

BONN ECON DISCUSSION PAPERS

Discussion Paper 29/2002

Imitators and Optimizers in Cournot Oligopoly

by

Burkhard Schipper

October 2002



Bonn Graduate School of Economics
Department of Economics
University of Bonn
Adenauerallee 24 - 42
D-53113 Bonn

The Bonn Graduate School of Economics is
sponsored by the

Deutsche Post  World Net
MAIL EXPRESS LOGISTICS FINANCE

Imitators and Optimizers in Cournot Oligopoly*

Burkhard C. Schipper

University of Bonn

October 24, 2002

*I wish to thank among others Carlos Alós-Ferrer, Itzhak Gilboa, and Jörg Oechssler as well as audiences of various seminars and conferences for very helpful comments. I thank the School of Economics at Tel Aviv University, especially Phyllis Klein, for hosting me within the European Doctoral Program of Quantitative Economics (EDP) during the year 2001-02 in which the paper was completed. Financial support by the DAAD and the DFG is gratefully acknowledged. A slightly different but also more extensive previous version had the title “Imitators and Optimizers in Symmetric n-Firm Cournot Oligopoly.”

Proposed title: Imitators and Optimizers in Cournot Oligopoly

Author: Burkhard C. Schipper, University of Bonn

Current address: Bonn Graduate School of Economics, Dept. of Economics, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany, Tel: +49-228-73 9473, Fax: +49-228-73 9221, Email: burkhard.schipper@wiwi.uni-bonn.de, Web: <http://www.bgse.uni-bonn.de/~burkhard>

Abstract

We present a formal model of symmetric n -firm Cournot oligopoly with a heterogeneous population of profit optimizers and imitators. Imitators mimic the output decision of the most successful firms of the previous round à la Vega-Redondo (1997). Optimizers play myopic best response to the opponents' previous output. The dynamics of the decision rules induce a Markov chain. As expression of bounded rationality, firms are allowed to make mistakes and deviate from the decision rules with a small probability. Applying stochastic stability analysis, we characterize the long run behavior of the oligopoly. 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: imitation, myopic best reply, bounded rationality, profit maximization hypothesis, stochastic stability, learning, Stackelberg.

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 n -firm symmetric Cournot oligopoly with a homogeneous population of 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 a homogeneous population of optimizers, which is 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 to Friedman, we should find that optimizers are better off than are imitators, and that consequently optimizers drive out imitators in any payoff monotone

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

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. 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. Moreover, today’s familiar sequential representation of the Stackelberg game is not due to von Stackelberg but was introduced in a different context as the “majorant game” by von Neumann and Morgenstern (1944), pp. 100.² Optimizers are “dependent” since by definition they play best reply. Although the imitators are “independent” because they do not perceive their influence on the price but take them as given, they do not conform exactly to the Stackelberg conjecture. Second, analogous to the profits in von Stackelberg’s (1934) independent and dependent firms, every imitator is better off than every optimizers. Finally, our analysis retains the important aspect of von Stackelberg’s (1934) idea: the modeling of asymmetries and behavioral heterogeneity of firms and in particular the modeling of independent and dependent firms.

Imitators and optimizers differ in 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 profits³, optimizers need to know the total output of their opponents as well as their own profit function,

²I thank Prof. Selten for pointing me to the “majorant game”.

³See Alós-Ferrer (2001) for a study of imitators in Cournot oligopoly who take a longer history into account.

which implies 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 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 a reason why one part of executives' remuneration-packages is often based on the firm's stock outperforming the market index or similar means of comparison.

Earlier experimental studies on Cournot oligopoly like the one by Sauer-
mann 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).

In the proofs of our results, we rely on two main concepts, submodularity of payoff functions and stochastic stability analysis. Submodularity (see Topkis, 1998, pp. 43) is closely related to strategic substitutes (see Bulow, Geanakoplos and Klemperer, 1985).⁴ The intuition for submodularity in our context is that a firm's payoff difference from an increase of its own output decreases in the total output.

⁴For submodularity and supermodularity in the context of a Cournot oligopoly see Vives (1990), Amir (1996) and Vives (2000), chapter 4. Formally, supermodularity is related to the single-crossing property and submodularity to the dual single-crossing property (see Topkis, 1998, Milgrom and Shannon, 1994).

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), as well as its characterization result (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 in a sense of experimenting with various output levels. Second, decision makers of the firms are assumed to be boundedly rational in the sense 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), Samuelson (1994) and Nöldeke and Samuelson (1997). 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 be contained in the minimal set of absorbing sets that is robust to a single mutation. Such a set is called recurrent set. In Theorem 1 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, in Theorem 2 we show that there are assumptions on the parameters of the game such that the entire set of Pseudo-Stackelberg states is the unique recurrent set and thus is the support of the unique limiting invariant distribution.

Apart from 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. His approach differs from the current one in that he takes the cost of optimizing into account, and this cost is a key for obtaining his results. Research on Friedman’s profit maximization hypothesis has been done for example by Blume and Easley (2002), who find support for it in a general equilibrium context. The present paper is also related to the literature on interdependent preferences. In particular, Koçkesen, Ok and Sethi (2000) 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. Similarly, biased perceptions as studied by Heifetz and Spiegel (2002), can yield a strategic advantage in our setting if they induce players to play more aggressively compared to unbiased profit maximizers.

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 works out 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 $N = \{1, 2, \dots, n\}$ and a market for a homogeneous good. Inverse demand is given by a function $p : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. 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 $i \in N$ possesses the same production technology. Hence the cost functions $c : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ are identical. For each firm it is a function of the quantity q_i it produces. Let the total output over all firms be $Q := \sum_{i \in N} q_i$. Profits are given by

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

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 $\Gamma = \{0, \delta, 2\delta, \dots, \nu\delta\}$, where both $\delta > 0$ and $\nu \in \mathbb{N}$ are arbitrary. This turns the strategic situation into a game with finite action space and allows us to focus on finite Markov chains later in the dynamic analysis.

We wish to show that each firm's payoff function is submodular in the firm's quantity and the total output. This observation will play a key role in the proofs of our results. It is closely related to strategic substitutes (see Bulow, Geanakoplos and Klemperer, 1985). To this end, a few definitions are needed.

A lattice is defined as a poset⁵ $\langle X, \leq \rangle$ whose least upper bound and great-

⁵A partially ordered set (poset) $\langle X, \leq \rangle$ is defined as a set X with a binary relation \leq such that $\forall x, y, z \in X$ the following conditions hold:

- (i) reflexivity: $\forall x \in X, x \leq x$,

est lower bound are given by $x' \vee x'' = \sup\{x', x''\}$ and $x' \wedge x'' = \inf\{x', x''\}$, for all $x', x'' \in X$ respectively. For example, if X is the product of several ordered sets, one may define $x' \vee x''$ (likewise $x' \wedge x''$) as the component-wise max (min) to define a lattice. In our context $\Gamma = \{0, \delta, 2\delta, \dots, n\nu\delta\}$ and products thereof endowed with component-wise max and min operations are lattices. A real valued function $f : X \longrightarrow \mathbb{R}$ on a lattice X is called submodular on X if $\forall x', x'' \in X$, $f(x' \wedge x'') + f(x' \vee x'') \leq f(x') + f(x'')$. The function f is called strictly submodular if the inequality holds strictly for all unordered $x', x'' \in X$.⁶

Assumption 1 (Strictly Decreasing Demand) $\forall Q, Q' \in \{0, \delta, 2\delta, \dots, n\nu\delta\}$, $Q' > Q \implies p(Q') < p(Q)$.

Lemma 1 *By Assumption 1, π_i is strictly submodular in (q_i, Q) on $\Gamma \times \{0, \delta, 2\delta, \dots, n\nu\delta\}$, $\forall i \in N$, i.e. $\forall (q'_i, Q'), (q''_i, Q'') \in \Gamma \times \{0, \delta, 2\delta, \dots, n\nu\delta\}$*

$$\pi_i((q'_i, Q') \wedge (q''_i, Q'')) + \pi_i((q'_i, Q') \vee (q''_i, Q'')) \leq \pi_i(q'_i, Q') + \pi_i(q''_i, Q''). \quad (2)$$

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 $\Gamma \times \{0, \delta, 2\delta, \dots, n\nu\delta\}$.

Remark 1 *(Strict) submodularity implies (strict) quasi-submodularity of π but not vice versa, i.e. inequality (2) implies⁷*

$$\pi((q', Q') \vee (q'', Q'')) \geq \pi(q'', Q'') \implies \pi(q', Q') > \pi((q', Q') \wedge (q'', Q'')), \quad (3)$$

$$\pi((q', Q') \wedge (q'', Q'')) \geq \pi(q'', Q'') \implies \pi(q', Q') > \pi((q', Q') \vee (q'', Q'')). \quad (4)$$

(ii) antisymmetry: if $x \leq y$ and $y \leq x \implies x = y$,

(iii) transitivity: if $x \leq y$ and $y \leq z \implies x \leq z$.

⁶A pair $x', x'' \in X$ is unordered iff none of the two statements holds: $x' \leq x''$, $x'' \leq x'$.

⁷For notational ease we drop subscripts and superscripts when no ambiguity is likely to arise.

The observation that payoffs are quasi-submodular⁸ in individual quantity and total output will be used in later proofs repeatedly, in particular in Lemma 4 and Remark 2. Note that this property follows directly from 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 $t = 0, 1, 2, \dots$. At each t the state of the system is identified by the current output schedule

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

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 $\pi(t) = (\pi_1(t), \pi_2(t), \dots, \pi_n(t))$ at t , defined as follows:

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

Assumption 2 (Inertia) *At every time $t = 1, 2, \dots$, 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 $\rho < 1$ 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.

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 $\theta = \frac{\#I}{\#N}$. Throughout the paper we assume that $\theta \in (0, 1)$, i.e. we have

⁸For quasisubmodularity see also Milgrom and Shannon (1994), p. 162 and Topkis (1998), pp. 58. It is the ordinal notion of submodularity.

a heterogeneous population of firms with at least one imitator and one optimizer. 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) 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...” (p. 165). Conventional economics focuses mainly on profit maximization. However, “benchmarking”, “best practices” and other imitation rules can be found in today’s management practice.

Definition 1 (Imitator) *An imitator $i \in I$ chooses with full support from the set*

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

Definition 2 (Optimizer) *An optimizer $i \in N \setminus I$ chooses from the set*

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

with $q_{-i} := \sum_{j \in N \setminus \{i\}} q_j$ and $b_i : \Gamma \longrightarrow \Gamma$ is defined to be firm i ’s best reply correspondence

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

It is assumed that initially in $t = 0$ both types of firms start 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 in Vega-Redondo (1997). Definition 2 means that an optimizer sets an output level that is a best reply to the opponents' total output in the previous period.

The process induced by the decision rules is a discrete time n -vector 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 $\text{prob}\{\omega(t+1)|\omega(t), \omega(t-1), \dots, \omega(t-k)\} = \text{prob}\{\omega(t+1)|\omega(t)\}$. 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 $\text{prob}\{\omega'(t+k+1)|\omega(t+k)\} = \text{prob}\{\omega'(t+1)|\omega(t)\}$, $k = 0, 1, \dots$

The Markov operator is defined in the standard way as the $\#\Gamma^n \times \#\Gamma^n$ -transition probability matrix $P = (p_{\omega\omega'})_{\omega, \omega' \in \Gamma^n}$ with $p_{\omega\omega'} = \text{prob}\{\omega'|\omega\}$, $p_{\omega\omega'} \geq 0$, $\omega, \omega' \in \Gamma^n$ and $\sum_{\omega' \in \Gamma^n} p_{\omega\omega'} = 1$, $\forall \omega \in \Gamma^n$. 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.

To complete the model of the decision process we make the 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 in t if it randomizes with full support in 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 $\varphi(\varepsilon)P(\varepsilon) = \varphi(\varepsilon)$, $\varphi^* := \lim_{\varepsilon \rightarrow 0} \varphi(\varepsilon)$ 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) $\forall \omega, \omega' \in A, \exists m \in \mathbb{N}, m \text{ finite, s.t. } p_{\omega\omega'}^{(m)} > 0, p_{\omega\omega'}^{(m)}$ being the m -step transition probability from ω to ω' .

Let Z be the collection of all absorbing sets in Γ^n .

Vega-Redondo (1997) showed that a homogeneous population of imitators in our framework converges to the competitive solution. Can the competitive solution be an absorbing state given a heterogeneous population of imitators and optimizers? Suppose the unique competitive solution exists 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 3 (Cournot Nash equilibrium) *A combination of output strategies $(q_1^\circ, q_2^\circ, \dots, q_n^\circ) \in \Gamma^n$ is a Cournot Nash equilibrium if $\forall i \in N$,*

$$q_i^\circ p(Q^\circ) - c(q_i^\circ) \geq q_i p(Q^\circ - q_i^\circ + q_i) - c(q_i), \forall q_i \in \Gamma. \quad (9)$$

It is known that in a homogeneous population of optimizers the Cournot Nash equilibrium is the solution of a sequential best response process under certain assumptions guaranteeing global convergence. 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 best response quantities. Thus it appears that the symmetric Cournot Nash

equilibrium is an absorbing state. However, is it the unique absorbing state?

Consider the following state:⁹

Definition 4 (Pseudo-Stackelberg Solution) *The Pseudo-Stackelberg solution is a state $\omega^S = (q_1, \dots, q_{\theta n}, q_{\theta n+1}, \dots, q_n)$ that satisfies the following conditions:*

$$(i) \quad \forall i \in I, q_i = q^S \text{ s.t.}$$

$$\begin{aligned} q^S p(\theta n q^S + (1 - \theta) n q^D) - c(q^S) > \\ q p(\theta n q^S + (1 - \theta) n q^D) - c(q), \forall q \in \Gamma, \end{aligned} \quad (10)$$

$$(ii) \quad \forall i \in N \setminus I, q_i = q^D,$$

$$q^D := b(\theta n q^S + ((1 - \theta)n - 1)q^D). \quad (11)$$

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. 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 every optimizer since inequality (10) holds for $\forall q \in \Gamma$, so also for q^D . I.e. it follows that¹⁰

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

⁹We assume that the best reply is unique. The uniqueness condition later in Assumption 5 ensures that the best reply to the opponents' output is indeed a singleton.

¹⁰For notational convenience we write $\pi_i(q, q', n, \theta)$ for $\pi_i(q, \theta n q + (1 - \theta) n q')$ if $i \in I$ or for $\pi_i(q', \theta n q + (1 - \theta) n q')$ if $i \in N \setminus I$.

Why is the Pseudo-Stackelberg solution a candidate for a solution? Assume that the Pseudo-Stackelberg solution exists in Γ^n and that is unique. Consider first the imitators: every imitator sets identical output and is strictly better off than any optimizer. Hence an imitator has no reason to deviate from its output. Also optimizers do not deviate from their output since they play the best response.

In the following, we will reserve q^S to denote the identical individual output of any imitator in the Pseudo-Stackelberg solution. q^I denotes that the individual quantity q is set by each imitator (superscript “I” stands for “independent” or all “imitators”). Analog, q^D means that the individual quantity q 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 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:¹¹

Definition 5 (Pseudo-Stackelberg States) *The set of Pseudo-Stackelberg states H consists of all states $\omega = (q_1, \dots, q_{\theta n}, q_{\theta n+1}, \dots, q_n) \in \Gamma^n$ that satisfy the following properties:*

- (i) $q_i = q^I, \forall i \in I$ and some $q^I \in \Gamma$,
- (ii) $q_i = q^D, \forall i \in N \setminus I, q^D := b(\theta n q^I + ((1 - \theta)n - 1)q^D)$,
- (iii) $\pi^I(q^I, q^D, n, \theta) \geq \pi^D(q^I, q^D, n, \theta)$,

¹¹Again, we assume that the best reply is unique. The uniqueness condition later in Assumption 5 ensures that the best reply to the opponents’ output is indeed a singleton.

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

Assume that each Pseudo-Stackelberg state is uniquely defined by the above conditions. If condition (i) is not satisfied, 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, not all optimizers are playing the best reply having an incentive to deviate. If condition (iii) is not satisfied, imitators will mimic optimizers. To understand the motivation of (iv) note that by identical costs, $q^I = q^D$ 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 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 the competitive solution is not a Pseudo-Stackelberg state since optimizers do not set a best reply in the competitive solution (unless $n \rightarrow \infty$). Finally, if the Pseudo-Stackelberg solution exists and 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 5 of Pseudo-Stackelberg states are satisfied. If c is not strictly convex condition (iv) might 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 The Result

Before we state and prove the results in this section, we need to introduce formally the assumptions. The first assumption concerns the existence of outcomes.

Assumption 4 (Existence) *The Cournot Nash equilibrium and at least another Pseudo-Stackelberg state exists in Γ^n .*

We like to focus on the dynamics in an oligopoly with a heterogeneous population of firms. We do not concern ourselves with questions of existence, which have been dealt with elsewhere. We avoid to state the above assumption in terms of primitives of the model since the existence of Cournot Nash equilibrium in a symmetric n -firm Cournot oligopoly can be shown under very general alternative assumptions even without concave profit functions (see for example Roberts and Sonnenschein, 1976, Amir, 1996, and Vives, 2000). In a non-degenerate Cournot oligopoly, whenever the Cournot Nash equilibrium exists, some other Pseudo-Stackelberg states also exist. The existence of another Pseudo-Stackelberg state makes the Cournot oligopoly interesting. The fact that those outputs are an element of Γ is not really an assumption since the grid Γ is arbitrary.

Assumption 5 (Uniqueness) *For $q'_{-i} < q_{-i}$, $q' \in b(q'_{-i})$, $q \in b(q_{-i})$, we have*

$$\frac{q' - q}{q'_{-i} - q_{-i}} > -1. \quad (12)$$

This second assumption is rather standard and 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 $q'_{-i} + q' < q_{-i} + q$. The requirement is that the best reply correspondence must have slopes strictly above -1 . It means that total output is strictly increasing in the opponents' output. The uniqueness condition implies that the unique Cournot Nash equilibrium is symmetric and that the best reply correspondence is in fact a function (see Vives, 2000, p. 43).

Assumption 6 (Generalized Ordinal Potential) *The Cournot game has a generalized ordinal potential.*¹²

A potential function monitors the drift towards the Nash equilibrium. In the context of a general Cournot oligopoly, a potential function is generally difficult to interpret. It raises the question what do the firms consciously or unconsciously try to jointly optimize? Important is that such a potential function implies the existence of a finite improvement path towards the unique Nash equilibrium (Monderer and Shapley, 1996, Lemma 2.5). Slade (1994) proves the existence of an exact potential, a much stronger version than the generalized ordinal potential function, for the case of a Cournot oligopoly with linear demand. An example for the existence of an ordinal potential is given for a Cournot oligopoly with linear costs by Monderer and

¹²A generalized ordinal potential (see Monderer and Shapley, 1996) is a function $P : \Gamma^n \rightarrow \mathbb{R}$ for the Cournot game such that for every $i \in N$ and for every $q_{-i} \in \{0, \delta, 2\delta, \dots, (n-1)\nu\delta\}$ and every $q, q' \in \Gamma$ it holds that

$$\pi_i(q, q_{-i}) - \pi_i(q', q_{-i}) > 0 \implies P(q, q_{-i}) - P(q', q_{-i}) > 0.$$

Shapley (1996), p. 124. We do not know of a full characterization of symmetric n -firm Cournot oligopoly satisfying the weaker requirement of the existence of a generalized ordinal potential. We conjecture that such characterization would be more general than the ones satisfying exact, weighted, or ordinal potentials. At this point of time, investigation of this conjecture is left to further research. The existence of a finite improvement path implies that the sequential best response dynamics converges to the Nash equilibrium. Huck, Normann, and Oechssler (1999) make use of this result to show that the best reply process with inertia converges to the Cournot Nash equilibrium in a quantity setting symmetric linear quadropoly. Note that there are of course alternative assumptions that would guarantee convergence to the Cournot Nash equilibrium, for instance dominance solvability (see Moulin, 1984).

Assumption 7 (Quasiconcavity) π_i is quasiconcave in q_i , $\forall i \in I$. That is $\forall q_i, q'_i \in \Gamma, \forall \lambda \in [0, 1]$

$$\pi(\lambda q_i + (1 - \lambda)q'_i, Q) \geq \min\{\pi(q_i, Q), \pi(q'_i, Q)\}, \forall Q \in \{0, \delta, 2\delta, \dots, n\nu\delta\}. \quad (13)$$

The intuition for quasiconcavity is single-peakedness. The set of maxima is convex. If the inequality holds strictly for $\lambda \in (0, 1)$, then π is strictly quasiconcave and the maximum is unique. Concavity and weak concavity imply quasiconcavity but not vice versa.

We are finally ready to state the main result.

Theorem 1 *Let S denote the support of the long run distribution φ^* . Under the above assumptions, we have $S \subseteq H$. Moreover, it is never true that $S = \{\omega^\circ\}$.*

The result is that given our assumptions the support S of the long run distribution φ^* is contained in the set of Pseudo-Stackelberg states $H \subset \Gamma^n$.

Moreover, the Cournot state ω° , the only state in which optimizers are as well off as imitators, is never the unique long run outcome. It follows that imitators are strictly better off than are optimizers in the long run.

The proof in the appendix is divided into three lemmata. First, in Lemma 2 we work out the absorbing sets of the unperturbed process. As we suggested in the previous section, the set of absorbing sets comprises exactly of the set of Pseudo-Stackelberg states, whereby each Pseudo-Stackelberg state is an absorbing state. Since by standard arguments the support of the long run distribution is a subset of the set of absorbing sets, we can conclude immediately that the support is the subset of the Pseudo-Stackelberg states. It implies that the imitators are weakly better off than are optimizers. Second, making use of the Lemma 3 by Nöldeke and Samuelson (1993), Samuelson (1994) and Nöldeke and Samuelson (1997), we show with Lemma 4 that the Cournot Nash equilibrium is not uniquely stochastically stable. That is, the Cournot Nash equilibrium is not robust against a single mutation. Since the Cournot Nash equilibrium is the only state in which optimizers are as well off as imitators, we can finally conclude that imitators are strictly better off than are optimizers in the long run.

In the proof of Theorem 1, particularly in Lemma 4, we show even more. We can order the Pseudo-Stackelberg states according to the identical output of imitators. We then show that, making use of quasi-submodularity and quasi-concavity, we can find a sequence of single mutations followed by the unperturbed decision dynamics by which one can move through the Pseudo-Stackelberg states step-wise upwards starting at the Cournot Nash equilibrium. If the Pseudo-Stackelberg solution does not exist within the set of Pseudo-Stackelberg states, this step-wise mutation-sequence terminates at the Pseudo-Stackelberg state with the highest output of imitators. If

the Pseudo-Stackelberg solution does exist in the set of Pseudo-Stackelberg states, then the sequence ends there. In this case one can also find a sequence of single mutations by which one can move step-wise downwards starting from the Pseudo-Stackelberg state with the highest output of imitators and ending at the Pseudo-Stackelberg solution. In fact, if the Pseudo-Stackelberg solution exists within the set of Pseudo-Stackelberg states, we show the following.

Remark 2 *Suppose that Assumptions 1 to 6 hold. If $\omega^S \in H$, then any state in H can be connected to the Pseudo-Stackelberg solution by a single suitable mutation followed by the unperturbed decision dynamics.*

We can connect any Pseudo-Stackelberg state to the Pseudo-Stackelberg solution by a single suitable mutation even without the assumption of quasi-concavity. It suggests that the Pseudo-Stackelberg solution - if it exists in H - is a candidate for a unique long run solution. It also establishes without the assumption of quasi-concavity that the Cournot Nash equilibrium is not uniquely stochastically stable. However, it is premature to conjecture that the Pseudo-Stackelberg solution is the unique long run outcome. In particular, we can destabilize any Pseudo-Stackelberg state including the Pseudo-Stackelberg solution by a single suitable mutation with a sufficient large output. In Theorem 2 we show that it is even possible to find plausible assumptions on the parameters of the model that are sufficient for the entire set of the Pseudo-Stackelberg states being in the support of the long run distribution.

Theorem 2 *Suppose that above assumptions hold. Then there exist p , c , θ , δ and finite n such that $S = H$.*

Indeed, one can find reasonable functions p , c and parameters θ , δ as well as a finite n that are sufficient for the entire set of Pseudo-Stackelberg states to be the support of the unique limiting invariant distribution. For example, consider $p(Q) = 10 - Q$, $c(q_i) = q_i^2$, $\theta = 0.6$, $\delta = 0.2$ and $n = 5$. Those assumptions are sufficient to show that any Pseudo-Stackelberg state can be connected to the Cournot Nash equilibrium as well as to the Pseudo-Stackelberg state with the largest output of imitators by a sequence of single suitable mutations followed by the unperturbed decision dynamics. This is done in Lemmata 6 and 7 respectively, making use of Lemma 5. Together with Lemma 4, in which we showed the existence of increasing and decreasing sequences of single mutations in the set of Pseudo-Stackelberg states, it implies that the entire set of Pseudo-Stackelberg states is the unique recurrent set. Using Lemma 3 by Nöldeke and Samuelson (1993), Samuelson (1994) and Nöldeke and Samuelson (1997) we can conclude that the entire set of Pseudo-Stackelberg states is the support of the long run distribution. Again, it follows that imitators are strictly better off than are optimizers.

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.¹³ 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. Consider what happens if we make the optimizers more and more sophisticated. Suppose first 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

¹³Alternatively, 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 decision rules is sufficiently lower than the probability of adjusting outputs, then the same result is likely to emerge.

be slower since the adjustment process becomes similar to fictitious play.¹⁴ Second, suppose that optimizers can forecast the behavior of the imitators. What does it help them if imitators set some large output, which they will do 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. Finally, assume that optimizers are so sophisticated that they can even forecast their effect of their own action on the imitators' behavior. For imitators to make lower profits than optimizers, the latter have to induce the former to lower quantities. Lower quantities in turn are played by imitators if the most successful firm of some previous round did play such low quantity. In Cournot games a lower quantity is more profitable than a larger one if the total market output is huge, e.g. some firms must play huge quantities making even higher losses than imitators in order to reduce the imitators quantities. This strategy is very costly to optimizers such that it is unlikely to make optimizers relatively better off than imitators. Moreover, any profitable "low-balling" of imitators by optimizers can happen by inertia just temporarily since imitators would seek to mimic optimizers immediately. It appears that no matter how sophisticated the optimization behavior of optimizers is, by the order structure of the Cournot game, optimizers can not beat imitators.

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 com-

¹⁴Regarding fictitious play refer for example to Fudenberg and Levine (1998), pp. 29.

mitment 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.

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 profit may not be the way to actually obtain the maximal 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 rather than a flaw of this approach.

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 necessary condition for a state being contained in the support of the long run distribution introduced by Nöldeke and Samuelson (1993) is that one can conclude immediately that just a single suitable mutations is required to trigger the long run outcome. That is, given our set of absorbing states, convergence to the long run outcome is comparatively rather fast. Note how-

ever, that this condition is not applicable generally and does not allow in general for conclusions about the speed of convergence. See Ellison (2000) for a more general method also characterizing the speed of convergence.

The key property driving the result is the observation that the payoff functions in a Cournot game are submodular in individual and total quantity. This is closely related to strategic substitutes. It suggests that the same result does not emerge in games with a different order structure such as games with strategic complements. It also suggests that the same result holds in other games with strategic substitutes and negative aggregate externalities. E.g. consider a repeated Nash demand game.¹⁵ Suppose now 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.") The mechanism is similar to the idea of spiteful behavior in evolutionary biology. A firm's action can be called spiteful if it forgoes profits in order to lower the profits of its competitors even further (see for example Schaffer, 1989). In our context, we face the "dual" to spite. The optimizers, while increasing their absolute profits, do not care about increasing the profits of the imitators even further.

¹⁵I thank Ariel Rubinstein for motivating this example.

Since both, imitation behavior as well as best response, is supported by experimental findings in Cournot markets depending on the amount of information provided to subjects (see Sauermann and Selten, 1959, Huck, Normann, and Oechssler, 1999, 2000, Offerman, Potters, and Sonnemans, 2002), it is only natural to test whether my results can be supported experimentally if different amounts of information are given to various firms in an oligopoly experiment. This shall be left to further research.

A Proofs

A.1 Proof of Lemma 1.

To show this Lemma for all $(q, Q) \in \Gamma \times \{0, \delta, 2\delta, \dots, n\nu\delta\}$, one has to show the following inequality

$$\pi((q', Q') \wedge (q'', Q'')) + \pi((q', Q') \vee (q'', Q'')) \leq \pi(q', Q') + \pi(q'', Q'') \quad (14)$$

for the cases (i) $q' \geq q''$ and $Q' \geq Q''$, (ii) $q' < q''$ and $Q' \geq Q''$, (iii) $q' \geq q''$ and $Q' < Q''$ as well as (iv) $q' < q''$ and $Q' < Q''$. In particular, strict submodularity requires to show a strict inequality for the cases (ii) and (iii) since neither $(q', Q') \geq (q'', Q'')$ nor $(q'', Q'') \geq (q', Q')$.

Case (i) and (iv):

$$\pi((q', Q') \wedge (q'', Q'')) + \pi((q', Q') \vee (q'', Q'')) = \pi(q', Q') + \pi(q'', Q'').$$

Case (iii): Since by Assumption 1, p is strictly decreasing

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

$$\begin{aligned}
\pi(q'', Q') + \pi(q', Q'') &< \pi(q', Q') + \pi(q'', Q'') \\
\pi((q', Q') \wedge (q'', Q'')) + \pi((q', Q') \vee (q'', Q'')) &< \pi(q', Q') + \pi(q'', Q'').
\end{aligned}$$

Case (ii): By Assumption 1, p is strictly decreasing. Thus we have analogous to previous steps

$$\begin{aligned}
p(Q'') &> p(Q') \\
p(Q'')(q' - q'') &< p(Q')(q' - q'') \\
\pi((q', Q') \wedge (q'', Q'')) + \pi((q', Q') \vee (q'', Q'')) &< \pi(q', Q') + \pi(q'', Q'').
\end{aligned}$$

This completes the proof of Lemma 1.

A.2 Proof of Theorem 1

The proof of the Theorem 1 follows from below lemmata.

Recall that Z is the collection of all absorbing sets of the unperturbed decision dynamics when $\varepsilon = 0$.

Lemma 2 (Absorbing Sets) *If Assumptions 2, 4, 5 and 6 hold, then $Z = H$ with $Z = \{\{\omega\} : \omega \in H\}$.*

PROOF. 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 5 is violated. Thus there will be an incentive for some imitators or some optimizers to deviate from their output in ω . By Assumptions 2 and 6 we can construct an unperturbed adjustment process based on the decision rules leading in the subsequent periods to a state $\omega' \in H$, noting that by Assumptions 4 and 5 such $\omega' \in H$ exists and is uniquely defined.

Second, we show that every absorbing set $A \subseteq H$ is a singleton. Suppose $\exists \omega', \omega \in A \subseteq H$ such that $\omega' \neq \omega$. Note that by Assumptions 4 and 5 at least two Pseudo-Stackelberg states exist and are uniquely defined. 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 5 (i), (iii), and (iv) as well as D_I that no imitator $i \in I$ wants to deviate from its output in $\omega \in H$. Now consider an optimizer $i \in N \setminus I$. Since $\omega \in H$, it follows by aforesaid Definition 5 (ii) that no optimizer $i \in N \setminus I$ wants to deviate from its best reply in $\omega \in H$, which is by Assumption 5 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 can conclude that $\bar{A}\omega \notin H$ s.t. $\omega \in A, A \in Z$. Hence $Z = \{\{\omega\} : \omega \in H\}$. This completes the proof of Lemma 2. Q.E.D.

In order to characterize the support of the unique limiting invariant distribution further, we consider perturbations introduced by Assumption 3. We show that the Cournot Nash equilibrium is not the unique stochastically stable state.

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 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 $\forall \omega \in A, M(\omega) \cap B(A') \neq \emptyset$. 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), Samuelson (1994) and Nöldeke and Samuelson (1997).

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, $\forall \omega' \in R$, $\varphi^*(\omega') > 0$.

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

Definition 6 Define $\bar{\omega} \in H$ to be the Pseudo-Stackelberg state with the largest possible identical output of imitators, that is

$$\begin{aligned} \bar{q}p(\theta n\bar{q} + (1-\theta)nq^D) - c(\bar{q}) &> \\ q^Dp(\theta n\bar{q} + (1-\theta)nq^D) - c(q^D), \end{aligned} \quad (15)$$

$$\text{with } q^D = b(\theta n\bar{q} + ((1-\theta)n - 1)q^D),$$

$$\begin{aligned} (\bar{q} + \delta)p(\theta n(\bar{q} + \delta) + (1-\theta)nq^{D\delta}) - c(\bar{q} + \delta) &\leq \\ q^{D\delta}p(\theta n(\bar{q} + \delta) + (1-\theta)nq^{D\delta}) - c(q^{D\delta}), \end{aligned} \quad (16)$$

$$\text{with } q^{D\delta} = b(\theta n(\bar{q} + \delta) + ((1-\theta)n - 1)q^{D\delta}).$$

It is easy to see that $\bar{\omega}$ is indeed in H , since by definition all imitators set identical output \bar{q} , all optimizer play best response q^D (which is unique by Assumption 5), and imitators make strictly higher profits than optimizers. That is, all conditions of Definition 5 are satisfied. \bar{q} is indeed the largest identical output that each imitator can set within the set of Pseudo-Stackelberg states since an increase by δ yields a state not in H .

Remark 3 Given Assumptions 4 and 5, 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 .

We call a sequence of Pseudo-Stackelberg states $\omega_1, \dots, \omega_m \in H$ increasing (decreasing) iff the identical output of each imitator in those Pseudo-Stackelberg states is such that $q_{\omega_j}^I < q_{\omega_{j+1}}^I$ ($q_{\omega_j}^I > q_{\omega_{j+1}}^I$), $j = 1, \dots, m-1$.

Lemma 4 *Under above assumptions we conclude:*

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

PROOF. By Lemma 2 we know that each absorbing set is a singleton in H . Moreover, we can enumerate the absorbing sets since Γ is a finite output grid.

(i): $\omega^S \notin H$ then $q^S > \bar{q}$. Define a sequence of absorbing states $\omega_1, \dots, \omega_m \in H \subset \Gamma^n$ such that $\omega_1 = \omega^\circ$, ω_2 s.t. $q_{\omega_2}^I = q^\circ + \delta$, ω_3 s.t. $q_{\omega_3}^I = q^\circ + 2\delta$, ..., ω_{m-1} s.t. $q_{\omega_{m-1}}^I = \bar{q} - \delta$ and $\omega_m = \bar{\omega}$. Clearly, this sequence is increasing. In order to show that $M(\omega_j) \cap B(\{\omega_{j+1}\}) \neq \emptyset$ for $j = 1, \dots, m-1$, we have to show for $k = 1$, $\forall q \in [q^\circ, \bar{q}) \subset \Gamma$,

$$(q + \delta)p((\theta n - k)q + k(q + \delta) + (1 - \theta)nq^D) - c(q + \delta) > qp((\theta n - k)q + k(q + \delta) + (1 - \theta)nq^D) - c(q), \quad (17)$$

with $q^D = b((\theta n - k)q + k(q + \delta) + ((1 - \theta)n - 1)q^D)$, which is uniquely defined by Assumption 5.

By Assumption 1, Lemma 1 and Remark 1, π is strictly quasi-submodular (formulas (3) and (4)). Define $q' = q + \delta$, $q'' = q$, $Q' = (\theta n - k)q + k(q + \delta) + (1 - \theta)nq^{D'}$ and $Q'' = \theta n(q + \delta) + (1 - \theta)nq^{D''}$, with $q^{D'} = b(\theta n - k)q + k(q + \delta) + ((1 - \theta)n - 1)q^{D'}$ and $q^{D''} = b(\theta n(q + \delta) + ((1 - \theta)n - 1)q^{D''})$ being uniquely defined by Assumption 5. Then the right hand side of “ \implies ” in the upper formula (3) is equivalent to inequality (17) if $q^S \geq \bar{q}$, i.e. it is equivalent to $\pi(q', Q') > \pi((q', Q') \wedge (q'', Q''))$.

What is left to show is that the left hand side of “ \implies ” in the upper formula (3) is implied by the Assumption 7 of quasiconcavity. To see this note that for each

$q \in [q^\circ, \bar{q}] \subset \Gamma$ there exists a $\lambda \in [0, 1]$ s.t. $q + \delta = \lambda q + (1 - \lambda)\bar{q}$. Since $q^S \geq \bar{q}$, we have $\min\{\pi(q, Q), \pi(\bar{q}, Q)\} = \pi(q, Q)$, $\forall q \in [q^\circ, \bar{q}]$, $\forall Q \in \{0, \delta, 2\delta, \dots, n\nu\delta\}$. Hence $\pi(q + \delta, Q'') \geq \pi(q, Q'')$, $\forall q \in [q^\circ, \bar{q}]$, which is equivalent to the left hand side of “ \implies ” in the upper formula (3), i.e. it is equivalent to $\pi((q', Q') \vee (q'', Q'')) \geq \pi(q'', Q'')$. Thus we have shown that if $q^S \geq \bar{q}$ there exists an increasing sequence of absorbing states through which we can move from the absorbing state with the lowest output of imitators ω° to the absorbing state with the highest output of imitators $\bar{\omega}$ by a sequence of single suitable mutations.

(ii) and (iii): If $\omega^S \in H$ then $q^S \leq \bar{q}$. Partition $[q^\circ, \bar{q}] \subset \Gamma$ into $\{[q^\circ, q^S], q^S, (q^S, \bar{q}]\}$. For any $q \in [q^\circ, q^S]$ we can show inequality (17) analogous to previous case (i). The same holds if $q^S = \bar{q}$. Hence (ii) is shown.

To show (iii) consider the interval $(q^S, \bar{q}]$. Define a sequence of absorbing states $\omega_1, \dots, \omega_m \in H \subset \Gamma^n$ such that $\omega_1 = \bar{\omega}$, ω_2 s.t. $q_{\omega_2}^I = \bar{q} - \delta$, ω_3 s.t. $q_{\omega_3}^I = \bar{q} - 2\delta$, ..., ω_{m-1} s.t. $q_{\omega_{m-1}}^I = q^S + \delta$ and $\omega_m = \omega^S$. Clearly, this sequence is decreasing. In order to show that $M(\omega_j) \cap B(\{\omega_{j+1}\}) \neq \emptyset$ for $j = 1, \dots, m-1$, we have to show for $k = 1, \forall q \in (q^S, \bar{q}] \subset \Gamma$,

$$(q - \delta)p((\theta n - k)q + k(q - \delta) + (1 - \theta)nq^D) - c(q - \delta) > qp((\theta n - k)q + k(q - \delta) + (1 - \theta)nq^D) - c(q), \quad (18)$$

with $q^D = b((\theta n - k)q + k(q - \delta) + ((1 - \theta)n - 1)q^D)$, which is uniquely defined by Assumption 5.

By Assumption 1, Lemma 1 and Remark 1, π is strictly quasi-submodular (formulas (3) and (4)). Define $q' = q - \delta$, $q'' = q$, $Q' = (\theta n - k)q + k(q - \delta) + (1 - \theta)nq^{D'}$ and $Q'' = \theta n(q - \delta) + (1 - \theta)nq^{D''}$, with $q^{D'} = b((\theta n - k)q + k(q - \delta) + ((1 - \theta)n - 1)q^{D'})$ and $q^{D''} = b(\theta n(q - \delta) + ((1 - \theta)n - 1)q^{D''})$ being uniquely defined by Assumption 5. Since $q^S < \bar{q}$, the right hand side of “ \implies ” in the lower formula (4) is equivalent to inequality (18).

What is left to show is that the left hand side of “ \implies ” in the lower formula (4) is implied by the Assumption 7 of quasiconcavity. To see this note that for each

$q \in [q^S, \bar{q}] \subset \Gamma$ there exists a $\lambda \in [0, 1]$ s.t. $q - \delta = \lambda q + (1 - \lambda)q^S$. Since $q^S < \bar{q}$, we have $\min\{\pi(q, Q), \pi(q^S, Q)\} = \pi(q, Q)$, $\forall q \in [q^S, \bar{q}]$, $\forall Q \in \{0, \delta, 2\delta, \dots, n\nu\delta\}$. Hence $\pi(q - \delta, Q'') \geq \pi(q, Q'')$, $\forall q \in [q^S, \bar{q}]$, which is equivalent to the left hand side of “ \implies ” in the lower formula (4). This completes the proof of Lemma 4. Q.E.D.

Corollary 1 *From previous Lemma 4 follows that $S \neq \{\omega^\circ\}$.*

PROOF. By Lemma 4 we know that $\exists \omega \in H$, $\omega \neq \omega^\circ$ s.t. $M(\omega^\circ) \cap B(\{\omega\}) \neq \emptyset$. Hence by the definition of a recurrent set we have $R \neq \{\omega^\circ\}$. Thus by Lemma 3 we can conclude that $S \neq \{\omega^\circ\}$. Q.E.D.

This completes the proof of the Theorem 1.

A.3 Remark 2

Suppose that Assumptions 1 to 6 hold. We have to show that if $\omega^S \in H$ then $M(\omega) \cap B(\{\omega^S\}) \neq \emptyset$, $\forall \omega \in H \setminus \{\omega^S\}$.

Assume $\omega^S \in H$. It is sufficient to show that $\forall q \in \Gamma$, q being a component of an arbitrary $\omega \in H$, $\omega \neq \omega^S$, $k \in \mathbb{N}$, $1 \leq k \leq \theta n$,

$$\begin{aligned} q^S p((\theta n - k)q + kq^S + (1 - \theta)nq^D) - c(q^S) &> \\ qp((\theta n - k)q + kq^S + (1 - \theta)nq^D) - c(q), \end{aligned} \quad (19)$$

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

By Assumption 1, Lemma 1 and Remark 1, π is strictly quasi-submodular (formulas (3) and (4)) in (q, Q) on $\Gamma \times \{0, \delta, 2\delta, \dots, n\nu\delta\}$. Define $q' := q^S$, $q'' := q$, $Q' := (\theta n - k)q + kq^S + (1 - \theta)nq^{D'}$ and $Q'' := \theta nq^S + (1 - \theta)nq^{D''}$ with $q^{D'} = b((\theta n - k)q + kq^S + ((1 - \theta)n - 1)q^{D'})$ and $q^{D''} = b(\theta nq^S + ((1 - \theta)n - 1)q^{D''})$ being uniquely defined by Assumption 5. If $q' > q''$ then $\theta nq' > (\theta n - k)q'' + kq'$. By Assumption 5, we get $q^{D''} \geq q^{D'}$. We conclude that $Q'' > Q'$. If $q' < q''$ then $\theta nq' < (\theta n - k)q'' + kq'$. By Assumption 5, we get $q^{D''} \leq q^{D'}$ and conclude that

$Q'' < Q'$. It follows that if $q < q^S$ the left hand side of “ \implies ” in formula (3) is given by inequality (10) of Definition 5 of the Pseudo-Stackelberg solution (i). In this case the right hand side of “ \implies ” in formula (3) yields above inequality (19). If $q > q^S$ the left hand side of “ \implies ” in formula (4) is given by inequality (10) of Definition 5 of the Pseudo-Stackelberg solution (i). In this case the right hand side of “ \implies ” in formula (4) yields above inequality (19). Finally, set $k = 1$ to see that one suitable mutation only is required to connect every $\omega \in H$ to $\omega^S \in H$. This completes the proof of Remark 2.¹⁶

A.4 Proof of Theorem 2

The proof of the Theorem 2 follows from below lemmata.

Lemma 5 *There exist p, c, θ, δ and finite n such that*

$$q^\circ p((2n-3)q^\circ) - c(q^\circ) \leq 0, \quad (20)$$

$$(n-1)q^\circ p((n-1)q^\circ) - c((n-1)q^\circ) \leq 0, \quad (21)$$

$$q^\bullet \equiv \bar{q}, \quad (22)$$

with $q^\bullet \in \Gamma$ being the monopoly output.¹⁷

PROOF. Consider the example $p(Q) = 10 - Q$, $c(q_i) = q_i^2$, $\theta = 0.6$, $\delta = 0.2$, and $n = 5$. It is straight forward to verify that formulas (20) to (22) hold. Q.E.D.

¹⁶An analog proof to the one of Remark 2 can be used to prove a result of Vega-Redondo (1997) for $\theta = 1$. Just erase any output by optimizers, replace q^S by the competitive output, and plug in the formula of quasi-submodularity the definition of competitive solution instead the imitator’s Pseudo-Stackelberg solution. Then one can conclude that the competitive solution can be reached by a single suitable mutation in a homogeneous population of imitators from any other monomorphic absorbing state. It is the submodularity of the Cournot game, which makes Vega-Redondo’s (1997) result work.

¹⁷To avoid stating an additional definition note that the monopoly-output is a special case of the Cournot Nash equilibrium in Definition 3 for $n = 1$.

We show that the properties of Lemma 5 together with the above assumptions are sufficient to show that we can connect any Pseudo-Stackelberg state to the Cournot Nash equilibrium by a single suitable mutation followed by the unperturbed decision dynamics.

Lemma 6 *Suppose that Assumptions 2 to 6 hold. Moreover, let p, c, θ, δ and n such that the properties of Lemma 5 hold. Then $M(\omega) \cap B(\{\omega^\circ\}) \neq \emptyset, \forall \omega \in H$.*

PROOF. Suppose in t any arbitrary state $\omega(t) \in H$. By Assumptions 4 and 5 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 Assumption 4 and 5 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 5, inequality (20), $\pi_j(t+1) < 0, \forall j \in N$. W.l.o.g. assume that by Assumption 2 a firm $k \in N \setminus I, k \neq i$ and only a firm k has the opportunity to adjust output in $t+2$. Since $D_O(t+1) = 0$, we have $q_k(t+2) = 0$. $Q(t+2) \geq (2n-3)q^\circ$. By Lemma 5, inequality (20), $\pi_j(t+1) < 0, \forall j \in N \setminus \{k\}$. W.l.o.g. assume that by Assumption 2 and $D_O(t+2) = D_I(t+2) = 0$ all $j \in N \setminus \{i\}$ 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 5, inequality (21), $\pi_i(t+3) \leq 0$. W.l.o.g. assume that by Assumption 2, 5 and $D_O(t+3) = b((n-1)q^\circ) = q^\circ$ another firm $k \in N \setminus I$ has the opportunity to adjust output in $t+4$. Since $\pi_k(t+4) > \pi_j(t+4), j \in N \setminus \{k\}$ 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 and 6 let all remaining optimizers adjust output in the subsequent periods such that with positive probability ω° is reached in finite time. 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 $M(A) \cap B(\{\omega^\circ\}) \neq \emptyset, \forall A \in Z$. Q.E.D.

We show that the properties of Lemma 5 together with the above assumptions are also sufficient to show that we can connect any Pseudo-Stackelberg state to

the Pseudo-Stackelberg solution with the largest output of imitators by a single suitable mutation followed by the unperturbed decision dynamics.

Lemma 7 *Suppose that Assumptions 2 to 6 hold. Moreover, let p, c, θ, δ and n such that the properties of Lemma 5 hold. Then $M(\omega) \cap B(\{\bar{\omega}\}) \neq \emptyset, \forall \omega \in H$.*

PROOF. By Assumption 4 and 5, \bar{q} exists and is uniquely defined. Suppose in t any arbitrary state $\omega(t) \in H$, which by Assumption 4 exists. 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 \hat{q}_i such that

$$qp(Q - q_i + \hat{q}) - c(q) < 0, \forall q > 0. \quad (23)$$

That is, $\pi_j(t+1) < 0, \forall j \in N$. W.l.o.g. assume that by Assumption 2 all optimizers in $N \setminus I$ 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 (23), 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 $t + 4$ an optimizer k and only an optimizer $k \in N \setminus I$ adjusts output such that $D_O(t + 3) = b(0) = q^\bullet(t + 4)$, which by Assumptions 4 and 5 exists and is uniquely defined. W.l.o.g. assume that by Assumptions 2 in $t + 5$ all imitators in I adjust output such that $D_I(t + 4) = q^\bullet(t + 5)$. Let all optimizers in $N \setminus I$ adjust output in the subsequent periods such that by Assumptions 2, 4, 5, and 6 a state $\omega^\bullet = (q_1^\bullet, \dots, q_{\theta n}^\bullet, q_{\theta n+1}^D, \dots, q_n^D)$ is reached in finite time. Since by Lemma 5, $\bar{q} = q^\bullet$, we can conclude that $\omega^\bullet = \bar{\omega}$. Q.E.D.

In Lemma 4 we showed already that we can connect the Pseudo-Stackelberg states by a increasing or decreasing sequence of single suitable mutations followed by the decision dynamics starting in the Cournot Nash equilibrium or the Pseudo-Stackelberg state with the largest identical output of imitators. In Lemma 6 and 7 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 5 hold. 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 Theorem 2.

References

- [1] Alchian, Armen A. (1950). Uncertainty, evolution, and economic theory, *Journal of Political Economy* **58**, 211-221.
- [2] Alós-Ferrer, Carlos (2001). Cournot versus Walras in dynamic oligopolies with memory, mimeo, University of Vienna.
- [3] Alós-Ferrer, Carlos, Ania, Ana B. and Fernando Vega-Redondo (1999). An evolutionary model of market structure, in: Herings, P., Talman, A., and G. van der Laan, *The theory of markets*, Amsterdam: North Holland.
- [4] Amir, Rabah (1996). Cournot oligopoly and the theory of supermodular games, *Games and Economic Behavior* **15**, 132-148.
- [5] Blume, Lawrence and David Easley (2002). Optimality and natural selection in markets, forthcoming *Journal of Economic Theory*.
- [6] Bulow, Jeremy I., Geanakoplos, John D. and Paul D. Klemperer (1985). Multimarket oligopoly: strategic substitutes and complements, *Journal of Political Economy* **93**, 488-511.
- [7] Conlisk, John (1980). Costly optimizers versus cheap imitators, *Journal of Economic Behavior and Organization* **1**, 275-293.
- [8] Cournot, Augustine (1839). *Researches into the mathematical principles of the theory of wealth*, (translated by Nathaniel T. Bacon) New York: The MacMillan Company, 1927.

- [9] Ellison, Glenn (2000). Basin of attraction, long-run stochastic stability, and the speed of step-by-step evolution, *Review of Economic Studies* **67**, 17-45.
- [10] Freidlin, Mark and Alexander Wentzell (1984). *Random perturbations of dynamical systems*, Berlin et al.: Springer-Verlag.
- [11] Friedman, Milton (1953). *Essays in positive economics*, Chicago: University of Chicago Press.
- [12] Fudenberg, Drew and David K. Levine (1998). *The theory of learning in games*, Cambridge, M.A. and London, England: The MIT Press.
- [13] Heifetz, Aviad and Yossi Spiegel (2002). The evolution of biased perceptions, mimeo, Tel Aviv University.
- [14] Huck, Steffen, Normann, Hans-Theo and Jörg Oechssler (2000). Does information about competitors' actions increase or decrease competition in experimental oligopoly markets?, *International Journal of Industrial Organization* **18**, 39-57.
- [15] Huck, Steffen, Normann, Hans-Theo and Jörg Oechssler (1999). Learning in Cournot oligopoly - An experiment, *Economic Journal* **109**, C80-C95.
- [16] Kandori, Michihiro, Mailath, George J. and Rafael Rob (1993). Learning, mutation and long run equilibria in games, *Econometrica* **61**, 29-56.
- [17] Koçkesen, Levent, Ok, Efe A. and Rajiv Sethi (2000). The strategic advantage of negatively interdependent preferences, *Journal of Economic Theory* **92**, 274-299.
- [18] Koopmans, T. (1957). *Three essays on the state of economic science*, McGraw-Hill Book Company.
- [19] Masaaki, Kijima (1997). *Markov processes for stochastic modeling*, London et al.: Chapman & Hall.

- [20] Milgrom, Paul and Chris Shannon (1994). Monotone comparative statics, *Econometrica* **62**, 157-180.
- [21] Monderer, Dov and Lloyd S. Shapley (1996). Potential games, *Games and Economic Behavior* **14**, 124-143.
- [22] Moulin, Herve (1984). Dominance solvability and Cournot stability, *Mathematical Social Sciences* **7**, 83-102.
- [23] Nelson, Richard R. and Sydney G. Winter (1982). *An evolutionary theory of economic change*, Cambridge, Mass.: Harvard University Press.
- [24] von Neumann, John and Oskar Morgenstern (1944). *Theory of games and economic behavior*, Science Editions, 1964, New York: John Wiley & Sons, Inc.
- [25] Nöldeke, Georg and Larry Samuelson (1997). A dynamic model of equilibrium selection in signaling markets, *Journal of Economic Theory* **73**, 118-156.
- [26] Nöldeke, Georg and Larry Samuelson (1993). An evolutionary analysis of backward and forward induction, *Games and Economic Behavior* **5**, 425-454.
- [27] Novshek, William (1985). Perfectly competitive markets as the limit of Cournot markets, *Journal of Economic Theory* **35**, 72-82.
- [28] Offerman, T., Potters, J. and Sonnemans, J. (2002). Imitation and belief learning in an oligopoly experiment, mimeo, forthcoming *Review of Economic Studies*.
- [29] Roberts, J. and Hugo Sonnenschein (1976). On the existence of Cournot equilibrium without concave profit functions, *Journal of Economic Theory* **22**, 112-117.

- [30] Rhode, P. and M. Stegeman (2001). Non-Nash equilibria of darwinian dynamics with applications to duopoly, *International Journal of Industrial Organization* **19**, 415-453.
- [31] Samuelson, Larry (1997). *Evolutionary games and equilibrium selection*, Cambridge, M.A. and London, England: The MIT Press.
- [32] Samuelson, Larry (1994). Stochastic stability in games with alternative best replies, *Journal of Economic Theory* **64**, 35-65.
- [33] Sauermann, Heinz and Reinhard Selten (1959). Ein Oligopolexperiment, *Zeitschrift für die gesamte Staatswissenschaft* **115**, 427-471.
- [34] Schaffer, Mark E. (1989). Are profit-maximizers the best survivors?, *Journal of Economic Behavior and Organization* **12**, 29-45.
- [35] Schipper, Burkhard C. (2002). Imitators and Optimizers in Symmetric n-Firm Cournot Oligopoly, mimeo, University of Bonn.
- [36] Slade, Margaret E. (1994). What does an oligopoly maximize?, *Journal of Industrial Economics* **17**, 45-61.
- [37] von Stackelberg, Heinrich (1934). *Marktform und Gleichgewicht*, Wien, Berlin: Verlag von Julius Springer.
- [38] Topkis, Donald M. (1998). *Supermodularity and complementarity*, Princeton, New Jersey: Princeton University Press.
- [39] Vega-Redondo, Fernando (1997). The evolution of Walrasian behavior, *Econometrica* **65**, 375-384.
- [40] Vives, Xavier (2000). *Oligopoly pricing: old ideas and new tools*, Cambridge, M.A., London, England: The MIT Press.
- [41] Vives, Xavier (1990). Nash equilibrium with strategic complementarities, *Journal of Mathematical Economics* **19**, 305-321.

- [42] Weibull, Jörgen W. (1995). *Evolutionary game theory*, Cambridge, M.A., London, England: The MIT Press.
- [43] Young, H. Peyton (1993). The evolution of conventions, *Econometrica* **61**, 57-84.