous population of imitators converges to the competitive solution.

Definition 4 (Competitive Solution) *The competitive solution $\omega^* = (q_1^*, ..., q_n^*)$ is defined by for all $i \in N$,*

$$q_i^* p(Q^*) - c(q_i^*) \geq q_i p(Q^*) - c(q_i), \forall q_i \in \Gamma, \tag{9}$$

with $Q^* = \sum_{i \in N} q_i^*$.

Can the competitive solution be an absorbing state given a heterogeneous population of imitators and optimizers? Suppose that the competitive solution exists uniquely in the grid. Consider first the imitators. Every firm plays its share of the competitive solution. By symmetry all firms make identical profits. Thus nobody is better off and imitators have no reason to deviate from their output. However, since n is finite, optimizers do not generally play a best reply. Each optimizer's share of the competitive output is larger than the best response. Hence they will deviate to the best response leading to a state different from the competitive solution. It follows that the competitive solution is not an absorbing state.

Consider now a state where every firm sets its symmetric Cournot Nash equilibrium output assuming that it exists in the grid Γ and that it is unique.

Definition 5 (Cournot Nash Equilibrium) *A combination of output strategies $\omega^\circ = (q_1^\circ, q_2^\circ, ..., q_n^\circ) \in \Gamma^n$ is a Cournot Nash equilibrium if for all $i \in N$,*

$$q_i^\circ p(Q^\circ) - c(q_i^\circ) \geq q_i p(Q^\circ - q_i^\circ + q_i) - c(q_i), \forall q_i \in \Gamma. \tag{10}$$

It is known that in a homogeneous population of optimizers the Cournot Nash equilibrium is under certain assumptions guaranteeing global convergence the solution of a sequential best response process. In a heterogeneous population, imitators do not deviate since all firms set identical outputs and earn identical profits. Optimizers do not deviate too since they anyway set their best response quantities. Thus the symmetric Cournot Nash equilibrium is an absorbing state. However, is it the unique absorbing state? Consider the following state:[3]

Definition 6 (Pseudo-Stackelberg Solution) *The Pseudo-Stackelberg solution is a state $\omega^S = (q_1, ..., q_{\theta n}, q_{\theta n+1}, ..., q_n)$ that satisfies the following conditions:*

(i) for all $i \in I$, $q_i = q^S$ s.t.

$$q^S p(\theta n q^S + (1 - \theta)n q^D) - c(q^S) > q p(\theta n q^S + (1 - \theta)n q^D) - c(q), \forall q \neq q^S, \tag{11}$$

(ii) for all $i \in N \backslash I$, $q_i = q^D$,

$$q^D := b(\theta n q^S + ((1 - \theta)n - 1)q^D). \tag{12}$$

In the Pseudo-Stackelberg solution all imitators set the identical output. This output maximizes profits of imitators given that they do not perceive any influence on the price and the optimizers set the identical best reply. Clearly, this outcome has features of the

---

[3]We assume here that the best reply is unique. The uniqueness condition later in Assumption 4 ensures that the best reply to the opponents' output is indeed a singleton (see Lemma 2).

competitive solution (for imitators) and the Cournot Nash equilibrium (for optimizers). If $\theta = 1$, then it is identical to the competitive solution since Inequality (12) becomes vacuous. If $\theta = 0$, then it is identical to the Cournot Nash equilibrium since Inequality (11) becomes vacuous. We call this outcome the *Pseudo-Stackelberg Solution* because of its obvious similarities and differences to the notion of Stackelberg solution in the literature. Analogous to the profits of von Stackelberg's (1934) independent and dependent firms, every imitator is strictly better off than is every optimizer since Inequality (11) holds for for all $q \in \Gamma$, $q \neq q^S$, so also for q^D. I.e., it follows that[4]

$$\pi_i(q^S, q^D, n, \theta) > \pi_j(q^S, q^D, n, \theta), \forall i \in I, \forall j \in N \backslash I.$$

Every imitator is strictly better off than every optimizer.

Why is the Pseudo-Stackelberg solution an absorbing state? Assume that the Pseudo-Stackelberg solution exists in Γ^n.[5] Consider first the imitators: every imitator sets the identical output and is strictly better off than is any optimizer. Hence an imitator has no reason to deviate from its output. Optimizers do not deviate too from their output since they play the best response. Thus the Pseudo-Stackelberg solution is an absorbing state.

In the following text, we will reserve q^S to denote the identical individual output of any imitator in the Pseudo-Stackelberg solution. q^I means that the individual quantity q^I is set by *each* imitator (superscript "I" stands for "independent" or all "imitators"). Analogously, q^D means that the individual quantity q^D is set by *each* optimizer (superscript "D" stands for "dependent"). The analogous notation applies to the profit functions π^I and π^D. Generally, a superscript indicates identical individual values for all firms within a subpopulation whereas a subscript indicates individual not necessarily

---

[4]For notational convenience we write $\pi_i(q, q', n, \theta)$ for $\pi_i(q, \theta nq + (1-\theta)nq')$ if $i \in I$, or for $\pi_i(q', \theta nq + (1-\theta)nq')$ if $i \in N \backslash I$.

[5]Existence of Pseudo-Stackelberg solution is analogous to existence of competitive solution in Vega-Redondo (1997). Standard assumptions on costs, i.e., strictly increasing marginal costs and small fixed costs, suffice. By the strict Inequality (11), quasi-submodularity, and Assumption 4 in the next section, the Pseudo-Stackelberg solution must be unique if it exists (see Lemma 2 (vi)).

identical values.

Previous arguments suggest already that the Cournot Nash equilibrium and the Pseudo-Stackelberg solution may not be the only candidates for solutions. To facilitate the analysis we define the following set of states:[6]

Definition 7 (Pseudo-Stackelberg States) *The set of Pseudo-Stackelberg states H consists of all states $\omega = (q_1, ..., q_{\theta n}, q_{\theta n+1}, ..., q_n) \in \Gamma^n$ that satisfy the following properties:*

(i) $q_i = q^I$, for all $i \in I$ and some $q^I \in \Gamma$,

(ii) $q_i = q^D$, for all $i \in N \setminus I$, $q^D := b(\theta n q^I + ((1 - \theta)n - 1)q^D)$,

(iii) $\pi^I(q^I, q^D, n, \theta) \geq \pi^D(q^I, q^D, n, \theta)$,

(iv) $\pi^I(q^I, q^D, n, \theta) = \pi^D(q^I, q^D, n, \theta)$ iff $q^I = q^D$.

If condition (i) is not satisfied, then an imitator may mimic a different output decision from another imitator if the latter happens to have higher profits. If condition (ii) is not satisfied, all optimizers that don't play a best reply will have an incentive to deviate. If condition (iii) is not satisfied, imitators will mimic optimizers. To understand the motivation of (iv) note that by symmetry, $q^I = q^D$ implies $\pi^I(q^D, q^I, n, \theta) = \pi^D(q^D, q^I, n, \theta)$. To see the purpose of the other direction note that if $\pi^I(q^D, q^I, n, \theta) = \pi^D(q^D, q^I, n, \theta)$ and $q^I \neq q^D$ then imitators would be indifferent between q^I and q^D thus adding a source of instability.

In each Pseudo-Stackelberg state, imitators are weakly better off than are optimizers. In fact, imitators are strictly better off in any Pseudo-Stackelberg state except the Cournot Nash equilibrium, the only state where optimizers are as well off as imitators.

It is clear that the set of Pseudo-Stackelberg states is nonempty since the Cournot Nash equilibrium - assume that it exists - belongs to it. Moreover, it is easy to see that

---

[6]Again, we assume here that the best reply is unique. The uniqueness condition later in Assumption 4 ensures that the best reply to the opponents' output is indeed a singleton (see Lemma 2).

the competitive solution is not a Pseudo-Stackelberg state since optimizers do not set a best reply in the competitive solution (unless $n \to \infty$ or $\theta = 1$). Finally, if c is strictly convex then the Pseudo-Stackelberg solution is a Pseudo-Stackelberg state since $q^S > q^D$ are such that $\pi^I(q^S, q^D, \theta, n) > \pi^D(q^S, q^D, \theta, n)$. Thus properties (i) to (iv) of Definition 7 of Pseudo-Stackelberg states are satisfied. If c is not strictly convex, then condition (iv) may be violated. To see this, assume that costs are linear (weakly convex). Imitators make zero profits when price equals marginal costs. The optimizers' best response is zero output. Then imitators are indifferent between zero output and q^S. If costs are strictly convex, then imitators make strict positive profits and optimizers set a positive output level which is lower than q^S. Thus each optimizer makes less profit than any imitator.

4   Results

Before we state and prove the results in this section, we need to introduce formally an assumption. As before let q_{-i} denote the total output of all firms but i.

Assumption 4 *For $q'_{-i} < q_{-i}$, $q' \in b(q'_{-i})$, $q \in b(q_{-i})$, we have*

$$0 > \frac{q' - q}{q'_{-i} - q_{-i}} > -1. \tag{13}$$

This assumptions states that the slopes of the best reply correspondence are strictly lower than 0 and strictly larger than -1. Former implies by Dubey, Haimanko, and Zapechelnyuk (2005) that the game is a pseudo-potential game and has a Cournot Nash equilibrium. Moreover, since it is a pseudo-potential game there exists a finite improvement path such that sequential best reply converges to the Cournot Nash equilibrium. Note, that the assumption of the existence of a pseudo-potential is weaker than of an exact, weighted or ordinal potential (Monderer and Shapley, 1996). Most Cournot games in the literature are games with strategic substitutes and thus pseudo-potential games.

The assumption that the slopes of the best reply correspondence are strictly larger than -1 is made in order to obtain a unique best reply. Vives (2000, Theorem 2.8) shows

in a simple proof that if a Cournot Nash equilibrium exists and the above assumption holds, then it must be unique. Note that the condition is equivalent to if $q'_{-i} < q_{-i}$ then $q'_{-i} + q' < q_{-i} + q$. It means that total output is strictly increasing in the opponents' output when the player sets best responses. Since we have a symmetric game, the uniqueness condition implies that the unique Cournot Nash equilibrium is symmetric (Vives, 2000, Remark 17) and that the best reply correspondence is in fact a function (see Vives, 2000, p. 43). In Lemma 2 we show that by Assumptions 4, total output is increasing in imitators output, and that if the Pseudo-Stackelberg solution exists, it must be unique.

We are finally ready to state our results. Let Z be the collection of all absorbing sets in Γ^n. Recall that H is the set of all Pseudo-Stackelberg states (Definition 7).

Proposition 1 *Under Assumptions 1, 2 and 4, $Z = H$, whereby each Pseudo-Stackelberg state is an absorbing state.*

Let S denote the support of the long run distribution φ^*. By standard results Proposition 1 implies that $S \subseteq H$.

Corollary 1 *Under previous assumptions, in the long run imitators are weakly better off than are optimizers.*

The question we answer next is whether the noise selects among absorbing states.

Theorem 1 *If $\theta \in (0, 1]$ and $\omega^S \in H$, then under above assumptions it is never true that $S = \{\omega^\circ\}$.*

Theorem 1 states that if the Pseudo-Stackelberg solution exists in H then the Cournot Nash equilibrium is not the unique long run outcome. Note that the Cournot Nash equilibrium is the only Pseudo-Stackelberg state in which optimizers are as well off as are imitators. It is worth to put following implication on record:

Corollary 2 *If $\omega^S \in H$ then in the long run imitators are strictly better off than are optimizers.*

Since the Cournot state ω°, the only state in which optimizers are as well off as imitators, is never the unique long run outcome, the long run distribution must put strict positive weight on some other state in the non-singleton set of Pseudo-Stackelberg states. Hence, in the long run imitators are strictly better off than are optimizers.

Theorem 1 does not exclude any absorbing states from the support of the long run distribution. In fact, we prove the following result:

Example 1 *Consider for example $p(Q) = 10 - Q$, $c(q_i) = \frac{1000}{501}q_i^2 + 1$, $\theta = 0.2$, $\delta = 0.001$, $n = 5$ and a sufficiently large ν. Then $S = H$, i.e., the support of the long run distribution comprises of the entire set of Pseudo-Stackelberg states (see appendix).*

Above example shows that one can find reasonable assumptions on functions p and c and parameters θ, δ, ν and n that are sufficient for the entire set of Pseudo-Stackelberg states to be the support of the unique limiting invariant distribution. While the example appears rather standard (i.e., linear demand, convex cost), it takes quite a bit of proof to obtain the result (see appendix). In the proof of the example, we show in particular that if profit functions are strictly quasi-concave and the Pseudo-Stackelberg solution exists, then we can find a sequence of single mutations by which we can move through the set of Pseudo-Stackelberg states, starting from the Cournot Nash equilibrium state up to the Pseudo-Stackelberg solution as well as starting from the Pseudo-Stackelberg state with the largest output of imitators down to the Pseudo-Stackelberg solution (Lemma 5). For sufficiently large ν we can also show that the Pseudo-Stackelberg solution can be destabilized by a sufficiently large mutation that leads subsequently back to either the Cournot Nash equilibrium or the Pseudo-Stackelberg state with the largest output of imitators (Lemma 7 and 8). Thus the assumptions in the example are sufficient to show that any Pseudo-Stackelberg state can be connected to any other Pseudo-Stackelberg state by a sequence of single suitable mutations. Hence, we conclude by a result by Nöldeke and Samuelson (1993, 1997) and Samuelson (1994) that the set of Pseudo-Stackelberg states is the unique recurrent set and the support of the long run distribution.

Note that we obtain known results for homogeneous populations of either imitators or optimizers as extreme cases. If there is a homogeneous population of optimizers ($\theta = 0$), then the Pseudo-Stackelberg solution is the Cournot Nash equilibrium. In this case, the set of Pseudo-Stackelberg states is a singleton containing the Cournot Nash equilibrium only. Hence Proposition 1 implies that the Cournot Nash equilibrium is the unique absorbing set. Mutations does not matter since mutations must also select the unique absorbing set. If there is a homogeneous population of imitators ($\theta = 1$), then the Pseudo-Stackelberg solution is equivalent to the competitive solution. The proof of Lemma 4 in the appendix implies Vega-Redondo's (1997) result.[7] I.e., the competitive solution is the unique long run outcome.

Following simple example illustrates a process with mutations and adjustments.

Example 2 *Consider for example three players, $n = 3$, two imitators $1, 2$ and one optimizer D. The inverse demand function is given by $p(Q) = 10 - q_D - q_1 - q_2$ and the cost function by $c(q_i) = \frac{1}{2} q_i$, $i = 1, 2, D$. Following Figure 1 illustrates an example of process with mutations and adjustments. The upper graph plots the quantities over time, the lower one the profits. At $t = 1$ we start in an arbitrary starting state. It happens that the optimizer makes the highest profit. Thus imitators mimic in period 2. The optimizer adjusts in period 3, and we reach an absorbing state, in which the optimizer is worse off than is any imitator. At period 4, imitator 1 innovates with the quantity of the Pseudo-Stackelberg solution. Imitator 2 mimics imitator 1 and the optimizer adjusts in period 5. The optimizer adjusts again in period 6 and we reach the Pseudo-Stackelberg solution, which is an absorbing state. In the following periods we illustrate that also the Pseudo-Stackelberg solution can be destabilized by single large mutation and how we reach another absorbing state in period 12.*

---

[7]In this case, Inequality (15) holds also for $k = n - 1$. Thus more than one mutation is needed to escape the competitive solution.

Figure 1: Process in Example 2

5 Discussion

The significance of the previous results stems from the following conclusion: If imitators are strictly better off than are optimizers, then any payoff monotone selection dynamics (see for example Weibull, 1995) on the long run profits selects imitators in favor of optimizers. That is, an evolutionary dynamics reflecting the paradigm of "survival of the fittest" will show that imitators drive optimizers out of the market. Thus Friedman's (1953) conjecture is false in the oldest formal model of market competition in economics, the Cournot oligopoly (Cournot, 1838). In a working paper version (Schipper, 2002) we make this argument precise by showing how imitators drive out optimizers in the example of the discrete time finite population replicator dynamics on the long run profits. The intuition there is that firms enter each market day with a fixed decision rule and the market day takes as long as the long run outcome of outputs to emerge. Before markets are reopened the next day, the "evolutionary hand" chooses for each firm the decision rule selecting effectively among firms. Alternatively, one can assume that at the end of each day, the management of every firm holds independently a strategy meeting to decide on its decision rule for the next day according to the relative performance of their current decision rule. The market sessions are repeated day for day. One can show that a homogeneous population of imitators is the unique asymptotically stable population state. From this evolutionary prospectus we can not assume in economics that firms behave as if they maximize absolute profits.[8] After all Vega-Redondo's (1997) imitators are supported by those evolutionary arguments. The same holds for Alchian's (1950) suggestion since imitators want to be as well off as others, which is closely related to relative profit maximization.

There are a few critiques we like to address. First, one may criticize the limitations of the optimizers. Playing myopic best response is not really sophisticated optimization.

---

[8]Alternatively, one may want to extend the Markov chain to a product set of the output space and the decision rule space. If we assume that the probability of revising the decision rule is sufficiently lower than the probability of adjusting the output, then the same result is likely to emerge.

Consider what happens if we make the optimizers more and more sophisticated. Suppose first that we would allow optimizers to take a longer history of output decisions into account when deciding which output level to set. Then results are not likely to change but convergence may be slower since the optimizers' adjustment process becomes similar to fictitious play.[9] Second, suppose that optimizers are able forecast the behavior of the imitators. What does it help them if imitators set some large output, which happens in finite time by the noise assumed? All the optimizers can do is playing best reply against their beliefs leading them to play a smaller output with smaller profits than imitators. Even if they could temporarily "low-ball" the imitators with some quantities, imitators would erase any profit advantage by mimicking those quantities such that it never shows up in the support of the long run distribution.

In this context it is natural to ask, why optimizers do not just mimic imitators? Suppose they do. Then all firms behave as if they are imitators and Vega-Redondo's (1997) result of a competitive solution would emerge. However, in the competitive solution every optimizer has an incentive to deviate to its lower best reply output since it would increase its profit although it increases the profits of imitators even more. The imitation rule is a commitment technology, which the optimizer does not like to adopt, not because it involves some investment cost but because the optimizer is worse off in absolute terms when adopting the technology although it can improve its relative standing.

Note that our result is likely to break down if we make imitators more smart. I.e., consider imitators that take a longer memory than just one period into account. Then they may remember that they decreased their absolute profits (although they may increased their relative profits) and return to their former output. Alós-Ferrer (2004) shows in Vega-Redondo's (1997) framework that the long run outputs of imitators with longer (but finite) memory converge to a set of monomorphic states between the Cournot Nash equilibrium and the competitive solution. An analogous result is likely to hold in our setting with a heterogeneous population of optimizers and imitators with a longer memory.

---

[9]Regarding fictitious play refer for example to Fudenberg and Levine (1998), pp. 29.

However, our arguments in this article suggest that similar to optimizers, such imitators with longer memory may do worse against imitators with just a single-period memory.

A second critique could aim at the semantics of profit optimization. Obviously in my setting the optimizers are absolute profit maximizers in regard to their objective but not in terms of the result. This highlights the ambiguity of profit maximization in Cournot oligopoly. Aiming to maximize absolute profit may not be the way to actually achieve the highest relative profit. We show that the standard text book understanding of profit maximizing firms can not be supported by evolutionary arguments in Cournot oligopoly. Our awareness of the ambiguity of "profit maximization" in a class of games is an insight gained from this analysis.

A third possible critique point is of a more technical nature. We use the concept of stochastic stability developed by Freidlin and Wentzel (1984) as well as Kandori, Rob, and Mailath (1993) and Young (1993). In many applications of this concept in literature (for a partial review of the increasing literature using this method see for example Fudenberg and Levine, 1998, chapter 5), the characterization of the long run distribution involves a comparison of a multiplicity of highly unlikely mutations. A meaningful application of this method must address the question about the speed of convergence. How long does it take for the long run outcome to emerge? The advantage of applying the concept of recurrent set by Nöldeke and Samuelson (1993, 1997) and Samuelson (1994) is that one can conclude immediately that just a single suitable mutations is required to trigger the long run outcome. That is, in our model convergence to the long run outcome is comparatively rather fast. In Theorem 1 we show that the Cournot Nash equilibrium is not robust against a single mutation. Instead using stochastic stability as a refinement tool, we use it as a robustness check. In Example 1 we show that there are reasonable assumptions on the game such that no proper subsets of absorbing states can be selected by stochastic stability. Thus we show that the concept of stochastic stability is of limited use for a refinement of our results in this model.

The key property driving the result is the observation that the payoff functions of

a Cournot game are quasi-submodular in the individual and the total quantity. This is closely related to strategic substitutes. It suggests that the same result holds in other games with strategic substitutes such as Cournot oligopoly with differentiated substitute products, Bertrand oligopoly with differentiated complementary products, some rent seeking games, common pool dilemmatas etc. For instance, consider a repeated Nash demand game[10] and suppose that the imitator demands a share larger than 50% of the pie. What can an optimizer do? It can optimize by demanding the highest share compatible to the claim of the imitator. If the optimizer demands less then it forgoes profits. If the optimizer demands more then both make zero profits. Assuming that the imitator mimics itself in such situation we can conclude that the optimizer can not manipulate the decision of the imitator in its favor. Hence it appears that also in this repeated Nash demand game the imitator is better off than is the optimizer. What is eventually wrong with Friedman's conjecture is that he does not consider a class of strategic situations in which "the wise one gives in" (a translated German proverb: "der Klügere gibt nach").[11] Our result is likely to be generalized to a class of aggregative quasi-submodular games (for such generalization of Vega-Redondo's result see Schipper, 2003).

Earlier experimental studies of Cournot oligopoly like the one by Sauermann and Selten (1959) found some support for the convergence to Cournot Nash equilibrium. Recent studies by Huck, Normann, and Oechssler (1999, 2000) found support for imitative behavior in experimental Cournot settings. Whereas in former experiments subjects had profit tables for easy calculation of the best reply available, in later studies subjects received feedback about the competitors' profits and output levels. The informational framework of these experimental designs corresponds closely to the information required by each of the two decision rules (see also Offerman, Potters, and Sonnemans, 2002, for further experimental evidence). Since both, imitation behavior as well as best response, is supported by experimental findings in Cournot markets depending on the information

---

[10]I thank Ariel Rubinstein for suggesting this example.

[11]To be fair, Friedman (1953) had probably only perfectly competitive situations in mind.

provided to subjects, it is only natural to test whether our results can be supported experimentally if different information is given to various firms in an oligopoly experiment. This shall be left to further research.

A   Proofs

Proof of Lemma 1.

Let $q'' > q'$ and $Q'' > Q'$. Since by Assumption 1, p is strictly decreasing

$$
\begin{aligned}
p(Q') &> p(Q'') \\
p(Q')(q'' - q') &> p(Q'')(q'' - q') \\
p(Q')(q'' - q') - c(q'') + c(q') &> p(Q'')(q'' - q') - c(q'') + c(q') \\
\pi(q'', Q') - \pi(q', Q') &> \pi(q'', Q'') - \pi(q', Q'')
\end{aligned}
$$

This completes the proof of Lemma 1. □

For the proofs of the following results it is useful to state following lemma:

Lemma 2 *If Assumptions 1 and 4 hold then we conclude the following:*

   (i) *The Cournot oligopoly has a Cournot Nash equilibrium.*

  (ii) *Sequential best reply converges to the Cournot Nash equilibrium in finite time.*

 (iii) *A best reply is unique.*

 (iv) *The Cournot Nash equilibrium is unique and symmetric.*

  (v) *Given $\theta \in (0,1)$, let x', x be total outputs of all imitators. If $x' < x$, then $x' + (1-\theta)nq^{D'} \leq x + (1-\theta)nq^D$, with $q^D = b(x + ((1-\theta)n - 1)q^D)$ and $q^{D'} = b(x' + ((1-\theta)n - 1)q^{D'})$.*

 (vi) *If ω^S exists, then it is unique.*

   Proof.   For (i) and (ii) see Dubey, Haimanko, and Zapechelnyuk (2005). For (iii) and (iv) see Vives (2000, p. 43).

(v) Suppose to the contrary that $x' < x$ and $x' + (1-\theta)nq^{D'} > x + (1-\theta)nq^D$. Last inequality is equivalent to $x' - x > (1-\theta)n(q^D - q^{D'})$. Since by assumption $0 > x' - x$ we conclude that $0 > (1-\theta)n(q^D - q^{D'})$. Since $\theta \in (0,1)$ and $n > 1$, the last equality is satisfied if and only if $0 > q^D - q^{D'}$. Define $q'_{-i} := x' + ((1-\theta)n - 1)q^{D'}$ and analogously for q_{-i}. Suppose $q'_{-i} > q_{-i}$ then by Assumption 4 (strictly decreasing best responses) $q^{D'} < q^D$, a contradiction to above. Hence we must have $q'_{-i} \leq q_{-i}$. Then by Assumption 4, $q'_{-i} + q^{D'} \leq q_{-i} + q^D$, a contradiction to $x' + (1-\theta)nq^{D'} > x + (1-\theta)nq^D$ above.

(vi) Let $\omega^{S'}$ and $\omega^{S''}$ be two Pseudo-Stackelberg solutions with $\omega^{S'} \neq \omega^{S''}$. Denote by $Q^{S'} = \theta nq^{S'} + (1-\theta)nq^{D'}$, $Q^{S''} = \theta nq^{S''} + (1-\theta)nq^{D''}$, $q^{D'} = b(\theta nq^{S'} + ((1-\theta)n - 1)q^{D'})$, and $q^{D''} = b(\theta nq^{S''} + ((1-\theta)n - 1)q^{D''})$. Inequality 11 implies $\pi_i(q^{S'}, Q^{S'}) > \pi_i(q^{S''}, Q^{S'})$ and $\pi_i(q^{S''}, Q^{S''}) > \pi_i(q^{S'}, Q^{S''})$. If $q^{S''} > q^{S'}$ then $Q^{S''} \geq Q^{S'}$ by (v). By Assumption 1 (Lemma 2 and Remark 15, quasi-submodularity (upper formula (3)) $\pi_i(q^{S''}, Q^{S''}) > \pi_i(q^{S'}, Q^{S''})$ implies $\pi_i(q^{S''}, Q^{S'}) > \pi_i(q^{S'}, Q^{S'})$, a contradiction. Likewise for $q^{S''} < q^{S'}$ (using lower formula (4)). □

Proof of Proposition 1.

Recall that Z is the collection of all absorbing sets of the unperturbed decision dynamics when $\varepsilon = 0$. We need to show the following: If Assumptions 1, 2 and 4 hold, then $Z = H$, with $Z = \{\{\omega\} : \omega \in H\}$.

First, suppose that some state $\omega \notin H$, $\omega \in \Gamma^n$ is an element of an absorbing set A. At least one condition of (i) to (iv) of Definition 7 is violated. Thus there will be an incentive for some imitators or some optimizers to deviate from their output in ω. By Assumptions 1, 2 (Lemma 2) and 4 we construct an unperturbed adjustment process based on the decision rules leading in the subsequent periods to a state $\omega' \in H$, noting that by Lemma 2 such $\omega' \in H$ exists.

Second, we show that every absorbing set $A \subseteq H$ is a singleton. Suppose there exists $\omega', \omega \in A \subseteq H$, $\omega' \neq \omega$. By the definition of absorbing set, $\exists m \in \mathbb{N}$, m finite s.t. $p_{\omega\omega'}^{(m)} > 0$. Consider any imitator $i \in I$. Since in $\omega \in H$ it follows by Definition 7 (i), (iii), and (iv) that no imitator $i \in I$ wants to deviate form its output in $\omega \in H$. Now consider an optimizer $i \in N \setminus I$. Since $\omega \in H$, it follows by aforesaid Definition 7 (ii) that no optimizer $i \in N \setminus I$ wants to

25

deviate from its best reply in $\omega \in H$, which is by Lemma 2 uniquely defined. Since both types of firms do not deviate in $\omega \in H$, no firm $i \in N$ deviates in any of the following periods. Thus $p_{\omega\omega'}^{(m)} = 0, \forall m \in \mathbb{N}$, which contradicts that $\omega', \omega \in A$, $\omega' \neq \omega$. It follows that $p_{\omega\omega} = 1$ for each $\omega \in H$ such that $\{\omega\} = A$, $\forall \omega \in H$. From the first part of the proof we conclude that there does not exist a state $\omega \notin H$ s.t. $\omega \in A$, $A \in Z$. Hence $Z = \{\{\omega\} : \omega \in H\}$. This completes the proof of Proposition 1. $\qquad\square$

In order to characterize the support of the unique limiting invariant distribution, we consider perturbations introduced by Assumption 3. We call states ω and ω' *adjacent* if exactly one mutation can change the state from ω to ω' (and vice versa), i.e., if exactly one firm's change of output changes the state ω to the state ω'. The set of all states adjacent to the state ω is the *single mutation neighborhood* of ω denoted by $M(\omega)$. The *basin of attraction* of an absorbing set A is the set $B(A) = \{\omega \in \Gamma^n | \exists m \in \mathbb{N}, \exists \omega' \in A \text{ s.t. } p_{\omega\omega'}^{(m)} > 0\}$. It is the collection of all states from which there is a strict positive probability that the (unperturbed) dynamics leads to the absorbing set A. A *recurrent set* R is a minimal collection of absorbing sets with the property that there do not exist absorbing sets $A \in R$ and $A' \notin R$ such that for all $\omega \in A$, $M(\omega) \cap B(A') \neq \emptyset$. That is, a recurrent set R is a minimal collection of absorbing sets for which it is impossible that a single mutation followed by the unperturbed dynamics leads to an absorbing set not contained in R. The importance of the recurrent set stems from below Lemma 3 by Nöldeke and Samuelson (1993, 1997) and Samuelson (1994).

Lemma 3 (Nöldeke and Samuelson) *Given a regularly perturbed finite Markov chain, then at least one recurrent set exists. Recurrent sets are disjoint. Let the state ω be contained in the support of the unique limiting invariant distribution φ^*. Then $\omega \in R$, R being a recurrent set. Moreover, for all $\omega' \in R$, $\varphi^*(\omega') > 0$.*

A proof of Lemma 3 is contained in Samuelson (1997), Lemma 7.1 and Proposition 7.7., proof pp. 236-238.

Proof of Theorem 1.

It is sufficient to show that $\{\omega^\circ\}$ is not a singleton recurrent set.

Lemma 4 *Given previous assumptions, if $\omega^S \in H$ then $M(\omega) \cap B(\{\omega^S\}) \neq \emptyset$, $\forall \omega \in H \setminus \{\omega^S\}$.*

Proof of Lemma. Assume $\omega^S \in H$. It is sufficient to show that $\forall q \in \Gamma$, q being a component of an arbitrary $\omega \in H$, $\omega \neq \omega^S$, $k \in \mathbb{N}$, $0 < k \leq \theta n$,

$$q^S p((\theta n - k)q + kq^S + (1-\theta)nq^D) - c(q^S) > qp((\theta n - k)q + kq^S + (1-\theta)nq^D) - c(q), \quad (14)$$

with $q^D = b((\theta n - k)q + kq^S + ((1-\theta)n - 1)q^D)$.

By Lemma 1 and Remark 1, π_i is quasi-submodular (formulas (3) and (4)) in (q, Q) on $\Gamma \times \{0, \delta, 2\delta, ..., n\nu\delta\}$. Set $q'' \equiv q^S$, $q' \equiv q$, $Q' = (\theta n - k)q + kq^S + (1 - \theta)nq^{D'}$ and $Q'' = \theta n q^S + (1-\theta)nq^{D''}$ with $q^{D'} \equiv q^D$ and $q^{D''} \equiv b(\theta n q^S + ((1-\theta)n - 1)q^{D''})$ being uniquely defined by Lemma 2. If $q'' > q'$ then $\theta n q'' > (\theta n - k)q' + kq''$. By Lemma 2, we conclude that $Q'' \geq Q'$. If $q'' < q'$ then $\theta n q'' < (\theta n - k)q' + kq''$. By Lemma 2, we conclude that $Q'' \leq Q'$. It follows that if $q^S > q$ then the left hand side of "$\Longrightarrow$" in formula (3) is given by Inequality (11) of Definition 7 of the Pseudo-Stackelberg solution (i). In this case the right hand side of "$\Longrightarrow$" in formula (3) yields above Inequality (14). If $q^S < q$ then the left hand side of "$\Longrightarrow$" in formula (4) is given by Inequality (11) of Definition 7 of the Pseudo-Stackelberg solution (i). In this case the right hand side of "$\Longrightarrow$" in formula (4) yields above Inequality (14). Finally, set $k = 1$ to see that one suitable mutation only is required to connect every $\omega \in H$ to $\omega^S \in H$. $\square$

Since Lemma 4 holds for any absorbing state except the Pseudo-Stackelberg solution, it holds also for the Cournot Nash equilibrium ω°. This implies by Lemma 3 the theorem. $\square$

Proof of Example 1.

We show that under certain assumptions on the parameters of the game, we connect all Pseudo-Stackelberg states by a sequence of single suitable mutations.

Remark 2 *The Cournot Nash equilibrium $\omega^\circ \in H$ is the state with the lowest identical output of imitators in the set of Pseudo-Stackelberg states H.*

Denote by $\bar{\omega} \in H$ the *Pseudo-Stackelberg state with the largest possible identical output of imitators.* I.e., if $\bar{q}^I$ is the identical output of imitators in $\bar{\omega}$, then there does not exist a state $\omega \in H$ such that if q^I is the identical output of imitators in ω we have $q^I > \bar{q}^I$.

Let $q^I_{\omega_j}$ be the identical output of imitators in the state $\omega_j \in H$. We call a sequence of Pseudo-Stackelberg states $\omega_1, ..., \omega_m \in H$ *increasing (decreasing)* iff the identical output of

each imitator in those Pseudo-Stackelberg states is ordered such that $q^I_{\omega_j} < q^I_{\omega_{j+1}}$ ($q^I_{\omega_j} > q^I_{\omega_{j+1}}$), $j = 1, ..., m-1$. Such order on H is the natural order on Γ.

We call π_i *strictly quasi-concave* in q_i if for all $q_i, q'_i \in \Gamma$, $q_i \neq q'_i$ and for all $\lambda \in (0,1)$ s.t. $\lambda q_i + (1-\lambda)q'_i \in \Gamma$,

$$\pi_i(\lambda q_i + (1-\lambda)q'_i, Q) > \min\{\pi_i(q_i, Q), \pi_i(q'_i, Q)\}, \forall Q \in \{0, \delta, 2\delta, ..., n\nu\delta\}. \tag{15}$$

Lemma 5 *Let π_i be strictly quasi-concave in q_i. Under previous assumptions we conclude:*[12]

(i) *If $\omega^S \in H$, then there exists an increasing sequence $\omega_1, ..., \omega_m \in H$ with $\omega_1 = \omega^\circ$ and $\omega_m = \omega^S$ s.t. $M(\omega_j) \cap B(\{\omega_{j+1}\}) \neq \emptyset$, $j = 1, ..., m-1$.*

(ii) *If $\omega^S \in H$, then there exists a decreasing sequence $\omega_1, ..., \omega_m \in H$ with $\omega_1 = \bar{\omega}$ and $\omega_m = \omega^S$ s.t. $M(\omega_j) \cap B(\{\omega_{j+1}\}) \neq \emptyset$, $j = 1, ..., m-1$.*

Proof of Lemma. (i): Let $\omega_1, ..., \omega_m \in H$ be an increasing sequence of absorbing states with $\omega_1 = \omega^\circ$ and $\omega_m = \omega^S$. In order to show that $M(\omega_j) \cap B(\{\omega_{j+1}\}) \neq \emptyset$ for $j = 1, ..., m-1$, it is sufficient to show for $0 < k < \theta n$

$$q^I_{\omega_{j+1}} p((\theta n - k)q^I_{\omega_j} + kq^I_{\omega_{j+1}} + (1-\theta)nq^D) - c(q^I_{\omega_{j+1}}) >$$
$$q^I_{\omega_j} p((\theta n - k)q^I_{\omega_j} + kq^I_{\omega_{j+1}} + (1-\theta)nq^D) - c(q^I_{\omega_j}), \tag{16}$$

with $q^D = b((\theta n - k)q^I_{\omega_j} + kq^I_{\omega_{j+1}} + ((1-\theta)n - 1)q^D)$, which is uniquely defined by Lemma 2.

By Lemma 1 and Remark 1, π is quasi-submodular. Set $q'' = q^I_{\omega_{j+1}}$, $q' = q^I_{\omega_j}$, $Q' = (\theta n - k)q^I_{\omega_j} + kq^I_{\omega_{j+1}} + (1-\theta)nq^{D'}$ and $Q'' = \theta n q^I_{\omega_{j+1}} + (1-\theta)nq^{D''}$, with $q^{D'} = q^D$ and $q^{D''} = b(\theta n q^I_{\omega_{j+1}} + ((1-\theta)n - 1)q^{D''})$ being uniquely defined by Lemma 2.

For each $q^I_{\omega_{j+1}} \in [q^\circ, q^S] \cap \Gamma$ there exists a $\lambda \in (0,1)$ s.t. $q^I_{\omega_{j+1}} = \lambda q^I_{\omega_j} + (1-\lambda)q^S$. We claim that $\pi(q^I_{\omega_j}, Q'') = \min\{\pi(q^S, Q''), \pi(q^I_{\omega_j}, Q'')\}$. We know that $\pi(q^S, Q^S) > \pi(q^I_{\omega_j}, Q^S)$ by definition. Since $q^I_{\omega_{j+1}} \leq q^S$ we know by Lemma 2 that $Q^S \geq Q''$. By quasi-submodularity (upper formula (3)) $\pi(q^S, Q'') > \pi(q^I_{\omega_j}, Q'')$ and the claim follows. By strict quasi-concavity $\pi(q^I_{\omega_{j+1}}, Q'') > \pi(q^I_{\omega_j}, Q'')$. It implies Inequality (16) by quasi-submodularity (upper formula (3)). (i) follows from setting $k = 1$ in Inequality (16).

_____

[12] Lemma 5 provides also an alternative proof of Theorem 1 under the additional assumption of strict quasi-concavity of π_i in q_i.

(ii) The proof is analogous to (i). Set $q' = q^I_{\omega_{j+1}}$, $q'' = q^I_{\omega_j}$, $Q'' = (\theta n - k)q^I_{\omega_j} + kq^I_{\omega_{j+1}} + (1 - \theta)nq^{D''}$ and $Q' = \theta nq^I_{\omega_{j+1}} + (1-\theta)nq^{D'}$, with $q^{D''} = q^D$ and $q^{D'} = b(\theta nq^I_{\omega_{j+1}} + ((1-\theta)n - 1)q^{D'})$ and use the lower formula (4) of quasi-submodularity.

This completes the proof of Lemma 5. $\qquad\square$

Lemma 6 *Let $q^{\circ 2} \in \Gamma$ be a firm's Cournot duopoly equilibrium output.*[13] *Given Assumptions 1 and 4 are satisfied, there exist p, c, θ, δ, ν, and finite n such that*

$$q^\circ p((2n - 3)q^\circ) - c(q^\circ) \leq 0, \qquad (17)$$

$$(n - 1)q^\circ p((n - 1)q^\circ) - c((n - 1)q^\circ) \leq 0, \qquad (18)$$

$$\pi(q, \nu\delta) < 0, \forall q > 0, \qquad (19)$$

$$q^{\circ 2} \equiv \bar{q}^I, \qquad (20)$$

π_i *is strictly quasi-concave in q_i, and $\omega^S \in \Gamma^n$ exists uniquely.*

Proof of Lemma. Consider p, c, θ, δ, ν, and n in Example 1, i.e., $p(Q) = 10 - Q$, $c(q_i) = \frac{1000}{501}q_i^2 + 1$, $\theta = 0.2$, $\delta = 0.001$, $n = 5$, and sufficiently large ν (e.g. $\nu = 30000$). Straight forward calculations verify that formulas (17) to (20) as well as Assumptions 1 and 4 hold, and that ω^S exists uniquely. Moreover, since π_i is strictly concave in q_i, it is strictly quasi-concave. $\qquad\square$

Lemma 7 *Given previous assumptions, let p, c, θ, δ, ν, and n such that the properties of Lemma 6 hold. Then $M(\omega) \cap B(\{\omega^\circ\}) \neq \emptyset$, $\forall \omega \in H$.*

Proof of Lemma. Suppose in t any arbitrary state $\omega(t) \in H$. By Lemma 2 such state exists and is uniquely defined. W.l.o.g. suppose by Assumptions 2 and 3 that in $t + 1$ a mutation by one firm $i \in N$ occurs such that $q_i(t + 1) = (n - 1)q^\circ$. Note that by Lemma 2 the Cournot Nash equilibrium output $q^\circ \in \Gamma$ exists and is unique. Since $\omega(t) \in H$, we have $Q(t + 1) \geq (n - 1)q^\circ + (n - 1)q^\circ = (2n - 2)q^\circ > (2n - 3)q^\circ$. By Lemma 6, Inequality (17), $\pi_j(t + 1) < 0$, $\forall j \in N$. W.l.o.g. assume that by Assumption 2 a firm $k \in N\backslash I$, $k \neq i$ and only a firm k has the opportunity to adjust output in $t + 2$. Since $D_O(t + 1) = 0$, we have $q_k(t + 2) = 0$.

---

[13]The Cournot duopoly equilibrium is a special case of the Cournot Nash equilibrium (Definition 5) for $n = 2$.

$Q(t + 2) \geq (2n - 3)q^\circ$. By Lemma 6, Inequality (17), $\pi_j(t + 1) < 0$, $\forall j \in N\backslash\{k\}$. W.l.o.g. assume that by Assumption 2 and $D_O(t+2) = D_I(t+2) = 0$ all $j \in N\backslash\{i\}$ adjust output in $t+3$ such that $Q(t+3) = q_i(t+3) = q_i(t+2) = q_i(t+1) = (n-1)q^\circ$. By Lemma 6, Inequality (18), $\pi_i(t + 3) \leq 0$. W.l.o.g. assume that by Assumption 2, 4, and $D_O(t + 3) = b((n - 1)q^\circ) = q^\circ$ another firm $k \in N\backslash I$ has the opportunity to adjust output in $t+4$. Since $\pi_k(t+4) > \pi_j(t+4)$, $j \in N\backslash\{k\}$ we can assume w.l.o.g. that by Assumption 2 and $D_I(t + 4) = q^\circ$ all $j \in I$ adjust output. By Assumptions 2 let all remaining optimizers adjust output in the subsequent periods such that with positive probability ω° is reached in finite time (by Lemma 2). Since we started in any arbitrary absorbing state $\omega(t) \in H$ (in particular it also includes ω^S if $\omega^S \in H$) we have shown that $M(A) \cap B(\{\omega^\circ\}) \neq \emptyset$, $\forall A \in Z$. $\qquad\square$

Lemma 8 *Given previous assumptions, let p, c, θ, δ, ν and n such that the properties of Lemma 6 hold. Then $M(\omega) \cap B(\{\bar{\omega}\}) \neq \emptyset$, $\forall \omega \in H$.*

Proof of Lemma. Suppose in t any arbitrary state $\omega(t) \in H$. W.l.o.g. assume that by Assumptions 2 and 3 in $t+1$ a mutation by an imitator $i \in I$ occurs setting a large quantity $\nu\delta$ such that by Lemma 6, Inequality (19) $\pi_j(t + 1) < 0$, $\forall j \in N$. W.l.o.g. assume that by Assumption 2 all optimizers in $N\backslash I$ have the opportunity to adjust output in $t+2$. Since $D_O(t+1) = 0$, we have $q^D(t + 2) = 0$ with $\pi^D(t + 2) = 0$. By Inequality (19), we have $\pi^I(t + 2) < \pi^D(t + 2)$. W.l.o.g. assume that by Assumption 2 all imitators in I adjust output in $t+3$ to $D_I(t+2) = q^I(t+3) = 0$. Hence, $Q(t + 3) = 0$. W.l.o.g. assume now that by Assumption 2 in $t + 4$ two optimizers and only two optimizers adjust output such that by Lemma 2 we reach a market output of $2q^{\circ 2}$ in finite time, i.e. at $t + k$. W.l.o.g. assume that by Assumptions 2 in the following period all imitators in I adjust output such that $D_I(t + k) = q^{\circ 2}(t + k + 1)$. Let all optimizers in $N\backslash I$ adjust output in the subsequent periods such that by Assumptions 2 and Lemma 2 a state $\omega^{\circ 2} = (q_1^{\circ 2}, ..., q_{\theta n}^{\circ 2}, q_{\theta n+1}^D, ..., q_n^D)$ is reached in finite time. Since by Lemma 6, $\bar{q}^I = q^{\circ 2}$, we have that $\omega^{\circ 2} = \bar{\omega}$. $\qquad\square$

In Lemma 5 we showed that we can connect the Pseudo-Stackelberg states by an increasing (starting with the Cournot Nash equilibrium) or decreasing sequence (starting with the Pseudo-Stackelberg state with the largest output of imitators) of single suitable mutations followed

by the decision dynamics to the Pseudo-Stackelberg solution if the profit functions are quasi-concave. In Lemma 7 and 8 we showed that we can connect by single suitable mutations followed by the decision dynamics any Pseudo-Stackelberg state to the Cournot Nash equilibrium and the Pseudo-Stackelberg state with the largest output of imitators if the properties of Lemma 6 hold. In particular we can also connect the Pseudo-Stackelberg solution $\omega^S \in H$ to the Cournot Nash equilibrium and the Pseudo-Stackelberg state with the largest output of imitators. Hence there exists a sequence of single suitable mutations by which we can move through the entire set of Pseudo-Stackelberg states. It follows that H is the unique recurrent set. By Lemma 3 it follows that $S = H$. This completes the proof of Example 1. $\qquad\square$

References

[1] Alchian, A. A., 1950. Uncertainty, evolution, and economic theory, Journal of Political Economy 58, 211-221.

[2] Alós-Ferrer, C., 2004. Cournot versus Walras in dynamic oligopolies with memory, International Journal of Industrial Organization 22, 193-217.

[3] Alós-Ferrer, C., Ania, A. B. and Vega-Redondo, F., 1999. An evolutionary model of market structure, in Herings, P., Talman, A. and van der Laan, G. (Eds.), The theory of markets. North Holland, Amsterdam.

[4] Amir, R., 1996. Cournot oligopoly and the theory of supermodular games, Games and Economic Behavior 15, 132-148.

[5] Amir, R. and Lambson, V.E., 2000. On the effect of entry in Cournot markets, Review of Economic Studies 67, 235-254.

[6] Banerjee, A. and Weibull, J., 1995. Evolutionary selection and rational behavior, in Kirmann, A. and Salmon, M. (Eds.), Learning and rationality in economics, Basil Blackwell, Oxford.

[7] Blume, L. and Easley, D., 2002. Optimality and natural selection in markets, Journal of Economic Theory 107, 95-135.

[8] Bulow, J. I., Geanakoplos, J. D. and Klemperer, P. D., 1985. Multimarket oligopoly: strategic substitutes and complements, Journal of Political Economy 93, 488-511.

[9] Conlisk, J., 1980. Costly optimizers versus cheap imitators, Journal of Economic Behavior and Organization 1, 275-293.

[10] Cournot, A., 1838. Researches into the mathematical principles of the theory of wealth, transl. by N. T. Bacon, MacMillan Company, New York, 1927.

[11] Dubey, P., Haimanko, O. and Zapechelnyuk, A., 2005. Strategic complements and substitutes, and potential games, forthcoming Games and Economic Behavior.

[12] Dutta, P.K. and Radner, R., 1999. Profit maximization and the market selection hypothesis, Review of Economic Studies 66, 769-798.

[13] Ellison, G., 2000. Basin of attraction, long-run stochastic stability, and the speed of step-by-step evolution, Review of Economic Studies 67, 17-45.

[14] Freidlin, M. and Wentzell, A., 1984. Random perturbations of dynamical systems, Springer-Verlag, Berlin.

[15] Friedman, M., 1953. Essays in positive economics, University of Chicago Press, Chicago.

[16] Fudenberg, D. and Levine, D. K., 1998. The theory of learning in games, The MIT Press, Cambridge, M.A.

[17] Gale, D. and Rosenthal, R. W., 1999. Experimentation, imitation, and stochastic stability Journal of Economic Theory 84, 1-40.

[18] Huck, S., Normann, H.-T. and Oechssler, J., 2000. Does information about competitors' actions increase or decrease competition in experimental oligopoly markets?, International Journal of Industrial Organization 18, 39-57.

[19] Huck, S., Normann, H.-T. and Oechssler, J., 1999. Learning in Cournot oligopoly - An experiment, Economic Journal 109, C80-C95.

[20] Kaarbøe, O. and Tieman, A., 1999. Equilibrium selection under different learning modes in supermodular games, mimeo., University of Bergen.

[21] Kandori, M., Mailath, G.J. and Rob, R., 1993. Learning, mutation and long run equilibria in games, Econometrica 61, 29-56.

[22] Koçkesen, L., Ok, E.A. and Sethi, R., 2000. The strategic advantage of negatively interdependent preferences, Journal of Economic Theory 92, 274-299.

[23] Koopmans, T., 1957. Three essays on the state of economic science, McGraw-Hill Book Company.

[24] Masaaki, K., 1997. Markov processes for stochastic modeling, Chapman & Hall, London.

[25] Milgrom, P. and Shannon, C., 1994. Monotone comparative statics, Econometrica 62, 157-180.

[26] Monderer, D. and Shapley, L. S., 1996. Potential games, Games and Economic Behavior 14, 124-143.

[27] Nelson, R. R. and Winter, S. G., 1992. An evolutionary theory of economic change, Harvard University Press, Cambridge, M.A.

[28] von Neumann, J. and Morgenstern, O., 1944. Theory of games and economic behavior, John Wiley & Sons, Science Editions, New York.

[29] Nöldeke, G. and Samuelson, L., 1997. A dynamic model of equilibrium selection in signaling markets, Journal of Economic Theory 73, 118-156.

[30] Nöldeke, G. and Samuelson, L., 1993. An evolutionary analysis of backward and forward induction, Games and Economic Behavior 5, 425-454.

[31] Offerman, T., Potters, J. and Sonnemans, J., 2002. Imitation and belief learning in an oligopoly experiment, Review of Economic Studies 69, 973-997.

[32] Rhode, P. and Stegeman, M., 2001. Non-Nash equilibria of darwinian dynamics with applications to duopoly, International Journal of Industrial Organization 19, 415-453.

[33] Samuelson, L., 1997. Evolutionary games and equilibrium selection, The MIT Press, Cambridge, M.A.

[34] Samuelson, L., 1994. Stochastic stability in games with alternative best replies, Journal of Economic Theory 64, 35-65.

[35] Sandroni, A., 2000. Do markets favor agents able to make accurate predictions?, Econometrica 68, 1303-1341.

[36] Sauermann, H. and Selten, R., 1959. Ein Oligopolexperiment, Zeitschrift für die gesamten Staatswissenschaften 115, 427-471.

[37] Schaffer, M. E., 1989. Are profit-maximizers the best survivors?, Journal of Economic Behavior and Organization 12, 29-45.

[38] Schipper, B. C., 2003. Submodularity and the evolution of Walrasian behavior, International Journal of Game Theory 32, 471-477.

[39] Schipper, B. C., 2002. Imitators and optimizers in symmetric n-firm Cournot oligopoly, mimeo., University of Bonn.

[40] Schlag, K., 1998. Why imitate, and if so, How? A bounded rational approach to multi-armed bandits, Journal of Economic Theory 78, 130-156.

[41] von Stackelberg, H., 1934. Marktform und Gleichgewicht, Verlag von Julius Springer, Wien.

[42] Topkis, D. M., 1998. Supermodularity and complementarity, Princeton University Press, Princeton, N.J.

[43] Vega-Redondo, F., 1997. The evolution of Walrasian behavior, Econometrica 65, 375-384.

[44] Vives, X., 2000. Oligopoly pricing: old ideas and new tools, The MIT Press, Cambridge, M.A.

[45] Weibull, J. W., 1995. Evolutionary game theory, The MIT Press, Cambridge, M.A.

[46] Young, H. P., 1993. The evolution of conventions, Econometrica 61, 57-84.